

The Professional Work Handbook

The craft of 454 professional skills, compiled from the open-source pm-claude-skills library — written to instruct AI, readable as a book for humans.

Part I — 11 chapters of craft (the Production-Ready tier, in full)

Part II — The Anti-Pattern Almanac: **2237 rules** of professional judgment

Edition 2026-07-06 · CC-licensed · regenerated from the library with every release

Chapter 1 — Engineering

10 skills

api-docs-writer

Write clear, developer-facing API documentation.

API Docs Writer Skill

This skill transforms raw API specs, endpoint descriptions, or Postman collections into clean, developer-facing documentation following OpenAPI-adjacent conventions. Output is ready for a developer portal, README, or Notion/Confluence page.

Required Inputs

Ask the user for these if not provided:

- **API or endpoint details** (raw spec, Postman export, or verbal description)
- **Auth method** (API key / Bearer token / OAuth 2.0 / None)
- **Base URL**
- **API version** (e.g. v1, v2.3, or "unversioned" — affects deprecation notes and versioning headers)
- **Rate limits** (requests per second/minute per token or IP, if known — or "unknown")
- **Audience** (internal developers / external partners / public)
- **Output format** (Markdown for developer portals and READMEs / Plain prose for Confluence or Notion — note: OpenAPI YAML is not produced by this skill)

Output Format

For each endpoint, produce the following:

[METHOD] **/path/to/endpoint**

Summary: [One line — what this endpoint does]

Description: [2–4 sentences. When to use this endpoint. What it returns. Any important behaviour to know (pagination, rate limits, async processing, etc.)]

Authentication: [Required / Optional — method]

Request

Headers:

Header	Required	Description
Authorization	Yes	Bearer <token>
Content-Type	Yes	application/json

Path Parameters:

Parameter	Type	Required	Description
id	string	Yes	Unique identifier for the resource

Query Parameters:

Parameter	Type	Required	Default	Description
limit	integer	No	20	Max results per page (1–100)
cursor	string	No	—	Pagination cursor from previous response

Request Body:

```
{
  "field_name": "value",
  "another_field": 42
}
```

Field	Type	Required	Description
field_name	string	Yes	[Plain description of what this field does]
another_field	integer	No	[Description. Include valid range or enum values if applicable]

Response

Success Response: 200 OK

```
{
  "id": "abc123",
  "status": "active",
}
```

```
"created_at": "2025-04-01T10:00:00Z"
}
```

Field	Type	Description
id	string	Unique identifier for the created/retrieved resource
status	string	Current status. Enum: <code>active</code> , <code>inactive</code> , <code>pending</code>
created_at	ISO 8601 string	Timestamp of creation in UTC

Error Codes

Status Code	Error Code	Description	How to Resolve
400	INVALID_REQUEST	Request body is malformed or missing required fields	Check request body against schema above
401	UNAUTHORIZED	Missing or invalid authentication token	Verify your API key or refresh your token
404	NOT_FOUND	The requested resource does not exist	Check the ID in the path parameter
429	RATE_LIMITED	Too many requests	Back off and retry after <code>Retry-After</code> header value
500	INTERNAL_ERROR	Unexpected server error	Retry with exponential backoff; contact support if persists

Code Examples

Produce examples in at least 2 languages relevant to the audience (default: cURL + Python):

cURL:

```
curl -X POST https://api.example.com/v1/endpoint \
  -H "Authorization: Bearer YOUR_TOKEN" \
  -H "Content-Type: application/json" \
  -d '{"field_name": "value"}'
```

Python:

```
import requests
```

```
response = requests.post(
    "https://api.example.com/v1/endpoint",
    headers={"Authorization": "Bearer YOUR_TOKEN"},
    json={"field_name": "value"}
)
data = response.json()
```

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/example-first-docs.md** — Example-First API Docs: the Rules That Make Docs Usable. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/endpoint-entry.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every parameter is documented (type, required/optional, description)
- ☐ Response fields are fully documented with types
- ☐ All relevant error codes are listed with resolution guidance
- ☐ Error codes cover at minimum: 400 (bad request), 401/403 (auth), 404 (not found), 429 (rate limited), 500 (server error) — or explicitly note which don't apply to this endpoint
- ☐ Code examples use the actual base URL and a realistic placeholder token — no examples reference undefined variables or "YOUR_ENDPOINT" outside the snippet
- ☐ Auth method is clearly stated at the top
- ☐ Enum values are listed where applicable
- ☐ Pagination documented if the endpoint is a list endpoint

Anti-Patterns

- ☐ Do not document only the happy path — every endpoint must have error codes for at least 400, 401/403, 404, 429, and 500
- ☐ Do not use placeholder values like "YOUR_ENDPOINT" or "INSERT_TOKEN" in code examples — use realistic-looking placeholders anchored to the actual base URL
- ☐ Do not skip enum values for fields with a fixed set of accepted values — undocumented enums cause integration bugs
- ☐ Do not omit pagination documentation on list endpoints — developers who miss this will build integrations that silently miss data

- ☐ Do not describe what a field "is" without describing what it "does" — "the ID" is not documentation; "the unique identifier used to retrieve or update this resource" is

Usage Examples

- "Document this API endpoint: [paste spec or description]"
- "Turn this Postman collection into developer docs"
- "Write API reference docs for [endpoint]"
- "Write a developer guide for our [product] API"

architecture-decision-record

Create an Architecture Decision Record (ADR) for any technical decision.

Architecture Decision Record (ADR) Skill

This skill produces a complete Architecture Decision Record (ADR) following the Nygard format — the most widely adopted standard. ADRs document the reasoning behind significant technical decisions so future team members understand not just *what* was decided, but *why*.

Required Inputs

Ask the user for these if not provided:

- **ADR number** (sequential number in your ADR registry — e.g. 012; or "next available" if unknown)
- **Decision title** (brief, e.g. "Use PostgreSQL as primary datastore")
- **Context** (what situation led to this decision needing to be made?)
- **Options considered** (at least 2; if only 1 is given, prompt for alternatives that were considered or ruled out)
- **Decision made** (which option was chosen)
- **Reason for choice**
- **Status** (Proposed / Accepted / Deprecated / Superseded)
- **Author and date**
- **Team context** (optional — team size, relevant experience, org constraints; helps calibrate formality and depth of the Context section)

Output Format

ADR-[NNN]: [Decision Title]

Date: [YYYY-MM-DD] **Status:** [Proposed / Accepted / Deprecated / Superseded by ADR-NNN] **Author(s):** [Name(s)] **Deciders:** [Who had final say — individual or team]

Context

[3–6 sentences. Describe the situation, constraints, and forces at play that made this decision necessary. Include: the problem being solved, relevant system state, team constraints, timeline pressures, or non-negotiable requirements. Write as if explaining to someone joining the team 18 months from now who has no prior context.]

Key constraints:

- [Constraint 1: e.g. "Must be deployable on-premise for enterprise customers"]
 - [Constraint 2: e.g. "Team has no prior Go experience"]
 - [Add as many as are relevant]
-

Options Considered

For each option, produce:

Option [N]: [Name]

Description: [What this option is — 1–3 sentences]

Pros:

- [Pro 1]
- [Pro 2]

Cons:

- [Con 1]
- [Con 2]

Why this was ruled out (if not chosen): [Honest reason]

Decision

We will [chosen option].

[2–4 sentences explaining the decision in plain language. This should be readable in isolation — someone should understand the decision from this paragraph alone without reading the full document.]

Consequences

Positive Consequences

- [What this decision enables or improves]
- [What risk it mitigates]

Negative Consequences / Accepted Tradeoffs

- [What we're giving up or taking on as a result of this decision]
- [Technical debt or limitations introduced]
- [What must now be true for this decision to remain valid]

Risks

- [What could cause this decision to be wrong in hindsight]
 - [What would trigger us to revisit this decision]
-

Implementation Notes

[Include if the decision has non-obvious implementation gotchas, or if there are related tickets/RFCs implementers will need. Skip only if the decision is purely tooling selection with no implementation ambiguity.]

Review Date

[Include unless the decision is permanent or self-evidently final. State a specific trigger condition — e.g. "Review if team grows beyond 20 engineers or traffic exceeds 10M requests/day" — not just "should be reviewed periodically".]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/decision-scoping.md** — What Deserves an ADR (and What "Context" Must Contain). Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/adr.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Context explains the *why* — not just the *what*

- ☐ At least 2 options are documented (including the rejected ones)
- ☐ Rejected options include honest reasons for rejection
- ☐ Consequences include *negative* consequences — no decision is consequence-free
- ☐ Decision is stated in plain language in the Decision section
- ☐ Risks section identifies what would invalidate this decision
- ☐ Context section states the problem explicitly in its first 1–2 sentences (does not assume the reader knows what problem the team was solving)
- ☐ Each rejected option's "Why ruled out" explanation names a specific constraint or trade-off (not a circular statement like "didn't meet our requirements")

Anti-Patterns

- ☐ Do not write an ADR after the decision has already been fully implemented and the team has moved on — ADRs written retrospectively often omit the real reasons and alternatives
- ☐ Do not list only the chosen option — rejected options with honest reasons are the most valuable part of an ADR for future readers
- ☐ Do not write consequences that are all positive — every architectural decision involves trade-offs; an ADR with no negative consequences was not scrutinised honestly
- ☐ Do not leave the status as "Proposed" indefinitely — an ADR that no one has approved is not guiding anyone's decisions
- ☐ Do not write context that assumes the reader already knows what problem was being solved — the context section exists precisely for readers who lack that background

Usage Examples

- "Write an ADR for using [technology]"
- "Document our decision to [architectural choice]"
- "Create an architecture decision record for [topic]"
- "Help me write up why we chose [option] over [alternative]"

changelog-generator

Convert a git log, commit list, or release notes into a polished, user-facing changelog.

Changelog Generator Skill

Converts raw git commits, a diff summary, or developer release notes into a polished changelog entry — categorised, user-facing, and following Keep a Changelog conventions.

Required Inputs

Ask for these if not provided:

- **Commits or release notes** (paste `git log --oneline`, raw commit messages, or a description of what changed)
- **Version number** (e.g. 2.4.0, v1.0.0-beta.2)
- **Release date** (or "today")
- **Audience** (developers using an API / end users of a product / internal team — affects language)
- **Any breaking changes** (flag these explicitly if known)
- **Previous version behaviour** (optional — paste the previous changelog entry or describe what is changing; needed for accurate "Changed" entries)
- **Scope** (whole product / specific package or module — e.g. "payments SDK only", "iOS app", "all services")

Output Format

Follow Keep a Changelog format:

[X.Y.Z] — YYYY-MM-DD

Breaking Changes

[Only include if there are breaking changes]

- **[Breaking change]:** [What changed and what it breaks]
- **Migration required:** [Specific action the user must take]

Added

- [New feature or capability, written from the user's perspective]
- [Another addition]

Changed

- [Changed behaviour — what it did before vs. what it does now]
- [Performance improvement with measurable impact if known]

Fixed

- [Bug fixed — describe what was broken, not the fix implementation]

- [Another fix]

Deprecated

- [Deprecated thing] — use [replacement] instead. Will be removed in [version].

Removed

- [Removed thing] — was deprecated in [version]

Security

- [Security fix — describe the vulnerability class, not exploit details]

Skill guidance — do not include the following section in the delivered changelog:

Formatting Rules Applied

Language: Write for the reader, not the committer. "Add dark mode support" not "implement ThemeProvider with dark palette variant".

Breaking changes: Always call these out first with $\triangle$. Include a migration path.

Bug fixes: Describe what was broken, not what was changed. "Fix crash when user has no profile picture" not "null-check avatar URL before rendering".

Granularity: Group related commits into one line. Don't list every micro-commit separately.

Tone: Active voice, imperative mood. "Add", "Fix", "Remove" — not "Added", "Fixed", "Removed".

Empty sections: Omit any section with no entries. Don't include empty `### Fixed` blocks.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/user-translation.md** — Commit-to-Changelog Translation: Writing for the People Affected. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/release-entry.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Breaking changes are at the top with migration instructions
- ☐ All entries are user-facing language (no internal variable names or implementation details)
- ☐ Related commits are grouped into single entries (not listed individually)
- ☐ Version and date header is correct
- ☐ Empty sections are omitted
- ☐ No entries start with past-tense verbs (no "Added", "Fixed", "Removed" — use "Add", "Fix", "Remove")
- ☐ Every breaking change entry includes a specific migration action (not just "update your code")

Anti-Patterns

- ☐ Do not include implementation details in changelog entries — users need to know what changed for them, not how the code was refactored internally
- ☐ Do not list every micro-commit as a separate entry — related commits should be grouped into one user-facing change
- ☐ Do not omit the migration path for breaking changes — a breaking change entry without a specific migration action forces users to read the source code
- ☐ Do not include empty sections — a "#### Fixed" section with no entries signals the template was filled in carelessly
- ☐ Do not write breaking changes in the same casual tone as minor additions — breaking changes must be visually prominent and call out migration requirements explicitly

Usage Examples

- "Write a changelog for version [X]" + [paste commits]
- "Generate release notes from these commits"
- "Turn this git log into a CHANGELOG entry"
- "Write the CHANGELOG.md update for this release"
- "What changed in this release?" + [paste commit list]

code-review-checklist

Generate a tailored code review checklist for any pull request based on the language, type of change, and risk level.

Code Review Checklist Skill

Produces a tailored code review checklist for a specific pull request — scaled to the language, type of change, and risk level. Not a generic template.

Required Inputs

Ask the user for these if not provided:

- **Language and framework** (e.g. TypeScript + React / Python + FastAPI / Go)
- **Type of change** (feature / bug fix / refactor / dependency upgrade / security patch / performance)
- **Risk level** (low / medium / high / critical)
- **PR description** (paste the description or link to the PR)
- **Code or diff** (optional — paste key changed files or a `git diff` ; significantly improves checklist specificity)
- **Author context** (new starter / experienced / external contributor)

Output Format

Code Review: [PR Title or Reference]

1. PR Overview

Scope assessment: [Small / Medium / Large / Too large — should be split] **Recommended review depth:** [Skim / Standard / Deep dive] **Estimated review time:** [e.g. 20–30 min — use 5 min per 50 lines of diff as a rough guide]

2. Correctness Checks

Language-specific correctness checks — choose based on the language stated:

For TypeScript/JavaScript:

- Type definitions match actual usage
- No implicit `any` in non-test code
- Async/await used consistently; no unhandled promises
- Null/undefined handling is explicit

For Python:

- Type hints present on public functions
- Exception handling is specific (no bare `except`)
- Resources are closed (context managers, `with` blocks)

For Go:

- Errors are handled or explicitly ignored with a comment
- Context propagation is correct
- Goroutine lifetimes are bounded

[Include only the section matching the stated language]

3. Change-Type-Specific Checks

For bug fixes:

- A test exists that would have caught this bug
- The fix addresses root cause, not symptom
- Related code paths checked for the same issue

For features:

- Acceptance criteria met
- Edge cases handled (empty, large, concurrent)
- Error paths tested, not just happy path
- Telemetry/logging added for debugging

For refactors:

- Behaviour unchanged (tests still pass)
- No scope creep — refactor only
- Complexity reduced, not just moved

For dependency upgrades:

- Breaking changes reviewed
- Security advisories checked
- License compatibility verified

[Include only the section matching the stated change type]

4. Risk-Appropriate Checks

Low risk: basic correctness, style conventions, test coverage **Medium risk:** above + rollback plan, monitoring updates, performance considerations **High risk:** above + security implications, data migration safety, feature flag/gradual rollout **Critical risk:** above + staging validation plan, incident response plan, post-deploy verification checklist

5. Testing Adequacy

- Unit tests cover new logic
- Integration tests cover the contract changes
- Edge cases tested

- Failure modes tested
- Performance tests if performance-sensitive

6. Review Decision Framework

Approve if: [2-3 specific conditions based on this PR] **Request changes if:** [Specific blockers] **Comment (non-blocking) if:** [Items worth discussing but not blocking merge]

7. Common Pitfalls for This Change Type

Based on the change type and language, flag 2-3 things reviewers typically miss for this combination.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/review-depth-calibration.md** — Calibrating Review Depth: Not Every PR Deserves the Same Eyes. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/review-record.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Checklist is tailored to the stated language (not generic)
- ☐ Change-type-specific section is included
- ☐ Risk-appropriate depth matches stated risk level
- ☐ Decision framework includes at least one named blocking condition and one named non-blocking comment condition
- ☐ Common pitfalls are specific to the stated language + change-type combo (not generic advice like "watch out for bugs")

Anti-Patterns

- ☐ Do not generate a generic checklist that ignores the stated language — a Python checklist and a Go checklist have fundamentally different correctness concerns
- ☐ Do not treat "looks fine" as a valid review outcome — the checklist exists to surface specific concerns, not validate a superficial read
- ☐ Do not scope a "high risk" review the same as a "low risk" review — depth must scale with the stated risk level
- ☐ Do not flag every stylistic preference as a blocking issue — distinguish between blocking correctness issues and non-blocking comments

- ☐ Do not skip the "common pitfalls" section for the stated language and change-type combination — this is where the most valuable knowledge lives

Usage Examples

- "Generate a code review checklist for [PR description]"
- "What should I check in this pull request?"
- "Give me a code review checklist for a [language] [change type]"
- "Review checklist for a high-risk PR in [language]"

incident-postmortem

Write a structured incident postmortem or post-incident review.

Incident Postmortem Skill

This skill produces a complete, blameless incident postmortem document following industry-standard format. Output enforces blameless framing throughout — system gaps over individual failures — and drives toward specific, closeable action items rather than vague process commitments.

Proposes Actions

The action items don't have to stay on the page: hand them to `action-runner`, which previews them (dry-run, risk-rated), runs only what you approve via the connected action MCP, and records what was done back to the brain. Typical: **file a follow-up issue per action item** (🟡), assigned to its owner with a due date. This skill proposes; action-runner gates and runs — never silently.

Required Inputs

Ask the user for these if not provided:

- **Incident title / ID**
- **Severity** (P1 / P2 / P3 or SEV1 / SEV2 / SEV3)
- **Date and duration** of the incident
- **What happened** (rough notes are fine — the skill will structure them)
- **Services or systems affected**
- **Customer impact** (how many users, what was degraded)
- **How it was detected**
- **How it was resolved**
- **Initial thoughts on root cause**

- **Action items already identified** (optional)
- **Responders** (who was on-call or responded — names or roles; used for the timeline, not for blame)
- **Customer or external communications sent** (optional — any status page updates, emails, or support messages with timestamps)

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it before asking:

- **Read first:** the affected system's `entities/` file and any related prior `decisions/` or past incidents (recurring root causes are the most important thing to surface).
- **Write after:** log the action items and decisions to `decisions/` , and the root-cause learning to `knowledge/` — tag a measured cause `[data]` and a suspected one `[hunch]` , never the reverse.

Deeper Materials

- `references/root-cause-digging.md` — five-whys done properly (stop at a changeable system property, branch into cause/detection/response chains), a contributing-factor taxonomy to sweep, and blame-shaped → systemic language rewrites. Use it while writing the Root Cause section and to reframe any blameful input notes.
- `templates/review-meeting-agenda.md` — a 45-minute, document-first agenda for the postmortem review meeting, with ground rules and an action-item quality gate. Offer it alongside the finished postmortem.

Output Format

Incident Postmortem: [Incident Title]

Incident ID: [ID] **Severity:** [P1/P2/P3] **Date:** [Date] **Duration:** [Start time → Resolution time — total duration] **Status:** [Resolved / Monitoring / Ongoing] **Author:** [Leave blank for user to fill] **Last updated:** [Date]

Executive Summary

[3–5 sentences. Describe what happened, who was affected, and what was done to resolve it. Written for a non-technical stakeholder. No jargon. No blame.]

Impact

Dimension	Details
Users affected	[Number or percentage]
Services degraded	[List affected services]
Business impact	[Revenue, SLA breach, support tickets, etc. if known]
Duration	[Total time from first detection to full resolution]

Timeline

List events in chronological order. Each entry: [HH:MM UTC] — [What happened. Who did what. What changed.]

Rules for timeline entries:

- Use passive or system-focused language — avoid "X made a mistake"
- Include: first symptom, detection, escalation, hypothesis tested, fix applied, confirmation of resolution
- Note time between key events (e.g. "22 minutes between detection and escalation")

Timeline, drawn — also render the incident timeline as a Mermaid Gantt so the gaps (e.g. detection → escalation) are visible at a glance (it renders live in the playground and exports as PNG). Use the incident phases as bars; keep it blameless and system-focused:

```
gantt
    title Incident timeline (UTC)
    dateFormat HH:mm
    axisFormat %H:%M
    section Phases
        Undetected impact    :22:00, 18m
        Detection             :milestone, 22:18, 0m
        Investigation         :22:18, 22m
        Mitigation            :22:40, 15m
        Resolved              :milestone, 22:55, 0m
```

Root Cause

Primary root cause: [One clear sentence. Technical but plain. "A misconfigured deployment config caused..."]

Contributing factors:

- [Factor 1 — e.g. lack of canary deployment meant change hit 100% of traffic immediately]

- [Factor 2 — e.g. alert threshold was set too high to catch the initial degradation]
- [Factor 3 — add as many as are relevant]

Why did our existing safeguards not prevent this? [Honest paragraph explaining why monitoring, tests, or processes didn't catch this earlier. This is where blameless analysis matters most — focus on system gaps, not individual failures.]

Detection

- **How was it first detected?** [Customer report / automated alert / internal monitoring / manual observation]
 - **Time from incident start to detection:** [X minutes]
 - **Should we have detected this faster?** [Yes / No — and why]
-

Resolution

What fixed it? [Clear description of the actual fix — one paragraph] **Why did this work?** [Brief technical explanation] **Was there a temporary mitigation before full resolution?** [Yes/No — describe if yes]

Action Items

#	Action	Owner	Due Date	Priority
1	[Specific, testable action]	[Team or person]	[Date]	P1/P2/P3

Rules for action items:

- Each action must be specific enough to close as "done" or "not done" — no vague items like "improve monitoring"
 - Distinguish between: **Prevent recurrence** (fix the root cause), **Improve detection** (catch it faster next time), **Improve response** (resolve it faster next time)
 - Assign a real owner — not "team" or "TBD" if avoidable
 - Flag P1 actions as items that block the incident from being marked fully closed
-

What Went Well

[3–5 honest observations about the response. Include: fast collaboration, good runbooks used, effective escalation, clear communication. This section builds team confidence and reinforces good habits.]

Lessons Learned

[3–5 key insights from this incident that are worth sharing beyond this team. Write these as transferable lessons — e.g. "Our runbook for database failover didn't account for read-replica lag. All runbooks involving database failover should be reviewed."]

Communication Log

[Optional — list external communications sent: status page updates, customer emails, support responses. Include timestamps.]

Quality Checks

- ☐ Timeline has no blame-focused language
- ☐ Root cause is specific (not "human error")
- ☐ Root cause answers "why did this happen?" not just "what happened?" — it names a system or process gap, not a symptom
- ☐ Contributing factors explain the systemic gaps
- ☐ Every action item has an owner and due date
- ☐ "What went well" section is genuine, not token
- ☐ No action item contains vague language like "improve monitoring", "increase resilience", or "better testing" — each must name a specific change
- ☐ Executive summary is readable by non-technical leadership

Anti-Patterns

- ☐ Do not assign blame to individuals — postmortems must focus on system and process failures
- ☐ Do not write action items with vague language like "improve monitoring" — each must name a specific, ownable change
- ☐ Do not skip the contributing factors — root cause alone misses the systemic issues that enable incidents
- ☐ Do not omit the detection timeline — how long it took to detect matters as much as how long it took to resolve
- ☐ Do not treat the postmortem as closed until all action items have named owners and due dates

Usage Examples

- "Write a postmortem for the [incident name] outage"
- "Help me write a P1 incident report"
- "Generate an RCA document for [service] going down on [date]"

- "Draft a blameless postmortem from these notes: [paste notes]"

pr-description-writer

Write a clear, structured pull request description from a git diff, branch summary, or commit list.

PR Description Writer Skill

Writes structured, reviewer-friendly pull request descriptions from a diff, commit list, or informal notes. Covers the what, why, and how-to-review so reviewers can start immediately.

Required Inputs

Ask for these if not provided:

- **What changed** (paste a git diff, `git log --oneline`, or describe the changes in plain English)
- **Why it was changed** (the problem being solved or feature being added)
- **How to test it** (any specific steps a reviewer needs to verify it works)
- **Risk level** (low / medium / high — affects how much reviewer guidance to include)
- **PR type** (feature / bug fix / refactor / dependency upgrade / config change / hotfix)
- **Target branch** (e.g. main / develop / release/2.4 — affects risk framing and reviewer guidance)
- **Linked issue or ticket** (e.g. JIRA-1234, GitHub #567 — or "none")

Output Format

Title

A clear, imperative-mood title under 72 characters: `[type]: [concise description of what changed]`

Examples:

- `feat: add rate limiting to the public API`
- `fix: resolve race condition in session expiry`
- `refactor: extract payment logic into PaymentService`

Summary

2–3 sentences covering:

- What this PR does (the change)

- Why it was needed (the problem or goal)
- The approach taken (at a high level)

Changes Made

Bullet list of specific changes — one bullet per logical change, not per file:

- Added [X] to handle [Y]
- Refactored [A] to reduce [B]
- Removed [C] as it was replaced by [D]
- Updated [E] to fix [F]

Screenshots / Demo

[If UI change: include before/after screenshots or a screen recording] [If API change: include example request/response] [If no visual change and no API contract change: omit this section entirely — do not leave it as a placeholder]

How to Test

Step-by-step instructions a reviewer can follow:

1. [Setup step if needed]
2. [Action to take]
3. [What to verify]
4. [Edge case to check]

Include any specific commands, test data, or environment flags needed.

Testing Checklist

- ☐ Unit tests added/updated
- ☐ Integration tests added/updated
- ☐ Edge cases covered
- ☐ Manual testing completed
- ☐ No regressions in existing tests

Reviewer Notes

Flag anything that warrants extra attention:

- Areas of uncertainty where a second opinion is welcome
- Deliberate trade-offs made (and why)
- Out-of-scope items noticed but not addressed
- Dependencies on other PRs (link them)

Related

- Closes #[issue number] (if applicable)
- Related to #[PR/issue number]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/reviewer-empathy.md** — PR Descriptions as Review Navigation. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/pr-template.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Title is imperative mood and under 72 characters
- ☐ Summary explains what AND why (not just what)
- ☐ Changes list describes logical changes (not file-by-file changes)
- ☐ Title starts with a valid type prefix (feat / fix / refactor / chore / deps / config / hotfix) and is under 72 characters
- ☐ Testing steps are reproducible by someone unfamiliar with the code
- ☐ For high-risk PRs, Reviewer Notes flags at least one specific area of concern or deliberate trade-off; for low-risk PRs, Reviewer Notes is either omitted or kept to one line

Anti-Patterns

- ☐ Do not write a description that only restates what changed — explain why the change was made
- ☐ Do not skip the testing steps — reviewers need to know how to verify the change works
- ☐ Do not omit the reviewer notes for high-risk PRs — flag deliberate trade-offs and areas needing careful review
- ☐ Do not describe implementation details that are obvious from the diff — add context that the diff cannot convey
- ☐ Do not produce a single paragraph — structure with headers so reviewers can navigate to what they need

Usage Examples

- "Write a PR description for these changes" + [paste diff or description]
- "Draft a pull request for [feature]"

- "I need a PR description — here's what I changed"
- "Summarise these commits into a PR description"
- "Write the PR body for this branch"

runbook-writer

Write an operational runbook for a service, incident type, or deployment procedure.

Runbook Writer Skill

Produces operational runbooks for services, incident types, and deployment procedures — structured so an on-call engineer who's never touched the system can follow them under pressure.

Required Inputs

Ask for these if not provided:

- **What the runbook is for** (e.g. deploying the payment service, responding to a database failover, rotating API keys)
- **Runbook type** (Deployment / Incident Response / Maintenance / Disaster Recovery)
- **System/service name and what it does** (brief description)
- **Audience** (new on-call engineers / experienced SREs / DevOps team)
- **Tech stack** (where relevant — e.g. Kubernetes, AWS RDS, Node.js)
- **Monitoring tools** (e.g. Grafana, Datadog, CloudWatch, Splunk — used to name specific dashboards and alert links in the steps)
- **Key environment details** (e.g. Kubernetes cluster name, AWS account/region, relevant namespaces or resource names — paste what's relevant for exact commands)

Output Format

Runbook: [Runbook Title] **Service:** [Service Name] **Type:** [Deployment / Incident Response / Maintenance / DR] **Last Updated:** [Insert today's date in YYYY-MM-DD format] **Owner:** [Team or person] **Severity:** [P1 / P2 / P3 — if incident-type]

Overview

What this runbook covers: [1–2 sentences on the scenario this runbook handles]

When to use this runbook:

- [Specific trigger condition 1 — e.g. PagerDuty alert: `high-error-rate-payment-service`]
- [Specific trigger condition 2 — e.g. Deploy needed after PR merged to `main`]

Estimated time to complete: [X minutes / X–Y minutes depending on outcome]

Impact if not completed correctly: [e.g. Payment processing degraded / Data loss risk / Users locked out]

Prerequisites

Access required:

- ☐ [System/tool access — e.g. AWS Console: `production-account`]
- ☐ [Credential — e.g. `vault read secret/payment-service`]
- ☐ [VPN / bastion access if needed]

Tools required:

- ☐ [Tool name and version — e.g. `kubectl v1.28+`]
- ☐ [CLI or dashboard name]

Before you start:

- ☐ [Prerequisite check — e.g. Verify current deployment is healthy in Grafana]
 - ☐ [Prerequisite action — e.g. Announce in `#ops-live` that you're starting]
-

Procedure

Number every step. Use exact commands. Do not paraphrase tool names or flags.

Step 1: [Action name] [What you're doing and why — one sentence]

```
# Exact command  
[command here]
```

Expected output: [what should appear if this worked] **If this fails:** [Exact error message to look for] → [What to do, or see Troubleshooting]

Step 2: [Action name] [Same structure as Step 1]

Step 3: Verify Always include a verification step after the main procedure:

```
[verification command]
```

Expected state: [What a healthy system looks like after this runbook completes]

Rollback

How to undo this procedure if something went wrong:

Step R1: [Rollback action]

[rollback command]

Verify rollback: [command to confirm rollback succeeded]

Troubleshooting

Symptom	Likely Cause	Resolution
[Error message or observable symptom]	[Why this happens]	[Exact fix or next step]
[Another symptom]	[Cause]	[Resolution]

Escalation

If this runbook does not resolve the issue:

Condition	Who to Contact	How
[e.g. DB unavailable after 10 min]	[DBA on-call]	[PagerDuty policy: db-oncall]
[e.g. Payment provider unresponsive]	[Vendor contact]	[Contact in 1Password: vendor-escalation]

Always update the incident timeline in [tool] before escalating.

Post-Procedure Checklist

After completing the runbook:

- ☐ Announce completion in #ops-live with outcome
 - ☐ Update the incident ticket / deploy log
 - ☐ Verify alerts have resolved in monitoring dashboard
 - ☐ If this revealed a gap in this runbook — update it now (link to edit process)
-

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/3am-usability.md** — The 3AM Test: Runbooks for Degraded Humans. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/runbook.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every step has an exact command (no "run the deploy script")
- ☐ Expected output is specified for each step so engineer knows if it worked
- ☐ Failure path is explicit for each step (not "if it fails, investigate")
- ☐ Rollback procedure is complete and independently testable
- ☐ Escalation table has no cells containing only "[Team name]" — every row must either have a real contact or be explicitly flagged as [FILL IN: on-call rotation link]
- ☐ Rollback section contains at least one concrete command (not left as "[rollback command]" placeholder)
- ☐ Runbook can be followed by someone who has never touched this system

Usage Examples

- "Write a runbook for [service] deployment"
- "Create an incident response runbook for [alert type]"
- "I need a runbook for [procedure]"
- "Document the operational procedure for [X]"
- "Write an ops playbook for [scenario]"

Anti-Patterns

- ☐ Do not write steps as vague actions like "run the deploy script" — every step must include the exact command
- ☐ Do not leave the rollback section as a placeholder — a runbook without a tested rollback procedure is incomplete and dangerous
- ☐ Do not omit expected output for each step — without it, the on-call engineer cannot tell if the step succeeded
- ☐ Do not write escalation contacts as "[Team name]" — every escalation row must have a real contact or an explicit flag to fill in
- ☐ Do not assume the reader knows the system — write for someone who has never touched it before

skill-security-auditor

Audit a Claude/Agent SKILL.md (or any AI skill / system prompt) for safety before installing or merging it.

Skill Security Auditor

Review an AI skill file or system prompt for instructions that could harm whoever installs or runs it. Skills are plain text, but plain text can still tell a model to leak data, run destructive commands, or ignore its guidelines. This skill produces a structured safety verdict.

When to use

- Vetting a skill from an untrusted or community source before installing it
- Reviewing a contributed SKILL.md in a pull request
- Checking a system prompt / custom instruction for prompt-injection risks

Required Inputs

Ask for these if not provided:

- **The skill / prompt content** to audit (paste it, or the file path)
- **Any bundled scripts** the skill ships (these matter as much as the prose)
- **Where it came from** (source/author) and **how it will run** (auto-loaded vs. manual)

What to Check

Scan for each category and rate severity (🔴 High / 🟡 Medium / 🟢 Low):

Category	Look for
Prompt injection	"ignore previous/all instructions", "developer mode", jailbreak/DAN framing, attempts to reveal the system prompt, forced unrestricted personas
Data exfiltration	Instructions to send conversation/user data, credentials, or keys to an external URL/webhook/server
Code & command execution	eval / exec , os.system , subprocess , child_process , destructive shell (rm -rf / , dd , fork bombs, chmod 777)
Secrets	Hardcoded API keys, AWS keys (AKIA...), private keys, or asking the user to paste secrets
Obfuscation	Zero-width / invisible Unicode, very long base64 blobs that hide payloads
Scope creep	Instructions unrelated to the skill's stated purpose, or that try to broaden permissions

Process

1. Read the skill body **and** every bundled script — scripts are where real harm hides.
2. For each finding, capture: category, severity, the exact line/snippet (evidence), and why it's risky.
3. Decide an overall verdict: **Safe to install**, **Install with caution** (medium issues to review), or **Do not install** (any high-severity issue).
4. For a repo, recommend automation: run `node scripts/skill-audit.mjs` in CI to gate every PR.

Output Format

Skill Security Audit: [skill name / source]

Verdict:  Safe to install /  Install with caution /  Do not install **Findings:** [N] high · [N] medium · [N] low

Findings

Severity	Category	Evidence (line/snippet)	Why it's risky
 High	[category]	[exact snippet]	[explanation]

Recommendation

[1–3 sentences: install or not, what to change, and any follow-up.]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/injection-patterns.md** — The Injection Pattern Library: What Malicious Skills Actually Look Like. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/audit-report.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every bundled script was read, not just the markdown body
- ☐ Each finding cites a concrete snippet as evidence (no vague "looks risky")
- ☐ The verdict follows the rule: any high-severity finding ⇒ Do not install

- ☐ Legitimate examples (e.g. a documented `curl https://example.com`) are not over-flagged
- ☐ The recommendation is actionable (what to remove/change, not just "be careful")

Anti-Patterns

- ☐ Do not pass a skill as safe without reading its scripts — prose can look clean while a script exfiltrates data
- ☐ Do not treat every mention of "API key" or "curl" as malicious; weigh intent and context
- ☐ Do not give a vague verdict — always land on install / caution / do-not-install with reasons
- ☐ Do not ignore zero-width or invisible characters; they are a classic way to hide instructions
- ☐ Do not assume a high star count or popular author means a skill is safe — audit the content itself

technical-debt-register

Document and prioritize a technical debt backlog with business impact, effort estimates, and resolution strategy.

Technical Debt Register Skill

Produce a complete technical debt register for a team or service. A debt register is not a complaint list — it is a prioritized, business-impact-aware inventory that lets an engineering team make deliberate choices about which debt to pay down, in what order, and with what expected return.

Good debt management is not eliminating all debt. It is ensuring debt is visible, owned, and resolved when the interest cost exceeds the cost of fixing it.

Required Inputs

Ask for these if not already provided:

- **Team or service name** — what team and/or service this register covers
- **Known debt items** — list of known technical debt, or ask Claude to elicit them by asking about: legacy code, missing tests, outdated dependencies, architectural shortcuts, manual processes, observability gaps, security backlogs
- **Tech stack** — language, frameworks, infrastructure (helps Claude categorise and score items correctly)

- **Team size and velocity** — number of engineers and approximate story points or days per sprint (needed for effort estimates)
- **Current quarter / planning period** — so the roadmap targets the right timeframe

Output Format

Technical Debt Register: [Team / Service Name]

Team: [Name] | **Service(s):** [Name(s)] **Author:** [Name] | **Last updated:** [Date] **Planning period:** [Q[X] [Year]] | **Review cadence:** [Monthly / Quarterly]

Overview

[2–3 sentences describing the team's current debt situation, the main categories of debt, and the business context — e.g. are they in a growth phase where velocity matters, or approaching a compliance deadline where security debt is critical?]

Total items in register: [X] **Unresolved items:** [X] **Critical/High priority items:** [X]
Estimated total resolution effort: [X story points / X engineer-weeks]

Debt Category Definitions

Category	Description	Examples
Code quality	Code that works but is hard to change safely	Duplicated logic, deeply nested conditionals, inconsistent error handling, missing abstraction
Architecture	Structural decisions that limit scalability or increase coupling	Monolith that should be decomposed, sync calls that should be async, missing domain boundaries
Testing	Gaps in test coverage that increase regression risk	Missing unit tests, no integration tests, flaky test suite, no test data management
Security	Known vulnerabilities or missing security controls	Outdated dependencies with CVEs, missing rate limiting, hard-coded secrets, insufficient auth
Dependencies	Outdated or risky external dependencies	End-of-life libraries, major version lag, abandoned packages

Category	Description	Examples
Infrastructure	Infrastructure that limits reliability or developer productivity	Manual deployment steps, no IaC, single-AZ, missing autoscaling
Observability	Gaps in visibility that slow incident response	Missing metrics, no distributed tracing, poor log structure, no alerting on key SLIs
Process	Manual or error-prone operational processes	Manual DB migrations, no runbooks, tribal knowledge not documented

Debt Register

Scoring Method

Business impact (1–5):

- 5 — Blocking growth, causing production incidents, or creating compliance risk
- 4 — Significantly slowing delivery or increasing incident likelihood
- 3 — Noticeable slowdown; manageable but accumulating
- 2 — Minor friction; low immediate risk
- 1 — Cosmetic or aspirational; no current business impact

Effort to resolve (1–5, lower = easier):

- 1 — <0.5 day; single engineer
- 2 — 0.5–2 days; single engineer
- 3 — 3–5 days; single engineer or small pair
- 4 — 1–2 weeks; team collaboration required
- 5 — >2 weeks; significant planning and coordination

Priority score = Business impact × (6 – Effort) (*rewards high-impact, low-effort items*)

ID	Item	Category	Business impact (1–5)	Effort (1–5)	Priority score	Status	Owner
TD-001	[e.g. No integration tests for payment flow]	Testing	5	3	15	Open	[Name]
TD-002	[e.g. Authentication library 3 major	Security	5	2	20	Open	[Name]

ID	Item	Category	Business impact (1–5)	Effort (1–5)	Priority score	Status	Owner
	versions behind]						
TD-003	[e.g. Database queries not using connection pooling]	Architecture	4	2	16	Open	[Name]
TD-004	[e.g. Manual deployment process for [service]]	Infrastructure	4	3	12	In progress	[Name]
TD-005	[e.g. 200-line God function in order processing]	Code quality	3	3	9	Open	[Name]
TD-006	[e.g. No structured logging — plain text only]	Observability	3	2	12	Open	[Name]
TD-007	[e.g. ORM version has known N+1 query issue]	Dependencies	3	3	9	Open	[Name]
TD-008	[e.g. No runbook for [critical operation]]	Process	3	1	15	Open	[Name]
TD-009	[e.g. Test coverage at 34% — no meaningful safety net]	Testing	4	4	8	Open	[Name]
TD-010	[e.g. Hard-coded config values in application code]	Code quality	2	1	10	Open	[Name]

ID	Item	Category	Business impact (1–5)	Effort (1–5)	Priority score	Status	Owner
TD-011	[e.g. Service deployed single-AZ with no failover]	Infrastructure	5	4	10	Open	[Name]
TD-012	[e.g. No alerting on P95 latency for [endpoint]]	Observability	4	1	20	Open	[Name]

Category Breakdown

Category distribution (by item count):			
Code quality		[X items]	([X]%)
Architecture		[X items]	([X]%)
Testing		[X items]	([X]%)
Security		[X items]	([X]%)
Dependencies		[X items]	([X]%)
Infrastructure		[X items]	([X]%)
Observability		[X items]	([X]%)
Process		[X items]	([X]%)
Priority distribution:			
Critical (score 20–25): [X items]			
High (score 12–19): [X items]			
Medium (score 6–11): [X items]			
Low (score 1–5): [X items]			

Top 5 Priority Items — Resolution Plans

TD-XXX: [Highest priority item name]

Priority score: [Score] | **Category:** [Category] | **Owner:** [Name]

Problem: [2–3 sentences describing what the debt is, how it manifests, and what pain it currently causes. Be specific — reference actual incidents, slowdowns, or risks.]

Business impact: [What happens if this is not resolved? Reference any incidents, near-misses, or growth blockers. E.g. "This caused 2 production incidents in the last quarter and adds ~30 minutes of debugging time to any change in this area."]

Resolution approach: [Clear description of the fix. Not "improve the code" — describe the actual work: "Extract the payment processing logic into a dedicated `PaymentService` class, write unit tests to 80% coverage, and update the 3 call sites."]

Steps:

1. [Specific, ticketable step]
2. [Specific, ticketable step]
3. [Specific, ticketable step]

Acceptance criteria:

- ☐ [Measurable criterion — e.g. "Zero hard-coded config values remain in application code"]
- ☐ [Measurable criterion — e.g. "CI pipeline passes with new tests"]
- ☐ [Measurable criterion]

Effort estimate: [X story points / X days] **Suggested sprint:** [Q[X] Sprint [Y] / When [dependency] is complete]

TD-XXX: [Second priority item name]

Priority score: [Score] | **Category:** [Category] | **Owner:** [Name]

Problem: [Description]

Business impact: [Impact description]

Resolution approach: [Approach description]

Steps:

1. [Step]
2. [Step]
3. [Step]

Acceptance criteria:

- ☐ [Criterion]
- ☐ [Criterion]

Effort estimate: [X story points / X days] **Suggested sprint:** [Sprint or timeframe]

TD-XXX: [Third priority item]

(Follow same format as above)

TD-XXX: [Fourth priority item]

(Follow same format as above)

TD-XXX: [Fifth priority item]

(Follow same format as above)

Debt Reduction Roadmap

Guiding principles

- Allocate [X%] of each sprint's capacity to debt resolution — recommended 15–20% for healthy teams
- Security and dependency debt is addressed on a fixed cadence regardless of priority score
- No new feature work in modules with Critical debt unless the debt is scheduled for the current sprint
- Debt items closed without a resolution (accepted/deferred) must have a named owner and a review date

Quarterly plan

Quarter	Focus area	Items targeted	Estimated capacity	Expected outcome
[Q1 Year] (current)	Security + observability	TD-002, TD-012, TD-006	[X] points / [Y] eng-days	Auth library current; latency alerting live; structured logging shipped
[Q2 Year]	Architecture + reliability	TD-003, TD-011, TD-004	[X] points / [Y] eng-days	Connection pooling fixed; multi-AZ deployed; deploy automation complete
[Q3 Year]	Testing coverage	TD-001, TD-009	[X] points / [Y] eng-days	Payment flow integration tests live; overall coverage ≥60%
[Q4 Year]	Code quality + process	TD-005, TD-008, TD-010	[X] points / [Y] eng-days	God functions refactored; runbooks complete; zero hard-coded config

Sprint allocation model

Sprint capacity: [X] story points

Allocation:

- └ Feature work: $[X * 0.75 = \sim Y]$ points (75%)
- └ Debt resolution: $[X * 0.15 = \sim Y]$ points (15%)
- └ Unplanned/bugs: $[X * 0.10 = \sim Y]$ points (10%)

Debt items that fit in one sprint ($[\leq Y]$ points each):

- ✓ TD-002 ($[X]$ points)
- ✓ TD-012 ($[X]$ points)
- ✓ TD-006 ($[X]$ points)
- ✓ TD-008 ($[X]$ points)

Multi-sprint debt items (break into phases):

- ~ TD-001: Phase 1 ($[X]$ pts) → Phase 2 ($[X]$ pts)
- ~ TD-009: Requires dedicated debt sprint or pairing

Accepted / Deferred Debt

Items where the cost of remediation currently exceeds the business value, accepted with explicit review dates.

ID	Item	Reason for deferral	Review date	Owner
TD-XXX	[Item]	[e.g. "Rewrite would require 3 weeks with no user-facing value at current scale; revisit at 10× traffic"]	[Date]	[Name]
TD-XXX	[Item]	[e.g. "Dependency has a CVE but no upgrade path exists until Q3; mitigated by WAF rule"]	[Date]	[Name]

Policy: No item may be deferred more than twice without escalation to the engineering manager.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/debt-pricing.md** — Pricing Debt: Turning "It's Bad" Into a Number Someone Can Rank. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/debt-entry.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every item has a named owner — no unowned debt

- ☐ Priority scores are calculated using the formula, not assigned arbitrarily
- ☐ Security and dependency items are not scored below their actual business impact because they feel "technical"
- ☐ Top-5 resolution plans include specific, ticketable steps — not vague descriptions like "improve test coverage"
- ☐ The quarterly roadmap allocates realistic capacity — debt allocation does not exceed actual sprint budget
- ☐ Accepted/deferred items have a review date and a named owner — no permanently deferred items
- ☐ The register distinguishes between debt (deliberate or accumulated shortcuts) and bugs (unintended defects)
- ☐ Items are closed as resolved only when acceptance criteria are met — not when the PR is merged

Anti-Patterns

- ☐ Do not score debt items arbitrarily — priority scores must be calculated using the documented formula
- ☐ Do not conflate technical debt (deliberate shortcuts) with bugs (unintended defects) — they require different remediation strategies
- ☐ Do not underrate security and dependency items because they feel abstract — score based on actual business impact
- ☐ Do not create "permanently deferred" items — every accepted item must have a review date and named owner
- ☐ Do not include resolution plans that are vague descriptions — each plan must have specific, ticketable steps

writing-great-skills

Author a high-quality Agent Skill (SKILL.md) that an AI reliably triggers and executes well — strong frontmatter, a sharp description with trigger phrases, a clear output contract, quality checks, and anti-patterns.

Writing Great Skills Skill

A skill is a promise: *given this kind of request, produce this kind of professional output, every time*. The best SKILL.md files win on two things — the model **triggers** them at the right moment, and once triggered it **produces the right artifact** without hand-holding. This skill helps you write one that does both.

Working from a brief

Given a rough idea ("a skill for writing changelogs"), **produce the full SKILL.md anyway** — infer the deliverable, inputs, and structure, and mark genuinely open choices. Never hand back a skeleton with `<!-- TODO -->` left in; fill them.

Required Inputs

Ask for (if not already provided), else infer and label:

- **What the skill should do** and the **concrete artifact** it produces
- **When it should trigger** (the phrasings a user would actually type)
- **The inputs** it needs from the user
- Any **framework or standard** it encodes (for attribution)

The anatomy of a great SKILL.md

1. Frontmatter (this is what gets your skill *found*)

```
---
name: kebab-case-name           # matches the folder; short, specific
description: "<one rich sentence>"
---
```

The **description is the most important line in the file** — it's all the model sees when deciding whether to load the skill (progressive disclosure: only names + descriptions are in context until one is invoked). A strong description has three parts:

- **What it does** + the concrete deliverable.
- A **"Use when ..." trigger clause** listing the real phrasings ("Use when asked to write a postmortem, do a root-cause analysis, or document an incident").
- A **"Produces ..." clause** naming the output ("Produces a blameless postmortem with timeline, root cause, and action items").

Write triggers the way users *speak*, not the way you'd categorise the skill. Cover synonyms.

2. One-line value statement

Open the body with a single sentence on the value, in the voice of a senior practitioner.

3. Working from a brief

State that the skill delivers a complete artifact even with thin input — infer and label assumptions, never leave bracketed placeholders, never refuse for missing context. This is what separates a skill that *works* from one that nags.

4. Required Inputs

A short list of what to ask for — and an instruction to proceed with labelled inferences if they're missing.

5. Output Format / Structure

The heart of the skill: a concrete template — real headings, tables, and sections — of the final artifact. Show the shape, don't describe it abstractly. This is where most of the quality lives.

6. Quality Checks

A short checklist the output must satisfy (the rubric a reviewer would apply). Make them *observable*.

7. Anti-Patterns

The specific failure modes to avoid — the lazy or generic outputs a weaker model would produce.

Process

1. Nail the **deliverable** in one sentence before writing anything else.
2. Write the **description** and stress-test the triggers ("would the model pick this over a neighbouring skill?").
3. Draft the **Output Format** as a real template.
4. Add **Quality Checks** and **Anti-Patterns** that target this skill's specific failure modes.
5. Validate: `npm run skillcheck` (structure) and run it against a thin brief to confirm it doesn't beg for inputs.

Output Format

Return:

1. The **complete SKILL.md** in a fenced block, ready to save to `skills/<name>/SKILL.md`.
2. A 3–5 bullet **"why this works"** note: the trigger phrases chosen, the deliverable, and the sharpest anti-pattern it guards against.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/description-engineering.md** — Description Engineering: the 300 Characters That Decide Everything. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.

- `templates/skill-scaffold.md` — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ `name` is kebab-case and matches the intended folder
- ☐ Description states what it does, has a "Use when ..." trigger clause, and names what it **Produces**
- ☐ Body has: value line, working-from-a-brief, inputs, a concrete Output Format template, Quality Checks, Anti-Patterns
- ☐ No `TODO` /placeholder text left in
- ☐ Triggers are distinct from neighbouring skills (won't mis-fire or get skipped)
- ☐ Would pass `npm run skillcheck` with no errors

Anti-Patterns

- A vague description with no trigger phrases — the skill never gets picked
- An Output Format that *describes* the artifact instead of *templating* it
- Quality Checks that aren't observable ("output should be good")
- Leaving `<!-- TODO -->` or `[bracketed]` placeholders in the final file
- Overlapping so heavily with an existing skill that the model can't choose between them

Chapter 2 — Delivery

7 skills

ab-test-planner

Design statistically rigorous A/B tests for product features, UI changes, onboarding flows, and pricing experiments.

A/B Test Planner Skill

Design experiments that produce trustworthy results — not just directional signals. Every test output includes hypothesis, success metrics, sample size, duration, and a results interpretation guide.

Required Inputs

Ask the user for these if not provided:

- **What is being tested** (feature, UI change, copy, pricing, onboarding step)
- **Hypothesis** (or ask to help formulate one)
- **Primary metric** (conversion rate, click-through, completion rate, etc.)
- **Baseline rate** and **minimum detectable effect** (MDE)
- **Daily eligible users** (to calculate duration)

Experiment Design Checklist

Before running any test, confirm:

- ☐ Clear hypothesis with predicted direction
- ☐ Single primary metric (plus up to 2 guardrail metrics)
- ☐ Minimum detectable effect (MDE) defined
- ☐ Sample size calculated
- ☐ Test duration estimated
- ☐ Segment isolated (no overlap with other running tests)
- ☐ Rollback plan defined

Hypothesis Template

"We believe that [change] will cause [primary metric] to [increase/decrease] by [X%] for [user segment], because [rationale based on data or insight]."

Never run a test without a directional hypothesis. "Let's just see what happens" is not a hypothesis.

Sample Size Calculator Logic

Use this formula (provide the output, not the formula, to the user):

- **Baseline conversion rate:** Current rate of primary metric
- **MDE:** Smallest change worth detecting (recommend 10–20% relative lift for most features)
- **Statistical power:** 80% (standard)
- **Significance level:** 95% ($p < 0.05$)

For common scenarios, provide pre-calculated estimates:

Baseline Rate	MDE (Relative)	Required Sample per Variant
5%	20%	~19,000
10%	15%	~14,000
20%	10%	~15,000
40%	10%	~9,500
60%	5%	~42,000

Always warn: "These are estimates. Use a tool like Evan Miller's calculator or Statsig for precision."

Test Duration Guidance

Minimum: 2 full weeks (to capture weekly seasonality) Maximum: 4 weeks (novelty effect distorts results beyond this)

$$\text{Duration} = \text{Required sample} \div (\text{Daily traffic} \times \% \text{ exposed})$$

Flag if traffic is too low to reach significance in under 8 weeks — recommend a different approach (e.g., holdout test, qualitative research).

Output Format

A/B Test Plan — [Test Name] — [Date]

Hypothesis:

[Filled hypothesis template]

Variants:

- Control (A): [Current experience]
- Treatment (B): [Changed experience — be specific]

Primary Metric: [Metric name + how measured] **Guardrail Metrics:** [Metrics that must not degrade]

Target Segment: [Who sees the test — % of traffic, user type] **Traffic Split:** [50/50 recommended unless ramp-up needed]

Sample Size Required: ~[N] users per variant **Estimated Duration:** [X] weeks (based on [Y] daily eligible users) **Significance Threshold:** 95% confidence, 80% power

Exclusions: [Any user segments to exclude and why]

Rollback Trigger: If [guardrail metric] degrades by [X%], stop the test immediately.

Results Interpretation Guide:

-  Ship if: Treatment shows [X%]+ lift on primary metric at 95% confidence AND guardrail metrics are stable
-  Iterate if: Direction is positive but not significant — consider extending or redesigning
-  Reject if: No lift or negative direction at significance
-  Inconclusive: Do not ship. Do not call it a win.

Guidelines

- Always recommend against peeking at results before the test reaches planned sample size — explain p-hacking risk
- If user wants to test multiple variants, explain the multiple comparisons problem and recommend a Bonferroni correction or a Bayesian approach
- If traffic is very low (<1,000 users/day), recommend qualitative alternatives: moderated testing, 5-second tests, or user interviews
- Never approve a test with no guardrail metrics — always protect revenue, retention, or core engagement

Anti-Patterns

- ☐ Do not run a test without a directional hypothesis — "let's see what happens" produces uninterpretable results

- ☐ Do not declare a winner before reaching the pre-planned sample size — peeking at results inflates false positive rates
- ☐ Do not test multiple independent changes in a single variant — you won't know which change caused the result
- ☐ Do not use engagement metrics (clicks, time-on-page) as the primary metric when the goal is revenue or retention — proxy metrics mislead
- ☐ Do not ignore guardrail metrics — a conversion lift that causes a support ticket spike is not a win

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/test-validity-traps.md** — The Validity Traps That Quietly Invalidate A/B Tests. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/test-plan.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Hypothesis is directional (predicts a specific direction and magnitude, not "let's see")
- ☐ Primary metric is singular (guardrail metrics are secondary)
- ☐ Sample size is calculated from actual MDE and baseline (not guessed)
- ☐ Test duration accounts for weekly seasonality (minimum 2 weeks)
- ☐ Guardrail metrics are defined (at least one to protect revenue or core engagement)
- ☐ Rollback trigger is specified with a concrete threshold

product-launch-checklist

Generate a comprehensive pre-launch, launch day, and post-launch checklist for any product release.

Product Launch Checklist Skill

Never launch without checking everything. Generate a complete, role-assigned checklist covering pre-launch readiness, launch day execution, and post-launch monitoring.

Proposes Actions

Once the checklist is approved, it can be *executed*: hand the items to `action-runner`, which previews them (dry-run, risk-rated), runs only what you approve via the connected action MCP (GitHub/Linear/Slack), and records what was done back to the brain. Typical: **open an issue per checklist item** in the named repo/project (🟡), and **post the launch summary to Slack** (🔴 — approved individually). This skill proposes; action-runner gates and runs — never silently.

Required Inputs

Ask the user for these if not provided:

- **Launch name** and planned launch date
- **Launch tier** (1 = major product launch, 2 = significant feature release, 3 = incremental update)
- **Team members and their roles** (engineering lead, PM, marketing, support, etc.)
- **Feature description** (what is being launched)
- **Rollback capability** (can this be feature-flagged or reverted quickly?)

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it before asking:

- **Read first:** the `entities/` feature being launched and related `decisions/` (scope, dates, owners).
- **Write after:** log launch decisions and owners to `decisions/`. This skill can also hand the checklist to `action-runner` to file the tickets — which records what was actually done back to the brain, closing the loop.

How to Use This Skill

Provide:

- Launch name and date
- Launch tier (1 = major, 2 = feature, 3 = incremental)
- Team members and their roles

The skill generates a tiered checklist. Tier 3 launches use only the Essentials section. Tier 2 adds Marketing & Comms. Tier 1 uses all sections.

Output Format

Launch Checklist — [Feature/Product Name] — Target Date: [Date]

Launch Tier: [1 / 2 / 3] **Launch Owner:** [PM Name] **Engineering Lead:** [Name] **Go/No-Go Decision By:** [Date and time — typically 24 hours before launch]

 PRE-LAUNCH — Engineering & Product (T-2 weeks)

- ☐ Feature flag created and tested in staging
- ☐ All acceptance criteria signed off by PM
- ☐ Code reviewed and merged to main
- ☐ QA sign-off completed (regression + new feature)
- ☐ Performance testing completed (load, latency)
- ☐ Security review completed (if data or auth changes)
- ☐ Rollback procedure documented and tested
- ☐ Monitoring and alerting configured
- ☐ Error logging in place with correct severity levels
- ☐ Database migrations tested on staging with production data volume

 PRE-LAUNCH — Marketing & Comms (T-1 week)

- ☐ Blog post written, reviewed, and scheduled
- ☐ In-app announcement or tooltip configured
- ☐ Email campaign drafted and QA'd
- ☐ Social media posts drafted and scheduled
- ☐ Landing page or feature page live in staging
- ☐ Press outreach sent (Tier 1 only)
- ☐ Product Hunt / community posts prepared (Tier 1 only)

 PRE-LAUNCH — Sales & Support (T-1 week)

- ☐ Sales enablement one-pager completed
- ☐ FAQ document shared with sales and support teams
- ☐ Help centre articles written and published
- ☐ Support team demo / training completed
- ☐ Customer success team briefed on top accounts
- ☐ Pricing updated (if applicable)
- ☐ Contracts / ToS updated (if applicable)

 PRE-LAUNCH — Analytics (T-1 week)

- ☐ Analytics events firing correctly in staging
 - ☐ Dashboard configured for launch metrics
 - ☐ Baseline metrics documented
 - ☐ Success criteria documented and shared with team
 - ☐ A/B test configured (if applicable)
-

✅ GO / NO-GO DECISION — T-24 hours

Criteria	Status	Owner
All critical bugs resolved	🟢 / 🔴	Eng Lead
QA sign-off complete	🟢 / 🔴	QA
Rollback tested	🟢 / 🔴	Eng Lead
Help centre articles live	🟢 / 🔴	Support
Monitoring active	🟢 / 🔴	Eng Lead
PM sign-off	🟢 / 🔴	PM

Go / No-Go Decision: [GO / NO-GO] **Decision Owner:** [PM + Eng Lead jointly]

🚀 LAUNCH DAY

- ☐ Feature flag enabled for [X%] of users (start low — 5–10%)
 - ☐ Launch confirmed in team Slack/channel
 - ☐ Metrics dashboard open and being monitored
 - ☐ Error rate checked at T+15 min, T+1 hr, T+4 hr
 - ☐ Blog post published / email sent
 - ☐ Social posts live
 - ☐ Support team on standby for first 4 hours
 - ☐ PM available and reachable all day
 - ☐ Feature flag expanded to 50% if T+2hr checks pass
 - ☐ Feature flag expanded to 100% if T+4hr checks pass
-

📊 POST-LAUNCH (D+7, D+30)

- ☐ D+7 metrics review: adoption, errors, support tickets
 - ☐ D+7 customer feedback synthesised
 - ☐ Retrospective scheduled
 - ☐ Learnings documented
 - ☐ D+30 success metrics reviewed against targets
 - ☐ Feature flag removed from codebase (clean up)
 - ☐ Follow-up features added to backlog based on feedback
-

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/launch-tiering.md** — Launch Tiering: Matching Ceremony to Stakes. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/launch-plan.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Launch tier confirmed before generating checklist (scope determines depth)
- ☐ Go/No-Go decision has a named owner and a specific decision time
- ☐ Rollback procedure is documented and tested (not just planned)
- ☐ Feature flag expansion is staged (5% → 50% → 100%), not all-at-once
- ☐ Post-launch retrospective is scheduled at launch time

Anti-Patterns

- ☐ Do not apply a Tier 1 checklist to an incremental update — tier the launch appropriately before generating the checklist
- ☐ Do not launch on a Friday without confirmed weekend engineering coverage
- ☐ Do not leave the Go/No-Go decision owner as "the team" — it must be a named individual
- ☐ Do not skip the rollback plan for Tier 1 and 2 launches — know the revert time before going live
- ☐ Do not close the launch without scheduling the post-launch retrospective — it must be booked at launch time, not after

Guidelines

- The Go/No-Go decision must have a named owner — "the team" is not an owner
- Never launch on a Friday unless you have weekend engineering coverage
- Recommend starting all launches at <10% traffic — even for simple features
- Document rollback time: "We can revert this in X minutes" should be known before launch

retro-analysis

Analyses sprint delivery data and produces a structured retrospective brief.

Retrospective Analysis Skill

Generate a data-grounded retrospective brief that separates facts from feelings, so the team spends retro time on solutions rather than debating what happened.

Required Inputs

Ask the user for these if not provided:

- **Sprint tickets: planned vs. completed**
- **Carry-over tickets and reasons** (if known)
- **Tickets reopened after closing** (quality signal)
- **Any incidents or unplanned work** (scope creep signal)
- **Sprint velocity vs. historical average** (trend context)

Process

1. Calculate: completion rate, carry-over rate, unplanned work percentage
2. Identify patterns: which ticket types were most likely to carry over? Which caused blockers?
3. Note any process or communication breakdowns visible in the data
4. Prepare 3 "Start / Stop / Continue" prompts based on the data — not generic, specific to this sprint
5. Suggest 1 concrete experiment for the next sprint based on the biggest friction point
6. **Validate** — Confirm each prompt is specific to this sprint (not a recycled generic prompt), and that the recommended experiment is concrete and measurable

Output Structure

Sprint [Number] Retrospective Brief

By the Numbers:

- Planned: [n] tickets | Completed: [n] | Carry-over: [n] | Completion rate: [%]
- Unplanned work: [n] tickets ([%] of capacity)
- Velocity: [points] vs. [average] average

What the Data Suggests: [2-3 observations grounded in the numbers above]

Discussion Prompts:

- Start: [specific prompt based on this sprint's data]
- Stop: [specific prompt based on this sprint's data]
- Continue: [specific prompt based on this sprint's data]

Suggested Experiment for Next Sprint: [One concrete, testable process change — with a specific success metric]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/root-cause-vs-symptom.md** — Retros That Change Things: Root Causes vs Symptoms. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/retro-board.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Each Start/Stop/Continue prompt names a specific behaviour, not a vague category
- ☐ The recommended experiment is testable in one sprint
- ☐ Carry-over analysis identifies the ticket type or cause, not just the count
- ☐ Data observations don't assign blame — they describe patterns
- ☐ Velocity trend is mentioned in context (is this a one-off or a pattern?)

Anti-Patterns

- ☐ Do not assign blame to individuals in the retrospective brief — observations must describe patterns, not people
- ☐ Do not produce Start/Stop/Continue prompts that are vague categories — each must name a specific behaviour
- ☐ Do not recommend an experiment that cannot be completed within one sprint — small, testable experiments only
- ☐ Do not treat carry-over tickets as a velocity problem without first identifying the root cause category
- ☐ Do not run the same retrospective format every sprint — vary the format to prevent engagement fatigue

sprint-brief

Generate a structured sprint brief from sprint data and goals.

Sprint Brief Skill

Produce a clear, scannable sprint brief that every team member — engineer, designer, PM — can read in under three minutes and understand exactly what we're doing and why.

Required Inputs

Ask the user for these if not provided:

- **Sprint name and number**
- **Sprint goal** (1-2 sentences — flag if too vague)
- **Ticket list with owners** (or a description of the work)
- **Known dependencies or blockers**
- **Carry-over items from previous sprint** (if any)

Process

1. Read sprint goal and check it's specific and measurable — flag if it's too vague
2. Group tickets by theme or feature area
3. Identify the critical path — which tickets must complete for the sprint goal to be met?
4. Flag risks: tickets with unclear acceptance criteria, missing designs, unresolved dependencies
5. Note carry-over items and whether they affect this sprint's goal
6. **Validate** — Confirm the sprint goal is achievable given the ticket scope and capacity.
If the critical path items alone would fill the sprint, flag it as overloaded.

Output Structure

Sprint [Number] Brief — [Dates]

Sprint Goal: [1-2 sentences — specific and measurable] **Why This Sprint Matters:**
[Connect to quarterly OKR in 2-3 sentences]

What We're Building:

- [Theme 1]: [tickets and owners]
- [Theme 2]: [tickets and owners]

Critical Path: [The 2-3 tickets everything else depends on]

Risks to Flag:

- [Risk 1 + mitigation]
- [Risk 2 + mitigation]

Carry-over from Last Sprint: [List + impact on current goal]

Definition of Done: [Specific, agreed criteria for sprint success]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/goal-writing.md** — Writing Sprint Goals That Steer. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/brief-one-pager.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Sprint goal is specific enough to score pass/fail at the end of the sprint
- ☐ Critical path items are named — not just "the important ones"
- ☐ Every risk has a mitigation or owner (not just "this is a risk")
- ☐ Carry-over items are connected to their impact on this sprint's goal
- ☐ Definition of Done is agreed criteria, not a task list

Anti-Patterns

- ☐ Do not write a sprint goal as a task list — the goal must be a single outcome-focused statement that can be scored pass/fail
- ☐ Do not leave the critical path unnamed — "the important tickets" is not a critical path
- ☐ Do not list risks without a mitigation or owner — a risk without a response is just a worry list
- ☐ Do not ignore carry-over items' impact on this sprint's capacity and goal
- ☐ Do not write a Definition of Done that mixes task completion with outcome criteria — they must be observable and agreed before the sprint starts

sprint-planning

Structure and facilitate sprint planning sessions.

Sprint Planning Skill

Transform raw backlog items into a structured, achievable sprint with clear goals, velocity-calibrated scope, and team-ready output.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** priority `decisions/` (what the team agreed matters), feature `entities/`, and open `hypotheses/` the sprint might test. Run `python3 ../professional-brain/scripts/brain_query.py ./brain "<sprint goal>"` and carry each fact's provenance tag through.
-  **Propose to the Brain:** after producing, propose logging the sprint commitment (goal + committed scope) as a `decisions/` record, provenance-tagged. Show it, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Proposes Actions

Once the sprint is agreed, hand it to `action-runner`: it previews (dry-run, risk-rated), runs only what you approve via the connected action MCP, and records what was done back to the brain. Typical: **create a ticket per committed backlog item** and **set the sprint milestone** (🟡). This skill proposes; action-runner gates and runs — never silently.

What This Skill Produces

- **Sprint Goal** — single, outcome-focused sentence the whole team can rally around
- **Sprint Backlog** — prioritised list of user stories with story point estimates and acceptance criteria
- **Capacity Plan** — team availability breakdown accounting for holidays, meetings, and focus time
- **Sprint Planning Agenda** — structured 2-hour meeting agenda with timings
- **Risk Flags** — blockers or dependencies that could derail the sprint

Required Inputs

Ask for (if not already provided):

- Sprint duration (1 or 2 weeks)
- Team size and velocity (average story points per sprint)
- Top 3–5 backlog items or epics to pull from
- Any known absences, holidays, or team events
- Previous sprint's incomplete items (carry-overs)

Sprint Goal Formula

Use this structure:

"This sprint we will [deliver X outcome] so that [user/business benefit], measured by [success indicator]."

Never write sprint goals as task lists. Always outcome-first.

Story Point Calibration

Complexity	Points	Description
Trivial	1	Clearly understood, no unknowns
Small	2	Straightforward, minor effort
Medium	3	Some complexity, clear path
Large	5	Complex, needs design or research
Very Large	8	High uncertainty, may need splitting
Epic	13+	Too large — must be split before sprint

Flag any item estimated at 8+ and recommend splitting.

Capacity Formula

```
Available capacity = (Team size × Sprint days × Focus hours/day) ×
Availability factor
Focus hours/day: 6 (accounting for meetings, Slack, admin)
Availability factor: 0.7–0.85 depending on holidays/events
Story points to commit = Historical velocity × Availability factor
```

Programmatic Helper

This skill ships with a stdlib-only Python script that computes capacity instead of estimating it by hand. Use it whenever the team's numbers are known — it applies the availability and 80% commit-ratio rules consistently.

```
# Quick estimate from flags
python3 scripts/capacity_calculator.py --team 5 --days 10 --velocity 30 --
availability 0.8 --carryover 5

# Detailed estimate from per-member availability (JSON via stdin or --input
file.json)
echo '{"sprint_days":10,"historical_velocity":40,"carryover_points":8,
      "members":[{"name":"Ada","available_days":10},
{"name":"Linus","available_days":7}]}' \
| python3 scripts/capacity_calculator.py --input -
```

The script returns available focus hours, a velocity figure adjusted for real availability, the **recommended commitment** (capped at 80% of velocity), and the remaining **capacity for new work** after carry-overs. Run it first, then build the sprint backlog to fit the recommended number. Add `--json` to pipe the result into other tooling.

Output Format

Sprint [N] — [Start Date] to [End Date]

Sprint Goal:

[Goal statement]

Team Capacity: [X] story points available (based on [Y] team members, [Z]% availability)

Sprint Backlog:

Priority	Story	Points	Owner	Acceptance Criteria
1	[Story title]	[N]	[Team member]	[When X then Y]

Carry-Overs from Previous Sprint:

- [Item] — Reason for carry-over: [brief explanation]

Risks & Dependencies:

- [Risk description] → Mitigation: [action]

Sprint Planning Agenda:

- 00:00–00:10 — Review sprint goal and team capacity
- 00:10–00:40 — Walk through backlog items, confirm estimates
- 00:40–01:20 — Assign stories, identify dependencies
- 01:20–01:50 — Review acceptance criteria per story
- 01:50–02:00 — Confirm sprint commitment and close

Guidelines

- Always challenge stories missing acceptance criteria — flag them explicitly
- Recommend the team commits to 80% of available capacity, not 100%
- If no velocity data is provided, assume 20–30 points for a 5-person team as a starting point
- Highlight any story with unclear ownership as a blocker

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/capacity-honesty.md** — Capacity Honesty — the numbers teams lie to themselves about. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.

- `templates/planning-worksheet.md` — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Sprint goal is outcome-focused (not "implement X" — something like "users can do Y")
- ☐ Team capacity is calculated using actual availability, not theoretical 100%
- ☐ Every story has an acceptance criterion (flag any that don't)
- ☐ Stories estimated at 8+ points are flagged for splitting
- ☐ Carry-overs from last sprint are accounted for in capacity

Anti-Patterns

- ☐ Do not write sprint goals as task lists — goals must be outcome-focused and scoreable pass/fail at sprint end
- ☐ Do not commit to 100% of available capacity — always recommend 80% to preserve slack for unplanned work
- ☐ Do not carry stories with no acceptance criteria into the sprint — flag them as blockers before committing
- ☐ Do not allow stories estimated at 8+ points into the sprint without splitting them first
- ☐ Do not ignore carry-over items when calculating capacity — they consume capacity and must be accounted for before new work is pulled in

Execution

For tool-using or computer-use agents that can reach the team's tracker (Jira, Linear, GitHub Projects). Runtimes without tool access ignore this section and deliver the document. See SKILLSPEC.md §5 for the rules this block follows.

Preconditions

- The sprint plan above has been produced and **explicitly approved by a human** — never build a sprint from an unreviewed draft.
- Tracker access is already authenticated in the agent's environment; the target board/project is named by the user.
- A dry-run listing of intended changes has been shown and confirmed.

Allowed actions

- Create the sprint/iteration container with the approved name and dates.

- Move the approved, already-existing backlog items into the sprint — only the items listed in the approved plan.
- Set story-point estimates on those items to the approved values.
- Post the sprint goal as the sprint description or a pinned comment.
- Nothing else: no creating new issues, no deleting or closing anything, no editing item descriptions, no touching other sprints.

Verification

- Re-read the sprint from the tracker: item count and total points equal the approved plan; every moved item is in the sprint; sprint dates match.
- Post the verification summary (items, points, dates) back to the user.

Rollback

- Undo = move the items back to the backlog and delete the empty sprint container.
- Stop and ask a human if: any item in the plan no longer exists or changed since approval, the tracker rejects an action, or the board contains an active sprint with overlapping dates.

technical-spec-template

Create structured technical specification documents that bridge product requirements and engineering implementation.

Technical Spec Template Skill

Write technical specifications that engineers actually read — clear problem framing, unambiguous requirements, explicit decisions, and documented trade-offs.

Required Inputs

Ask the user for these if not provided:

- **Feature or system description** (what needs to be specced)
- **Related PRD or product brief** (if available)
- **Engineering reviewers** (whose sign-off is needed)
- **Known constraints** (technical limitations, security requirements, performance targets)

When to Write a Tech Spec

Write a tech spec when:

- The feature requires changes to 2+ systems
- There are significant architectural decisions to make
- More than one engineer will work on the implementation
- The feature has security, privacy, or compliance implications
- Estimated effort is >5 story points

Skip the spec for trivial bug fixes or 1-2 hour changes.

Technical Spec Output Format

Technical Specification — [Feature Name]

Author: [Name] **Status:** Draft | In Review | Approved | Implemented **Created:** [Date] | **Last Updated:** [Date] **Reviewers:** [Eng Lead, Architect, PM, Security if needed] **Related PRD:** [Link] | **Jira Epic:** [Link]

1. Problem Statement

[2–3 sentences. What problem are we solving and why now? No solution language here.]

2. Goals & Non-Goals

Goals (in scope):

- [Specific, measurable outcome]
- [Specific, measurable outcome]

Non-Goals (explicitly out of scope):

- [What this spec does NOT cover]
- [Common assumption to shut down early]

3. Background & Context

[Any prior art, related systems, or context engineers need to understand the decision space. Link to previous specs, ADRs, or research.]

4. Proposed Solution

High-Level Approach: [2–4 sentences describing the chosen solution. Why this approach vs alternatives?]

System Architecture Diagram: [Describe or embed: which services are involved, how data flows, what APIs are called]

Data Model Changes:

```
-- New tables or schema changes
[Include DDL or schema definition]
```

API Design:

```
[Endpoint] [Method]
Request: { [fields and types] }
Response: { [fields and types] }
Error codes: [list]
```

Key Implementation Details:

- [Important technical constraint or approach]
- [Edge case handling]
- [Third-party dependency and version]

5. Alternative Approaches Considered

Option	Pros	Cons	Why Rejected
[Alt 1]	[Benefits]	[Drawbacks]	[Reason not chosen]
[Alt 2]	[Benefits]	[Drawbacks]	[Reason not chosen]

6. Security & Privacy Considerations

- Data stored: [What PII or sensitive data is involved]
- Authentication: [How is access controlled]
- Authorisation: [What permissions are required]
- Encryption: [At rest / in transit requirements]
- Compliance implications: [GDPR, SOC2, etc. if relevant]

7. Performance & Scalability

- Expected load: [Requests/second, data volume]
- Latency requirements: [P50 / P95 targets]
- Caching strategy: [If applicable]
- Database indexing: [New indexes required]
- Known bottlenecks: [Where to watch]

8. Testing Plan

- Unit tests: [Key scenarios to cover]
- Integration tests: [System boundaries to test]
- Load tests: [If performance-critical]
- Edge cases: [Known tricky scenarios]
- Rollback plan: [How to revert if something goes wrong]

9. Rollout Plan

- Feature flag: [Yes / No — name of flag]
- Rollout stages: [% of users at each stage]
- Monitoring: [Metrics and alerts to set up]
- Success criteria to progress rollout: [What needs to be true]
- Rollback trigger: [What would cause immediate rollback]

10. Open Questions

Question	Owner	Due Date	Resolution
[Unresolved question]	[Name]	[Date]	[Pending]

11. Implementation Timeline (Rough)

Phase	Work	Estimated Effort
[Phase 1]	[What gets built]	[X days/points]
[Phase 2]	[What gets built]	[X days/points]
Total		[X story points]

Guidelines

- The spec is a decision record, not a task list — document *why* decisions were made
- All open questions must have an owner and due date
- Security and privacy sections are never optional for features that touch user data
- Recommend async review: engineers read first, then a 30-minute sync to resolve questions
- Keep the spec updated as implementation progresses — stale specs are worse than no specs

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/spec-decisions.md** — What a Spec Is For: Decisions, Alternatives, and the Blast Radius. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/spec-skeleton.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Problem statement contains no solution language
- ☐ Non-goals explicitly list at least 2 things that might be assumed in scope
- ☐ At least 2 alternative approaches are documented with reasons for rejection
- ☐ Security and privacy section is completed for any feature touching user data
- ☐ All open questions have a named owner and due date (not "TBD")

Anti-Patterns

- ☐ Do not include solution language in the problem statement — the problem must be described independently of the proposed solution
- ☐ Do not omit alternatives considered — a spec that considers only one approach has not been properly evaluated
- ☐ Do not leave open questions as "TBD" without a named owner and due date — unresolved questions are blockers
- ☐ Do not skip security and privacy sections for any feature that touches user data
- ☐ Do not write a non-goals section that is empty — always list at least two things that might be assumed in scope

user-story-writer

Write well-structured user stories with acceptance criteria and edge cases.

User Story Writer Skill

This skill produces production-ready user stories from a feature brief, PRD section, or verbal description. Each story follows the standard format with a clear who/what/why, behavioural acceptance criteria in Given/When/Then format, edge cases, and definition of done. Output is ready to paste into Jira, Linear, or your planning tool.

Required Inputs

Ask the user for these if not provided:

- **Feature or change** to break into stories — paste the brief, PRD section, or describe the feature
- **User types / personas** involved (e.g. admin, end user, guest, API consumer)
- **Scope** — are we writing one story or decomposing an epic into a full set of stories?
- **Acceptance criteria format preference** — Given/When/Then, bullet checklist, or both?
- **Technical constraints or notes** — anything the engineering team has flagged that should shape the stories

Output Structure

For each story:

Story: [Short title — verb + noun, e.g. "Filter search results by date range"]

Epic: [Parent epic name — e.g. "Advanced Search"] **Story ID:** [Jira/Linear ID — leave blank if not yet created] **Priority:** [P1 / P2 / P3] **Story points:** [Leave blank — for engineering to estimate]

User Story

As a [specific user type — not "user"], **I want to** [concrete action they want to take], **So that** [the outcome they achieve — business value, not feature description].

Example:

As an **account manager**, I want to **filter my client list by last contact date**, so that I **can quickly identify clients I haven't spoken to in over 30 days and prioritise outreach**.

Context

[1–3 sentences of context that aren't in the user story itself: when does this story matter, what triggers the need, how does it fit into a larger flow. This helps engineers understand why before they ask.]

Acceptance Criteria

Format: Given / When / Then

Each criterion tests one specific behaviour. Write one GWT per observable outcome — not one GWT for the whole feature.

AC1: [Short name for this criterion]

Given [starting state or context]
When [user action]
Then [observable system behaviour]

AC2: [Short name]

Given [...]
When [...]
Then [...]

AC3: [Short name]

Given [...]
When [...]
Then [...]

Edge Cases

[List scenarios that are non-obvious but must be handled. These become additional ACs or notes to engineering.]

- ☐ **[Edge case 1]:** [e.g. User applies a date filter that returns 0 results — show empty state with clear messaging and a "clear filters" action]
- ☐ **[Edge case 2]:** [e.g. User has >10,000 clients — filter must not degrade load time >200ms]
- ☐ **[Edge case 3]:** [e.g. Date filter persists across page refresh — or explicitly should not if that's the decision]
- ☐ **[Permission edge case]:** [e.g. Read-only users can see the filter but cannot save filter presets]

Out of Scope

[Explicitly state what this story does NOT cover — prevents scope creep and clarifies where the next story begins.]

- Saving and sharing filter presets (separate story — see [Story X])
- Bulk actions on filtered results
- Exporting filtered client list to CSV

Definition of Done

- ☐ Acceptance criteria all pass
- ☐ Edge cases handled (or explicitly deferred with a new ticket raised)
- ☐ Unit tests written for each AC
- ☐ Works on mobile viewport (if applicable)
- ☐ Accessibility: keyboard navigable and screen-reader compatible
- ☐ Error states are handled and copy approved
- ☐ Product and design have reviewed in staging
- ☐ No console errors in production build

Epic Decomposition Template

If the user provides an epic or feature brief, decompose it into a full set of stories before writing them:

Epic: [Name] **Goal:** [What outcome does completing this epic achieve?] **Stories:**

#	Story	Notes	Dependencies
1	[Core happy path story — the simplest version of the feature that delivers value]		
2	[Validation / error handling story]		Depends on #1
3	[Edge case or power user story]		Depends on #1
4	[Admin or configuration story]		
5	[Performance or scale story — if applicable]		Depends on #1

Suggested sprint order: [Which stories are P1 for MVP? Which can follow in a later sprint?]

Common Story Anti-Patterns — and Fixes

Use these to review stories before handing to engineering:

Anti-pattern	Example	Fix
Solution in the story	"As a user I want a dropdown filter"	Remove the UI decision — "As a user I want to filter by date range"
Vague "so that"	"so that it's easier to use"	Make it specific — "so that I can prioritise outreach without opening each record manually"
Too big	Story covers 5 distinct user flows	Split into separate stories per flow

Anti-pattern	Example	Fix
No acceptance criteria	Story has description only	Add at least 3 GWT criteria before engineering starts
ACs that test the solution, not the behaviour	"Given the dropdown is open, When I select an option"	Test the outcome — "Given I have applied a date filter, When I view my results, Then only clients last contacted in that date range appear"
Missing empty state	No AC for what happens with 0 results	Add it — empty states are part of the feature
Missing error state	No AC for network failure or invalid input	Add error handling ACs explicitly

Example: Full Story Set for a Feature

Feature brief: "Allow users to export their invoice history as a PDF or CSV"

Story 1: Export invoice list as CSV

As a **finance admin**, I want to **export my invoice history as a CSV file**, so that I can **import it into our accounting software without manual data entry**.

AC1: Successful export

Given I am on the Invoices page with at least one invoice
 When I click "Export" and select "CSV"
 Then a CSV file is downloaded containing all visible invoices with columns:
 Invoice ID, Date, Amount, Status, Customer Name

AC2: Empty state

Given I am on the Invoices page with no invoices
 When I click "Export"
 Then the export button is disabled and a tooltip reads "No invoices to export"

AC3: Filtered export

Given I have applied a date filter showing invoices from Jan 2026 only
 When I click "Export" and select "CSV"
 Then the export contains only invoices from Jan 2026 — not all invoices

Edge cases:

- ☐ Export with >10,000 invoices — must complete in <30s or show a progress indicator
- ☐ Export triggered on mobile — downloads to device's default download location

Out of scope: PDF export (Story 2), scheduled exports (future epic)

Story 2: Export invoice list as PDF

As a **finance admin**, I want to **export my invoice history as a formatted PDF**, so that I can **share a professional summary with our accountant**.

[... ACs follow same pattern ...]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/acceptance-criteria-craft.md** — Acceptance Criteria That Actually Gate. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/story-card.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every story has a specific user type — not "a user" or "the system"
- ☐ The "so that" explains business value — not just feature description
- ☐ Each AC tests one observable outcome — not a bundle of behaviours
- ☐ Empty states, error states, and edge cases are explicitly handled
- ☐ Out of scope is documented — not assumed
- ☐ Stories are independent — they can be shipped individually without depending on unreleased work (except where explicitly noted)

Anti-Patterns

- ☐ Do not write user stories from a technical perspective — every story must be from the user's point of view and state their goal
- ☐ Do not write acceptance criteria that are untestable — every criterion must have a clear pass/fail condition
- ☐ Do not create stories that are too large to complete in a single sprint — break epics into estimable, independently deliverable stories

- ☐ Do not omit edge cases — unhappy paths and error states are required, not optional
- ☐ Do not skip the Definition of Done — without it, "done" means different things to different people

Example Trigger Phrases

- "Write user stories for [feature] from this brief"
- "Break this PRD section into user stories with acceptance criteria"
- "Convert these feature requirements into Jira tickets"
- "Write the user stories and ACs for [feature name]"
- "Decompose this epic into individual stories ready for sprint planning"

Chapter 3 — Cs

6 skills

churn-analysis

Produce a structured churn analysis that separates avoidable from unavoidable churn.

Churn Analysis Skill

Produce a structured churn analysis that goes beyond the headline rate — identifying why customers leave, which segments are most at risk, and what interventions will have the highest impact on retention.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** `context.md` (metric definitions — what "churn" means here), `knowledge/` , and related segment `entities/` . Run `python3 ../professional-brain/scripts/brain_query.py ./brain "churn"` and carry each fact's provenance tag through.
-  **Propose to the Brain:** after producing, propose recording the headline retention finding to `knowledge/` (`[data]`), any retention decision to `decisions/` , and at-risk drivers as `hypotheses/` . Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Required Inputs

Ask for these if not already provided:

- **Time period** being analysed (e.g. Q1, last 12 months)
- **Total customers at start of period** and **customers churned**
- **ARR or revenue lost** to churn
- **Churn reasons data** — exit survey results, CSM notes, support data, or sales loss reasons
- **Customer segments** — by tier, industry, cohort, or product line
- **Current retention rate** if known

- **Any recent changes** — pricing, product, support model — that may have affected churn

Churn Categories

Always classify churn before analysing it:

Category	Definition
Voluntary — avoidable	Customer left due to a problem we could have addressed (product gaps, poor onboarding, relationship failures)
Voluntary — unavoidable	Customer left for reasons outside our control (budget cuts, acquisition, company shutdown)
Involuntary	Payment failure, contract non-renewal by mistake, admin error

The interventions for each category are different. Conflating them leads to wrong conclusions.

Output Format

Churn Analysis: [Product / Segment / Company]

Period: [Start date] — [End date] **Prepared by:** [Name] | **Date:** [Date]

Headline Numbers

Metric	Value
Customers at start of period	[N]
Customers churned	[N]
Customer churn rate	[X]%
ARR at start of period	£/\$/€[X]
ARR lost to churn	£/\$/€[X]
Revenue churn rate (gross)	[X]%
ARR from expansions (same period)	£/\$/€[X]
Net revenue retention (NRR)	[X]%

Benchmark context:

- Customer churn rate: [X]% vs. industry benchmark [Y]% — [above / below / in line]
- NRR: [X]% — [What this means: above 100% = expansion offsets churn; below 100% = shrinking base]

Churn Breakdown by Category

Category	Customers	% of churn	ARR lost
Voluntary — avoidable	[N]	[X]%	£/\$/€[X]
Voluntary — unavoidable	[N]	[X]%	£/\$/€[X]
Involuntary	[N]	[X]%	£/\$/€[X]
Total	[N]	100%	£/\$/€[X]

Avoidable churn as % of total churn: [X]% — this is the number we can actually influence.

Churn Reasons — Avoidable Churn Only

Rank by frequency. Include ARR weight where data allows.

Reason	Count	% of avoidable churn	ARR lost	Representative quote
[Reason 1 — e.g. "Product missing key feature"]	[N]	[X]%	£/\$/€[X]	"[Quote]"
[Reason 2]	[N]	[X]%	£/\$/€[X]	"[Quote]"
[Reason 3]	[N]	[X]%	£/\$/€[X]	"[Quote]"
[Reason 4]	[N]	[X]%	£/\$/€[X]	"[Quote]"
Other	[N]	[X]%	£/\$/€[X]	—

Theme synthesis: [2–3 sentences grouping the top reasons into 2–3 themes. E.g. "The top three reasons cluster around two themes: product gaps in [area] (affecting X% of avoidable churn) and onboarding failures where customers never achieved value (Y%)."]

Churn by Segment

Identify which segments over- or under-index for churn.

By Tier

Tier	Churn rate	vs. Overall	Notes
Enterprise	[X]%	+/-[X]pp	
Mid-Market	[X]%	+/-[X]pp	
SMB	[X]%	+/-[X]pp	

By Cohort (Acquisition Year)

Cohort	Churn rate	Notes
[Year 1]	[X]%	
[Year 2]	[X]%	
[Year 3]	[X]%	

By Industry / Use Case (if data available)

Segment	Churn rate	Notes
[Segment 1]	[X]%	
[Segment 2]	[X]%	

Key pattern: [Which segment has the highest churn rate and what likely explains it]

Timing Analysis

- **Average contract length before churn:** [X months]
- **Highest-risk moment:** [e.g. "Month 3 — when trial value has worn off but full adoption hasn't happened"]
- **Churn timing distribution:**

When churn occurred	% of churned accounts
0–3 months	[X]%
3–6 months	[X]%
6–12 months	[X]%
12+ months	[X]%

Early Warning Signals

Based on the churned accounts, identify the signals that preceded churn (and could have triggered earlier intervention):

Signal	Lead time before churn	How to detect
[Signal 1 — e.g. "DAU/MAU dropped below 15%"]	[~X weeks]	[Usage dashboard / alert]
[Signal 2 — e.g. "No QBR in 90+ days"]	[~X weeks]	[CRM flag]
[Signal 3 — e.g. "Champion left the account"]	[~X weeks]	[LinkedIn alert / CSM tracking]
[Signal 4]	[~X weeks]	[Detection method]

Intervention Recommendations

Ranked by estimated impact × feasibility.

Intervention	Addresses	Est. churn reduction	Effort	Owner
[Intervention 1 — e.g. "Improve onboarding for [segment] with dedicated 30-day check-in"]	[Reason 1]	[X accounts / £X ARR]	Low / Med / High	[Team]
[Intervention 2]	[Reason 2]	[X accounts / £X ARR]	Low / Med / High	[Team]
[Intervention 3]	[Reason 3]	[X accounts / £X ARR]	Low / Med / High	[Team]

Priority call: [Which one intervention, if implemented this quarter, would have the biggest impact and why]

What We Don't Know (Data Gaps)

- [Data gap 1 — e.g. "Exit survey response rate is only 30% — the reasons data may not be representative"]
- [Data gap 2 — e.g. "No product usage data for SMB tier — can't confirm usage signal correlation"]
- [Data gap 3]

Anti-Patterns

- ☐ Do not mix avoidable and unavoidable churn in intervention plans — recommending product fixes for customers who churned due to company shutdown wastes resources
- ☐ Do not calculate churn rate using end-of-period customer count as the denominator — this understates churn; always divide churned customers by the starting cohort
- ☐ Do not rely solely on exit survey data for churn reasons — response rates are typically low and self-selection biases the sample toward customers who are engaged enough to complete a survey
- ☐ Do not recommend interventions without linking them to a specific churn reason — interventions disconnected from root causes will not move retention
- ☐ Do not report only gross revenue churn — without net revenue retention (NRR), a healthy-looking retention number can hide a shrinking revenue base

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/avoidability-calls.md** — Avoidable or Not? The Judgment Calls in Churn Classification. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/churn-report.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Churn rate is correctly calculated (churned ÷ starting cohort, not end-of-period total)
- ☐ Avoidable and unavoidable churn are separated — interventions target avoidable churn only
- ☐ Churn reasons are customer-reported, not internally assumed
- ☐ Segment analysis identifies which segments over-index — not just averages
- ☐ Early warning signals are specific and detectable, not generic ("low engagement")
- ☐ Interventions link directly to the top churn reasons — no recommendations without a root cause match

cs-escalation-brief

Write a structured escalation brief for an at-risk customer account.

Customer Escalation Brief Skill

Produce a clear, concise escalation brief that gives internal stakeholders — VP CS, CCO, product leadership, or the CEO — everything they need to understand the situation, make decisions, and act fast.

A good escalation brief is not a complaint. It is a professional document that states the facts, assigns accountability honestly, and proposes a specific resolution plan.

Required Inputs

Ask for these if not already provided:

- **Account name**, tier, and ARR
- **CSM name** and account owner
- **Nature of the escalation** — what happened, what the customer is saying
- **Timeline** of events leading to escalation
- **Customer contact** who escalated (name, role, influence level)
- **What the customer wants** — their stated ask
- **What we believe the root cause is**
- **What has already been done** to address the situation
- **Renewal date** and current renewal risk assessment

Escalation Levels

Calibrate urgency and audience based on escalation level:

Level	Trigger	Audience	Response time
L1 — Account Risk	Customer expressing dissatisfaction; renewal at risk	CSM + CS Manager	24 hours
L2 — Executive Escalation	Customer escalated to their exec; requesting vendor exec involvement	VP CS + Account Exec	4 hours
L3 — Churn Risk	Customer has issued notice or is in active churn conversation	CCO / CEO + Revenue leadership	1 hour
L4 — Public Risk	Customer threatening public escalation, legal, or press	CCO / Legal / Comms	Immediate

Output Format

Escalation Brief: [Account Name]

Escalation level: L[1/2/3/4] — [Label] **Date raised:** [Date] **Raised by:** [CSM name]

Escalation owner: [Name of exec or senior stakeholder now leading response]

Account at a Glance

Field	Detail
ARR	£/\$/€[X]
Tier	Enterprise / Mid-Market / SMB
Customer since	[Date]
Renewal date	[Date] — [N] days away
Renewal risk (pre-escalation)	Green / Amber / Red
Renewal risk (current)	Green / Amber / Red
Customer contact who escalated	[Name, role, seniority]
Executive sponsor (customer)	[Name, role — active / passive / vacant]
Executive sponsor (vendor)	[Name, role]

What Happened — Summary

[3–5 sentences. State the facts plainly. What the customer experienced, how they reacted, and how we learned about the escalation. No editorialising. No blame.]

Timeline

List in chronological order. Each entry: [Date / time] — [What happened. Who did what.]

Include:

- When the original issue or trigger event occurred
 - When the customer first raised concerns (informally)
 - When it escalated (formal escalation or exec involvement)
 - Actions taken since escalation
-

Root Cause

Primary cause: [One clear sentence. What specifically went wrong.]

Contributing factors:

- [Factor 1 — be honest about internal failures as well as external ones]
- [Factor 2]

Is this a systemic issue or isolated? [] Isolated to this account [] Pattern seen in other accounts — details: [_____] [] Product or process gap that needs fixing

Customer's Stated Position

What the customer says happened: [Their version of events — fair and unfiltered]

What they are asking for: [Their explicit ask — compensation, fix by date, exec call, SLA credit, exit clause]

Sentiment of escalating contact: [Frustrated but constructive / Angry / Seeking exit / Unknown]

Risk of public escalation: Low / Medium / High — [evidence if Medium or High]

Business Impact

Impact type	Detail
ARR at risk	£/\$/€[X]
Potential churn probability	[X]%
Reputational risk	Low / Medium / High
Reference / case study status	[Was a reference — now at risk / Not a reference]
Expansion pipeline at risk	£/\$/€[X]

What Has Been Done So Far

1. [Action taken — by whom — date — outcome]
2. [Action taken — by whom — date — outcome]
3. [Action taken — by whom — date — outcome]

Has a formal apology or acknowledgement been issued? Yes / No

Proposed Resolution Plan

Immediate actions (next 24–48 hours):

Action	Owner	By when
[Action]	[Name]	[Date]
[Action]	[Name]	[Date]

Medium-term actions (next 2–4 weeks):

Action	Owner	By when
[Action]	[Name]	[Date]

What we are NOT offering: [Be explicit about what is not on the table — avoids misaligned expectations]

Success criteria: [How will we know the escalation is resolved? What does the customer need to confirm they are satisfied?]

Decision Required from Escalation Owner

[State clearly what decision or resource the escalation owner needs to provide. Be specific — do not make them ask. E.g.: "We need approval to offer a 20% service credit for Q2" or "We need an exec call with [name] within 48 hours."]

Communication Plan

Audience	Message	Channel	Owner	By when
Escalating customer contact	[Summary of message]	Email / Call	[Name]	[Date]
Customer exec sponsor	[Summary]	Call	[Name]	[Date]
Internal CS team	[Summary]	Slack / Meeting	CS Manager	[Date]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/deescalation-sequencing.md** — De-escalation Sequencing: the Order of Operations When an Account Is on Fire. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.

- `templates/escalation-brief.md` — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Root cause is specific — not "communication breakdown" or "product gap" without detail
- ☐ Customer's position is stated fairly — not minimised or dismissed
- ☐ A clear decision is requested from the escalation owner — brief does not end with "what do you think?"
- ☐ ARR at risk is quantified
- ☐ Communication plan has owners and dates — not "TBD"
- ☐ Language is professional and blameless toward individuals

Anti-Patterns

- ☐ Do not assign blame to individuals — focus on system failures and process gaps
- ☐ Do not downplay ARR at risk or describe churn risk vaguely without a number
- ☐ Do not leave resolution plan ownership as "TBD" or unassigned
- ☐ Do not write the brief without a clear ask from the escalation owner
- ☐ Do not omit the customer's own stated position — their perspective must be represented fairly

cs-health-scorecard

Build a customer health scorecard for a specific account.

Customer Health Scorecard Skill

Produce a structured, data-driven health scorecard for a customer account — giving the CSM and leadership a clear view of renewal risk, expansion potential, and the actions needed to move the account in the right direction.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** the account's `entities/` file, its `stakeholders/` (champion, economic buyer, detractors), and `knowledge/`. Run `python3 ../professional-`

brain/scripts/brain_query.py ./brain "<account name>" and carry each fact's provenance tag through.

-  **Propose to the Brain:** after producing, propose recording the health verdict + key risks to the account `entities/` file, and a renewal-risk entry to `decisions/` if a call is made, each provenance-tagged. Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Required Inputs

Ask for these if not already provided:

- **Account name** and tier (enterprise / mid-market / SMB)
- **Contract value** (ARR) and **renewal date**
- **Product usage data** — logins, DAU/MAU ratio, key feature adoption
- **Support data** — open tickets, CSAT or NPS score, recent escalations
- **Engagement data** — last QBR date, executive sponsor status, champion name
- **Commercial data** — payment history, expansion conversations, seats used vs. licensed
- **Any known risks or recent changes** at the account

Scoring Framework

Score each dimension 1–5. Weight as shown. Calculate weighted total out of 100.

Dimension	Weight	What to Score
Product Adoption	30%	DAU/MAU ratio, breadth of features used, power users identified
Engagement	20%	QBR cadence, executive sponsor active, champion strength
Outcomes	20%	Customer hitting their stated goals / success metrics
Support Health	15%	Ticket volume trend, unresolved escalations, CSAT
Commercial	15%	On-time payments, seats utilised, expansion signals

Score → RAG conversion:

- 80–100: Green (healthy, renew likely)
- 60–79: Amber (at risk, needs attention)
- 0–59: Red (high churn risk, escalate)

Programmatic Helper

This skill ships with a stdlib-only Python script that applies the weights above and converts the weighted total to a RAG status — so the headline score is computed identically every

time and weights always sum to 100%.

```
# Five scores 1-5 in order: adoption engagement outcomes support commercial
python3 scripts/health_score.py --scores 4 3 4 2 5 --account "Acme Corp"

# Or from JSON (lets you override the default weights per account/segment)
python3 scripts/health_score.py --input account.json
```

It returns the per-dimension weighted points, the **total out of 100**, and the **RAG band** (Green ≥80, Amber 60–79, Red <60) with a one-line next step. Run it to set the headline number, then write the dimension detail and actions below around it. Add `--json` for downstream tooling.

Output Format

Customer Health Scorecard: [Account Name]

CSM: [Name] | **Tier:** [Enterprise / Mid-Market / SMB] **ARR:** £/\$/€[X] | **Renewal date:** [Date] | **Days to renewal:** [N] **Overall health:** [Green / Amber / Red] — [Score]/100 **Last updated:** [Date]

Health Score Summary

Dimension	Score (1–5)	Weight	Weighted Score	Trend
Product Adoption	[1–5]	30%	[X]	↑ / → / ↓
Engagement	[1–5]	20%	[X]	↑ / → / ↓
Outcomes	[1–5]	20%	[X]	↑ / → / ↓
Support Health	[1–5]	15%	[X]	↑ / → / ↓
Commercial	[1–5]	15%	[X]	↑ / → / ↓
Total	—	100%	[X]/100	

Dimension Detail

Product Adoption — [Score]/5

- **DAU/MAU ratio:** [X]% (benchmark: >25% = healthy)
- **Key features adopted:** [List features in use]

- **Features not adopted:** [List unused high-value features]
- **Power users identified:** [Yes / No — how many]
- **Assessment:** [1–2 sentences on adoption health]

Engagement — [Score]/5

- **Last QBR:** [Date] — [Outcome summary]
- **Next QBR:** [Scheduled / Overdue]
- **Executive sponsor:** [Active / Passive / Vacant]
- **Champion:** [Name, role, strength: strong / moderate / weak]
- **Assessment:** [1–2 sentences]

Outcomes — [Score]/5

- **Customer's stated goals:** [List 2–3 goals from onboarding or last QBR]
- **Progress against goals:** [On track / Partial / Off track]
- **Evidence of value:** [Metric or quote that demonstrates ROI]
- **Assessment:** [1–2 sentences]

Support Health — [Score]/5

- **Open tickets:** [N] (priority breakdown: P1: X, P2: X, P3: X)
- **CSAT / NPS:** [Score] (benchmark: >8 CSAT / >30 NPS = healthy)
- **Unresolved escalations:** [Yes / No — details if yes]
- **Ticket trend (last 90 days):** Increasing / Stable / Decreasing
- **Assessment:** [1–2 sentences]

Commercial — [Score]/5

- **Seats licensed:** [N] | **Seats active:** [N] ([X]% utilisation)
- **Payment history:** [On time / Late — details]
- **Expansion signals:** [Yes — describe / No]
- **Downgrade or cancellation signals:** [Yes — describe / No]
- **Assessment:** [1–2 sentences]

Top Risks

Risk	Severity	Mitigation
[Risk description]	High / Medium / Low	[Specific action to mitigate]

Recommended Actions

Immediate (this week):

1. [Action — owner — deadline]

This month:

1. [Action — owner — deadline]

Before renewal:

1. [Action — owner — deadline]

Renewal Forecast

Scenario	Probability	ARR at risk
Full renewal at current ARR	[X]%	£/\$/€0
Renewal with contraction	[X]%	£/\$/€[X]
Churn	[X]%	£/\$/€[full ARR]

Recommended renewal play: [Expand / Hold / Save / Manage out]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/leading-signals.md** — Health Signals That Lead (Instead of Eulogise). Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/account-scorecard.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Score is based on data, not gut feel — each dimension has evidence
- ☐ Risks are specific (not "low engagement" — something like "executive sponsor left in March, no replacement identified")
- ☐ Actions have owners and deadlines
- ☐ Renewal probability is calibrated against pipeline reality
- ☐ Trend arrows reflect direction of change vs. last scorecard, not just current state

Anti-Patterns

- ☐ Do not score health dimensions on gut feel — every score needs specific supporting evidence

- ☐ Do not give a Green status to accounts with unresolved P1 issues or missed milestones
- ☐ Do not list risks vaguely — "low engagement" without specifics is not actionable
- ☐ Do not leave recommended actions without named owners and deadlines
- ☐ Do not conflate product usage frequency with product value delivery

customer-success-plan

Build a joint customer success plan for a specific account.

Customer Success Plan Skill

This skill produces a joint customer success plan — a living document shared between the CSM and the customer that aligns on outcomes, milestones, and mutual commitments. Output is ready to co-author with the customer in a kickoff call or QBR.

Required Inputs

Ask the user for these if not provided:

- **Account name** and industry
- **Product / plan purchased**
- **Key stakeholders** — customer champion and economic buyer
- **Customer's stated business goals** — why did they buy? What problem are they solving?
- **Contract term and renewal date**
- **Current onboarding stage** (new customer / expanding / post-QBR / pre-renewal)
- **Seats / licenses / usage purchased**
- **Any known risks** — adoption gaps, champion uncertainty, competing priorities

Output Structure

Customer Success Plan: [Account Name]

Product: [Product name / plan tier] **Contract term:** [Start date → Renewal date] **CSM:** [Name] **Customer champion:** [Name, Title] **Customer executive sponsor:** [Name, Title — if known] **Last updated:** [Date] **Status:** [Active / Under review / Completed]

1. Partnership Objectives

What does success look like for [Account Name] at contract end?

[Write 2–3 sentences describing the customer's core objective in plain English — what they are trying to achieve in their business, not what features they are using.]

Primary business goal: [e.g. Reduce time-to-hire by 30% across engineering teams]

Secondary goal: [e.g. Consolidate three legacy tools into one platform, saving £X/year]

Success statement (customer's words): "[Direct quote from champion about what success looks like — ask for this in kickoff]"

2. Success Metrics

Define how both parties will measure success. Agreed in the kickoff call and tracked in QBRs.

Metric	Baseline (today)	Target	By when	Data source
[e.g. Seat utilisation]	[X%]	[≥ 80%]	[Month 3]	[Product analytics]
[e.g. Time to hire]	[X days]	[< Y days]	[Month 6]	[Customer's ATS]
[e.g. Reports produced/month]	[X]	[≥ Y]	[Month 3]	[Product analytics]
[e.g. NPS]	[X]	[≥ 8]	[Month 6]	[Quarterly survey]

Leading indicators (early signs the plan is on track):

- [e.g. 5+ users log in within the first 2 weeks]
- [e.g. First workflow automated within 30 days]
- [e.g. Champion presents the tool to their team by end of Month 1]

3. Milestone Roadmap

Break the success journey into phases with clear milestones and owners:

Phase 1: Onboard (Month 1)

Milestone	Owner	Due date	Status
Admin setup complete (SSO, permissions, data integration)	[IT contact]	[Date]	[]

Milestone	Owner	Due date	Status
All purchased seats activated and users invited	[Champion]	[Date]	[]
Core workflow [X] configured and tested	[CSM + Champion]	[Date]	[]
First training session delivered (all teams)	[CSM]	[Date]	[]
Kickoff call completed and success plan co-signed	[CSM + Champion]	[Date]	[]

Phase 2: Adopt (Months 2–3)

Milestone	Owner	Due date	Status
[Core feature] in active daily use by $\geq X$ users	[Champion]	[Date]	[]
First business outcome achieved and documented	[Champion + CSM]	[Date]	[]
30-day check-in completed	[CSM]	[Date]	[]
[Power user workflow] enabled for advanced users	[CSM]	[Date]	[]

Phase 3: Value (Months 4–6)

Milestone	Owner	Due date	Status
QBR 1 delivered — ROI evidence presented	[CSM + AE]	[Date]	[]
Success metric [X] hit target	[Champion]	[Date]	[]
Expansion use case identified and introduced	[AE]	[Date]	[]
Reference call or case study agreed	[Champion]	[Date]	[]

Phase 4: Renew & Expand (Months 7–12)

Milestone	Owner	Due date	Status
QBR 2 delivered — renewal conversation started	[CSM + AE]	[Date]	[]
Renewal proposal sent	[AE]	[Date]	[]
Expansion or flat renewal signed	[AE]	[Date]	[]

4. Mutual Commitments

Success plans work when both parties commit. Document what each side will do:

[Vendor] commits to:

- Dedicated CSM available [X days/week / by email within 24 hours]
- Monthly [call / check-in / async update] with champion
- QBR every [90 days] with executive summary and ROI report
- Priority support for [Account] — response SLA of [X hours] for P1 issues
- Roadmap preview for relevant upcoming features
- [Any other specific commitment made in sales cycle]

[Account Name] commits to:

- Champion available for [30-min monthly] check-in
- Users complete onboarding training by [date]
- Feedback on product experience shared monthly (async or sync)
- Executive sponsor participates in QBR 1 and renewal discussion
- Provide outcome data to CSM quarterly for ROI tracking

5. Stakeholder Engagement Plan

Stakeholder	Role	Engagement frequency	Format	Owner
[Champion]	Day-to-day owner	Weekly (async) + Monthly (call)	Slack / Email + Zoom	CSM
[Economic buyer]	Budget holder	Quarterly	QBR (in-person or video)	CSM + AE
[IT contact]	Integration owner	As needed	Email	CSM
[End users]	Active users	Training only	Group session	CSM

6. Risk & Mitigation

Risk	Likelihood	Impact	Mitigation plan
Low adoption in first 30 days	[M]	[H]	CSM hosts live onboarding; champion sends internal comms day 1
Champion changes role	[L]	[H]	Multi-thread: introduce CSM to 2 additional stakeholders by Month 2
Budget pressure at renewal	[M]	[H]	Build ROI case monthly; document value continuously
Competing priorities delay rollout	[H]	[M]	Agree minimum viable adoption path with champion; don't require perfection to declare value

7. Communication Plan

Communication	Audience	Frequency	Format	Owner
Health update	Champion	Monthly	Email summary (3 bullets: what's good, what needs attention, one ask)	CSM
QBR	Champion + Exec	Quarterly	45-min video call with slide deck	CSM + AE
Product updates	Champion	As released	Release notes email	CSM
Support status	Champion	When open tickets exist	Email / Slack	Support + CSM

8. Escalation Path

If the success plan falls off track:

Trigger	Action	Owner	Timeline
Health drops to Amber	Internal review + champion call within 5 days	CSM	Immediate
Health drops to Red	CS leadership + AE looped in; escalation brief drafted	CS Manager	Within 24 hours
Champion is unresponsive for >10 days	AE attempts exec sponsor contact	AE	After CSM attempt fails
Adoption <40% at Month 3	Emergency enablement session + revised milestone plan	CSM	Within 1 week of flag

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/outcome-contracting.md** — Outcome Contracting: Success Plans That Bind Both Sides. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/success-plan.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Success metrics are the customer's metrics — not just product usage metrics

- ☐ Milestones have specific owners and due dates — not "TBD"
- ☐ Mutual commitments section is genuinely mutual — not just what the vendor will do
- ☐ Risk register includes champion departure and low adoption
- ☐ Plan is written to be shared with the customer — no internal-only commentary in this document
- ☐ Executive sponsor is identified and has an engagement role

Anti-Patterns

- ☐ Do not define success metrics that the vendor controls — metrics must reflect the customer's business outcomes
- ☐ Do not set milestone dates without customer confirmation — unilateral timelines undermine joint ownership
- ☐ Do not create a plan the customer hasn't agreed to — it must be mutual, not a CSM's internal plan
- ☐ Do not leave ownership fields blank or assigned to "CS team" — every action needs a named owner
- ☐ Do not confuse product adoption milestones with customer business outcomes — both are needed but are not the same

Example Trigger Phrases

- "Build a success plan for [Account Name] who just signed"
- "Create a joint success plan for our new enterprise customer"
- "Write a 6-month customer success roadmap for [Company]"
- "I need a mutual action plan for our QBR with [Account]"
- "Generate a customer success plan for an at-risk account"

qbr-deck

Build a Quarterly Business Review (QBR) deck structure and narrative for a customer account.

QBR Deck Skill

Produce a complete Quarterly Business Review deck — structured, data-backed, and customer-focused. A good QBR demonstrates value delivered, aligns on goals for the next quarter, and strengthens the executive relationship. It should never feel like a product demo or a vendor update.

Required Inputs

Ask for these if not already provided:

- **Account name**, CSM name, and customer stakeholders attending
- **Contract details** — ARR, contract start date, renewal date
- **Last quarter's goals** (from previous QBR or kickoff)
- **Usage and adoption data** — key metrics for the quarter
- **Support summary** — tickets raised, resolution time, any escalations
- **Business outcomes the customer cares about** — what success looks like for them
- **Product updates or new features** relevant to this customer
- **Goals for next quarter**
- **Any open commercial conversations** (expansion, renewal, at-risk signals)

QBR Principles

- Lead with customer outcomes, not product features
- Every metric should connect to a business result the customer cares about
- The agenda is a conversation, not a presentation — build in time for customer input at every stage
- Close with mutual commitments, not just vendor actions

Output Format

QBR: [Account Name] × [Your Company]

[Quarter] [Year] Business Review

Date: [Date] | **Location / Call link:** [TBC] **Customer attendees:** [Names and roles] **[Your company] attendees:** [Names and roles]

Slide 1: Agenda (5 min)

Time	Topic	Owner
0:00	Welcome and introductions	CSM
0:05	[Last quarter] — how did we do?	CSM + Customer
0:20	Value delivered — business impact	CSM
0:35	What's coming — roadmap preview	CSM / Product
0:45	[Next quarter] — goals and priorities	Customer
0:55	Actions and mutual commitments	CSM
1:00	Close	

Talking point: "We've kept today to 60 minutes. We want as much of this to be a conversation as possible — please push back, redirect, and ask questions throughout."

Slide 2: Where We Are Together (2 min)

Partnership snapshot:

- **Customer since:** [Date]
- **Contract value:** £/\$/€[ARR]/year
- **Renewal date:** [Date]
- **Active users:** [N] of [N] licensed seats ([X]% adoption)
- **Products / modules active:** [List]

Talking point: "Before we dive in — a quick picture of where we are. [X] months in, [Y] active users, and this is our [Nth] QBR together."

Slide 3: Last Quarter — Goals We Set Together (5 min)

Goal	Set in [Last QBR / Kickoff]	Status
[Goal 1]	[What we committed to]	✅ Achieved / ⚠️ Partial / ❌ Missed
[Goal 2]	[What we committed to]	✅ Achieved / ⚠️ Partial / ❌ Missed
[Goal 3]	[What we committed to]	✅ Achieved / ⚠️ Partial / ❌ Missed

For any partial or missed goal: state what happened and what changes next quarter.

Talking point: "Let's start with accountability. Here's what we said we'd achieve last quarter — let's be honest about where we landed."

Slide 4: Usage and Adoption (5 min)

Quarter-over-quarter trend:

Metric	[Q-1]	[Q]	Change
Monthly active users	[N]	[N]	+/-X%
Sessions per user per week	[N]	[N]	+/-X%
[Key feature 1] adoption	[X]%	[X]%	+/-X%
[Key feature 2] adoption	[X]%	[X]%	+/-X%

Highlights:

- [Positive adoption trend to call out]
- [Feature or workflow with strongest engagement]

Opportunity:

- [Feature with low adoption that could drive more value — link to their goals]

Talking point: "Usage is [up / stable / something we want to talk about]. The area I'd like to focus on is [feature] — we're not seeing the adoption we'd expect given [their goal], and I want to understand why."

Slide 5: Business Impact — Value Delivered (10 min)

Lead with outcomes, not activity.

[Outcome 1: customer's primary success metric]

- Before: [baseline]
- Now: [current state]
- Impact: [quantified business result — time saved, revenue influenced, cost reduced, risk mitigated]

[Outcome 2]

- [Same structure]

[Outcome 3]

- [Same structure]

Customer evidence (use if available):

"[Quote from champion or user about value experienced]"

Talking point: "This is the section I most want your input on. Are these the outcomes that matter to your business? Are there other ways you're measuring success that we should be tracking?"

Slide 6: Support Summary (3 min)

Metric	This quarter	Last quarter	Trend
Tickets raised	[N]	[N]	↑ / → / ↓
Average resolution time	[X hrs]	[X hrs]	↑ / → / ↓
P1 / critical issues	[N]	[N]	↑ / → / ↓

Metric	This quarter	Last quarter	Trend
CSAT score	[X/10]	[X/10]	↑ / → / ↓

Notable issues this quarter:

- [Any escalation or major ticket — brief summary and resolution]

What we're doing differently:

- [Any process change or improvement based on support patterns]

Slide 7: What's Coming — Roadmap Preview (5 min)

Focus only on what's relevant to this customer's goals. Do not dump the full roadmap.

Feature / Improvement	Expected	Why it matters to [Account Name]
[Feature 1]	[Q+1]	[Direct link to their goal or pain point]
[Feature 2]	[Q+1 / Q+2]	[Direct link]
[Feature 3]	[H2]	[Direct link]

Talking point: "I've filtered the roadmap to what I think matters most to your team. I'd love your reaction — are these the right priorities from your perspective?"

Slide 8: Next Quarter — Your Goals (10 min)

Customer input section — facilitate, don't present.

Prompt questions:

- "What does success look like for your team in [next quarter]?"
- "What's the biggest challenge you're trying to solve in the next 90 days?"
- "Is there anything about the way you're using [product] you want to change?"

Capture live:

Goal for next quarter	Owner (customer)	How we'll support it	How we'll measure it
[Goal 1]	[Name]	[CSM / product action]	[Metric]
[Goal 2]	[Name]	[CSM / product action]	[Metric]

Slide 9: Mutual Commitments (5 min)

[Your company] commits to:

1. [Specific action — owner — by when]
2. [Specific action — owner — by when]
3. [Specific action — owner — by when]

[Account Name] commits to:

1. [Specific action — owner — by when]
2. [Specific action — owner — by when]

Next touchpoint: [Date of next check-in or mid-quarter review]

Slide 10: Thank You + Open Q&A (5 min)

- Recap the one headline from today: [The single most important thing you want them to remember]
 - Confirm actions are captured and shared after the call
 - Ask: "Is there anything we didn't cover today that you wanted to raise?"
-

Preparation Checklist

- ☐ Usage data pulled and QoQ comparison calculated
- ☐ Last QBR goals reviewed — status confirmed before the meeting
- ☐ Business outcomes framed in customer language (not product language)
- ☐ Roadmap filtered to this account's specific use cases
- ☐ Customer's goals for next quarter researched or pre-confirmed with champion
- ☐ Executive sponsor briefed on any sensitive topics before the call
- ☐ Actions from previous QBR reviewed — any outstanding items addressed

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/value-narrative.md** — The QBR Value Narrative: Their Numbers, Not Your Features. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/qbr-outline.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every slide has a talking point, not just a title

- ☐ Value slide leads with business outcomes, not product activity
- ☐ Roadmap preview links each item to a customer goal
- ☐ Mutual commitments section has real owners on both sides
- ☐ Customer has at least 20 minutes of airtime in the agenda

Anti-Patterns

- ☐ Do not fill the QBR with product activity metrics — lead with business outcomes the customer cares about
- ☐ Do not present a roadmap without linking each item to a customer goal — vendor priorities are not a QBR agenda
- ☐ Do not run a QBR as a one-sided presentation — it must include structured time for the customer to speak
- ☐ Do not close a QBR without documented mutual commitments with named owners on both sides
- ☐ Do not skip the "what's not working" slide — suppressing problems erodes trust and misses renewal risks

renewal-playbook

Build a structured renewal playbook for a customer account.

Renewal Playbook Skill

This skill produces a complete renewal playbook for a specific customer account, covering health assessment, commercial strategy, negotiation preparation, expansion opportunity mapping, and a step-by-step timeline. Output is ready for the CSM or account team to execute 90–180 days before renewal.

Required Inputs

Ask the user for these if not provided:

- **Account name**
- **Renewal date**
- **Current ARR** and proposed renewal ARR (if different)
- **Account health** — RAG status and main reasons (or describe the account situation)
- **Key stakeholders** — economic buyer, champion, and any detractors
- **Renewal risk factors** — budget pressure, low adoption, competitive threat, champion departure, etc.
- **Expansion opportunity** — any upsell or cross-sell potential?
- **Contract terms** — current plan, duration, and any terms up for renegotiation

Renewal Playbook: [Account Name]

Renewal date: [Date] **Current ARR:** [£/\$/€ X] **Target renewal ARR:** [£/\$/€ X — flat / +X% expansion / contraction risk] **Health status:** [Green / Amber / Red] **CSM:** [Name] **Account executive:** [Name] **Days to renewal:** [X days]

1. Account Health Snapshot

Dimension	Score (1–5)	Evidence
Product adoption	[X/5]	[e.g. 3 of 5 purchased seats active; core feature used weekly]
Business outcomes	[X/5]	[e.g. Customer reports X% improvement in [metric]; no formal ROI review done]
Relationship depth	[X/5]	[e.g. Strong champion in [name/role]; limited exec sponsorship]
Support & satisfaction	[X/5]	[e.g. 2 open P2 tickets; last NPS 7; no escalations in 6 months]
Commercial engagement	[X/5]	[e.g. Invoice paid on time; no discount pressure raised yet]
Overall health	[X/5 — weighted]	[Green / Amber / Red]

Renewal thesis: [One sentence: why this account will renew — or what must change for it to renew.]

2. Stakeholder Map

Stakeholder	Role	Influence	Sentiment	Our relationship
[Name]	Economic buyer	High	[Positive / Neutral / Negative]	[Warm / Cold / Unknown]
[Name]	Champion	High	[Positive]	[Warm]
[Name]	End user	Low	[Neutral]	[Limited]
[Name]	IT / procurement	Medium	[Neutral]	[Transactional]

Champion risk: [Is our champion secure in their role? Any signals of departure or reorganisation?]

Multi-thread plan: [Who else do we need relationships with before renewal? How do we get there?]

3. Risk Register

Risk	Likelihood (H/M/L)	Impact (H/M/L)	Mitigation
[Budget pressure / cost-cutting]	[H]	[H]	[Build ROI case 90 days out; identify budget holder's priorities]
[Low adoption in [department]]	[M]	[H]	[Run targeted enablement session; tie to champion's OKRs]
[Competitor evaluation]	[M]	[M]	[Request competitive intelligence; schedule exec-level call]
[Champion departure]	[L]	[H]	[Map two additional stakeholders; executive intro call]

4. Value Story

Build the ROI narrative for the renewal conversation:

Headline result: [e.g. "[Account] saved X hours/week or reduced [metric] by X% using [product]"]

Evidence sources:

- ☐ Product usage data (logins, features used, seat utilisation)
- ☐ Business metric improvement (pull from QBR deck or success plan)
- ☐ Support resolution time improvement
- ☐ Customer-provided testimonial or case study quotes

Value gaps to close before renewal: [Are there outcomes the customer expected but hasn't seen yet? What's the plan to close these?]

5. Expansion Opportunity

Map upside beyond flat renewal:

Opportunity	Type	Estimated value	Likelihood	Timing
[Seat expansion — [dept] wants to add 10 users]	Upsell	[+£X ARR]	[High]	[Renewal or +3M]
[Cross-sell — [Product B] use case identified]	Cross-sell	[+£X ARR]	[Medium]	[+6M]
[Multi-year commitment]	Discount for term	[+£X TCV / -X% discount]	[Low]	[At renewal]

Expansion play: [Which opportunity to lead with, and the sequence for raising it in the renewal conversation]

6. Commercial Strategy

Renewal scenario planning:

Scenario	Probability	ARR outcome	Response strategy
Flat renewal	[X%]	[£X — same as current]	[Accept; plant seeds for +6M expansion]
Expansion	[X%]	[£X]	[Lead with ROI evidence; pitch seat or feature expansion]
Contraction risk	[X%]	[£X — downgrade to lower tier]	[Propose phased commitment; demonstrate path to full adoption]
Churn risk	[X%]	[£0]	[Escalate to leadership; executive sponsor engagement]

Discount guardrails:

- Floor discount: [X% — do not go below without VP approval]
- Triggers for discount: [Multi-year / volume / reference customer commitment]
- What to ask for in return: [Reference case study / G2 review / executive intro / case study participation]

Pricing flexibility:

- [e.g. Can offer monthly billing in exchange for 24-month commit]
- [e.g. Can offer X seats free in exchange for expansion commitment]

7. Objection Responses

Prepare for the most likely objections:

"The price is too high"

Anchor on value delivered: "[Customer] achieved [X outcome] — at [£X ARR], that's [£Y per outcome / hour saved / user]. What would it cost to deliver that outcome without us?" If budget is genuinely constrained, explore: phased payment, reduction in scope rather than full churn, multi-year pricing.

"We're not seeing enough adoption"

Acknowledge, then commit: "You're right — [X seats] are actively using [core feature] out of [Y]. We want to fix this. Here's our 60-day plan: [exec sponsor on enablement call / training session / in-product nudge campaign]."

"We're evaluating [Competitor]"

Don't panic. Ask: "What's driving the evaluation — is it specific features, pricing, or something else?" Then map gaps honestly. Offer a feature roadmap preview if relevant. Get clarity on their criteria and timeline before responding defensively.

"We need to reduce spend this quarter"

Separate the commercial conversation from the value conversation. Offer to protect the relationship with a reduced scope today with a committed expansion trigger at a business milestone. Avoid discounting without a reason.

8. Renewal Timeline

Week	Action	Owner	Notes
W-16 (4 months out)	Internal renewal review — health, expansion opportunity, risk	CSM	Flag to leadership if Red
W-12	QBR / executive business review — ROI evidence delivered	CSM + AE	Book 45–60 min with economic buyer
W-10	Champion 1:1 — pulse check on satisfaction and upcoming priorities	CSM	Uncover internal dynamics before commercial discussion
W-8	Expansion conversation — plant seeds, share roadmap	AE	Do not lead with pricing
W-6	Send renewal proposal — pricing, terms, options	AE	Include multi-year option
W-4	Negotiation — address objections, finalise commercial terms	AE + CSM	Escalate to VP if >X% discount required

Week	Action	Owner	Notes
W-2	Legal / procurement — contract redlines, signature process	AE + Legal	
W-0	Signed. Handoff to post-renewal success plan	CSM	Thank the champion; begin next cycle

9. Success Criteria

- ☐ Renewal signed before deadline
- ☐ ARR outcome within target range
- ☐ Champion relationship maintained or improved
- ☐ At least one expansion conversation started
- ☐ ROI evidence documented and accepted by customer

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/risk-timeline.md** — The Renewal Clock: What Happens at T-minus-When. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/renewal-plan.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Stakeholder map includes the economic buyer — not just the champion
- ☐ Risk register has a mitigation for every H/H risk
- ☐ Value story uses product data and business outcomes, not just feature lists
- ☐ Commercial strategy includes a floor discount and a reason-to-discount framework
- ☐ Timeline starts at least 90 days before renewal date
- ☐ Objection responses are specific to this account, not generic

Anti-Patterns

- ☐ Do not start renewal conversations less than 90 days before the renewal date for accounts over \$50K ARR
- ☐ Do not build a renewal strategy without first honestly assessing account health — wishful thinking leads to last-minute churn

- ☐ Do not treat all renewal objections as negotiating tactics — some objections signal genuine dissatisfaction that requires resolution first
- ☐ Do not offer discounts as the first response to price objections — explore value gaps before reducing price
- ☐ Do not close the renewal without confirming the expansion opportunity — every renewal is also an expansion conversation

Example Trigger Phrases

- "Build a renewal playbook for [Account Name] renewing in [Month]"
- "Help me plan the renewal strategy for an at-risk customer"
- "Prepare a renewal brief for my QBR with [Company]"
- "What's my renewal strategy for a Red account coming up in 60 days?"
- "Create a renewal and expansion plan for [Account]"

Chapter 4 — Discovery

5 skills

assumption-mapper

Extract and risk-rate hidden assumptions in a product brief or PRD.

Assumption Mapper Skill

Surface and prioritize the untested assumptions embedded in any product plan before development begins.

Required Inputs

Ask the user for these if not provided:

- **Product brief, PRD, or concept description** (even rough notes work)
- **Stage** (concept / discovery / pre-build / post-launch — affects which assumptions matter most)

Process

1. Read the provided brief, PRD, or concept description
2. Extract assumptions across four categories:
 - **Desirability** (do users want this?)
 - **Feasibility** (can we build it?)
 - **Viability** (will it sustain the business?)
 - **Usability** (can users actually use it?)
3. Score each assumption:
 - Confidence (1-5): How sure are we this is true?
 - Impact (1-5): How badly does the plan fail if this assumption is wrong?
 - Priority = Impact – Confidence (higher = test first)
4. **Validate completeness** — Ensure at least one assumption per category. If a category is empty, re-read the brief looking specifically for that type.
5. Output a ranked list with recommended validation methods

Output Structure

Assumption Map: [Feature/Product Name]

Assumption	Category	Confidence	Impact	Priority	Validation Method
[assumption]	[type]	[1-5]	[1-5]	[score]	[method]

Critical Assumptions (Impact 4+ and Confidence 2 or below)

[Flagged items with detailed validation recommendations]

Top 3 Assumptions to Validate First

[Detailed recommendations including specific research method, estimated effort, and what the result would change]

Example (Partial)

Input: "We're building a self-serve onboarding flow to reduce time-to-value for SMB customers."

Assumption	Category	Confidence	Impact	Priority	Validation Method
SMB users can complete onboarding without human help	Usability	2	5	3	Unmoderated usability test (n=8)
Faster onboarding correlates with higher retention	Viability	3	4	1	Cohort analysis of current onboarding times vs. 90-day retention
The current onboarding is the primary reason for slow time-to-value	Desirability	2	4	2	User interviews with recent churned SMB accounts

Anti-Patterns

- ☐ Do not only surface desirability assumptions — feasibility and viability assumptions are equally likely to kill a product and are often overlooked
- ☐ Do not assign high confidence to an assumption just because it hasn't been challenged yet — absence of evidence is not evidence
- ☐ Do not recommend "user interviews" as the validation method for every assumption — some assumptions require quantitative data, competitive analysis, or technical spikes
- ☐ Do not list assumptions that cannot be tested — every assumption in the map must have a plausible validation method, or it should be flagged as unknowable and treated as a risk

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/cheap-tests.md** — The Cheap-Test Catalog: Right-Sizing Validation. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/assumption-board.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ At least one assumption per category (Desirability, Feasibility, Viability, Usability)
- ☐ All Impact 4+ / Confidence 2– assumptions flagged as CRITICAL
- ☐ Each validation method is specific (not just "do research" — name the method and sample size)
- ☐ Priority scores are consistent (Impact – Confidence, higher = more urgent)

customer-journey-map

Build a customer journey map for a product, service, or experience.

Customer Journey Map Skill

This skill produces a complete customer journey map covering every stage from awareness through advocacy. Each stage includes touchpoints, customer actions, emotions, pain points, and specific improvement opportunities. Output is ready for use in product discovery, UX design, or cross-functional alignment workshops.

Required Inputs

Ask the user for these if not provided:

- **Product or service** being mapped
- **Customer persona** — which customer segment is this map for? (be specific — one persona per map)
- **Journey scope** — full end-to-end (awareness → advocacy), or a specific phase (e.g. onboarding only)?
- **Current state or future state?** — mapping how it works today, or designing how it should work?
- **Data sources** — any research, user interviews, support tickets, NPS comments, analytics available?

- **Goal of the map** — what decision will this inform? (redesign, prioritisation, stakeholder alignment, new feature)

Output Structure

Customer Journey Map: [Product / Service]

Persona: [Name — e.g. "Sarah, the overwhelmed HR manager"] **Journey scope:** [Full end-to-end / Onboarding / Purchase / Renewal] **Current or future state:** [Current state / Desired future state] **Prepared by:** [Name / Team] **Date:** [Date] **Based on:** [Research sources — interviews, analytics, support data, assumed/hypothetical]

Persona Summary

Name	[Sarah]
Role	[HR Manager at a 200-person professional services firm]
Goal	[Reduce time spent on manual employee data management]
Frustrations	[Too many tools that don't talk to each other; always chasing approvals]
Tech comfort	[Moderate — comfortable with SaaS tools but not a power user]
Decision power	[Recommends tools; budget approved by CHRO]

Journey Overview

AWARENESS → CONSIDERATION → DECISION → ONBOARDING → ADOPTION → ADVOCACY

[Stage 1] [Stage 2] [Stage 3] [Stage 4] [Stage 5]

[Stage 6]

Overall experience rating (current state): [🤬 Frustrating / 😐 Neutral / 😊 Positive]

Stage 1: Awareness

How does the customer first discover the product exists?

Customer goal at this stage: [e.g. Realise they have a problem worth solving — or find a solution to a specific pain]

Element	Detail
Trigger	[What event makes them start looking? — e.g. Manual process breaks down / peer recommendation / saw ad]
Where they are	[Google search / LinkedIn / conference / colleague conversation / email newsletter]
What they do	[e.g. Searches "automate employee onboarding" / asks peers in HR community / clicks LinkedIn ad]
Emotion	[😞 Frustrated — overwhelmed by manual processes and hoping for a better way]
Pain points	[Overwhelming number of options / hard to know which tools are credible / can't tell what's B2B vs B2C from homepage]
Opportunities	[SEO content targeting the trigger keyword / LinkedIn thought leadership / peer community presence]

Stage 2: Consideration

The customer is actively evaluating options. What do they do to decide?

Element	Detail
Customer goal	[Narrow down from many options to a shortlist of 2–3]
What they do	[Reads G2/Capterra reviews / watches demo video / downloads comparison guide / asks peers who use something similar]
Touchpoints	[Website / review sites / social proof / demo request flow / sales email]
Emotion	[😟 Anxious — worried about making the wrong choice; past tool purchases haven't delivered]
Pain points	[Pricing not visible on website / demo requires a call before seeing the product / unclear if it works with their existing stack]
Opportunities	[Self-serve demo or interactive product tour / transparent pricing page / ROI calculator / case studies from similar company size]

Stage 3: Decision

The customer is ready to buy — or not. What makes them commit?

Element	Detail
Customer goal	[Get sign-off from CHRO and justify the decision with a business case]

Element	Detail
What they do	[Books sales call / requests security questionnaire / builds internal business case / negotiates contract]
Touchpoints	[AE / sales call / security review / contract / procurement process]
Emotion	[😬 Cautious — doesn't want to be wrong; presenting to leadership adds pressure]
Pain points	[Sales process is slow / security questionnaire takes weeks / contract terms are non-standard and require legal]
Opportunities	[Security FAQ self-serve / standard contract with predictable terms / champion toolkit (slides, business case template) to help them sell internally]

Stage 4: Onboarding

The customer has bought. Now they need to get value fast.

Element	Detail
Customer goal	[Get the product working and show their CHRO it was a good decision]
What they do	[Receives welcome email / attends kickoff call / configures integrations / invites team]
Touchpoints	[Onboarding email sequence / in-product onboarding checklist / CSM / help centre / integrations marketplace]
Emotion	[😬 Anxious but hopeful — excited about potential but stressed about the setup work]
Pain points	[Setup is more complex than expected / IT required for SSO but IT is slow to respond / generic onboarding doesn't match their use case]
Opportunities	[Role-specific onboarding paths / IT connector with pre-filled request template / quick win email at day 3 (show them one thing that already works)]

Key moment of truth: [What single moment in this stage determines whether they'll become an active user or ghost? — e.g. "First time the product saves them 30 minutes on a task they used to do manually"]

Stage 5: Adoption

The customer is using the product. Are they getting consistent value?

Element	Detail
Customer goal	[Make the product a regular part of their workflow; demonstrate ROI to leadership]
What they do	[Uses core features daily / discovers new features / hits a limitation / contacts support / attends webinar]
Touchpoints	[Product UI / in-app notifications / email / support / community / customer success manager]
Emotion	[Variable — some days 😊 when the product works well; some days 😞 when hitting a gap or bug]
Pain points	[Feature they expected isn't there / reporting doesn't show the metric leadership wants / power features are too complex / feels like they're underutilising what they're paying for]
Opportunities	[Proactive CSM check-in at day 30 / in-product feature discovery / usage dashboard for the customer to see their own ROI / community for peer learning]

Adoption health indicators:

- [DAU/MAU ratio — what does healthy look like?]
- [Feature X used by Y% of seats within Z weeks]
- [First NPS survey at 60 days — target score]

Stage 6: Advocacy


The customer loves the product. How do you turn them into a referral engine?

Element	Detail
Customer goal	[Solve problems faster; feel like an expert; feel valued as a customer]
What they do	[Refers a peer / writes a G2 review / participates in case study / speaks at event / becomes a power user / joins community]
Touchpoints	[CSM / community / review request email / referral programme / case study outreach / conference sponsorship]
Emotion	[😊 Proud — the tool is part of their professional identity; they feel smart for choosing it]
Pain points	[Referral programme is clunky / no structured way to connect with peers / case study process is slow and effortful for them]
Opportunities	[One-click G2 review request at high-satisfaction moment / peer community / referral programme with meaningful reward / case study process that does]

Element	Detail
	most of the work for them]

Emotion Curve

Plot the customer's emotional experience across the journey:



Lowest point: [Which stage has the worst experience — and why?] **Highest point:** [When is the customer most delighted — what drove it?] **Biggest drop:** [Where does sentiment fall most sharply — this is usually the biggest opportunity]

Prioritised Opportunities

Opportunity	Stage	Impact on customer	Effort to fix	Priority
[Self-serve product tour before sales call]	Consideration	[High — removes top buying barrier]	[Medium]	P1
[Quick win email at day 3]	Onboarding	[High — builds early habit]	[Low]	P1
[IT SSO setup template]	Onboarding	[Medium — removes specific blocker]	[Low]	P2
[30-day proactive CSM check-in]	Adoption	[Medium — catches churn signals early]	[Medium]	P2
[Peer referral programme]	Advocacy	[High for growth — reduces CAC]	[High]	P3

What We Don't Know (Research Gaps)

Gap	How to close it	Priority
[What actually triggers the decision to start looking?]	[5 JTBD interviews with recent buyers]	[High]

Gap	How to close it	Priority
[What causes customers to stall in onboarding?]	[Drop-off analysis in onboarding funnel + 3 interviews with churned customers]	[High]
[What % of customers have reached the advocacy stage?]	[Product analytics — identify power users; NPS by cohort]	[Medium]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/evidence-based-mapping.md** — Journey Maps Built on Evidence (Not Conference-Room Fiction). Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/journey-canvas.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Map covers one specific persona — not "all customers"
- ☐ Each stage includes the customer's emotional state — not just actions
- ☐ Pain points are the customer's pain — not the company's pain
- ☐ Opportunities are specific enough to become backlog items or design prompts
- ☐ Emotion curve shows the real experience — not an aspirationally positive version
- ☐ Research gaps are documented — the map reflects what is known, not assumed

Anti-Patterns

- ☐ Do not build the map from assumptions alone — ground at least the pain points in real customer data or research
- ☐ Do not treat all journey stages as equally weighted — identify the highest-friction moments explicitly
- ☐ Do not omit the emotional layer — a journey map without emotions is a process flow, not a customer map
- ☐ Do not create generic touchpoints that apply to any product — each touchpoint must be specific to this product and customer
- ☐ Do not leave opportunities unranked — prioritise by impact and feasibility

Example Trigger Phrases

- "Map the customer journey for [product]"
- "Build a user journey from awareness to advocacy"

- "Create a journey map for our onboarding experience"
- "Map out the touchpoints and pain points for [customer type]"
- "Design an experience map for [process or product]"

discovery-interview-guide

Create a structured user discovery interview guide with screener questions, a discussion guide, and a synthesis framework.

Discovery Interview Guide Skill

Design interviews that surface genuine insight — not validation of what you already believe. Every guide follows a story-based, past-behaviour-focused structure.

Core Principles

1. **Never ask about the future.** "Would you use X?" tells you nothing. "Tell me about the last time you did X" tells you everything.
2. **Interview for behaviour, not opinion.** Opinions are cheap. Behaviour is evidence.
3. **The 5 Whys.** Every surface answer is a door. Keep opening doors.
4. **Confirm the problem before exploring the solution.** Never show a prototype until you've confirmed the pain exists unprompted.

Interview Structure (60 minutes standard)

1. Warm-Up (5 min)

Build rapport. Get them talking. Don't discuss the topic yet.

- "Tell me a bit about your role and what a typical week looks like for you."
- "What tools do you rely on most day-to-day?"

2. Context Setting (10 min)

Understand their world before diving into the problem space.

- "Walk me through how you currently [handle the domain area]."
- "What does that process look like from start to finish?"
- "Who else is involved when you do this?"

3. Problem Exploration (25 min) — THE CORE

Surface pain without leading.

- "Tell me about the last time you had to [relevant task]. What happened?"

- "What was the hardest part of that?"
- "How did you handle it?"
- "What did you try before settling on that approach?"
- "What does it cost you when this goes wrong?" (time, money, stress, reputation)
- "If you could wave a magic wand and change one thing about this process, what would it be?"

⚠ **Do not mention your product or feature during this phase.**

4. Current Solutions (10 min)

Understand the competitive landscape from their perspective.

- "What tools or workarounds do you use today for this?"
- "What do you like about [current solution]? What frustrates you?"
- "Have you tried other approaches? What happened?"

5. Wrap-Up (10 min)

- "Is there anything about this topic we haven't covered that you think I should know?"
- "Is there anyone else you'd recommend I speak to?"
- "Would you be open to a follow-up if I have more questions?"

Output Format

Discovery Interview Guide — [Topic] — [Date]

Research Goal: [One sentence: what decision will this research inform?] **Target**

Participant Profile: [Role, company size, behaviour qualifier]

Screening Questions (for recruiting):

1. [Question] → Must answer: [Y/N or specific]
2. [Question] → Must answer: [Y/N or specific]
3. [Disqualifier question] → Disqualify if: [answer]

Interview Guide:

[Full structured guide using the format above, customised to the specific research topic]

Synthesis Template (fill after each interview):

- Key quote: "[verbatim]"
- Core pain: [1 sentence]
- Current workaround: [what they're doing today]
- Intensity (1–5): [how painful is this?]

- Surprise/unexpected finding: [anything that challenged your assumptions]

Pattern Detection (after 5+ interviews):

- Pain mentioned by [X/N] participants: [theme]
 - Workaround used by [X/N] participants: [theme]
 - Most emotionally charged moment in interviews: [observation]
-

Required Inputs

Ask the user for these if not provided:

- **Research topic or question** (what decision will this inform?)
- **Target participant profile** (role, behaviour, company type)
- **Session length** (30 / 45 / 60 / 90 minutes)
- **Number of interviews planned**
- **Known hypotheses to test or avoid confirming prematurely** (optional)

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/question-craft.md** — Question Craft: Getting Truth Instead of Politeness. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/guide-skeleton.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ No future-tense questions ("would you...") — only past-behaviour questions
- ☐ Product or solution not mentioned until after pain is confirmed
- ☐ Questions open-ended (cannot be answered yes/no)
- ☐ Synthesis template included for per-session notes
- ☐ Screener questions identify and disqualify wrong participants

Guidelines

- Recommend 5–8 interviews to reach thematic saturation for most discovery questions
- Always record with permission — transcripts beat notes
- If user is new to interviewing: remind them to stay silent after asking a question (aim for 80/20 participant-to-interviewer talking ratio)

- Never synthesise during the interview — do it after, when you can look across sessions
- Flag confirmation bias: if user writes questions that lead toward a predetermined answer, rewrite them as open-ended alternatives

Anti-Patterns

- ☐ Do not use future-tense questions ("Would you use this?") — hypothetical responses do not predict real behaviour and produce false confidence in an idea
- ☐ Do not mention your product or solution before problem exploration is complete — doing so anchors the participant's responses and invalidates the discovery
- ☐ Do not synthesise across fewer than 5 interviews — themes from 2–3 interviews reflect anecdote, not pattern; wait for saturation
- ☐ Do not write screener questions that are too easy to pass — if participants can guess the "right" answer, you will recruit the wrong people
- ☐ Do not treat participant opinions as evidence of future behaviour — what people say they will do consistently diverges from what they actually do

job-story-mapper

Write Jobs-to-be-Done (JTBD) job stories and map customer jobs across functional, social, and emotional dimensions.

Job Story Mapper Skill

Stop writing features. Start understanding jobs. This skill translates product requirements and user interviews into precise job stories that keep the team focused on outcomes — not outputs.

Jobs-to-be-Done Fundamentals

A "job" is the progress a customer is trying to make in a given situation. People don't buy products — they hire them to get a job done.

Three dimensions of every job:

- **Functional job:** The practical task ("get from A to B")
- **Emotional job:** How they want to feel ("feel confident I made the right choice")
- **Social job:** How they want to be perceived ("look like a competent professional to my team")

Great products address all three. Most roadmaps only address the functional one.

Job Story Format

Template:

When [situation/trigger], I want to [motivation/goal], so I can [expected outcome].

Not a user story: User stories focus on roles and features: "As a [role] I want [feature] so that [benefit]." Job stories focus on situations and motivations: "When [I'm in this specific situation] I want [this capability] so I can [achieve this outcome]."

The situation is the most important part. "When I'm in the middle of a sprint and my PM asks for an update" is a much richer trigger than "As a developer."

Mapping Process

Step 1: Identify the main job

One sentence: What is the core job your product is hired for?

"Help [user type] [accomplish outcome] when [context]."

Step 2: Break into job steps

What are all the sub-tasks within the main job? (Use a job map: Define → Locate → Prepare → Confirm → Execute → Monitor → Modify → Conclude)

Step 3: Identify pain points per step

Where does the job fall down today? Where do customers use workarounds?

Step 4: Write job stories for each pain point

One job story per distinct situation-motivation pair.

Step 5: Map to product opportunities

Which job stories are underserved? Which have existing solutions? Where is your differentiation?

Output Format

Job Story Map — [Product/Feature Area] — [Date]

Core Job Statement:

When [context], [user type] wants to [main job outcome], so they can [ultimate goal].

Job Map:

Step	Sub-Job	Current Solution	Pain Points	Underserved?
Define	[What user does]	[Tool/method used]	[Frustration]	H/M/L
Locate				
Prepare				
Confirm				
Execute				
Monitor				
Modify				
Conclude				

Job Stories (prioritised by underservice):

Job Story 1 — [Situation label]

When [specific situation], I want to [motivation], so I can [outcome].

Functional dimension: [What they need to get done] Emotional dimension: [How they want to feel] Social dimension: [How they want to be perceived]

Current workaround: [What they do today] Pain intensity: [High / Medium / Low]

Frequency: [How often this situation occurs] Product opportunity: [What we could build to address this]

Repeat for each major job story.

Opportunity Scoring: Rate each job story on:

- Importance to customer (1–10)
 - Satisfaction with current solution (1–10)
 - Opportunity score = Importance + max(Importance – Satisfaction, 0)
 - Prioritise: Opportunity score > 10
-

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/situation-mining.md** — Situation Mining — the "When" Is the Whole Method. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/job-story-canvas.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Job stories use the "When / I want to / So I can" format (not user story format)
- ☐ Situation is specific (not "as a user" — a real moment or trigger)
- ☐ All three dimensions covered: functional, emotional, social
- ☐ Opportunity score calculated for each job story
- ☐ Current workaround identified for each high-opportunity story
- ☐ Product opportunity is distinct from "build the feature" (it's an outcome)

Required Inputs

Ask the user for these if not provided:

- **Product or feature area** to map (e.g. onboarding, checkout, dashboard)
- **User type or persona** (who are we mapping jobs for?)
- **Source material** (user interview notes, support tickets, discovery findings, or describe from memory)
- **Scope** (full product job map vs. a single feature area)

Anti-Patterns

- ☐ Do not write job stories that describe a feature rather than a situation-motivation pair
- ☐ Do not skip the social and emotional dimensions — mapping only functional jobs misses the most defensible differentiation opportunities
- ☐ Do not define situations too broadly ("as a user who wants to manage their work") — the situation must be a specific moment or trigger
- ☐ Do not conflate opportunity scoring with priority — a high opportunity score still requires feasibility and strategic fit assessment
- ☐ Do not produce a job map without identifying current workarounds — the workaround reveals what the job is worth to the customer

Guidelines

- Never write a job story for a feature — write it for the situation that makes the feature valuable

- If you can't identify the situation, you don't understand the job yet — go back to user research
- Social and emotional jobs are harder to surface but often the most defensible differentiators
- Recommend sharing job stories with engineering — they make better technical decisions when they understand the "why"

user-interview-synthesis

Synthesises user interview transcripts into structured research findings.

User Interview Synthesis Skill

Transform raw interview transcripts into a structured synthesis document that surfaces themes, pain points, and actionable insights.

Required Inputs

Ask the user for these if not provided:

- **Interview transcripts or notes** (even rough notes work)
- **Number of participants and their profiles** (role, company size, context)
- **Research questions** (what was the study trying to answer?)
- **Date range** of research (for context)

Process

1. Read all provided transcripts fully before drawing conclusions
2. Identify recurring themes (minimum 3 mentions to qualify as a theme)
3. Categorize findings into: Pain Points, Workflow Insights, Feature Requests, Delight Moments
4. Select 2-3 verbatim quotes per theme that best represent the pattern
5. Draft "So What" implications for each theme — what does this mean for the product?
6. **Validate** — Confirm every theme has quotes from at least 3 participants. Flag any insight resting on fewer as low-confidence.

Output Structure

Research Synthesis: [Study Name]

Participants: [n] **Date Range:** [dates] **Research Questions:** [list]

Theme 1: [Theme Name]

- Summary (2-3 sentences)
- Supporting quotes (from at least 3 participants)
- Implication for product

[Repeat for each theme]

Low-Confidence Signals (1-2 participants only)

[Findings worth tracking but not acting on yet — note what further research would confirm or deny]

Recommended Next Steps

[Specific, actionable recommendations based on findings]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/coding-transcripts.md** — Coding Interview Transcripts Without Losing the Signal. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/per-session-capture.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every theme is supported by quotes from at least 3 participants
- ☐ Implications connect to specific product decisions, not just observations
- ☐ Researcher bias check: no leading language, findings don't all support one hypothesis
- ☐ Single-source signals are flagged separately, not mixed into main themes
- ☐ Research questions from the study brief are each addressed (even if the answer is "inconclusive")

Anti-Patterns

- ☐ Do not mix single-source signals into main themes — insights cited by only one participant must be flagged separately
- ☐ Do not write implications that are observations restated rather than product decisions enabled
- ☐ Do not include themes that only support the project hypothesis — contradictory findings must be surfaced, not omitted

- ☐ Do not present findings without quotes — every theme requires verbatim evidence from at least 3 participants
- ☐ Do not leave research questions unanswered — each question from the study brief must be explicitly addressed, even if the answer is inconclusive

Chapter 5 — Essentials

5 skills

competitive-analysis

Analyze competitors and create competitive landscape documentation with feature matrices, positioning maps, and strategic recommendations.

Competitive Analysis Skill

Create structured competitive analyses for product decision-making.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** `knowledge/` (market + positioning) and competitor `entities/` . Run `python3 ../professional-brain/scripts/brain_query.py ./brain "<competitor or market>"` and carry each fact's provenance tag through — a competitor claim from a press release is `[external]` , not `[data]` .
-  **Propose to the Brain:** after producing, propose recording new competitor facts to `knowledge/` (`[external]`) and creating/updating competitor `entities/` . Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Required Inputs

Ask the user for these if not provided:

- **Your product or company** (what you're comparing against)
- **Competitors to analyze** (or ask to identify the top 3-5)
- **Analysis focus** (full landscape / feature comparison / pricing / positioning / win-loss)
- **Audience** (product team / leadership / sales / board)

Process

1. Gather competitor information from provided inputs and available context
2. Build profiles for each competitor

- 3. Create feature comparison matrix on dimensions that matter to the user's customers
- 4. Analyze pricing and positioning
- 5. Identify win/loss patterns and strategic implications
- 6. **Validate** — Confirm all claims reference a specific source or are flagged as assumptions. Verify feature comparisons note quality differences, not just presence/absence.

Output Structure

1. Executive Summary

- **Market Position:** Where we stand relative to competitors
- **Key Findings:** Top 3-5 insights
- **Strategic Implications:** What this means for the roadmap

2. Competitor Profiles

For each competitor:

- **Company Overview:** Size, funding, market position
- **Target Customer:** Who they serve
- **Value Proposition:** Core positioning
- **Strengths / Weaknesses:** What they do well and where they fall short
- **Recent Activity:** Major updates, funding, announcements

3. Feature Comparison Matrix

Feature	Us	Competitor A	Competitor B	Competitor C
[Feature]	✔ Full	△ Limited	✖ None	✔ Full

Legend: ✔ Full (production-ready) · △ Limited/Beta · ✖ None

Include notes on quality and implementation differences where significant.

4. Pricing Comparison

Plan	Us	Competitor A	Competitor B
Free/Trial	[price]	[price]	[price]
Pro	[price]	[price]	[price]
Enterprise	[price]	[price]	[price]

5. Market Positioning Map

Position competitors on two key dimensions relevant to the market:

- Y-Axis: [e.g., Enterprise vs. SMB]
- X-Axis: [e.g., Simple vs. Comprehensive]

Whitespace Opportunities: [Underserved segments]

6. Win/Loss Analysis

Why We Win:

- Better at: [specific capabilities]
- Customers who value: [what matters to them]

Why We Lose:

- When customers need: [specific requirements]
- Their advantage: [what tips the decision]

7. Strategic Recommendations

Immediate Actions (0-3 months):

1. [Action] — [Rationale]

Medium-term (3-12 months):

1. [Action] — [Rationale]

Anti-Patterns

- ☐ Do not present competitor feature claims as facts without citing a source or flagging them as assumptions — outdated or incorrect feature data misleads sales and product decisions
- ☐ Do not build a competitive analysis that only covers features — pricing, messaging, go-to-market motion, and who they hire for are equally strategic signals
- ☐ Do not treat all buyers as identical — the same product may win against Competitor A in the enterprise segment and lose in SMB; segment-specific win/loss matters
- ☐ Do not soften weaknesses and threats in the SWOT to avoid internal discomfort — an honest SWOT is only useful if the negatives are real

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/feature-matrix-honesty.md** — Feature Matrices That Don't Lie. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.

- `templates/landscape-doc.md` — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ All competitor claims cite a source or are flagged as assumptions
- ☐ Feature comparison notes quality differences, not just feature presence
- ☐ Strategic recommendations are specific actions, not generic advice
- ☐ Win/loss analysis reflects customer perspective, not internal assumptions
- ☐ Different customer segments are considered (not all buyers value the same things)

meeting-notes

Structure and format meeting notes following PM best practices.

Meeting Notes Skill

This skill structures meeting notes to maximize value and ensure follow-through.

Required Inputs

Ask the user for these if not provided:

- **Meeting title and date**
- **Attendees** (names and roles)
- **Raw notes or transcript** (paste discussion notes, a transcript, or describe what was discussed)
- **Meeting type** (1:1 / sprint planning / product review / stakeholder sync / other) — determines which template to use

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, this is where notes become durable memory:

- **Read first:** the relevant `stakeholders/` files (so you arrive knowing each attendee's open asks and concerns) and any `decisions/` the meeting revisits.
- **Write after:** append each **decision** (with its rationale and a `reopen-when`) to `decisions/` , add new **asks/concerns** to the right `stakeholders/` file, and flag any new **assumption** into `hypotheses/` . Tag every captured fact with its provenance — most meeting statements are `[verbal]` until independently confirmed. Save the raw notes to `source/` .

Standard Meeting Notes Template

Meeting Header

Meeting: [Meeting Title]

Date: [Date]

Attendees: [Names/Roles]

Note Taker: [Name]

Duration: [Actual duration]

Agenda

- ☐ Topic 1
- ☐ Topic 2
- ☐ Topic 3

(Check off items as discussed)

Decisions Made

Clear documentation of decisions:

Decision: [What was decided]

Context: [Why this decision]

Owner: [Who's responsible for executing]

Deadline: [When if applicable]

Use this format for each decision made.

Action Items

All action items should be:

- ☐ **[Action item]** - @Owner - Due: [Date]
- ☐ **[Action item]** - @Owner - Due: [Date]

Format:

- Clear, specific action
- Single owner (no "team" ownership)
- Concrete deadline
- Checkbox for tracking

Discussion Notes

Key points discussed organized by topic:

Topic 1: [Name]

- Key point or discussion highlight
- Important context or concern raised
- Any data or information shared

Topic 2: [Name]

- Key discussion points
- Decisions or conclusions reached

Open Questions / Follow-Up

Questions that couldn't be answered:

- **Question:** [What we need to know]
- **Owner:** [Who will find out]
- **By When:** [Deadline]

Next Steps

Clear summary of what happens next:

1. [Immediate next action]
2. [Follow-up meeting if needed]
3. [Any broader process to start]

Best Practices

During the meeting:

- Focus on decisions and action items over dialogue
- Capture specific commitments, not general discussion
- Note dissenting opinions on important decisions
- Ask for clarity on vague commitments ("I'll look into it" → "I'll analyze the data and share findings by Friday")

After the meeting:

- Send notes within 2 hours while fresh
- Tag action item owners (@mention them)
- Include links to relevant documents
- Follow up on overdue action items

What to capture: ✅ Decisions made ✅ Action items with owners and deadlines ✅ Key points of discussion ✅ Open questions ✅ Next steps

What to skip: ❌ Verbatim transcripts ❌ Off-topic tangents ❌ Preliminary discussion before decisions ❌ Redundant information

Meeting Types & Adaptations

1:1 Meetings

Focus on:

- Career development discussions
- Feedback (both directions)
- Current challenges
- Action items for both parties

Template additions:

- **Recent Wins:** What's going well
- **Challenges:** What's not going well
- **Career Discussion:** Development topics
- **Feedback:** For both parties

Sprint Planning

Focus on:

- Story acceptance criteria
- Sizing/estimation decisions
- Dependency identification
- Sprint commitment

Template additions:

- **Sprint Goal:** What we're committing to
- **Story Points:** Capacity and estimates
- **Dependencies:** External blockers
- **Definition of Done:** Acceptance criteria

Product Reviews

Focus on:

- Design decisions
- User feedback discussed
- Changes requested
- Launch readiness assessment

Template additions:

- **Design Decisions:** What was approved/rejected
- **User Feedback:** Key insights discussed

- **Open Design Questions:** What needs iteration
- **Launch Criteria:** Remaining requirements

Stakeholder Sync

Focus on:

- Status updates delivered
- Concerns raised
- Approvals given
- Escalation needs

Template additions:

- **Status Overview:** High-level progress
- **Approvals Obtained:** Sign-offs received
- **Escalations:** Issues raised to stakeholders
- **Next Sync:** When and what to cover

Example Meeting Notes

Product Roadmap Review - Q1 2026

Date: January 20, 2026

Attendees: Sarah (CPO), Mike (Eng Lead), Jennifer (Design), Tom (PM)

Note Taker: Tom

Duration: 45 minutes

Agenda

- [x] Review Q1 planned features
- [x] Discuss resource constraints
- [x] Prioritization discussion
- [x] Timeline alignment

Decisions Made

Decision: Move multi-channel dashboard to Q2, prioritize mobile app improvements for Q1

Context: Customer feedback shows mobile experience is significantly impacting retention (65% of users primarily mobile). Engineering team can only tackle one major initiative this quarter.

Owner: Tom (PM) to communicate to stakeholders

Deadline: January 22

Decision: Allocate 20% of engineering time to technical debt

Context: Accumulated tech debt is slowing feature development. Team velocity dropped 30% last quarter.

Owner: Mike (Eng Lead) to create tech debt backlog

****Deadline**:** January 27

****Decision**:** Run mobile beta with 100 users before full launch

****Context**:** Need to validate improvements on diverse devices

****Owner**:** Jennifer (Design) to coordinate with QA

****Deadline**:** February 10

Action Items

- [] ****Update Q1 roadmap deck with new prioritization**** - @Tom - Due: Jan 22
- [] ****Schedule alignment meeting with support team about dashboard delay**** - @Tom - Due: Jan 24
- [] ****Create tech debt prioritization rubric**** - @Mike - Due: Jan 27
- [] ****Run user testing on mobile designs**** - @Jennifer - Due: Feb 3
- [] ****Document decision rationale for executives**** - @Sarah - Due: Jan 23
- [] ****Identify 100 beta users for mobile**** - @Tom - Due: Feb 1

Discussion Notes

****Q1 Feature Prioritization****

- Customer retention is #1 company priority this quarter
- Mobile app NPS score is 6.2 (vs 8.1 for web)
- Mobile accounts for 65% of daily active users
- Multi-channel dashboard would take 8 engineering weeks
- Mobile improvements estimated at 6 engineering weeks with higher ROI
- Sales has 3 enterprise deals waiting on dashboard feature

****Resource Constraints****

- Currently 4 engineers available (down from 6 last quarter due to attrition)
- Design team can support both initiatives but at reduced capacity
- QA team needs 2 weeks for thorough testing on mobile
- One engineer on loan to security team through February

****Risk Discussion****

- Delaying dashboard may impact enterprise sales (3 deals waiting)
- Sarah noted: "We can position mobile improvements as foundation for enterprise features"
- Mike raised concern about mobile tech stack stability - addressed through tech debt allocation
- Need to communicate clearly with Sales about timeline change

****Mobile Implementation Plan****

- Week 1-2: Design refinements based on user feedback
- Week 3-4: Engineering implementation
- Week 5: Internal testing
- Week 6: Beta with 100 users
- Week 7: Full rollout

Open Questions

- **Question**: What's the impact on enterprise pipeline if we delay dashboard?

Owner: Sarah will check with Sales leadership

By When: January 23

- **Question**: Can we do a limited beta of dashboard for enterprise customers?

Owner: Tom will explore MVP scope with Mike

By When: January 25

- **Question**: What's our plan if mobile improvements don't hit target metrics?

Owner: Tom will create contingency plan

By When: January 27

Next Steps

1. Tom to send updated roadmap to leadership by EOD Wednesday (Jan 22)
2. Team to begin sprint planning for mobile improvements next Monday (Jan 27)
3. Follow-up meeting on Feb 1 to review progress and validate prioritization
4. Sarah to present decision rationale to executive team on Jan 24

Next Meeting: February 1, 2026 - Progress Check-in

Notes Sent: January 20, 2026 5:30 PM

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/decisions-vs-discussion.md** — Separating Decisions from Discussion. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/notes-skeleton.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every action item has a single named owner (not "team")
- ☐ Every action item has a concrete deadline
- ☐ Decisions include context (why the decision was made)
- ☐ Open questions have an owner and a "by when"
- ☐ No verbatim transcripts — synthesis only

Anti-Patterns

- ☐ Do not assign action items to "the team" or "everyone" — every action item must have exactly one named owner or it will not be completed
- ☐ Do not capture verbatim transcript content — meeting notes record decisions and commitments, not the full conversational path to get there
- ☐ Do not omit the context for decisions — a decision without its rationale is useless when someone asks "why did we do that?" six months later
- ☐ Do not leave open questions without an owner and deadline — an unanswered question with no follow-up assigned is a blocked decision
- ☐ Do not delay sending notes beyond 2 hours after the meeting — notes sent the next day miss the window when action item owners can act on commitments while fresh

Notes Distribution

Subject Line Format: "[Meeting Type] Notes - [Date] - [Key Topic]"

Example: "Product Roadmap Review Notes - Jan 20 - Q1 Prioritization"

Recipients:

- All attendees
- Anyone mentioned in action items
- Anyone who requested notes

Follow-Up:

- Send reminder 3 days before action item due dates
- Weekly summary of all open action items
- Mark action items as complete and share updates

Execution

For tool-using agents with connected MCP servers (Notion, Linear/Jira, Slack). Runtimes without tool access ignore this section and deliver the document. See SKILLSPEC.md §5 and connectors/mcp-pairings.md.

Preconditions

- The structured notes above have been shown to the human and **explicitly approved**, including the destination (which Notion database/page, which tracker project).
- The MCP servers are already connected and authenticated in the agent's environment.

- Action items each have a named owner — unowned items are resolved with the human first, never assigned by guess.

Allowed actions

- Create ONE page in the approved Notion database (or equivalent docs tool) containing the approved notes, verbatim.
- Create one tracker issue per approved action item (title, owner, due date from the notes) in the approved project.
- Post the page link (only the link and a one-line summary) to the approved channel, if the human named one.
- Nothing else: no editing existing pages/issues, no inviting or notifying people beyond the named channel, no calendar writes.

Verification

- Fetch the created page and each created issue; confirm titles, owners, and dates match the approved notes.
- Report every created URL back to the human in one list.

Rollback

- Undo = archive/delete the just-created page and issues, only on explicit human instruction.
- Stop and ask a human if: the destination database/project is not found, any issue creation fails partway (report what WAS created), or an action-item owner does not exist in the tracker.

prd-template

Create a Product Requirements Document following proven PM template structure.

PRD Template Skill

This skill helps create professional Product Requirements Documents following industry best practices.

Required Inputs

Ask the user for these if not provided:

- **Feature or product name**
- **Problem being solved** (from the user's perspective)
- **Target user** (role, context, what they're trying to accomplish)

- **Success metrics** (how will you know it worked?)
- **Scope** (MVP vs full vision — what's in and out of scope)
- **Key stakeholders** (who needs to review and approve)

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it instead of asking for context you already have:

- **Read first:** `context.md` (product, metrics definitions, voice), `knowledge/strategy.md` (where the product is going), any related `hypotheses/` and the matching `entities/` feature file. Run `python3 ../professional-brain/scripts/brain_query.py ./brain "<feature>"` to pull grounded facts, and carry their provenance tags into the PRD (don't present a `[hunch]` as a settled requirement).
- **Write after:** save the feature as `into entities/<feature>.md` , log any scoping decision to `decisions/` , and add new assumptions to `hypotheses/` . Tag each with its provenance.

Deeper Materials

This skill ships with two support files — use them when they're available:

- `templates/prd-skeleton.md` — a fill-in PRD skeleton with a "what good looks like" hint per section. Start from it when the user wants a document to complete themselves rather than a generated draft.
- `references/success-metrics-guide.md` — calibration for the Success Metrics section: the four-part metric test, the standard adoption/outcome/business/guardrail set, and the common traps. Consult it whenever writing or reviewing the metrics table.

Template Structure

Every PRD should include these sections in order:

1. Overview

- **Problem Statement:** What problem are we solving? (2-3 sentences)
- **Proposed Solution:** High-level description of what we're building (2-3 sentences)
- **Success Metrics:** How we'll measure success (3-5 key metrics)

2. Context & Background

- **Why Now:** Why is this the right time?
- **Strategic Alignment:** How does this align with company objectives?

- **User Research Summary:** Key insights from research (if applicable)

3. User Stories & Use Cases

Format: "As a [user type], I want to [action] so that [benefit]"

- Include 3-7 primary user stories
- Add acceptance criteria for each

4. Requirements

Functional Requirements:

- Must-have features (P0)
- Should-have features (P1)
- Nice-to-have features (P2)

Non-Functional Requirements:

- Performance expectations
- Security considerations
- Accessibility requirements

5. Design & User Experience

- Link to design mocks or wireframes
- Key user flows
- Edge cases and error states

6. Technical Considerations

- Architecture implications
- Dependencies on other systems
- Technical risks and mitigations

7. Implementation Plan

- **Phase 1 (MVP):** What goes in first version
- **Phase 2:** What comes next
- **Phase 3:** Future enhancements

8. Open Questions

- Decisions that still need to be made
- Stakeholders to consult
- Research needed

9. Appendix

- Research links
- Related documents
- Competitive analysis

Writing Guidelines

Tone: Clear, concise, actionable **Audience:** Engineers, designers, stakeholders **Length:** Aim for 3-6 pages for features, 8-12 for products

Best Practices:

- Use concrete examples over abstractions
- Include "why" not just "what"
- Make requirements testable
- Link to supporting materials
- Update as decisions are made

What Makes a Good PRD

✅ Do:

- Write from the user's perspective
- Include specific success metrics
- Address edge cases
- Link to research and data
- Make trade-offs explicit

❌ Don't:

- Write implementation details (that's tech spec)
- Assume everyone has context
- Leave requirements ambiguous
- Skip the "why"
- Forget about accessibility

Quality Checks

- ☐ Problem statement is written from the user's perspective (not the company's)
- ☐ Success metrics are specific and measurable
- ☐ User stories include acceptance criteria
- ☐ Requirements are testable (not vague)
- ☐ Open questions are listed explicitly
- ☐ Implementation plan distinguishes MVP from future phases

Anti-Patterns

- ☐ Do not write requirements from the company's perspective — every requirement must trace back to a user need
- ☐ Do not include vague requirements like "the system should be fast" — every requirement must be testable
- ☐ Do not conflate MVP with future phases — be explicit about what is and is not in scope for the first release
- ☐ Do not leave success metrics as percentages without baselines — specify the current state and the target
- ☐ Do not skip open questions — unresolved assumptions are risks; surfacing them is the PM's job

Example PRD Opening

PRD: Multi-Channel Customer Support Dashboard

Overview

****Problem Statement**:** Support teams are currently managing customer inquiries across email, chat, and social media using three separate tools, leading to delayed responses, duplicated work, and inconsistent customer experiences. On average, support agents waste 2.3 hours per day switching between tools and manually tracking conversation history.

****Proposed Solution**:** Build a unified dashboard that aggregates customer inquiries from all channels into a single interface, maintains conversation history across channels, and provides intelligent routing based on agent expertise and availability.

****Success Metrics**:**

- Reduce average response time from 4 hours to 1 hour
- Decrease tool-switching time by 80% (from 2.3 to <0.5 hours)
- Improve customer satisfaction score from 3.8 to 4.5 (out of 5)
- Increase support agent productivity by 35%

Context & Background

****Why Now**:** Customer satisfaction has declined 15% over the past 6 months, primarily due to slow response times. Our top competitor launched a unified support dashboard last quarter, and we're hearing about it in sales calls. Support team turnover is at 45% annually, with "tool complexity" cited as a top frustration.

****Strategic Alignment**:** This aligns with our Q1 company objective to "Improve customer retention by 10%" and our support team's OKR to "Reduce average handle time by 25%."

****User Research Summary****: We conducted interviews with 12 support agents and observed 20 hours of support sessions. Key findings:

- Agents spend 35% of their time finding context from previous interactions
- 65% of escalations are due to lack of conversation history
- Agents rated tool-switching as their #1 daily frustration (9.2/10 pain)
- Current NPS for support experience is -12

User Stories & Use Cases

****US1: Unified Inbox****

As a support agent, I want to see all customer inquiries in one place so that I don't miss urgent requests and can prioritize effectively.

Acceptance Criteria:

- Inbox shows inquiries from email, chat, and social media
- Inquiries are sorted by priority (urgent, high, normal, low)
- Agent can filter by channel, customer, or status
- Real-time updates when new inquiries arrive

****US2: Cross-Channel Context****

As a support agent, I want to see the full conversation history regardless of channel so that I can provide consistent, informed responses without asking customers to repeat themselves.

Acceptance Criteria:

- Timeline view shows all interactions chronologically
- Each interaction displays channel, timestamp, and content
- Customer profile shows demographics and account information
- Previous issues and resolutions are accessible

[Continue with 5-7 total user stories...]

stakeholder-update

Create concise executive stakeholder updates using the BLUF (Bottom Line Up Front) framework.

Stakeholder Update Skill

This skill creates effective status updates for executives and stakeholders following the BLUF (Bottom Line Up Front) principle.

Required Inputs

Ask the user for these if not provided:

- **Project or product being reported on**
- **Audience** (CEO, board, cross-functional leads, investors — changes depth and format)
- **Period** (this week / this sprint / this month)
- **Current status** (on track / at risk / blocked)
- **Key metrics** and their current values vs. targets

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it before asking:

- **Read first:** the relevant `stakeholders/` files (what each person cares about and their prior asks), `context.md` (voice/tone), and recent `decisions/` for what's changed since the last update.
- **Write after:** append any new ask, concern, or commitment surfaced to the relevant `stakeholders/` file, provenance-tagged (`[verbal]` for something said in a meeting, not yet documented).

Deeper Materials

- `references/status-honesty-guide.md` — calibration for the 🟢/🟡/🔴 call (the watermelon problem, the consecutive-🟡 rule, re-baselining honestly) and fact → impact → action → ask phrasing for bad news. Apply it whenever the status is 🟡/🔴 or the input notes feel rosier than the metrics.
- `templates/update-skeleton.md` — a one-page fill-in update with the quality gates inline and a pre-send checklist. Offer it to users who want to write updates themselves.

Update Structure

1. BLUF (Bottom Line Up Front)

Start with the most important information:

- **Status:** 🟢 On track / 🟡 At risk / 🔴 Blocked / ✅ Complete
- **Key Takeaway:** One sentence summary of current state
- **Action Needed:** What you need from stakeholders (if anything)

2. Progress Summary

Brief overview of accomplishments:

- What shipped this period
- Milestones achieved
- Key metrics movement

Keep to 3-5 bullet points maximum.

3. Metrics Dashboard

Key Metrics

Metric	Current	Target	Trend	Status
[Metric name]	[Value]	[Target]	↑/→/↓	  

Include 3-5 most important metrics only.

4. Risks & Blockers

High Priority Issues:

- **Issue:** Brief description
- **Impact:** What's at stake
- **Mitigation:** What you're doing about it
- **Help Needed:** What stakeholders can do (if applicable)

Only include issues that matter at executive level.

5. Upcoming Milestones

Next 30 Days:

- Milestone (expected date)
- Milestone (expected date)

Next 90 Days:

- Major milestone (month)
- Major milestone (month)

6. Decisions Needed (if applicable)

- **Decision:** Clear description
- **Options:** 2-3 options with pros/cons
- **Recommendation:** What you recommend and why
- **Timeline:** When decision is needed

Writing Guidelines

Tone: Professional, concise, action-oriented **Length:** Keep under 1 page (or 2 minutes reading time) **Frequency:** Weekly for active projects, bi-weekly for maintenance

Executive Communication Principles:

1. Lead with conclusions, not process

- ❌ "We ran 5 experiments this week and analyzed the data..."
- ✅ "Conversion rate increased 15% from optimization work"

2. Focus on impact, not activities

- ❌ "Held 12 customer interviews"
- ✅ "Identified #1 barrier to adoption (complexity of setup)"

3. Make problems visible early

- Don't sugarcoat risks
- Propose solutions, not just problems
- Be specific about help needed

4. Use data to tell story

- Quantify whenever possible
- Show trends, not just snapshots
- Connect metrics to business outcomes

5. Make it scannable

- Use headers and bullet points
- Bold key information
- Use visual indicators (🟢🟡🔴, ↑→↓)

Status Guidelines

🟢 **On Track:** Meeting all targets, no significant risks 🟡 **At Risk:** Potential issues that could impact delivery 🔴 **Blocked:** Critical issues preventing progress, needs intervention

Example Update

Product Update: Customer Onboarding Redesign

Week of Jan 20, 2026

BLUF

Status: 🟡 At Risk

Key Takeaway: New onboarding flow is performing well in tests (+35% completion), but launch delayed one week due to integration issues with billing system.

Action Needed: Decision needed on whether to launch onboarding separately or wait for billing integration fix.

Progress Summary

- Completed user testing with 24 participants (94% positive feedback)
- Implemented first-time user experience improvements
- Resolved 12 of 15 bugs identified in QA
- Engineering allocated resources to billing integration fix

Key Metrics

Metric	Current	Target	Trend	Status
Onboarding Completion	45%	60%	→	🟡
Time to First Value	4.2 min	3.0 min	↓	🟢
Setup Support Tickets	45/week	<30/week	↓	🟢
User Activation Rate	52%	65%	→	🟡

Risks & Blockers

HIGH: Billing System Integration Delay

- ****Impact****: Prevents users from completing onboarding flow; delays launch by 1-2 weeks
- ****Root Cause****: API deprecation by payment processor, requires code rewrite
- ****Mitigation****: Engineering team reallocated resources, fix ETA Feb 3
- ****Decision Needed****: Launch onboarding without payment integration or wait for fix? (See below)

MEDIUM: Mobile Testing Coverage

- ****Impact****: Some edge cases on older Android devices not tested
- ****Mitigation****: Partnering with QA to expand test matrix; running beta with internal users on diverse devices

Upcoming Milestones

Next 30 Days:

- Resolve billing integration (Feb 3)
- Launch onboarding redesign (Feb 5 or Feb 12 depending on decision)
- Begin measuring impact on conversion (Feb 12)

Next 90 Days:

- Iterate based on production data (March)
- Extend to mobile app (April)
- Launch advanced features (May)

Decision Needed

Should we launch onboarding separately from billing integration?

Option A: Launch Now (Recommended)

- Pros: Get 35% completion rate improvement to users immediately, gather production data, maintain momentum
- Cons: Users need to complete payment in old flow, slightly disjointed

experience

- Timeline: Launch Feb 5

****Option B: Wait for Billing Fix****

- Pros: Fully integrated experience from day one, no technical debt

- Cons: Delays benefits by 2 weeks, Q1 metric targets at risk, team momentum lost

- Timeline: Launch Feb 12

****Recommendation****: Option A. The onboarding improvements are valuable independently, and the old payment flow works fine. Waiting risks missing Q1 targets and delays validated improvements from reaching users.

****Timeline****: Need decision by Jan 22 for Feb 5 launch.

****Questions?**** Reply to this email or ping me on Slack.

Frequency Guidance

Daily standups:

- Ultra-brief (3 bullets)
- What shipped yesterday
- What's shipping today
- Blockers

Weekly updates:

- Use full template above
- Focus on progress and risks
- Keep to 1 page

Monthly reviews:

- Deeper metrics analysis
- Strategic reflections
- Quarterly goal progress
- Longer format (2-3 pages) acceptable

Quarterly business reviews:

- Comprehensive analysis
- Trends over time
- Strategic recommendations

- Presentation format

Adaptation by Audience

For C-Suite

- Lead with business impact
- Connect to company OKRs
- Focus on strategy and outcomes
- Minimize technical details

For Product/Engineering Leadership

- Include technical context
- Show sprint/milestone progress
- Discuss architecture implications
- Reference technical debt

For Cross-Functional Teams

- Balance technical and business context
- Highlight dependencies
- Call out collaboration needs
- Make asks explicit

For Board/Investors

- Focus on metrics and traction
- Competitive positioning
- Market opportunities
- Financial implications

Quality Checks

- ☐ Update leads with BLUF — status, key takeaway, and action needed before any detail
- ☐ Every metric has a target comparison (not just a raw number)
- ☐ Every risk has a mitigation and a "help needed" flag if stakeholder action is required
- ☐ Decisions needed have specific options and a clear recommendation
- ☐ Total length is under 1 page / 2 minutes reading time

Anti-Patterns

- ☐ Do not bury the status assessment at the bottom — BLUF means the most important information comes first
- ☐ Do not report metrics without a target or prior-period comparison — raw numbers without context are not useful
- ☐ Do not list risks without mitigation actions and clear flags for stakeholder help needed
- ☐ Do not write decisions needed as questions without providing a clear recommendation — executives need options, not open-ended questions
- ☐ Do not allow the update to exceed one page — if it requires more, the message needs editing, not expanding

Execution

For tool-using agents that can reach the team's communication channels (Slack, email). Sending an update is **outward-facing**: it is never automatic. Runtimes without tool access ignore this section. See SKILLSPEC.md §5.

Preconditions

- The final update text has been shown to the human **verbatim** and explicitly approved — including the exact channel/recipient list.
- The channel or recipient list is named by the user, not inferred from history.
- If the status is  or contains a Decision Needed, confirm the named decision-maker is among the recipients.

Allowed actions

- Post the approved text, unmodified, to the one approved channel — or send it as one email to the approved recipients with the approved subject line.
- Save a copy to the location the user names (doc, Brain, repo file).
- Nothing else: no scheduling recurring sends (see `schedule-recipe` for that, with its own gates), no @-mentions not present in the approved text, no cross-posting.

Verification

- Confirm the message exists in the channel/thread (fetch its permalink) and report the link back.
- Confirm the sent text is byte-identical to the approved text.

Rollback

- If the platform allows it, deletion of a just-posted message is permitted **only** on explicit human instruction — otherwise post a correction reply.
- Stop and ask a human if: the channel is not found, posting partially fails, or the approved text no longer matches what is about to be sent.

user-research-synthesis

Analyze and synthesize user research findings into structured, actionable insights.

User Research Synthesis Skill

This skill helps analyze user research data and transform it into actionable insights following a structured methodology.

Required Inputs

Ask the user for these if not provided:

- **Research data** (transcripts, notes, survey results, or summary bullets)
- **Research method** (interviews, surveys, usability tests, etc.)
- **Number of participants** and their profiles (role, context)
- **Research questions** the study aimed to answer

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it before asking:

- **Read first:** open `hypotheses/` (which assumptions this research can validate or invalidate) and `context.md` (who the users are).
- **Write after:** update each touched hypothesis's status, add durable insights to `knowledge/users.md` , and keep the raw notes in `source/` . Tag interview-derived claims `[interview]` — never launder them into `[data]` .

Synthesis Framework

1. Data Collection Overview

- **Research Type:** Interviews, surveys, usability tests, etc.
- **Participant Profile:** Demographics, segments, sample size
- **Research Questions:** What we sought to learn
- **Methodology:** How data was collected

2. Key Themes Identification

Organize findings into themes using this structure:

Theme Name

- **Description:** What this theme represents
- **Prevalence:** How many participants mentioned this (e.g., "8 out of 12 participants")

- **Supporting Quotes:** 2-3 representative quotes
- **Implication:** What this means for our product

Aim for 4-8 major themes per research effort.

3. Pain Points Analysis

For each identified pain point:

- **Pain Point:** Clear description
- **Severity:** High/Medium/Low (based on impact and frequency)
- **Current Workaround:** How users deal with it today
- **Evidence:** Specific examples from research

4. Feature Requests

Categorize requests:

- **Must-Have:** Critical needs blocking user success
- **High Value:** Would significantly improve experience
- **Nice-to-Have:** Incremental improvements

For each request:

- **Request:** What users asked for
- **Frequency:** How often it came up
- **User Quote:** Representative example
- **Underlying Need:** Why they want this (dig deeper than surface request)

5. User Workflow Insights

Document actual workflows observed:

- **Current State:** How users accomplish tasks today
- **Pain Points:** Where they struggle
- **Ideal State:** What they wish they could do
- **Opportunities:** Where we can add value

6. Segmentation Insights

If research reveals distinct user segments:

- **Segment Name:** Descriptive label
- **Characteristics:** What defines this segment
- **Unique Needs:** How their needs differ
- **Size/Importance:** Relative weight for prioritization

7. Competitive Insights

If users mentioned competitors or alternatives:

- **Competitor/Alternative:** What they use
- **Why They Use It:** What it does well
- **Gaps:** What it doesn't do
- **Switching Barriers:** Why they don't switch fully

8. Recommendations

Prioritized recommendations based on insights:

High Priority

- Recommendation with supporting evidence
- Expected impact

Medium Priority

- Recommendation with supporting evidence
- Expected impact

Low Priority / Future Consideration

- Recommendation with supporting evidence
- Expected impact

9. Open Questions

Research gaps identified:

- What we still need to understand
- Suggested follow-up research
- Uncertainties requiring validation

Analysis Guidelines

When synthesizing interviews:

- Look for patterns across multiple participants
- Note both what users say AND what they do
- Pay attention to emotional reactions
- Identify jobs-to-be-done, not just feature requests

When analyzing quotes:

- Use verbatim quotes in "quotation marks"

- Attribute quotes: [Participant ID, Role, Context]
- Select quotes that illustrate patterns, not outliers
- Include both positive and negative feedback

When identifying themes:

- Use descriptive names, not generic labels
- Provide evidence for each theme
- Quantify when possible ("7 out of 10 users...")
- Connect themes to business objectives

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/theme-validity.md** — When Is a Theme Real? Synthesis Validity Rules. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/synthesis-report.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Themes identify patterns across multiple participants, not individual responses
- ☐ Insights connect to specific product decisions, not just observations
- ☐ Each claim includes supporting evidence (quotes, counts, or examples)
- ☐ Observations and interpretations are clearly separated
- ☐ Findings are prioritised by impact, not just listed

Anti-Patterns

- ☐ Do not list every individual comment — synthesis must identify patterns across participants
- ☐ Do not make interpretive leaps without supporting evidence from the data
- ☐ Do not focus on feature requests before understanding the underlying problem — always identify the job-to-be-done first
- ☐ Do not ignore contradictory data — conflicting findings must be surfaced and noted
- ☐ Do not present results without quantifying prevalence — state how many participants held each view

Example Theme

****Theme: Information Overload During Onboarding****

****Description**:** Users consistently expressed feeling overwhelmed by the amount of information presented during initial setup, leading to incomplete onboarding and delayed time-to-value.

****Prevalence**:** 9 out of 12 participants mentioned this issue unprompted

****Supporting Quotes**:**

- "I just wanted to get started, but it felt like I needed to read a manual first" [P3, Marketing Manager]
- "By the third screen of instructions, I started clicking 'Next' without reading" [P7, Sales Rep]
- "I wish there was a 'quick start' option for people like me who just want to try it" [P11, Product Designer]

****Implication**:** Our current onboarding flow prioritizes completeness over engagement. We should consider a progressive disclosure approach where users can start using the product quickly and learn advanced features contextually.

****Recommended Action**:**

- Design a "Quick Start" path that gets users to first value in <3 minutes
- Move advanced configuration to contextual help within the app
- Test with 5-10 new users before full rollout
- Expected impact: +20-30% activation rate improvement

Template Output Structure

When synthesizing research, use this structure:

User Research Synthesis: [Research Topic]

Research Overview

- ****Date**:** [Date range]
- ****Methodology**:** [Interview/Survey/Testing]
- ****Participants**:** [Number] [User types]
- ****Research Questions**:**
 1. [Question 1]
 2. [Question 2]
 3. [Question 3]

Executive Summary

[2-3 sentence overview of key findings and implications]

Key Themes

Theme 1: [Theme Name]
[Full theme documentation as shown in example above]

Theme 2: [Theme Name]
[Full theme documentation]

[Continue with 4-8 themes]

Pain Points Summary

Pain Point	Severity	Frequency	Current Workaround
[Pain 1]	High	10/12 users	[How they cope]
[Pain 2]	Medium	7/12 users	[How they cope]

Feature Requests

Must-Have

1. **[Request]** - Mentioned by [X] participants
 - Quote: "[Representative quote]"
 - Underlying need: [Why they want this]

High Value

[Similar structure]

Nice-to-Have

[Similar structure]

Recommendations

High Priority (0-3 months)

1. **[Recommendation]**
 - Supporting evidence: [Data from research]
 - Expected impact: [What will improve]
 - Effort estimate: [Rough sizing]

Medium Priority (3-6 months)

[Similar structure]

Future Consideration (6+ months)

[Similar structure]

Open Questions

1. [Question requiring more research]
2. [Uncertainty to validate]
3. [Follow-up study needed]

Appendix

- Interview guide used
- Full participant demographics
- Raw notes/transcripts ([link](#))

Chapter 6 — Planning

5 skills

feature-prioritisation

Apply prioritisation frameworks (RICE, MoSCoW, Kano, ICE, Opportunity Scoring) to rank features and backlog items.

Feature Prioritisation Skill

Apply the right prioritisation framework to any backlog and produce a clear, defensible ranking with rationale — not just a sorted list.

Required Inputs

Ask the user for these if not provided:

- **List of features or initiatives to prioritise**
- **Goal or metric** being prioritised against (OKR, launch, sprint)
- **Preferred framework** (or recommend based on context below)
- **Team data:** reach estimates, effort estimates, velocity (for RICE)

Framework Selection Guide

Ask the user which framework they prefer, or recommend based on context:

Situation	Recommended Framework
Need a quick, data-driven score	RICE
Stakeholder alignment meeting	MoSCoW
Understanding customer delight vs expectations	Kano
Early-stage startup, fast decisions	ICE
Identifying underserved customer needs	Opportunity Scoring
Strategic portfolio decisions	Value vs Effort Matrix

RICE Scoring

Formula: $(\text{Reach} \times \text{Impact} \times \text{Confidence}) \div \text{Effort}$

Factor	Definition	Scale
Reach	Users impacted per quarter	Actual number
Impact	Effect on goal per user	0.25 / 0.5 / 1 / 2 / 3
Confidence	How certain are you?	50% / 80% / 100%
Effort	Person-months required	Actual number

Output table:

Feature	Reach	Impact	Confidence	Effort	RICE Score	Priority
---------	-------	--------	------------	--------	------------	----------

MoSCoW Method

Categorise each feature as:

- **Must Have** — non-negotiable for launch/sprint; product fails without it
- **Should Have** — important but not critical; workarounds exist
- **Could Have** — nice to have; include only if time allows
- **Won't Have (this time)** — explicitly out of scope now; may revisit

Always ask: "Must have for *what?*" — define the scope (launch, sprint, quarter) before categorising.

ICE Scoring (Startup/fast mode)

Formula: $\text{Impact} + \text{Confidence} + \text{Ease}$ (each 1–10)

Quick, subjective — good for early decisions before data exists.

Kano Model

Classify features into:

- **Basic (Must-be):** Expected; absence causes dissatisfaction
- **Performance:** More = better satisfaction; linear relationship
- **Excitement (Delighters):** Unexpected; creates delight; absence is neutral
- **Indifferent:** Users don't care either way
- **Reverse:** Some users want it, others don't

Recommend building: all Basic features first → Performance features for key use cases → 1–2 Excitement features per release.

Programmatic Helper

This skill ships with a stdlib-only Python script that computes ranking for the math-based frameworks (RICE, ICE) so feature scoring is consistent across sessions.

```
# RICE from JSON
python3 scripts/feature_prioritisation.py initiatives.json --framework rice

# RICE from CSV
python3 scripts/feature_prioritisation.py initiatives.csv --framework rice --
format csv

# ICE from JSON
python3 scripts/feature_prioritisation.py features.json --framework ice

# Pipe into it
printf '%s\n' '["name":"API refactor","impact":8,"confidence":80,"ease":5}]'
\
| python3 scripts/feature_prioritisation.py --framework ice -
```

Use `--json` to produce machine-readable output for downstream tooling.

Output Format

Feature Prioritisation — [Product/Team] — [Date]

Framework Used: [RICE / MoSCoW / ICE / Kano / Custom] **Scope:** [Sprint / Quarter / Release] **Goal being prioritised against:** [Metric or objective]

[Scored table using selected framework]

Recommended Build Order:

1. [Feature] — [1-line rationale]
2. [Feature] — [1-line rationale]
3. ...

Explicitly Deprioritised:

- [Feature] — Reason: [brief]

Assumptions Made:

- [Any estimates or judgements used in scoring]
-

Guidelines

- Always anchor prioritisation to a specific goal or metric — never prioritise in a vacuum
- Flag when two features have similar scores but very different risk profiles
- If stakeholder politics are influencing prioritisation, name it explicitly and suggest separating the framework score from the final decision
- Recommend revisiting priorities every 2 weeks minimum
- Never produce a single-column ranked list without rationale — explain the top 3 and bottom 3 decisions

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/framework-selection.md** — Picking the Prioritisation Framework (Instead of Defaulting to RICE). Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/prioritisation-session.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every item is scored against the same goal or metric (not different goals per item)
- ☐ Deprioritised items are explicitly listed with reasons (not just absent from the ranked list)
- ☐ Assumptions used in scoring are documented
- ☐ Stakeholder politics or personal preferences are separated from framework score
- ☐ Prioritisation is anchored to a specific scope (sprint / quarter / launch)

Anti-Patterns

- ☐ Do not score items against different goals — every item in a prioritisation session must be scored against the same objective
- ☐ Do not omit deprioritised items — explicitly listing what was cut and why is as important as the ranked list
- ☐ Do not let stakeholder politics override framework scores without documenting the override and reason
- ☐ Do not mix RICE, ICE, or MoSCoW scores across frameworks in a single session — pick one framework per prioritisation exercise
- ☐ Do not treat the output as final without documenting the assumptions used in scoring — assumptions change, and the list must be revisitable

okr-builder

Create well-structured OKRs (Objectives and Key Results) for product teams, startups, and individuals.

OKR Builder Skill

Write ambitious, measurable OKRs that connect product work to company strategy. Avoid vanity metrics, output-focused key results, and objectives that sound like task lists.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** `context.md` (metric definitions), `knowledge/strategy.md` (where the product is going), and any open `hypotheses/`. Run `python3 ../professional-brain/scripts/brain_query.py ./brain "<objective theme>"` and carry each fact's provenance tag through — don't set a key result off a `[hunch]` as if it were `[data]`.
-  **Propose to the Brain:** after producing, propose logging the chosen objectives + KR targets as a `decisions/` record (the period's bet) and any new metric definitions to `knowledge/`, each provenance-tagged. Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Working from a brief

You will often get a short brief without every detail (no baselines, no exact numbers).

Always deliver a complete, specific OKR set anyway — do not stop to ask questions and do not leave bracketed placeholders like `[target]`. Where a baseline or number is missing, infer a realistic value from the brief and the domain, and mark it (*assumed* — *confirm*). A clearly-labelled assumed baseline (e.g. "activation 40% (*assumed*) → 60%") is always better than a blank or an invented-as-fact figure.

Deeper Materials

- `references/bad-okr-gallery.md` — six realistic bad OKRs with diagnosis and rewrite (disguised roadmap, unfalsifiable objective, sandbagging, uncontrollable KR, metric zoo, missing guardrail), ending in a 5-question diagnostic. Use it when *reviewing* existing OKRs — match against the gallery before writing feedback.
- `templates/okr-worksheet.md` — a fill-in worksheet whose columns enforce the quality gates (baseline source, drift test, control test, guardrail) plus a pre-

committed quarter-end scoring rubric. Offer it when a team wants to draft OKRs themselves.

OKR Fundamentals

Objective: Qualitative, inspiring, time-bound. Answers "where are we going?" **Key Result:** Quantitative, specific, measurable. Answers "how will we know we've arrived?"

The Test for a Good KR

- Can it be scored 0.0–1.0 at the end of the period?
- Does it measure outcome, not output? ("Revenue from new customers increased by 30%" not "Launch 3 features")
- Is it ambitious but achievable? (Aim for 70% attainment as the gold standard)
- Is it within the team's control?

Common OKR Anti-Patterns to Flag and Fix

Anti-Pattern	Example	Better Version
Task masquerading as KR	"Launch onboarding redesign"	"New user activation rate increases from 42% to 65%"
Vanity metric	"Get 10,000 app downloads"	"30-day retention for new users reaches 40%"
Binary KR	"Ship API v2"	"API v2 adopted by 80% of active integrations"
Too many KRs	6+ per objective	Max 3–4 KRs per objective
No baseline	"Improve NPS"	"NPS increases from 32 to 50"

Always flag anti-patterns and offer a rewrite.

Output Format

[Quarter] OKRs — [Team/Product Area]

Objective 1: [Inspiring, qualitative statement]

Why this matters: [1–2 sentence strategic context]

#	Key Result	Baseline	Target	Measurement Method
KR1	[Measurable outcome]	[Current state]	[Target]	[How measured]
KR2	[Measurable outcome]	[Current state]	[Target]	[How measured]

#	Key Result	Baseline	Target	Measurement Method
KR3	[Measurable outcome]	[Current state]	[Target]	[How measured]

Owner: [Name/Role] Check-in cadence: Weekly

Repeat for each objective. Recommend 2–4 objectives per team per quarter.

Scoring Guide to Include

At quarter end, score each KR:

- 0.7–1.0 = Excellent (0.7 is the "sweet spot" — if all KRs score 1.0, they weren't ambitious enough)
- 0.4–0.6 = Made progress but missed
- 0.0–0.3 = Missed — needs retrospective discussion

Inputs (infer any not provided — label assumptions)

- **Team or individual** the OKRs are for
- **Quarter and year**
- **Company or product North Star metric** (OKRs should connect to this — if not given, infer a plausible one and label it *(assumed)*)
- **Top 3 priorities or goals for this quarter** (rough notes are fine)
- **Any existing OKRs to review or improve** (optional)

Guidelines

- Connect OKRs to the company/product North Star; if it isn't given, infer a plausible one and label it *(assumed)* rather than asking
- Recommend no more than 3 objectives per team per quarter
- If user provides output-based goals, always reframe as outcomes
- Include a "health check" section flagging which KRs have no current baseline data
- Remind user: OKRs are not performance reviews — they should be ambitious enough that missing them is okay

Quality Checks

- ☐ Each KR is measurable with a baseline and target
- ☐ No output-based KRs (no "launch X" or "complete Y")
- ☐ Maximum 4 KRs per objective
- ☐ OKRs connect to the company or product North Star
- ☐ Ambitious enough that 0.7 attainment is the expected score

Anti-Patterns

- ☐ Do not accept output-based key results — any KR phrased as "launch X" or "complete Y" must be rewritten as an outcome with a baseline and target
- ☐ Do not write OKRs without asking for the company or product North Star — OKRs disconnected from the strategic context are just a goal-setting exercise
- ☐ Do not write more than 4 KRs per objective — too many KRs dilute focus and make scoring ambiguous at quarter end
- ☐ Do not use binary KRs (ship/don't ship) — every KR must be scorable on a 0.0–1.0 scale based on degree of achievement
- ☐ Do not skip the health check section on baselines — OKRs without current baselines cannot be scored objectively at quarter end

rice-impact-matrix

Scores features using both RICE and strategic alignment for nuanced prioritisation.

RICE + Strategic Alignment Skill

Produce a prioritisation output that balances quantitative RICE scoring with qualitative strategic fit — because the highest RICE score isn't always the right next bet.

Required Inputs

Ask the user for these if not provided:

- **List of initiatives or features to prioritise** (names and brief descriptions)
- **Current strategic priorities or OKRs** (needed to rate strategic alignment)
- **Reach estimates** (users affected per quarter — even rough estimates work)
- **Effort estimates** (person-months — from engineering if available)
- **Quarter or planning period**

Two-Stage Process

Stage 1: RICE Scoring

- Reach: Users affected per quarter
- Impact: 3/2/1/0.5/0.25 scale
- Confidence: 100% / 80% / 50%
- Effort: Person-months
- $RICE = (R \times I \times C) / E$

Stage 2: Strategic Alignment Score

Rate each initiative against your current strategic priorities (provided as input):

- Directly supports top OKR: +3
- Supports secondary OKR: +2
- Neutral: +1
- Contradicts strategic direction: -1

Final Priority Score

Combined Score = RICE Score + (Strategic Alignment × 10)

Validate — Flag any initiative where RICE score and strategic alignment conflict sharply (e.g., high RICE, low alignment). These require an explicit team conversation before sequencing.

Output Structure

Priority Matrix — [Quarter]

Initiative	RICE Score	Strategic Alignment	Combined Score	Quadrant	Recommendation
[name]	[score]	[score]	[combined]	[Now/Next/Later/Drop]	[action]

Quadrant Definitions

- **Now:** High RICE + High Strategic Alignment → Build this quarter
- **Next:** High RICE + Lower Alignment → Queue for next quarter
- **Later:** Lower RICE + High Alignment → Revisit when capacity allows
- **Drop:** Low RICE + Low Alignment → Remove from backlog

Recommendations

[Top 5 initiatives with rationale for sequencing]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/strategic-weighting.md** — Blending RICE with Strategic Fit — Without Cooking the Books. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/matrix-worksheet.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ All RICE components have an estimate (even if low confidence — flag those)
- ☐ Strategic alignment is rated against specific OKRs, not general "feels strategic"
- ☐ Conflicts between RICE rank and strategic alignment are explicitly flagged
- ☐ "Drop" recommendations are specific — not just "low priority, deprioritise"
- ☐ Confidence levels on estimates are noted where weak (drives the 50% confidence flag)

Anti-Patterns

- ☐ Do not treat the combined score as a definitive ranking — use it to structure a conversation, not replace one
- ☐ Do not rate strategic alignment as "high" because an initiative feels important without mapping it to a specific OKR
- ☐ Do not place all initiatives in the "Now" quadrant — a matrix with no "Drop" recommendations is not credible
- ☐ Do not ignore the conflict flag when RICE rank and strategic alignment sharply diverge
- ☐ Do not accept 100% confidence on estimates that have not been validated with data

rice-prioritisation

Scores and ranks product initiatives using the RICE framework.

RICE Prioritisation Skill

Apply consistent, criteria-based RICE scoring to a list of features or initiatives to produce an objective prioritisation ranking.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** `knowledge/strategy.md` (so the ranking serves the direction), the items as `entities/` , and impact `hypotheses/` . Run `python3 ../professional-brain/scripts/brain_query.py ./brain "<initiative theme>"` and carry each fact's provenance tag through — an impact estimate is usually a `[hunch]` , not `[data]` .

- 🧠 **Propose to the Brain:** after producing, propose recording the ranking decision to decisions/ and the reach/impact estimates as hypotheses/ tagged by evidence strength. Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Required Inputs

Ask the user for these if not provided:

- **List of initiatives or features to score** (names and brief descriptions)
- **Reach estimates** (users affected per quarter — from analytics if available)
- **Impact estimates** (use the standard scale below)
- **Effort estimates** (person-months — from engineering if available)
- **Quarter or planning period**

RICE Definitions (adapt to your context)

- **Reach:** Number of users affected per quarter (use actual DAU/MAU data where available)
- **Impact:** Effect on your primary metric — use scale: 3=massive, 2=high, 1=medium, 0.5=low, 0.25=minimal
- **Confidence:** How certain are we about R and I estimates? 100%=high, 80%=medium, 50%=low
- **Effort:** Person-months required across all functions

RICE Formula

RICE Score = (Reach × Impact × Confidence) / Effort

Programmatic Helper

This skill ships with a stdlib-only Python script that calculates and ranks RICE scores so the maths is consistent and the quick-win / moonshot flags are applied by rule, not by feel. Feed it the initiatives once R, I, C, and E are gathered.

```
# From a JSON file (confidence accepts 0.8 or 80)
python3 scripts/rice_calculator.py initiatives.json

# Or from a CSV with header: name,reach,impact,confidence,effort
python3 scripts/rice_calculator.py initiatives.csv --format csv

# Or piped in
echo
'[{ "name": "Onboarding", "reach": 5000, "impact": 2, "confidence": 0.8, "effort": 3 } ]'
```

```
\
| python3 scripts/rice_calculator.py -
```

It outputs a ranked table with computed RICE scores and auto-flags **quick-win** (strong score, low relative effort), **moonshot** (high impact, high effort), and **low-confidence** ($\leq 50\%$) items. Use the computed ranking as the starting point, then apply the validation step below — never accept a surprising top rank without checking the estimates behind it.

Deeper Materials

- **references/estimate-calibration.md** — how to anchor each of the four estimates (reach sources, the impact scale with reserve-it-for examples, evidence-based confidence, cross-functional effort) and the cross-checks to run on the finished ranking. Apply it when challenging the user's inputs.
- **templates/scoring-worksheet.md** — a fill-in worksheet whose evidence columns force each score to name its source. Offer it when a team wants to score together rather than have the ranking generated.

Process

1. For each initiative provided, gather or estimate R, I, C, E values
2. Flag where estimates are weak and note what data would improve them
3. Calculate RICE score for each
4. Rank highest to lowest
5. Flag any "quick wins" (high RICE score, low effort) and "moonshots" (high impact, high effort)
6. Note dependencies between items that affect sequencing
7. **Validate** — Cross-check: if the top-ranked item surprises the team, investigate whether an estimate is inflated. RICE is a tool, not a verdict.

Output Structure

RICE Prioritisation: [Backlog/Quarter]

Initiative	Reach	Impact	Confidence	Effort	RICE Score	Notes
[name]	[n]	[score]	[%]	[months]	[score]	[flags]

Recommended Sequence

[Top 5 initiatives with rationale]

Quick Wins (high score, low effort)

[Items to pick up alongside bigger bets]

Data Gaps to Address

[What information would most improve scoring accuracy]

Quality Checks

- ☐ Every initiative has all four RICE components estimated (even roughly)
- ☐ Confidence is 50% for anything without data backing (not 100% as a default)
- ☐ Quick wins and moonshots are explicitly called out
- ☐ Dependencies that affect sequencing are noted
- ☐ Any surprising ranking is investigated before accepting it

Anti-Patterns

- ☐ Do not default to 100% confidence on estimates that lack supporting data — this inflates scores and misleads planning
- ☐ Do not treat RICE scores as a final decision — a ranking that surprises the team must be investigated before it is accepted
- ☐ Do not omit effort estimates from engineering — PM-only effort estimates are frequently optimistic and skew results
- ☐ Do not forget to note dependencies that would change the sequencing even if RICE scores suggest otherwise
- ☐ Do not score every initiative at the same impact level — if everything is "high impact," the framework produces no useful signal

roadmap-narrative

Transform a prioritised initiative list into a compelling strategic roadmap narrative.

Roadmap Narrative Skill

Convert a ranked list of product initiatives into a clear, strategic narrative that connects individual items to company goals and communicates a coherent product direction.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, ground in it instead of re-asking for what you already know:

- **Read first:** `knowledge/strategy.md` (the direction the narrative must ladder to), `priority decisions/`, and `feature entities/`. Run `python3 ../professional-`

brain/scripts/brain_query.py ./brain "<roadmap theme>" and carry each fact's provenance tag through.

-  **Propose to the Brain:** after producing, propose logging the sequencing/priority decisions to `decisions/` and updating the relevant feature `entities/`, each provenance-tagged. Show them, get a yes, then write with `../professional-brain/scripts/brain_write.py ... --commit` (append-only, dry-run by default).

Working from a brief

You will often get a short brief (a few themes, an audience) without a full initiative list or OKRs. **Always deliver the complete narrative anyway** — do not stop to ask questions and do not leave bracketed placeholders like `[Theme Name]`. Where detail is missing, infer specific, realistic themes, initiatives, and metrics from the brief and the domain, and mark any inferred fact or number as (*assumed* — *confirm*). Fill every section with concrete content, not template brackets.

Inputs (infer any not provided — label assumptions)

- **Prioritised initiative list** (with rough timelines or quarters)
- **Company OKRs or strategic priorities** (to connect roadmap to company goals)
- **Audience** (all-hands, board, investors, sales team — changes tone and depth)
- **Items explicitly NOT on the roadmap** (optional but strengthens credibility)

Process

1. Review the prioritised initiative list and company OKRs provided
2. Identify 2-3 strategic themes that group the initiatives naturally
3. For each theme, articulate: the problem it addresses, the customer it serves, the metric it moves
4. Write a quarter-level narrative that shows progression — how does H1 set up H2?
5. Draft an executive summary (3-4 sentences max) that non-technical stakeholders can repeat
6. **Validate** — Confirm every initiative maps to a theme. If an initiative is orphaned, either create a theme or flag it as a narrative gap to address

Output Structure

Product Roadmap: [Quarter/Half/Year]

Strategic Context: [1 paragraph: market moment, key challenge, our response]

Theme 1: [Theme Name]

- Strategic rationale

- Initiatives included
- Primary metric impacted
- Dependencies

[Repeat for each theme]

What's Not on the Roadmap (and Why): [2-3 items with rationale — shows strategic discipline, not just prioritisation]

Executive Summary (shareable): [3-4 sentences that could be shared in an all-hands or board update]

Tone Guidelines

- Write for a CFO, not an engineer
- Lead with customer outcomes, not features
- Be honest about what's NOT on the roadmap and why

Timeline, drawn

When the themes have a sequence or dates, also render the roadmap as a Mermaid Gantt chart so the shape of the plan is visible (it renders live in the playground; with real ISO dates it also exports to a calendar .ics). Use `section` per theme/quarter and mark key checkpoints as milestones.

```
gantt
  title Roadmap
  dateFormat YYYY-MM-DD
  section Theme 1
    Initiative      :2026-07-01, 30d
    Checkpoint      :milestone, 2026-07-31, 0d
  section Theme 2
    Initiative      :2026-08-01, 45d
```

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/now-next-later.md** — Now/Next/Later Done Right: Commitment Gradients, Not Date Camouflage. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/roadmap-onepager.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Every initiative in the input maps to a strategic theme
- ☐ The executive summary can stand alone and be repeated correctly after one reading
- ☐ Progression narrative shows causal links between quarters (not just chronological listing)
- ☐ "What's not on the roadmap" section includes at least 2 items with clear rationale
- ☐ Language throughout is free of engineering jargon — tested by asking: "could a CFO repeat this?"

Anti-Patterns

- ☐ Do not produce a list of features with dates and call it a narrative — every initiative must connect to a strategic theme
- ☐ Do not omit the "what's not on the roadmap" section — without it, the narrative lacks strategic discipline
- ☐ Do not write progression as a chronological list — show causal links between quarters (Q1 enables Q2 because...)
- ☐ Do not write the executive summary last and treat it as a summary — write it as the version stakeholders will repeat
- ☐ Do not let orphaned initiatives appear without a theme — either create a theme or flag the gap explicitly

Chapter 7 — Gtm

3 skills

competitor-teardown

Produce a structured competitive analysis for any product or market.

Competitor Teardown Skill

This skill produces a complete competitive analysis document — structured for use in strategy decks, investor materials, sales enablement, or product planning sessions.

Required Inputs

Ask the user for these if not provided:

- **Your product** (name + one-line description)
- **Competitors to analyse** (list 2–5 names; if not provided, ask)
- **Analysis depth** (quick overview / detailed teardown)
- **Primary use case for this analysis** (e.g. sales enablement, investor deck, internal strategy, product planning)

Deeper Materials

- **references/intel-sourcing-guide.md** — where competitive facts come from (four source tiers), which source to use per teardown section, the [verified]/[reported]/[assumed] confidence labels, and the ethics line. Apply its labelling to every substantive claim in the output.
- **templates/teardown-skeleton.md** — a fill-in teardown with the confidence labels and a verification queue built in. Offer it when the user wants to gather the intel themselves.

Output Structure

1. Competitive Landscape Overview

One paragraph summarising the market dynamic: who the key players are, how the market is segmented, and where the white space sits. Keep this under 150 words — it's the exec summary.

2. Positioning Map

Describe a 2x2 positioning map in text form (since you can't render images):

- Define the two axes relevant to this market (e.g. "Ease of Use vs. Depth of Features" or "Price vs. Enterprise Readiness")
- Place each competitor in one quadrant with a one-sentence rationale
- Place the user's product and highlight the strategic implication

3. Feature Comparison Table

Feature / Capability	[Your Product]	[Competitor A]	[Competitor B]	[Competitor C]
[Feature]	✅ / ❌ / 🟡 Partial			

Use ✅ (has it), ❌ (doesn't have it), 🟡 (partial/limited). Add a "Strategic Notes" column for features where the difference is a significant selling point or risk.

Include 10–15 rows. If user hasn't provided feature details, note which cells need to be verified.

4. Messaging Analysis

For each competitor, analyse their public-facing messaging (website headline, tagline, primary value prop):

[Competitor Name]

- **Their primary claim:** [what they say they do]
- **Target audience signal:** [who they seem to be targeting based on language/imagery]
- **Emotional hook:** [fear / aspiration / authority / speed / simplicity]
- **Gap or weakness in their messaging:** [what they don't address that your product could own]

5. SWOT Summary

Produce a clean SWOT for the user's product in the context of this competitive landscape:

- **Strengths:** [2–3 genuine differentiators]
- **Weaknesses:** [2–3 honest gaps or vulnerabilities]
- **Opportunities:** [2–3 market gaps or competitor weaknesses to exploit]
- **Threats:** [2–3 competitor moves or market shifts to watch]

6. Strategic Recommendations

3–5 actionable recommendations based on the analysis. Frame each as: "**Given [observation], [your product] should [action] to [outcome].**"

Quality Checks

- ☐ Axes on positioning map are meaningful and specific to this market
- ☐ Feature table includes strategic notes on key differentiators
- ☐ Messaging analysis covers all named competitors
- ☐ SWOT is honest — Weaknesses and Threats should not be softened
- ☐ Recommendations are specific and actionable, not generic strategy advice

Anti-Patterns

- ☐ Do not mark feature presence as equivalent across competitors without noting quality differences — both products may have "reporting" while one's is meaningfully better
- ☐ Do not position the user's product in the most favourable quadrant without justification — a self-serving positioning map that ignores real competitive pressure provides no strategic value
- ☐ Do not soften Weaknesses or Threats in the SWOT — a SWOT that only celebrates strengths is a marketing document, not a strategy tool
- ☐ Do not include unverifiable claims about competitor capabilities without flagging them as assumptions — presenting rumours as facts damages analytical credibility

Example Trigger Phrases

- "Do a competitor analysis of [Product] vs [Competitor A] and [Competitor B]"
- "Tear down [Competitor]'s positioning"
- "Give me a competitive landscape for [market]"
- "Build a SWOT for our product against [competitor]"

go-to-market

Create go-to-market assets for any product or feature.

Go-To-Market Skill

This skill produces a complete go-to-market asset pack for a product, feature, or initiative. It follows Geoffrey Moore's positioning framework and structures all outputs for use in sales decks, landing pages, launch emails, and internal alignment docs.

Working from a brief

You will often get a short brief without every detail. **Always deliver the full GTM pack anyway** — do not stop to ask questions and do not leave bracketed placeholders like [ADD PROOF POINT] OR [Technical capability]. Where a detail is missing (differentiators, proof points, features), infer specific, realistic ones from the product description and the target customer, and mark anything inferred as (*assumed* — *confirm*). A concrete, labelled assumption is always better than a blank.

Inputs (infer any not provided — label assumptions)

- **Product/feature name**
- **One-line description** (what it does, technically)
- **Target customer** (role, company size, industry if relevant)
- **Primary problem it solves**
- **Key competitor or alternative** (what people do today without this)
- **Top 3 differentiators**

Reads from / Writes to the Brain

If a professional-brain (brain/) exists, use it before asking:

- **Read first:** context.md (product, ICP, voice), knowledge/market.md and knowledge/strategy.md , and the matching entities/ feature being launched.
- **Write after:** save the launch plan to entities/ , and any positioning or channel decision to decisions/ , each provenance-tagged.

Output Structure

Always produce all four sections below in order.

1. Positioning Statement

Use the Geoffrey Moore format exactly:

For [target customer] who [has this problem or need], [Product Name] is a [product category] that [key benefit/outcome]. Unlike [primary alternative or competitor], our product [key differentiator].

Write one primary positioning statement, then offer a shorter tagline version (10 words or fewer) suitable for a hero headline.

2. Messaging Pillars

Generate 3–5 messaging pillars. Each pillar must include:

- **Pillar name** (2–4 words, bold)
- **One-sentence summary** of what this pillar claims
- **2–3 proof points** (specific and evidence-backed; if no data was provided, infer a realistic proof point and mark it *(assumed)* — never leave a bare placeholder)
- **Example use in copy** (one sentence as it would appear in a landing page or deck)

Pillars should be distinct — avoid overlap. Each pillar should be defensible against the primary competitor.

3. Feature & Functionality List

Produce a two-column table:

Feature / Functionality	Buyer Benefit (what it means for the user)
[Technical capability]	[Outcome in plain language — start with a verb: "Reduces...", "Enables...", "Eliminates..."]

Rules:

- Never list a feature without a corresponding benefit
- Benefits should reference the target customer's workflow or pain point
- Aim for 6–12 rows; if only 1–2 features were given, infer the rest plausibly from the product description
- Avoid jargon in the benefit column — write as if explaining to a buyer, not an engineer

4. Use Cases

Generate 3–5 role-specific use cases. Each use case must follow this format:

Use Case [N]: [Role] — [Scenario Title]

- **Who:** [Job title / role]
- **Situation:** [The specific moment or trigger that leads them to use the product]
- **Before:** [What they had to do without this product — be specific about time, friction, or risk]
- **With [Product Name]:** [What they do now — concrete action, not vague benefit]
- **Outcome:** [Measurable or tangible result]

Use cases should cover different buyer personas if possible (e.g. end user, manager, admin).

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/messaging-hierarchy.md** — The Messaging Hierarchy: One Claim, Then Everything Else. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/gtm-pack.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

Before delivering output, verify:

- ☐ Positioning statement follows Moore format exactly
- ☐ Tagline is 10 words or fewer
- ☐ Each pillar has at least 2 proof points (or flagged placeholders)
- ☐ Every feature has a benefit — no orphaned features
- ☐ Benefits start with action verbs
- ☐ Use cases include a Before/After structure
- ☐ Language is consistent with the target customer's vocabulary (not internal engineering terms)

Anti-Patterns

- ☐ Do not write feature descriptions instead of benefits — the GTM pack must translate features into customer value
- ☐ Do not use the same messaging across all buyer personas — each role has different priorities and language
- ☐ Do not create a positioning statement that could apply to any competitor — differentiation must be specific and defensible
- ☐ Do not skip the "not for" section — defining who this is not for sharpens positioning and prevents misdirected sales effort
- ☐ Do not list use cases without tying them to specific job titles or buyer roles

Example Trigger Phrases

- "Create a positioning statement for [product]"
- "Write a GTM plan for [feature]"
- "Give me key pillars for [product name]"
- "Build a feature and use case list for [product]"
- "We're launching [X] — help me with the messaging"

product-positioning-doc

Write a product positioning document and messaging framework.

Product Positioning Doc Skill

This skill produces a complete product positioning document following the April Dunford positioning methodology. Output covers category definition, target customer, unique attributes, proof points, and a messaging hierarchy — ready to align GTM, marketing, sales, and product teams.

Required Inputs

Ask the user for these if not provided:

- **Product name** and what it does
- **Target customer** — who is it for? (role, company type, size)
- **Problem it solves** — what pain or goal does it address?
- **Key alternatives** — what do customers use today instead? (not just direct competitors — include status quo, spreadsheets, DIY)
- **Differentiation** — what does this product do that alternatives cannot? (not features — capabilities that produce different outcomes)
- **Proof points** — any customer data, case studies, metrics, or validation?
- **Business goal** — is positioning for a new category, expansion into new segment, or repositioning away from a declining category?

Output Structure

Positioning Document: [Product Name]

Version: [1.0] **Owner:** [PMM / Founder / Marketing lead] **Date:** [Date] **Status:** [Draft / Reviewed / Approved] **Approved by:** [Names — this document must be signed off by product, marketing, and sales leadership before use]

1. Background & Context

[2–3 sentences describing why positioning is being done now. Is this a new product, a pivot, a segment expansion, or a rebrand? What triggered this work?]

Positioning objective: [e.g. Move from being perceived as a reporting tool to being the category leader in revenue intelligence for mid-market SaaS]

2. Market Category

What category does this product compete in?

This is the frame of reference your customer uses to understand what the product is. Choose the wrong category and everything downstream — competitors, value, messaging — is wrong.

Category: [e.g. Customer data platform / Revenue intelligence / No-code automation / Modern data stack]

Why this category, not [alternative category]? [1–2 sentences on why this framing serves the customer's understanding better than adjacent categories]

Category maturity:

- ☐ New category (we are creating it — high education burden, high upside if it works)
- ☐ Growing category (fast-growing segment — compete on differentiation)
- ☐ Mature category (well-understood — must disrupt with clear superiority or narrower niche)

3. Target Customer

Be precise. Vague targeting produces vague positioning.

Dimension	Description
Primary buyer / decision-maker	[e.g. VP of Revenue Operations at B2B SaaS companies with 100–500 employees]
Primary user	[e.g. Revenue operations analysts and sales ops managers]
Company profile	[Industry, size, growth stage, technology stack]
Business context	[What is happening in their world that makes them a buyer right now?]
Trigger event	[What just happened that makes them start looking for a solution? — e.g. Sales team grew past 20 reps, forecast accuracy became a board question]

Who this is NOT for: [Be explicit about who to exclude — this sharpens the positioning for those who are a fit]

4. Competitive Alternatives

What do buyers use today when they don't have your product? List all real alternatives — not just direct competitors.

Alternative	Who uses it	Why buyers choose it	What they sacrifice
[Direct competitor — e.g. Gong]	[Enterprise sales teams]	[Market leader, strong brand, sales coaching features]	[Price, complexity, implementation time]
[Adjacent tool — e.g. Salesforce reports]	[CRM-native users]	[Already have it, no additional cost]	[No AI analysis, manual reporting, siloed data]
[Status quo — e.g. spreadsheets + manual tracking]	[SMB, early-stage]	[Free, flexible, no change management]	[Time-consuming, error-prone, not scalable]
[Build in-house]	[Tech companies with data teams]	[Custom to their exact needs]	[Engineering cost, maintenance burden, 12+ month timeline]

Key insight: [What does this competitive landscape tell you about what your positioning must emphasise? e.g. "Every alternative either costs too much or requires too much manual work — positioning must nail 'fast time to value' and 'right-sized for mid-market'"]

5. Unique Differentiated Attributes

These are the features or capabilities your product has that alternatives genuinely cannot match — or cannot match at the same level. Do not list features that competitors also have.

Attribute	What it is	What it enables (outcome)	Why competitors can't match it
[e.g. Real-time CRM sync]	[Bidirectional sync with any CRM in <5 min]	[Reps see clean data in the tools they already use — no toggle between systems]	[Legacy competitors require 3-month integration projects; Salesforce-native tools only work in SFDC]
[e.g. Natural language querying]	[Ask questions in plain English, get data visualisations]	[Anyone on the revenue team can answer their own questions without SQL or waiting for an analyst]	[BI tools require analyst training; direct competitors have rigid dashboards]
[...]	[...]	[...]	[...]

The core differentiation thesis: [1–2 sentences that unite the above attributes into a single "why we win" statement — this is internal language, not customer-facing yet]

6. Value Proof Points

Back up the differentiation claims with evidence:

Claim	Proof point	Source
[Fastest time to value]	[Average customer is live in 4 hours vs 3 months for legacy alternatives]	[Customer data — average across [X] accounts]
[Better forecast accuracy]	[Customers achieve X% improvement in forecast accuracy within 90 days]	[Case study: [Company Name] — link]
[Loved by operators, not just managers]	[NPS of X among end users; 4.8/5 on G2 for ease of use]	[G2 reviews, internal NPS survey]

Proof gaps: [Are there claims you're making that you don't yet have evidence for? List them — they are either research projects or risks to the positioning]

7. Positioning Statement

The classic positioning template — internal only, never used verbatim in marketing:

For [target customer] **who** [trigger event or problem statement], **[Product name]** is a **[category]** **that** [primary differentiated value — the outcome, not the feature]. **Unlike** [primary alternative], **[Product name]** [the key thing that makes you different and better].

Draft positioning statement:

For [VP Revenue Ops at B2B SaaS companies with 50–500 reps] who [struggle to forecast accurately as the sales team scales], [Product Name] is a [revenue intelligence platform] that [gives every rep and manager accurate, real-time pipeline visibility without any analyst overhead]. Unlike [Salesforce dashboards and manual reporting], [Product Name] [syncs automatically, surfaces risks before they become missed quarters, and needs no configuration by IT or data teams].

8. Messaging Hierarchy

Translate the positioning into customer-facing language at three levels:

Tagline (5–8 words)

[The simplest possible statement of what you do and for whom. Used in ads, hero sections, email signatures.]

Options to test:

- 1. [e.g. "Revenue intelligence for scaling sales teams"]
- 2. [e.g. "Forecast with confidence. Close with clarity."]
- 3. [e.g. "The revenue platform your whole team will actually use"]

Value Proposition (1–2 sentences)

[Used in the hero section of the website, email subject lines, and sales decks. Must be instantly clear.]

[e.g. "[Product Name] gives revenue teams real-time pipeline visibility and accurate forecasting — without spreadsheets, custom reports, or waiting for an analyst. Get live in 4 hours, not 4 months."]

Full Description (3–5 sentences)

[Used in PR, partnership briefs, longer sales emails, and About Us pages.]

[e.g. "[Product Name] is the revenue intelligence platform built for mid-market SaaS teams. Unlike legacy BI tools that require analyst configuration or CRM dashboards that only show what's already happened, [Product Name] automatically syncs your entire revenue stack, surfaces AI-driven risk signals, and lets any rep or manager ask questions in plain English. [X] customers use [Product Name] to call their quarters with confidence. Average time to live: 4 hours."]

9. Persona-Specific Messaging

The core positioning is the same, but different buyers care about different aspects:

Persona	Their primary concern	Lead message	Proof point to use
VP Revenue Operations	Forecast accuracy, board credibility	"Call your quarter with confidence"	[X% improvement in forecast accuracy across N customers]
Head of Sales	Rep productivity, pipeline visibility	"Your reps close more, not admin more"	[X hours/week saved per rep]
CEO / CFO	Revenue predictability, cost	"Stop being surprised by quarters"	[ROI: £X saved vs X headcount required to replicate manually]

Persona	Their primary concern	Lead message	Proof point to use
Sales Rep	Ease of use, not adding to workload	"It works in the tools you already use"	[Ease of use NPS, G2 reviews]

10. Messaging Do's and Don'ts

Do say:

- [Specific, outcome-focused language — what the customer achieves]
- [Comparative language grounded in evidence]
- [Language your target buyer uses to describe their problem — not language you invented]

Don't say:

- ["Best-in-class", "innovative", "cutting-edge", "game-changing" — unless followed by evidence]
- [Feature lists without outcome context]
- [Jargon your buyer doesn't use themselves]
- [Claims your competitors could also make]

11. Distribution Plan

Positioning only works if it's implemented consistently:

Team	What they need	Format	Owner	When
Marketing	Tagline, value prop, messaging hierarchy	This doc + messaging playbook	PMM	[Date]
Sales	Competitive positioning, objection responses	One-pager + deck	Sales enablement	[Date]
Product	Category definition, target customer	Shared doc + roadmap input	PMM + PM	[Date]
Leadership	Full positioning narrative	This doc	PMM	[Date]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/category-choice.md** — Choosing Your Category: the Highest-Leverage Positioning Decision. Apply it while producing the output; it carries the calibration

and judgment calls the method summary above compresses.

- **templates/positioning-canvas.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Positioning statement has exactly one A — the product is accountable to exactly one primary differentiated claim
- ☐ Competitive alternatives include the status quo — not just named competitors
- ☐ Differentiated attributes describe outcomes, not features
- ☐ Every proof point cites a source — not "customers say..."
- ☐ Persona messaging uses the buyer's language, not the company's
- ☐ At least two people from product, marketing, and sales have reviewed and approved

Anti-Patterns

- ☐ Do not write positioning that could describe any competitor — differentiation must be specific, provable, and hard to copy
- ☐ Do not mix category design with category entry — know whether you are creating a new category or competing in an existing one
- ☐ Do not create persona messaging that uses the same headline for all personas — each persona has different priorities
- ☐ Do not include proof points that are claims without evidence — every proof point needs a supporting data point or reference
- ☐ Do not skip the "not for" section — defining who this is not for sharpens targeting and prevents off-persona deals

Example Trigger Phrases

- "Write a positioning document for [product]"
- "Build a messaging framework for our B2B SaaS tool"
- "Define our product positioning — who is this for and why should they care?"
- "Create a positioning statement and messaging hierarchy for [launch]"
- "Help me articulate our differentiation vs [Competitor]"

Chapter 8 — Analytics

3 skills

data-analysis-standard

Structure a product data analysis, metric deep-dive, funnel analysis, or cohort study.

Data Analysis Standard Skill

Turn raw numbers into product decisions. Structure every analysis with a clear question, methodology, finding, and recommended action.

Analysis Framework: The 4-Question Method

Every analysis starts here:

1. **What changed?** (describe the metric and its movement)
2. **Why did it change?** (root cause — segment, funnel step, cohort, channel)
3. **So what?** (business or product impact)
4. **Now what?** (recommended action with confidence level)

Never deliver data without answering all four. A chart with no narrative is not an analysis.

Metric Triage Template

Use when a metric has moved unexpectedly:

METRIC: [Name]

MOVEMENT: [X% change over Y period]

BASELINE: [What was normal]

SEGMENTATION CHECK:

- By platform (iOS / Android / Web)?
- By user cohort (new / returning / power users)?
- By acquisition channel?
- By geography?
- By plan/tier?

ROOT CAUSE HYPOTHESIS:

1. [Most likely explanation] — Evidence: [data point]
2. [Alternative explanation] — Evidence: [data point]
3. [Ruling out] — Eliminated because: [reason]

CONCLUSION: [Single sentence answer to "why did this change?"]

CONFIDENCE: [High / Medium / Low] — based on [data available]

Funnel Analysis Structure

Stage	Metric	Current	Benchmark/Target	Drop-off %	Notes
[Top of funnel]	[Users]	[N]	[N]	—	
[Step 2]	[Users]	[N]	[N]	[X%]	
[Step 3]	[Users]	[N]	[N]	[X%]	
[Conversion]	[Users]	[N]	[N]	[X%]	

Biggest drop-off: [Step X → Step Y] — Hypothesis: [reason] **Recommended investigation:** [specific query or test]

Cohort Analysis Guidelines

Always define:

- **Cohort definition:** [What groups users — signup week, first action, plan type]
- **Retention metric:** [What counts as retained — login, core action, revenue]
- **Retention window:** [D1, D7, D30, W4, M3, etc.]

Output a cohort retention table and annotate:

- Baseline retention for each cohort
- Cohorts that over/underperform and why (feature launch? campaign? seasonal?)
- Trend direction across cohorts (improving / declining / stable)

Stakeholder Analysis Output Format

[Analysis Title] — [Date]

Question being answered: [Specific question in plain English] **Time period:** [Date range]

Data source: [Where data comes from]

Finding:

[1–2 sentence plain-English summary of what the data shows]

Key chart / table: [Include or describe]

Root cause: [Best explanation with evidence]

Confidence level: [High / Medium / Low] — [reason]

Recommended action:

1. [Immediate action — owner, timeline]
2. [Investigation needed — what to check next]
3. [Monitoring — what metric to watch and at what cadence]

What this analysis does NOT tell us: [Important caveat — what data is missing or what can't be concluded]

Required Inputs

Ask the user for these if not provided:

- **Metric or question** being investigated
- **Time period** (what changed, from when to when)
- **Data available** (which segments, sources, or queries you have access to)
- **Business context** (what decision this analysis informs)
- **Audience** (who will read this — exec / team / data team)

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/analysis-integrity.md** — Analysis Integrity: the Checks Between Query and Conclusion. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/analysis-writeup.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Analysis answers all 4 questions: what changed, why, so what, now what
- ☐ Root cause has evidence (not just hypothesis)
- ☐ Confidence level is stated and justified
- ☐ What the data cannot tell us is explicitly named
- ☐ Recommended action includes an owner and timeline

Anti-Patterns

- ☐ Do not present correlations as causation — always state the distinction explicitly
- ☐ Do not report a metric movement without stating the time window and comparison baseline
- ☐ Do not skip the "so what" — raw observations without recommended actions are incomplete analysis
- ☐ Do not overstate confidence — label hypotheses clearly and note what data would be needed to confirm them
- ☐ Do not ignore segment breakdowns — aggregate metrics can mask opposing trends in sub-segments

Guidelines

- Always state what the data *cannot* tell you — never oversell confidence
- Correlations are not causation — flag this every time
- If the user has no baseline, recommend establishing one before drawing conclusions
- Recommend the simplest chart for each finding: bar for comparison, line for trends, scatter for correlation, table for detailed breakdowns
- Always specify the time window — "conversion dropped" is meaningless without "from X to Y over Z period"

product-health-analysis

Interpret product metrics against goals and surface actionable signals.

Product Health Analysis Skill

Transform raw metrics data into a clear health narrative — what's working, what's not, and what needs immediate attention.

Required Inputs

Ask the user for these if not provided:

- **Metrics data** (current values for key metrics — even rough numbers work)
- **Targets or benchmarks** (OKR targets, historical baselines, or industry benchmarks)
- **Period** (week / month / quarter being analysed)
- **Product area or segment** (are we looking at the whole product or a specific feature?)

Metrics Framework

Analyse across four layers:

1. **Acquisition** — new users, source quality, CAC trends
2. **Activation** — time to first value, onboarding completion rates
3. **Engagement** — DAU/MAU, feature adoption, session depth
4. **Retention** — D1/D7/D30 retention, churn rate, resurrection rate

Process

1. For each metric, compare: current period vs. previous period, current vs. target
2. Flag anything more than 10% off target as requiring investigation
3. Look for correlations — does a drop in activation explain a retention dip 2 weeks later?
4. Write a plain-English health summary (no jargon) suitable for sharing with non-data stakeholders
5. Recommend top 3 areas for immediate investigation with suggested diagnostic steps
6. **Validate** — Confirm every flagged metric has a plausible root cause hypothesis, not just a raw number, and every recommended action has a specific owner or team

Output Structure

Product Health Report — [Period]

Overall Health:  On Track /  Watch /  Action Required

Metric	Current	Target	vs. Last Period	Status
[metric]	[value]	[target]	[+/-%]	  

Key Observations: [3-5 bullet observations written in plain English]

Areas Requiring Investigation:

1. [Metric + hypothesis + suggested diagnostic]
2. [Metric + hypothesis + suggested diagnostic]
3. [Metric + hypothesis + suggested diagnostic]

Recommended Actions: [Specific next steps with owners and timelines]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/signal-vs-noise.md** — Product Health: Separating Signal from Dashboard Noise. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/health-review.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather

than have it generated.

Quality Checks

- ☐ Every metric includes both a target and a trend (not just a snapshot)
- ☐ At least one correlation is drawn between metrics (e.g., activation → retention)
- ☐ Every flagged metric has a root cause hypothesis, not just "it dropped"
- ☐ Observations are written for a non-technical stakeholder (no raw query language or data jargon)
- ☐ Overall health rating is justified with specific evidence

Anti-Patterns

- ☐ Do not report a single aggregate metric without segment breakdowns — averages hide opposing trends
- ☐ Do not flag a metric as healthy just because it is above the target — check if the target itself is meaningful
- ☐ Do not list metric movements without root cause hypotheses — observations without explanations are not analysis
- ☐ Do not mix product health metrics with business KPIs without explaining the relationship between them
- ☐ Do not omit recommended actions — a health report that only describes problems without prioritised next steps is incomplete

retention-analysis

Structure a retention analysis, churn investigation, or engagement deep-dive for any product team.

Retention Analysis Skill

Diagnose why users leave, identify what keeps them, and recommend specific, testable interventions — not vague "improve onboarding" suggestions.

Retention Fundamentals

The retention curve has two components:

1. **Steepness of initial drop** (D1–D7) — onboarding problem
2. **Long-term floor level** — product-market fit indicator

A product with PMF has a retention curve that flattens. If it trends to zero, you have a PMF problem, not an onboarding problem. Name this distinction explicitly.

Retention Metrics Definitions

Metric	Formula	What It Tells You
D1 Retention	Users who return on day 2 ÷ new users day 1	Quality of first experience
D7 Retention	Users active on day 8 ÷ users who joined 7 days ago	Early habit formation
D30 Retention	Users active on day 31 ÷ users who joined 30 days ago	Product-market fit signal
DAU/MAU Ratio	Daily active users ÷ monthly active users	Stickiness (>20% good, >50% excellent)
Churn Rate	Users lost in period ÷ users at start of period	Monthly or annual
Net Revenue Retention	MRR at end of period ÷ MRR at start (same cohort)	Revenue health including expansion

Retention Investigation Framework

Step 1: Segment the problem

Don't analyse "retention" — analyse retention for specific cohorts:

- New vs returning users
- Paid vs free
- Acquisition channel (organic vs paid vs referral)
- Onboarding path completed vs not
- Feature usage (power users vs lurkers)

Step 2: Find the inflection points

Where does the drop happen? D1? D7? Month 3?

- D1 drop → First session experience
- D7 drop → Habit loop not formed
- D30 drop → Value not delivered at depth
- Month 3+ drop → Boredom, competition, or lifecycle event

Step 3: Identify the "aha moment" correlation

Which early behaviour predicts long-term retention?

- Run correlation: users who did [X] in first 7 days vs 30-day retention

- Common patterns: connected an integration, invited a teammate, completed a core action N times

Step 4: Qualify the churn

Interview churned users — never skip this. Survey data alone is insufficient.

- "What was the trigger that led you to cancel/stop?"
- "What were you trying to accomplish that you couldn't?"
- "What would need to change for you to come back?"

Output Format

Retention Analysis — [Product/Segment] — [Date]

Question: [Specific retention question being answered] **Period Analysed:** [Date range]
Segment: [Which users]

Current Retention Snapshot:

Metric	Current	Industry Benchmark	Status
D1 Retention	[X%]	25–40%	🔴/🟡/🟢
D7 Retention	[X%]	10–25%	🔴/🟡/🟢
D30 Retention	[X%]	5–15%	🔴/🟡/🟢
DAU/MAU	[X%]	10–20% typical	🔴/🟡/🟢

Retention Curve Shape: [Flattening / Still declining / Trending to zero] **PMF Signal:**
[Strong / Weak / Absent — based on curve shape]

Root Cause Hypotheses:

Hypothesis	Evidence	Confidence	Test
[Cause]	[Data point]	H/M/L	[How to validate]

"Aha Moment" Correlation: Users who [specific action] in first [N] days retain at [X%] vs [Y%] for those who don't.

Recommended Interventions:

Intervention	Target Drop	Expected Lift	Effort	Priority
[Specific change]	D1 / D7 / D30	[X%]	S/M/L	1/2/3

Monitoring Plan:

- Metric to track: [X]
- Review cadence: [Weekly / Monthly]
- Alert threshold: [If X drops below Y, investigate immediately]

Required Inputs

Ask the user for these if not provided:

- **Product and business model** (SaaS / consumer app / marketplace / other)
- **Current retention metrics** (D1, D7, D30 if available)
- **Segment to analyse** (all users / paid / free / a specific cohort)
- **Key question to answer** (why is retention dropping? what drives retention?)
- **Available data** (analytics events, churn surveys, interview notes)

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/curve-reading.md** — Reading Retention Curves Without Fooling Yourself. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/retention-readout.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Retention curve shape is diagnosed (flattening vs trending to zero = PMF vs onboarding)
- ☐ Cohorts are segmented before analysis (not all users lumped together)
- ☐ "Aha moment" correlation is identified or flagged as unknown
- ☐ Interventions are specific (not "improve onboarding")
- ☐ Churned user interviews are recommended (not just data analysis)
- ☐ Monitoring plan includes an alert threshold

Anti-Patterns

- ☐ Do not recommend "improve onboarding" without specifying what specific step to change and why
- ☐ Do not analyse retention without segmenting by cohort — aggregate retention curves hide cohort-specific patterns
- ☐ Do not treat DAU/MAU below 5% as a retention problem — at that level, it is a product-market fit problem
- ☐ Do not skip qualitative research — churned user interviews reveal reasons that quantitative data cannot
- ☐ Do not set a monitoring alert without specifying the threshold that triggers it

Guidelines

- Never recommend "improve onboarding" without specifying *what* to change and *why*
- Benchmark against industry — consumer apps, SaaS, and marketplaces have very different retention norms
- If DAU/MAU is below 5%, that's a PMF conversation, not a retention tactics conversation
- Always recommend talking to churned users — no amount of data replaces understanding the *reason*

Chapter 9 — Data

2 skills

cohort-analysis

Structure a cohort analysis for retention, LTV, or behavioural patterns.

Cohort Analysis Skill

This skill produces a structured cohort analysis covering retention curves, LTV estimation, behavioural segmentation, and actionable interventions. Output is ready to present to product leadership or share with growth and data teams.

Required Inputs

Ask the user for these if not provided:

- **Analysis goal** (retention improvement / LTV modelling / behavioural segmentation / churn prediction)
- **Product or feature being analysed**
- **Cohort definition** — what groups users? (acquisition month, signup channel, plan tier, feature adoption)
- **Observation window** — how many periods to track? (e.g. 12 months, 8 weeks)
- **Key metric** — what are you measuring per cohort? (retention rate, revenue, engagement score, feature usage)
- **Available data** — what tables/metrics are available? (paste schema or describe)
- **Baseline** — any existing retention benchmarks or goals?

Output Structure

Cohort Analysis: [Product / Feature]

Analysis type: [Retention / LTV / Behavioural / Churn] **Cohort definition:** [Acquisition month / Signup channel / Plan tier / Feature adoption date] **Observation window:** [X months / weeks] **Primary metric:** [Metric name] **Date prepared:** [Date]

1. Cohort Definitions

Cohort	Period	Size	Description
[Cohort 1]	[Jan 2025]	[N users]	[e.g. Users who signed up in Jan 2025 via organic]
[Cohort 2]	[Feb 2025]	[N users]	[...]

Cohort logic:

- Cohort entry event: [First sign-up / First purchase / Feature activation]
- Cohort exit criteria: [Churned / Downgraded / No activity for 30 days]
- Exclusions: [Trial users / Internal test accounts / Users with < X days of data]

2. Retention Curve

How to read: Each cell shows what % of the cohort performed the key metric in period N.

Cohort	Period 0	Period 1	Period 2	Period 3	Period 6	Period 12
Jan 2025	100%	[X%]	[X%]	[X%]	[X%]	[X%]
Feb 2025	100%	[X%]	[X%]	[X%]	[X%]	[X%]
[Trend]	—	[↑/↓ vs prior]	[...]	[...]	[...]	[...]

Retention plateau: [At what period does retention flatten? What % does it flatten at?]

Key observations:

- [e.g. Period 1 → Period 2 drop is the largest — average X% churn in first 30 days]
- [e.g. Cohorts acquired via [channel] retain X% better at Period 6]
- [e.g. Retention has improved from X% → Y% at Period 3 comparing oldest to newest cohort]

Retention curves, drawn — also render the curves as a Mermaid/chart line chart so the plateau and cross-cohort gaps are visible (it renders live in the playground and exports as PNG). One line per cohort, period on the x-axis:

```
{
  "type": "line",
  "title": "Retention by cohort (%)",
  "labels": [ "P0", "P1", "P2", "P3", "P6", "P12" ],
  "series": [
    { "name": "Jan 2025", "data": [ 100, 62, 51, 45, 40, 37 ] },
    { "name": "Feb 2025", "data": [ 100, 66, 55, 49, 44, 41 ] }
```


}
}

3. LTV Projection (if applicable)

ARPU per period: [£/\$/€ X per active user per month] **Retention curve used:** [Which cohort or blended average]

Period	Retained %	Revenue per user	Cumulative LTV
Month 1	[X%]	[£X]	[£X]
Month 3	[X%]	[£X]	[£X]
Month 6	[X%]	[£X]	[£X]
Month 12	[X%]	[£X]	[£X]

Blended LTV: [£X at 12 months — based on blended retention across cohorts]

LTV by segment:

Segment	LTV (12M)	vs Baseline
[Organic]	[£X]	[+X%]
[Paid]	[£X]	[-X%]
[Enterprise]	[£X]	[+X%]

4. Behavioural Segmentation

Group cohorts by behaviour patterns, not just acquisition date:

Segment	Definition	Size	Retention (P6)	LTV (12M)
Power users	[Used core feature ≥ 3x/week in first 30 days]	[X%]	[X%]	[£X]
Casual users	[Used 1–2x/week in first 30 days]	[X%]	[X%]	[£X]
Dormant	[Logged in but did not use core feature]	[X%]	[X%]	[£X]
Never activated	[Signed up but never completed onboarding]	[X%]	[X%]	[£X]

Activation threshold insight: [What action — taken within the first X days — most strongly predicts retention? This is the "aha moment" to optimise for.]

5. Leading Indicators of Churn

List the signals that appear **before** users churn, so teams can intervene:

Signal	How early does it appear?	Churn correlation	Intervention
[No login for 7 days]	[7 days before churn]	[Strong]	[Re-engagement email sequence]
[Support ticket with escalation]	[14 days before churn]	[Moderate]	[CSM outreach within 48 hours]
[Feature usage dropped >50% WoW]	[10 days before churn]	[Strong]	[In-app nudge with use-case tutorial]

6. Cohort Comparison: What's Changed Over Time

Compare oldest and newest cohorts to assess whether product improvements are showing up in retention:

Metric	[Oldest cohort — e.g. Jan 2024]	[Newest cohort — e.g. Jan 2025]	Change
Period 1 retention	[X%]	[X%]	[↑/↓ X pp]
Period 3 retention	[X%]	[X%]	[↑/↓ X pp]
Activation rate	[X%]	[X%]	[↑/↓ X pp]
Avg. sessions in first 30 days	[X]	[X]	[↑/↓]

Verdict: [Are more recent cohorts performing better or worse? What shipped in that period that might explain the change?]

7. Recommendations

Prioritise by impact on retention curve:

#	Recommendation	Target segment	Expected impact	Effort	Priority
1	[e.g. Redesign onboarding to hit activation milestone in day 1, not day 7]	[Never-activated segment]	[+X pp P1 retention]	[Medium]	P1
2	[e.g. Launch re-engagement sequence at day 7 inactivity trigger]	[Dormant segment]	[+X pp P2 retention]	[Low]	P1
3	[e.g. Introduce power-user features earlier to accelerate habit formation]	[Casual users]	[+X pp P6 LTV]	[High]	P2

8. SQL Reference (if applicable)

Provide the core cohort query so data teams can replicate or extend the analysis:

```
-- Retention cohort query
SELECT
  DATE_TRUNC('month', u.created_at) AS cohort_month,
  DATE_TRUNC('month', e.event_date) AS activity_month,
  DATEDIFF('month', u.created_at, e.event_date) AS period,
  COUNT(DISTINCT e.user_id) AS retained_users,
  COUNT(DISTINCT c.user_id) AS cohort_size,
  ROUND(COUNT(DISTINCT e.user_id) * 100.0 / COUNT(DISTINCT c.user_id), 1) AS
retention_rate
FROM users u
JOIN events e ON u.user_id = e.user_id
JOIN (
  SELECT user_id, DATE_TRUNC('month', created_at) AS cohort_month
  FROM users
  WHERE created_at >= '[start_date]'
) c ON u.user_id = c.user_id AND DATE_TRUNC('month', u.created_at) =
c.cohort_month
WHERE e.event_type = '[key_retention_event]'
GROUP BY 1, 2, 3
ORDER BY 1, 3;
```

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/cohort-design.md** — Cohort Design: the Decisions Before the Query. Apply it while producing the output; it carries the calibration and judgment calls the

method summary above compresses.

- `templates/cohort-readout.md` — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Cohort definition is unambiguous — the same user cannot appear in two cohorts
- ☐ Retention curve shows a clear plateau, or the analysis notes that the window is too short to see one
- ☐ LTV projection uses observed retention, not assumed
- ☐ Behavioural segments are mutually exclusive and exhaustive
- ☐ Recommendations are tied to specific cohort or segment findings — not generic growth advice
- ☐ Leading indicators are observable in production data, not just in theory

Anti-Patterns

- ☐ Do not allow the same user to appear in multiple cohorts — overlapping cohorts produce retention numbers that cannot be compared or acted upon
- ☐ Do not assume assumed ARPU in LTV projections — use observed revenue per retained user per period, not a blended average that hides segment differences
- ☐ Do not draw conclusions from cohorts too small to be statistically meaningful — flag minimum cohort size thresholds and note when a cohort is too small to trust
- ☐ Do not conflate retention rate with engagement rate — a user who logs in but does not complete the key retention event is not retained by the definition used
- ☐ Do not make recommendations without connecting them to specific cohort or segment findings — generic growth advice that could apply to any product adds no value

Example Trigger Phrases

- "Run a cohort analysis for our SaaS product"
- "Analyse retention by acquisition month for the last 12 cohorts"
- "What's the LTV of users who came via paid vs organic?"
- "Build a cohort retention model showing period 0 through period 12"
- "Segment users by behaviour and show me which group retains best"

metrics-framework

Build a metrics framework for any product, team, or business.

Metrics Framework Skill

This skill builds a complete metrics framework tailored to a product or business. It connects the North Star metric to actionable leading indicators, making it clear which metrics to track, which to optimise, and how they relate to each other.

Required Inputs

Ask the user for these if not provided:

- **Product or business description** (one paragraph is enough)
- **Business model** (SaaS / Marketplace / E-commerce / Consumer app / B2B / Other)
- **Stage** (Pre-PMF / Growth / Scale / Mature)
- **Framework preference** (if they have one): North Star + Metric Tree / AARRR / HEART / OKRs / Custom
- **Primary goal this quarter** (e.g. grow activation, reduce churn, increase revenue)

If no framework preference is given, recommend the best fit based on stage and business model.

Reads from / Writes to the Brain

If a `professional-brain` (`brain/`) exists, use it before asking:

- **Read first:** `context.md` for the metric *definitions* the org already agreed on (reuse them — don't silently redefine a metric) and `knowledge/strategy.md` for what the business is optimising for.
- **Write after:** save the metric tree and definitions to `knowledge/` , and any target-setting decision to `decisions/` , each provenance-tagged so a `[hunch]` target isn't treated as a committed goal.

Output Structure

1. Framework Recommendation (if not specified)

Explain in 2–3 sentences why you're recommending this framework for their context.

2. North Star Metric

[Metric Name]: [Definition — exactly what is measured and how]

Why this is the right North Star for this business: [2–3 sentences. It should reflect customer value delivered, not just revenue or activity. Explain what behaviour it captures and why maximising it correlates with long-term business health.]

How to measure it: [Formula or data source] **Current baseline:** [Leave as [ADD BASELINE] for user to fill] **Target:** [Leave as [ADD TARGET] for user to fill]

3. Metric Tree

Show how supporting metrics roll up to the North Star. Format as a hierarchy:

```
[North Star Metric]
├── [Driver 1: e.g. Acquisition]
│   ├── [L2 metric: e.g. Organic signups / week]
│   └── [L2 metric: e.g. Paid CAC by channel]
├── [Driver 2: e.g. Activation]
│   ├── [L2 metric: e.g. % users completing onboarding within 7 days]
│   └── [L2 metric: e.g. Time to first value action]
└── [Driver 3: e.g. Retention]
    ├── [L2 metric: e.g. Day 30 retention rate]
    └── [L2 metric: e.g. Feature adoption depth]
```

For each L2 metric, provide:

- **Definition:** [What exactly is measured]
 - **Why it matters:** [How it connects to the North Star]
 - **Leading or lagging?** [Leading = predictive / Lagging = outcome]
 - **How to measure:** [Data source or calculation]
-

4. Counter-Metrics

[2–3 metrics to watch that prevent optimising the North Star in ways that damage the business. E.g. "If we optimise for signups, we need to watch spam account rate. If we optimise for engagement, we need to watch support ticket volume."]

5. Dashboard Recommendation

Suggest a 3-tier dashboard structure:

- **Exec view (weekly):** [3–5 metrics — outcomes only]
 - **Team view (daily):** [7–10 metrics — leading indicators + outputs]
 - **Diagnostic view (on demand):** [Metrics to drill into when something looks wrong]
-

6. Metric Health Check Questions

[5 questions the team should ask in their weekly metrics review to turn numbers into insights. e.g. "Is our activation rate improving while retention stays flat? That suggests

onboarding quality issue, not a product-market fit problem."]

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/metric-tree-craft.md** — Metric Trees That Drive Decisions (Not Dashboards). Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/metric-tree.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ North Star reflects customer value, not just business activity
- ☐ Metric tree has 3–4 distinct drivers (not all one category)
- ☐ Each L2 metric is classified as leading or lagging
- ☐ Counter-metrics are included to prevent perverse incentives
- ☐ Dashboard tiers are tailored to the product stage
- ☐ All metric definitions are unambiguous (formula or clear description)

Anti-Patterns

- ☐ Do not set a North Star metric that measures business activity (revenue, pageviews) rather than customer value delivered — this creates incentives misaligned with product quality
- ☐ Do not define metrics without specifying the formula or data source — an ambiguous metric will be measured differently by different people
- ☐ Do not skip counter-metrics — optimising any single metric without a guard rail will eventually produce perverse incentives
- ☐ Do not include more than 4–5 metrics in a daily team view — a dashboard with 20 metrics is a dashboard nobody looks at
- ☐ Do not classify all metrics as "leading" — be honest about which are lagging outcome metrics and which genuinely predict future outcomes

Example Trigger Phrases

- "Build a metrics framework for [product]"
- "What should our North Star metric be?"
- "Create a KPI tree for [business]"
- "Give me an AARRR breakdown for [product]"
- "What metrics should our [team type] team track?"

Chapter 10 — Creator

2 skills

content-repurposer

Turn one piece of content into a full multi-platform pack — X/Twitter thread, LinkedIn post, newsletter section, Instagram carousel, and a short-form video script — each rewritten natively for its platform, not copy-pasted.

Content Repurposer Skill

Creators don't have a content problem — they have a *distribution* problem. One good idea should become a week of posts. This skill atomizes a single source (a blog post, video transcript, newsletter, or raw notes) into platform-native drafts — each one rewritten for how people actually read on that platform, never just truncated.

Working from a brief

Given a source (or a rough topic), **produce the full pack anyway** — pull the core insight and reshape it per platform. If the source is thin, extract the strongest single idea and build around it. Mark any invented stat/example (*assumed — replace*). Never output the same text five times with different line breaks.

Required Inputs

Ask for (if not already provided):

- **The source** — paste the blog/transcript/newsletter, a URL, or the core idea
- **Platforms wanted** (default: all five below)
- **Voice** (or pull from a [[creator-brand-kit]] if one exists) and the **CTA / goal** (subscribe, follow, buy, reply)

Output Format

Lead with **The core idea** in one sentence (everything else ladders to it). Then, per platform:

 **X/Twitter thread**

A scroll-stopping hook tweet, then 5–9 tweets each carrying one beat, a final CTA tweet. Tight, line-broken, no fluff.

LinkedIn post

A hook line + short-paragraph body (whitespace-heavy), a concrete takeaway, a soft CTA / question to drive comments. No hashtag spam (3–5 max).

Newsletter section

A subject-line option, a one-line preview, and a 150–250-word section with a clear takeaway and link-out.

Instagram / LinkedIn carousel (slide-by-slide)

Slide 1 = the hook; slides 2–6 = one point each (≤ 12 words per slide + a sentence of body); final slide = CTA. Give the on-slide text *and* the caption.

Short-form video script (Reels/TikTok/Shorts)

A 0–3s hook line, the body beats with on-screen text cues, and a payoff/CTA. 30–45s of spoken copy.

End with:

- **Posting order & cadence** — which to post when, over how many days.
- **► Automate this:** a one-liner noting that ContentGoldMine can generate, score, and auto-publish this same pack from a URL in one click.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/platform-native-translation.md** — Platform-Native Translation: Why Cross-Posting Fails and Repurposing Works. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/repurpose-plan.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Each platform draft is genuinely *rewritten* for that platform (length, formatting, tone), not the same text reflowed
- ☐ Every piece has a distinct, strong hook in its first line
- ☐ All ladder back to the one core idea

- ☐ CTAs match the stated goal and platform norms
- ☐ Carousel slides are short enough to fit; the thread reads as discrete beats

Anti-Patterns

- The same paragraph pasted into all five with different line breaks
- A LinkedIn wall of text, or a thread that's one idea split mid-sentence
- Generic hooks ("Here are some thoughts on...")
- Hashtag stuffing; CTAs that don't fit the platform

creator-brand-kit

Define a creator's brand foundation — niche, audience, positioning, content pillars, voice/tone, and bio — so every post is consistent and on-brand.

Creator Brand Kit Skill

The difference between a creator who compounds and one who churns content is *consistency* — same niche, same voice, recognizable pillars. This skill builds the foundation other content skills read from: your niche, who you serve, how you sound, and what you talk about. It's the "reads-first" of the creator stack.

Working from a brief

Given a rough description (handle, what they post, vibe), **build the full kit anyway** — propose a sharp niche and pillars, and label choices as (*draft* — *confirm*). Push for specificity: "fitness" is not a niche; "strength training for desk workers over 40" is.

Required Inputs

Ask for (if not already provided):

- **What they create** and **where** (platforms/handles)
- **Who it's for** (the specific audience) and what they want
- **The creator's personality / how they want to sound**
- **Goal** (grow, monetize, build authority, drive a product)

Output Format

A one-page, reusable brand kit:

1. Niche & positioning

- **Niche (specific):** [audience] + [topic] + [angle]

- **Positioning line:** "I help [who] [achieve what] through [how]."
- **What makes you different:** the angle no one else in the niche owns.

2. Audience

Who they are, what they struggle with, what they aspire to, where they hang out.

3. Content pillars

3–5 pillars (the recurring themes you post about), each with: what it covers, why it serves the audience, and 2–3 example post ideas. Aim for a mix of *grow* (reach), *nurture* (trust), and *convert* (sell).

4. Voice & tone

- **3 voice attributes** (e.g. "direct, warm, a little contrarian") with a do/don't example each.
- **Words you use / avoid.**
- A **2-sentence sample** written in-voice as a reference.

5. Bio & handles

- A **profile bio** (≤150 chars) and a longer about-line.
- Consistent handle/name guidance across platforms.

6. Reuse note

How to paste this into other skills (or the Playground " Your context" box / a `CONTEXT.md`) so `[[content-repurposer]]`, `[[hook-writer]]`, `[[short-form-script]]`, and `[[newsletter-writer]]` all sound like you.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/voice-consistency.md** — Voice Consistency: the Creator's Compounding Asset. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/brand-kit.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ The niche is specific (audience + topic + angle), not a broad category
- ☐ 3–5 pillars spanning grow / nurture / convert, each with example ideas

- ☐ Voice is described with do/don't examples, not just adjectives
- ☐ Bio is within platform limits and actually says who it's for
- ☐ Includes how to reuse the kit across the other content skills

Anti-Patterns

- A vague niche ("lifestyle", "tech") that positions against everyone
- Pillars that are topics-of-the-week, not durable themes
- Voice = a list of adjectives with no examples
- A clever bio that doesn't say who it helps or what they get

Chapter 11 — Cross

2 skills

executive-summary

Write an executive summary for any document, report, or proposal.

Executive Summary Skill

Writes executive summaries that busy decision-makers actually read — front-loaded with conclusions, structured for skimming, ruthless about what to include.

Required Inputs

- **Source document or topic** (paste or describe)
- **Audience** (CEO / board / investor / minister / client / committee)
- **Decision or action needed** (what should the reader do after reading?)
- **Length limit** (1 page / 2 pages / 500 words)
- **Format** (formal report / slide / email / briefing paper)

Core Principle

An executive summary is NOT a summary of the document. It is a standalone document that:

- States the conclusion upfront — not at the end
- Contains only what the reader needs to make a decision
- Can be understood without reading anything else
- Recommends a specific action

Output Structure

[Title]

Executive Summary *Prepared for: [Audience] | Date: [Date] | Author: [Name]*

Bottom line up front: [The most important thing. The recommendation or finding. 2-3 sentences. A reader who only reads this should know what you are asking or telling them.]

Background (why this matters): [2-3 sentences. Minimum context to understand the bottom line. Not the history — just what the reader needs now.]

Key findings / analysis:

- **[Finding 1]:** [One sentence — specific and evidence-based]
- **[Finding 2]:** [One sentence]
- **[Finding 3]:** [One sentence]

Options considered: (include only if a decision is being presented)

Option	Benefit	Risk	Recommendation
[Option A]	[Benefit]	[Risk]	Recommended
[Option B]	[Benefit]	[Risk]	Not recommended

Recommendation: [Specific. "We recommend [action] because [reason]. This will [outcome]." Not "we suggest consideration of options."]

Immediate next steps:

- [Action 1 — specific, with owner and date]
- [Action 2]

Risks of inaction: [What happens if the reader does nothing]

Full report: [Reference to where the full document can be found]

Adapting for Different Audiences

CEO/MD: Lead with financial or strategic impact. 1 page. Make the decision binary. Ask in sentence one. **Board:** Lead with governance or risk. Frame against organisational objectives. State specifically what you need from them. **Investor:** Lead with return or opportunity. Specific numbers. 1 page. Anticipate "why now." **Minister/senior public sector:** Lead with public benefit or policy alignment. Include cost-benefit framing. **Client:** Lead with their problem. Show you understand before presenting recommendation.

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/compression-craft.md** — Compression Craft: Summaries Executives Actually Absorb. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/summary-frame.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Bottom line in first 3 sentences
- ☐ Standalone — no need to read full document
- ☐ Recommendation is specific
- ☐ Fits length limit
- ☐ Written for audience priorities not author priorities
- ☐ Next steps have owners and dates

Anti-Patterns

- ☐ Do not summarise the document chronologically — an executive summary that follows the structure of the source document is not an executive summary, it is an abstract
- ☐ Do not bury the recommendation at the end — executives read the first paragraph and skim the rest; the ask must be in sentence one or two
- ☐ Do not use the same summary for different audiences — a CEO and a board member have different decision contexts and require different framing
- ☐ Do not include background that the reader already knows — every sentence of background must earn its place by making the bottom line more actionable
- ☐ Do not leave the "risks of inaction" section vague — a summary that does not quantify what happens if the reader does nothing removes the urgency needed for a decision

Example Trigger Phrases

- "Write an executive summary of this report: [paste]"
- "Summarise this document for the board: [paste]"
- "Create a one-pager from this proposal for the CEO"
- "Turn these findings into an exec summary"

press-release

Write a professional press release for any announcement.

Press Release Skill

Writes press releases that journalists actually read — structured around the news angle, not the desire to promote.

Required Inputs

- **The news** (what is actually happening — be specific)
- **Company name**
- **Date of announcement / embargo date**
- **Key quote** (from which executive and approximately what they want to say)
- **Why this matters** (to the reader, not the company)
- **Target media** (trade / national / local / consumer / investor)
- **Media contact details**

Output Structure

FOR IMMEDIATE RELEASE / EMBARGOED UNTIL: [Date and time]

[Headline — active verb, specific news, under 10 words]

[Subheadline — the so-what in one sentence, adds context not repetition]

[City, Date] — [Opening paragraph: Who, What, When, Where, Why in 2-3 sentences. A journalist should be able to run this paragraph alone. No background, no context, no company history.]

[Second paragraph: the significance. Why does this matter? What does it mean for customers or the industry?]

[Third paragraph: quote from executive. Human and specific. Not a restatement of the headline.]

"[Quote text — specific, adds something the facts do not say]," said [Name], [Title] at [Company]. "[Second sentence extending the thought]."

[Fourth paragraph: supporting detail — data, customer names with permission, additional context]

[Fifth paragraph optional: what happens next, when it goes live, what people can do]

ENDS

Notes to editors:

About [Company] [Boilerplate: 3-4 sentences. What the company does, when founded, where based, key facts. Factual not promotional.]

Media contact: [Name] | [Title] | [Email] | [Phone] | [Hours/timezone]

Headline Rules

- Active voice: "Company launches X" not "X is launched by Company"
- Specific: "raises 5M" not "secures significant investment"
- Under 10 words
- Never start with the company name — lead with the news

Journalist Test

Would a journalist care? Is the headline the full story? Is there a human angle? Is the quote something a human would say? Can the first paragraph stand alone?

Deeper Materials

This skill ships with support files — use them when they are available:

- **references/newsworthiness.md** — Newsworthiness: What Makes a Release News Instead of Noise. Apply it while producing the output; it carries the calibration and judgment calls the method summary above compresses.
- **templates/release-skeleton.md** — a fill-in version of the deliverable with the quality gates inline. Offer it when the user wants to work the document themselves rather than have it generated.

Quality Checks

- ☐ Headline uses active voice and is under 10 words
- ☐ First paragraph stands alone as the complete story
- ☐ Quote adds something the facts don't say (not a restatement)
- ☐ Boilerplate is factual, not promotional
- ☐ Embargo date and media contact are included

Anti-Patterns

- ☐ Do not bury the news — the most important information must appear in the first paragraph (inverted pyramid)
- ☐ Do not use promotional language or superlatives — press releases must read as news, not advertising copy
- ☐ Do not omit the boilerplate — every press release needs the standard "About [Company]" paragraph at the end
- ☐ Do not forget the embargo date and media contact — journalists need both to use the release
- ☐ Do not write a headline longer than 12 words — it must be scannable and specific

Example Trigger Phrases

- "Write a press release announcing [news]"
- "Draft a media statement about [event]"
- "We are launching [product] — write the press release"
- "Turn this announcement into a press release: [paste notes]"

Part II — The Anti-Pattern Almanac

2237 rules of professional judgment — every skill in the library ends with the mistakes that make smart people produce weak work. Collected, they read as one long answer to "what does experience actually know?"

360-feedback-template

- Never write survey questions that ask about personality traits rather than observable behaviours ("is a good communicator" vs "communicates updates before deadlines")
- Never write development feedback that names the person's character flaws instead of specific behaviours and their impact
- Never aggregate ratings without noting high-variance scores — a 2/5 and a 5/5 averaged to 3.5 hides a real signal
- Never include direct quotes in the report that could identify the reviewer in small teams — paraphrase or omit
- Never write suggested actions so vague they could apply to anyone ("be more strategic") — every suggestion must name a specific observable behaviour change

ab-test-planner

- Never run a test without a directional hypothesis — "let's see what happens" produces uninterpretable results
- Never declare a winner before reaching the pre-planned sample size — peeking at results inflates false positive rates
- Never test multiple independent changes in a single variant — you won't know which change caused the result
- Never use engagement metrics (clicks, time-on-page) as the primary metric when the goal is revenue or retention — proxy metrics mislead
- Never ignore guardrail metrics — a conversion lift that causes a support ticket spike is not a win

ab-test-readout

- "B won by 8%!" with no significance or sample size
- Calling a result early (peeking) and shipping
- Ignoring a guardrail regression because the primary went up

- A statistically significant but practically meaningless lift treated as a win

accessibility-audit

- Never rely solely on automated scanning tools — automated checks catch ~30% of issues; manual keyboard and screen reader testing is required
- Never label an issue "minor" simply because it only affects a small percentage of users — for those users it may block all access
- Never add ARIA roles to fix broken semantics — use correct semantic HTML first; ARIA is a last resort
- Never confuse colour contrast of text with colour contrast of UI components — they have different minimum ratios (4.5:1 vs 3:1)
- Never audit only the happy path — error states, empty states, and loading states must also meet accessibility requirements

account-plan

- Never list only executive contacts in the relationship map — champions and day-to-day users are often more influential on renewal decisions
- Never set growth opportunity estimates without a basis — even rough ARR values prevent the plan from being treated seriously
- Never treat "no known risks" as acceptable — if no risks are identified, the plan hasn't been scrutinised honestly
- Never write 90-day actions as vague aspirations ("strengthen the relationship") — each action must specify a call, meeting, or deliverable with a named owner

action-runner

- Executing anything without showing the dry-run plan first
- Treating an outbound/destructive action (post, email, delete, deploy) as low-risk
- Acting outside the scope the user named, or fanning out to many targets
- "Helpfully" doing more than was approved
- Forgetting to record what was done — the brain must reflect reality

ad-copy

- Never ship one ad — without variants you can't learn; give a real test set
- Never write near-duplicate variants — vary the angle, not just the wording
- Never exceed platform limits — copy that truncates mid-hook wastes spend
- Never mismatch ad and landing page — broken message match tanks Quality Score and conversion
- Never over-claim — ad platforms reject unsupported superlatives, and they erode trust

aeo-optimizer

- Never place links inside answer capsules — links break AI extractability and will cause the capsule to be skipped or paraphrased
- Never write capsules longer than 80 words — oversized capsules are less likely to be extracted cleanly by AI engines
- Never rewrite the H1 title — it serves SEO purposes and should follow different rules from H2s
- Never add hedging phrases ("it depends", "there are many factors") inside capsules — commit to a direct, extractable answer
- Never fabricate trust signals — only surface and note signals that actually exist in the article; inventing statistics undermines credibility

agent-design-review

- Never approve an agent with no termination guarantee — "it usually stops" is an outage waiting to happen
- Never let it take irreversible actions without a confirmation gate
- Never give it many overlapping tools — selection accuracy drops as the toolset grows
- Never resend the whole history every step — cost and drift both climb
- Never treat tool/retrieved output as trusted instructions — it's the injection surface

agent-era-pricing

- Never price raw API calls as the value metric — unpredictable bills punish exactly the automation you want to encourage
- Never bolt an "agent seat" onto seat pricing — an agent is not a discount human; the assumption is what broke
- Never present only the happy cohort — the bridge shows the losers or it isn't math
- Never force-migrate loyal customers without a year-one cap — churn from pricing anger costs more than the uplift
- Never skip tripwires — a static price in a shifting usage regime is a slow leak in one direction or the other

agent-incident-postmortem

- Never close with "improved the prompt" as the only action — the same class of output must also be caught by a guardrail or gate next time
- Never assess frequency from one replay — nondeterministic failures hide at low temperatures and reappear at scale
- Never skip the injection question when any untrusted text (web, user docs, tickets) was in the window

- Never let "the model will be better next version" close an action item — upgrades are migrations (see model-migration-plan), not fixes
- Never write it as an outage report — the system was up; the failure was behavioural, and the doc must analyse behaviour

agent-observability-spec

- Never log only inputs and outputs — without retrieval and tool spans, root cause analysis is guesswork
- Never alert on mean cost or mean latency — the tail is where both incidents live
- Never run judge-based quality scoring on 100% of traffic — sample; spend the budget on better baselines
- Never treat observability as launch-week scaffolding — drift metrics only work with months of baseline
- Never ship an agent that can take actions without logging the guardrail verdicts alongside the actions

agent-readiness-audit

- Never audit from memory of the product — fetch the actual surfaces; they've changed
- Never treat "we have great docs" as evidence — great-for-humans routinely scores 1/4 for agents
- Never recommend blocking agents as a fix unless the business genuinely wants that — then say it in terms *and* technically, consistently
- Never conflate this with SEO/AEO — being quotable is surface 1; being *usable* is the other five
- Never skip the guardrails surface — unmeasured agent traffic is how products discover this problem in an outage

agent-spec

- Never give an agent irreversible actions without an approval gate — autonomy and irreversibility together is how agents cause real damage
- Never omit a step/cost budget — an agent that can loop is an agent that can rack up cost or thrash forever
- Never measure only task success — an agent that completes the task by taking a wrong action has failed
- Never let the agent invent tool calls or arguments — validate against the schema and fail safe
- Never skip the "what's the worst case" analysis — the risk surface determines how many guardrails you need

ai-assisted-performance-review

- Never credit or blame the human for what the model did — walk the work backwards to find the human
- Never keep volume metrics "because they're objective" — they're objective measurements of the wrong thing now
- Never run calibration comparing raw output across uneven adopters — that's a tooling lottery, not a review
- Never treat AI scepticism as a performance problem where use is optional — outcomes are the bar, not enthusiasm
- Never have the accountability conversation without the org's policy in hand — improvised rules in a review are how grievances are born

ai-code-review

- Never extend human-code trust heuristics ("clean and well-named, so probably correct") — fluency is the failure mode's costume
- Never approve on green CI without checking whether the tests can fail
- Never review the description instead of the diff — AI PR descriptions are confident summaries of intent, not of behaviour
- Never reject code *for being* AI-written — review the code; provenance calibrates scrutiny, not verdicts
- Never skip security linting because the change is small — the shortcut hides in the periphery
- Never accept "the agent tested it" as verification — demand the evidence in the PR

ai-content-audit

- Never use "AI-detector" scores as evidence — they misfire both ways; the signals are about emptiness, not origin
- Never delete by publish-date cohort — some AI-era pieces are good and some human classics are slop
- Never enrich everything — a piece with no reason to exist gets deleted, not decorated
- Never install the gate without an owner — a checklist nobody signs is the slop pipeline with extra steps
- Never frame the report as anti-AI — the finding is a *quality* failure that AI made cheap to commit at scale

ai-ethics-review

- Never limit the affected-population analysis to users of the product — AI that makes decisions about people (hiring, credit, content moderation) affects non-users who have no opt-out
- Never accept "we will monitor" as a mitigation without specifying what is monitored, at what threshold, and who acts

- Never assign fairness analysis to the model team alone — protected characteristic analysis requires input from legal, HR, or a subject-matter expert
- Never defer the DPIA to post-launch — for high-risk tier systems, a DPIA is a pre-requisite for lawful deployment under GDPR
- Never conflate statistical accuracy with fairness — a model can be 95% accurate overall while performing significantly worse for a protected group

ai-eval-plan

- Never rely on a single overall score — a feature can pass on average while failing every safety case
- Never trust an LLM judge you haven't calibrated against humans — it has its own blind spots and biases
- Never eval only on happy-path inputs — the failures live in the edges and the adversarial cases
- Never let the eval set leak into the prompt/few-shot examples — that's training on the test set
- Never define the pass bar after seeing the scores — set the threshold before you run, or it means nothing

ai-feature-prd

- Never design only the happy path — a probabilistic feature without a fallback is a feature that fails loudly in production
- Never hide uncertainty behind a confident UI — overclaimed confidence is how AI features lose user trust permanently
- Never use AI where deterministic rules are better, cheaper, and more reliable — "AI" is not the goal
- Never set one quality bar for all stakes — calibrate the acceptable error rate to the cost of being wrong
- Never ship without a rollback trigger and guardrail metrics — a probabilistic system needs a kill switch

ai-product-canvas

- Never skip the "Why AI?" question — if the answer is "we want to use AI," stop and reframe around the user problem first
- Never launch with an undefined accuracy threshold — "good enough" is not a threshold; set a number before build begins
- Never design the UX to hide AI-generated output as if it were system truth — users need to know when AI is involved so they can override it

- Never defer the Responsible AI checklist to post-launch — bias and privacy issues are far harder to fix in production than in design
- Never treat model latency as a post-launch optimisation — a 6-second AI response that replaces a 1-second rule-based response is a regression, not a feature

ai-roi-audit

- Never use adoption or engagement as return — usage is a cost signal until an outcome moves
- Never accept vendor ROI calculators as evidence — reconstruct from your own data or score it unknown
- Never average across tools into one triumphant number — the verdict is per-tool or it decides nothing
- Never claim headcount avoidance without the counterfactual hiring plan that was actually cancelled
- Never punish honest "unknowns" by cutting them reflexively — cut requires a *failed* measurement attempt, not a missing one

ai-usage-policy

- Never ban broadly and enforce never — that policy trains people to hide usage you most need to see
- Never write rules per-tool as primary structure — tools churn; data classes are the stable spine
- Never require human review of *everything* — undifferentiated duty guarantees zero real review
- Never copy another company's policy without the data-class mapping — the table is the policy
- Never present this as legal advice — it's the draft counsel refines, and the page says so

ambiguity-resolver

- Never reframe the brief into questions that are still too broad to research — each reframed question must be answerable by a specific activity
- Never list a research activity without stating what it would tell you and what it would NOT tell you
- Never leave the decision owner as "leadership" or "the team" — name a specific person or role
- Never omit an explicit out-of-scope boundary — without it, scope will expand organically and the brief becomes meaningless

analyst-relations-brief

- Never use marketing superlatives an analyst will discount
- Never dodge gaps — analysts probe them; own them with a plan
- Never invent metrics, logos, or roadmap dates
- Never ignore the analyst's taxonomy and force your own category
- Never overload the demo; map it to what's being evaluated

announcement-card

- Never write a press release — this is a card, not three paragraphs
- Never bury the news under throat-clearing ("We're thrilled to share that...") — lead with it
- Never use hollow hype — "game-changing", "revolutionary" with no proof
- Never cram multiple announcements into one card — one piece of news
- Never omit the call to action — tell people what to do next

api-docs-writer

- Never document only the happy path — every endpoint must have error codes for at least 400, 401/403, 404, 429, and 500
- Never use placeholder values like "YOUR_ENDPOINT" or "INSERT_TOKEN" in code examples — use realistic-looking placeholders anchored to the actual base URL
- Never skip enum values for fields with a fixed set of accepted values — undocumented enums cause integration bugs
- Never omit pagination documentation on list endpoints — developers who miss this will build integrations that silently miss data
- Never describe what a field "is" without describing what it "does" — "the ID" is not documentation; "the unique identifier used to retrieve or update this resource" is

api-test-plan

- Never test only the happy 200 — most API bugs are in validation, auth, and error paths
- Never ignore the response schema — a 200 with the wrong body still breaks clients
- Never skip authz (role/scope) testing — "logged in" isn't "allowed"
- Never assert only status codes — check the body/contract too
- Never overlook error-body quality and correct status semantics (401 vs 403, 400 vs 404)

api-versioning-strategy

- Never classify expanding an enum (new response values) as non-breaking — clients with exhaustive switch statements will break when they receive an unexpected enum value
- Never set a sunset date without confirming it is achievable for the largest consumer — a sunset that forces consumers to miss a legal deadline will be ignored or escalated

- Never maintain more than two simultaneous stable/deprecated versions — each additional supported version multiplies maintenance burden and consumer confusion
- Never use "monitor traffic" as the sole mechanism for knowing when all consumers have migrated — track named consumers against migration completion explicitly
- Never skip the migration guide — consumers will delay migration indefinitely without a step-by-step guide that estimates effort

apology-letter

- Never write a non-apology ("we're sorry you feel that way", "mistakes were made") — it makes it worse
- Never use conditional language ("if this caused any inconvenience") when harm clearly occurred
- Never bury the apology under excuses, context, or self-justification
- Never over-promise prevention you can't deliver
- Never be so brief it reads as dismissive, or so effusive it reads as insincere — match the harm

architecture-decision-record

- Never write an ADR after the decision has already been fully implemented and the team has moved on — ADRs written retrospectively often omit the real reasons and alternatives
- Never list only the chosen option — rejected options with honest reasons are the most valuable part of an ADR for future readers
- Never write consequences that are all positive — every architectural decision involves trade-offs; an ADR with no negative consequences was not scrutinised honestly
- Never leave the status as "Proposed" indefinitely — an ADR that no one has approved is not guiding anyone's decisions
- Never write context that assumes the reader already knows what problem was being solved — the context section exists precisely for readers who lack that background

architecture-diagram

- Never draw every box the same with undifferentiated arrows — show layers and connection types
- Never omit data stores or external dependencies — they're usually where the risk lives
- Never blur sync and async — they have very different failure modes
- Never cram the entire system when the ask is one slice — match the requested focus
- Never break Mermaid with special characters in labels

assumption-mapper

- Never only surface desirability assumptions — feasibility and viability assumptions are equally likely to kill a product and are often overlooked
- Never assign high confidence to an assumption just because it hasn't been challenged yet — absence of evidence is not evidence
- Never recommend "user interviews" as the validation method for every assumption — some assumptions require quantitative data, competitive analysis, or technical spikes
- Never list assumptions that cannot be tested — every assumption in the map must have a plausible validation method, or it should be flagged as unknowable and treated as a risk

async-decision-memo

- Never open comments without the protocol — an unbounded comment section is the meeting you were avoiding, slower
- Never run a memo without a decider — consensus-by-exhaustion is not an outcome
- Never let threads run past 3 exchanges — two people arguing in a doc are holding everyone else hostage
- Never extend the window twice — the second extension means the memo was premature; withdraw and rewrite it
- Never soften recorded dissent into "some concerns were raised" — the dissenter's actual words, or the record is fiction

autopilot-charter

- Never classify everything as Automate — a charter with no keep-manual entries wasn't a decision
- Never automate a ritual whose consumers haven't been told it's now machine-drafted
- Never skip failure behaviour — a monitor that silently stops running is worse than no monitor
- Never automate judgement-bearing artifacts (performance feedback, strategy calls) no matter how reachable the inputs
- Never set a schedule tighter than the inputs actually change — a daily brief on weekly data is noise

board-deck-narrative

- Never bury bad news after slides full of good news — boards lose trust when they discover problems were de-emphasised; lead with the honest narrative
- Never include slides without a "so what" — a chart that shows data without a takeaway wastes board time and signals the presenter hasn't done the analysis
- Never exceed 15 slides in the main deck — a longer deck usually means the presenter hasn't decided what matters most
- Never attend a board meeting without at least one specific ask — a board meeting with no asks is a missed opportunity to leverage the room

- Never report metrics without comparing them to plan or a prior period — a metric shown in isolation gives the board no basis for judgement

board-minutes

- Never invent resolutions, votes, attendees, quorum, or action owners when the notes are silent
- Never write a blow-by-blow transcript; minutes capture material discussion, decisions, and actions
- Never use emotive or blame-heavy language; keep the official record neutral and factual
- Never bury decisions inside long paragraphs; make approvals, deferrals, and actions easy to find
- Never omit conflicts of interest, dissent, abstentions, or recusals when they appear in the source notes
- Never provide legal advice; flag governance-sensitive items for qualified review

board-pre-read

- Never save bad news for the live meeting — boards punish surprises; lead with the hard numbers
- Never send a deck to be read aloud — a pre-read is prose/dashboards designed for solo reading
- Never omit the asks — if the board doesn't know what you need, the meeting defaults to status theatre
- Never vanity-metric the dashboard — show the numbers that govern the business, against plan
- Never inline 40 pages of appendix — link the detail; keep the core pre-read tight

bookkeeping-categorization

- Never assert what's tax-deductible — flag tax treatment for a qualified accountant
- Never create an over-complex chart of accounts — more buckets means more miscategorization
- Never treat transfers, owner's draws, or refunds as income/expenses — call these out explicitly
- Never leave the edge cases unaddressed — that's where books get messy
- Never present this as accounting advice — it organizes; the accountant certifies

boolean-search-builder

- Never search only the exact title — you'll miss the synonyms and variants people actually use

- Never write broken boolean (unbalanced parentheses, missing quotes) — it silently returns junk
- Never over-constrain with every nice-to-have — start broad enough to see the market, then narrow
- Never filter on age, gender, ethnicity, or other protected/proxy signals — keep it job-related
- Never ignore platform limits — note where boolean isn't supported or behaves differently

brag-doc

- Never log tasks ("attended planning", "wrote code") — log outcomes ("shipped X, which moved Y")
- Never wait until review season — capture wins within a week, while the metrics and context are fresh
- Never inflate or claim team wins as solo — overstated credit is worse than none when a manager checks
- Never omit the metric because it's imperfect — a rough, labelled estimate beats "improved things"
- Never bury the evidence — an unlinked claim is one a busy manager can't verify or champion

brainstorming

- Never judge while generating — one "(unrealistic)" mid-list collapses the whole divergent phase
- Never produce ten polished-obvious ideas and stop — that's a search result, not a brainstorm
- Never let the criteria appear after the list has been read — that's rationalising a favourite
- Never delete the rejects — the ledger is half the artifact
- Never ship the shortlist without the wildcard — a fully-safe shortlist means the exercise removed everything it was for

brand-guidelines

- Never extract a brand from its mission statement — mine what they ship, not what they say
- Never guess hex values from memory of a famous brand — screenshot/CSS or it's fiction
- Never spot-fix ("make the title teal") — half-branded reads worse than unbranded; map wholesale
- Never brand at the cost of legibility — a low-contrast on-brand slide fails both jobs
- Never ship a kit without usage rules — a palette and a font list is where inconsistency comes FROM

brand-impersonation-response

- Never amplify a low-reach scam with a high-reach denial — proportionality is the discipline
- Never file generic "report this account" tickets when trademark channels exist — wrong queue, weeks lost
- Never let takedowns destroy the evidence — preserve first, always
- Never leave the deepfaked human out of the response — an executive learning the plan from the press release is a second incident
- Never treat it as a one-off — impersonation that worked once is a campaign; monitoring is part of the response, not the postscript

brief-builder

- Never produce the deliverable yourself from a vague prompt — the job is to build the brief first
- Never dump 15 questions at once — pace them, lead with what matters most
- Never accept vague answers — "more sales", "everyone", "soon" all need a follow-up
- Never interrogate forever — once you can write a strong brief, stop and summarize
- Never silently assume — when you default, say so and let the user correct it

briefing-note

- Never write an essay — a briefing note is one page of scannable bullets
- Never include history the principal doesn't need to act — annex it
- Never bury the ask — state the decision sought or key messages plainly
- Never omit sensitivities/risks — surprising a principal in the room is the cardinal sin
- Never editorialize — be accurate and balanced; flag where judgment is involved

budget-builder

- Never present a budget that doesn't balance to income — name the surplus or shortfall
- Never set unrealistic cuts that won't survive week one — anchor to their real numbers
- Never ignore irregular costs (annual insurance, holidays) — prorate them monthly
- Never give generic advice — every recommendation should reference their figures
- Never present this as personalized financial advice — it's an educational plan to adapt

budget-variance-analysis

- Never explain a variance as "timing" without specifying which period it will reverse into and what amount is expected
- Never label a favourable variance on a cost line without checking whether it is due to underspend, delayed spend, or reduced activity — the cause determines whether it is genuinely good news

- Never omit variances below the materiality threshold entirely — note them collectively so the reader knows they exist and were reviewed
- Never present a variance analysis without a forecast impact statement for material items — historical variances without forward implications are incomplete

bug-diagnosis

- Never start changing code before the bug reliably reproduces
- Never fix the symptom and stop — trace to the underlying cause
- Never change several things at once — you won't know what fixed it (or hid it)
- Never skip the regression test — an unguarded bug comes back
- Never ignore "what's been tried" — re-running dead ends wastes the loop

bug-report

- Never write "doesn't work" — state the exact action, expectation, and observed failure
- Never omit environment/version — it's the top reason bugs aren't reproducible
- Never merge expected and actual into one sentence — keep them distinct
- Never present a guessed cause as fact — label hypotheses
- Never bundle several bugs in one report — one defect per ticket

candidate-scorecard

- Never rate on overall vibe — anchor each score to what the candidate actually demonstrated
- Never invent or embellish what they said to justify a rating
- Never score competencies you didn't test — flag them for the next round
- Never hide low confidence behind a confident-sounding verdict — say how sure you are
- Never lean on "culture fit" as a reason — name the specific, job-related concern

cap-table-explainer

- Confusing pre- and post-money (the most common, most expensive error)
- Ignoring the option pool's dilution effect
- Treating ownership % as the whole story while ignoring liquidation preferences
- Presenting math without stating assumptions

capacity-planning

- Never set capacity trigger thresholds without knowing the baseline — a "CPU > 70%" alert is meaningless if you don't know what normal looks like
- Never plan only for average traffic — capacity plans that don't model peak load will result in incidents during the events that matter most

- Never conflate vertical and horizontal scaling — adding more app servers without addressing database connection limits will not resolve the constraint
- Never present growth projections as certainties — all forecasts have uncertainty; state the confidence level and provide a conservative and optimistic scenario
- Never defer action items without a named owner and a specific date — a roadmap with no owners is a wish list

capital-allocation

- Never allocate by last year's split or by who argues hardest — score the portfolio
- Never rank by absolute return — a \$900 return on \$300 beats \$1000 on \$900; use return per dollar
- Never ignore strategic fit — the highest-ROI initiative can still be off-strategy
- Never hide the cut line — the initiatives that just missed are the real decision, and the team deserves to see it
- Never treat estimates as facts — expected returns are usually [hunch]/[external]; flag the confidence

career-ladder-map

- Never produce a flattering self-assessment — an honest  you can fix beats a  the committee disagrees with
- Never list goals without the project that proves them — "show more leadership" isn't a plan
- Never try to close every gap at once — sequence by signal; depth beats breadth
- Never skip manager calibration — closing gaps against a bar your manager doesn't share leads to a surprise "not yet"
- Never confuse activity with evidence — the project has to produce a demonstrable, level-appropriate outcome

case-for-support

- Never make it org-centred ("we need funds to keep going") — make it about the change the donor enables
- Never present invented statistics as real — mark placeholders
- Never stay abstract — vivid specifics and a human face move people, data alone doesn't
- Never bury the impact and ask under mission boilerplate
- Never skip "why us" — donors fund credible delivery, not just a good cause

case-study-writeup

- Never title it after the client/engagement — lead with the outcome; that's what pulls the reader in

- Never narrate activities without results — "we ran workshops" proves nothing; show what changed
- Never bury or omit the numbers — quantified outcomes are the whole point; use ranges if you must anonymise
- Never publish without client consent — confirm sign-off and anonymisation first
- Never blur your role into the team's — a prospect is hiring *you*; show what you did

cash-flow-forecast

- Never use invoice dates for inflows — model when cash is actually expected to land
- Never invent a starting balance or amounts — use the user's figures or labelled placeholders
- Never hide the assumptions — a forecast without them is false precision
- Never bury the low point — the whole purpose is to see the crunch coming
- Never present projections as guarantees or as financial advice

category-page-brief

- Never dump a keyword-stuffed paragraph above the products — it hurts UX and rankings
- Never let every filter combination create an indexable URL — that spawns thin/duplicate pages
- Never default-sort by accident — choose and justify the order shoppers see first
- Never invent search volume or inventory — flag for validation
- Never ignore out-of-stock handling — dead-end products lose the click and the ranking

change-management-plan

- Never treat communication as a one-time announcement — people need to hear a message multiple times before they internalise it; plan for repeated touchpoints
- Never assign change management to a single owner without involving line managers — managers are the most effective cascade channel and must be briefed before their teams
- Never schedule training after go-live — people who learn a new system on the day they need to use it will revert to the old process
- Never ignore resistors in the stakeholder analysis — resistors who are not explicitly engaged will undermine adoption, especially informal leaders
- Never measure adoption only at go-live — the real test is sustained adoption at 90 days, when novelty has worn off

changelog-generator

- Never include implementation details in changelog entries — users need to know what changed for them, not how the code was refactored internally

- Never list every micro-commit as a separate entry — related commits should be grouped into one user-facing change
- Never omit the migration path for breaking changes — a breaking change entry without a specific migration action forces users to read the source code
- Never include empty sections — a "### Fixed" section with no entries signals the template was filled in carelessly
- Never write breaking changes in the same casual tone as minor additions — breaking changes must be visually prominent and call out migration requirements explicitly

changelog-writer

- Never paste raw commit messages — translate to what the user gains or must do
- Never bury breaking changes among the features — they go first, with migration steps
- Never include internal-only noise (refactors, CI tweaks) the user doesn't care about
- Never mix change types into one list — group them
- Never misclassify the version bump — a breaking change is a major, not a patch

chart

- Never use a pie chart for more than ~6 slices or for trends — pies show composition, not change
- Never put units or text inside the numeric data array — it breaks the chart
- Never emit invalid JSON (trailing commas, single quotes, comments) — it won't render
- Never mismatch lengths — a series shorter/longer than the labels misaligns the chart
- Never chart numbers you weren't given — flag gaps instead of inventing data points

chart-data-extractor

- Never silently include low-confidence data points in the main table — flag them separately so the user knows which values to verify
- Never assume a linear scale without confirming it — logarithmic axes make extracted values incorrect by orders of magnitude if misread
- Never report extracted values with false precision — if the chart's Y-axis only shows gridlines every 10 units, a reported value of 37 is invented, not extracted
- Never omit the assumptions and caveats section — partial image quality, overlapping bars, or unlabelled axes must be disclosed

churn-analysis

- Never mix avoidable and unavoidable churn in intervention plans — recommending product fixes for customers who churned due to company shutdown wastes resources

- Never calculate churn rate using end-of-period customer count as the denominator — this understates churn; always divide churned customers by the starting cohort
- Never rely solely on exit survey data for churn reasons — response rates are typically low and self-selection biases the sample toward customers who are engaged enough to complete a survey
- Never recommend interventions without linking them to a specific churn reason — interventions disconnected from root causes will not move retention
- Never report only gross revenue churn — without net revenue retention (NRR), a healthy-looking retention number can hide a shrinking revenue base

cicd-playbook

- Never describe a rollback procedure that has never been tested — a theoretical rollback is not a rollback plan; test it in staging before production
- Never allow deploys on Fridays or before holidays without an explicit on-call engineer who will monitor through the weekend
- Never commit secrets to source control even in non-production branches — secret scanning in the pipeline catches this, but prevention is the standard
- Never skip post-deploy monitoring after a production deploy — the deploying engineer must watch error rates and latency for the specified observation window
- Never suppress a security scan finding without a linked ticket and a named owner — suppressions without accountability accumulate into unmanaged risk

claude-superpowers

- Never proceed to Stage 2 without explicit user confirmation of the plan — coding before confirmation defeats the entire purpose of the planning stage
- Never write tests after the implementation and call it "test-first" — tests must be written and confirmed failing before the implementation starts
- Never skip the Double Review when time is tight — the review is most valuable precisely when speed is the priority, because that is when errors are most likely
- Never expand scope during Stage 2 without surfacing it — silent scope expansion produces code the user did not approve and may not want
- Never mark both reviews as clean without actually performing them — a rubber-stamp review produces false confidence and defeats the framework

clause-explainer

- Re-stating the clause in slightly different legalese instead of explaining it
- "It depends" with no actual read
- Risk ratings with no scenario behind them
- Suggesting changes with no example of the better wording

client-discovery

- Never pitch during discovery — listen and diagnose; the proposal is where you prescribe
- Never accept the stated problem at face value — the real one (and real scope) is usually underneath
- Never skip qualifying budget/authority — a beautiful proposal to someone who can't buy is wasted
- Never ignore red flags to win work — a bad-fit client costs more than the fee
- Never end without a confirmed next step and date — momentum dies in the gap

clinical-case-summary

- Never include any identifiable patient information — full names, dates of birth, NHS or MRN numbers, or specific addresses must be anonymised or replaced with generic identifiers
- Never omit the clinical disclaimer — this output is for documentation and educational purposes only and must not be presented as clinical advice
- Never confuse the SBAR Recommendation with a treatment plan — R is what you need from the recipient, not a full management plan
- Never list differential diagnoses without noting the reasoning for ranking — an unranked list of differentials is not clinically useful

clinical-trial-protocol

- Never invent prior efficacy/safety data or effect sizes — mark them as placeholders to be set
- Never list multiple "primary" endpoints — pick one and demote the rest to secondary
- Never let objective, endpoint, and analysis drift apart — they must answer the same question
- Never present this as regulatory/statistical sign-off — it's a draft for expert review
- Never omit safety monitoring and stopping rules — a protocol without them isn't approvable

co-marketing

- Never propose a partner with no real audience overlap — "big brand" ≠ "right brand"
- Never partner with a competitor or design a lopsided deal — fairness sustains partnerships
- Never leave lead-sharing vague — agree it before the campaign, not after
- Never pitch by leading with what *you* want — lead with the partner's gain
- Never skip metrics — "we did a thing together" isn't a result

code-explainer

- Never narrate line-by-line in English ("this line sets x to 5") — explain intent and structure
- Never skip the gotchas — the value is in the non-obvious parts
- Never assume expert level if the question reads like a beginner's (or vice-versa)

- Never ignore a bug you can see just because you weren't asked to review it

code-review-checklist

- Never generate a generic checklist that ignores the stated language — a Python checklist and a Go checklist have fundamentally different correctness concerns
- Never treat "looks fine" as a valid review outcome — the checklist exists to surface specific concerns, not validate a superficial read
- Never scope a "high risk" review the same as a "low risk" review — depth must scale with the stated risk level
- Never flag every stylistic preference as a blocking issue — distinguish between blocking correctness issues and non-blocking comments
- Never skip the "common pitfalls" section for the stated language and change-type combination — this is where the most valuable knowledge lives

code-review-guide

- Never nitpick style while missing a correctness or security problem — priority first
- Never dump ungraded comments — rank them so the author knows what's blocking
- Never say "this is wrong" without why and a suggested fix
- Never rewrite it your way for taste — respect working approaches; flag real issues
- Never be a jerk — review the code, acknowledge good work, keep the author motivated

code-simplification

- Never simplify and change behaviour in one pass — the moment behaviour shifts, this became a rewrite without a spec
- Never delete weirdness without checking why it's weird — some of it is a production incident's scar tissue
- Never confuse terse with simple — code golf raises the reading tax this skill exists to cut
- Never remove flexibility that's actually on the roadmap — YAGNI applies to imagined futures, not planned ones
- Never skip the ledger — invisible simplification is indistinguishable from unexplained deletion in review

cohort-analysis

- Never allow the same user to appear in multiple cohorts — overlapping cohorts produce retention numbers that cannot be compared or acted upon
- Never assume assumed ARPU in LTV projections — use observed revenue per retained user per period, not a blended average that hides segment differences

- Never draw conclusions from cohorts too small to be statistically meaningful — flag minimum cohort size thresholds and note when a cohort is too small to trust
- Never conflate retention rate with engagement rate — a user who logs in but does not complete the key retention event is not retained by the definition used
- Never make recommendations without connecting them to specific cohort or segment findings — generic growth advice that could apply to any product adds no value

cohort-curve-model

- Never fit fewer than 4 periods — two points always fit a power law and mean nothing
- Never project a poor fit silently — a beautiful curve through bad residuals is how LTV fictions get funded
- Never quote LTV without the horizon — "lifetime" hides the assumption that matters
- Never average incomplete cohorts into the input (young cohorts drag the tail down mechanically — survivorship in reverse)
- Never present the fitted floor as a promise — it is an extrapolation, and the honest phrasing is "if the current shape holds"

cold-email

- Never open about yourself ("We're a leading platform...") — lead with them
- Never feature-dump — one outcome + one proof beats a capability list
- Never stack asks or ask for too much ("30-min demo" cold) — make the first yes tiny
- Never use fake personalisation ("loved your post!") — be specifically, verifiably relevant or don't claim it
- Never skip follow-ups or make them "just checking in" — each must add a reason to reply

collections-email

- Never open with hostility — most late payments are oversight; start friendly
- Never make it hard to pay — repeat the payment link/details in every message
- Never threaten legal action or fees you can't or won't enforce — keep it factual and lawful
- Never wait until 60 days to send the first chase — a pre-due/just-due nudge gets paid fastest
- Never present this as legal advice — flag interest/escalation for professional confirmation

community-management-playbook

- Never delete genuine customer complaints to silence negative feedback — deletion damages trust more than the original complaint and can escalate a minor issue to a viral one
- Never respond to competitor comparison comments publicly — engaging publicly with competitive comparisons amplifies them; redirect to DMs or ignore

- Never use the same template response for every complaint — copy-paste responses on visible complaints are noticed by other users and undermine brand authenticity
- Never leave a crisis without pausing scheduled content — queued posts published during an active brand crisis appear tone-deaf and make the situation worse
- Never set response time SLAs that cannot be met with the available team size — an SLA that is consistently missed is worse than no SLA

company-brief

- Never fabricate funding, metrics, or news — work from provided info and label inferences
- Never produce a generic company overview — focus on what matters for this role and interview
- Never list culture platitudes — read real signals (JD tone, reviews, how they describe the work)
- Never suggest generic questions ("what's a typical day?") — make them business-specific
- Never skip "likely challenges" — it's the section that makes you sound like a hire, not a tourist

comparative-market-analysis

- Never invent comparable sales or prices — use provided data or say what to pull
- Never give a single exact value with false precision — give a supported range
- Never skip adjustments — raw comp prices ignore the differences that matter
- Never ignore the market trend — a stale comp in a moving market misleads
- Never present this as a formal appraisal — flag for professional review

competitive-analysis

- Never present competitor feature claims as facts without citing a source or flagging them as assumptions — outdated or incorrect feature data misleads sales and product decisions
- Never build a competitive analysis that only covers features — pricing, messaging, go-to-market motion, and who they hire for are equally strategic signals
- Never treat all buyers as identical — the same product may win against Competitor A in the enterprise segment and lose in SMB; segment-specific win/loss matters
- Never soften weaknesses and threats in the SWOT to avoid internal discomfort — an honest SWOT is only useful if the negatives are real

competitive-intelligence-monitor

- Never mark a signal as Low priority simply because it is new and unfamiliar — unknown competitive moves often deserve investigation before dismissal

- Never provide "monitor" as the recommended response for a High-priority signal — High signals require a specific action with a named owner
- Never include signals from competitors that are not relevant to the stated roadmap or strategic priorities — noise reduces the brief's usefulness and trains the team to ignore it
- Never produce a diff-mode brief that is longer than the full report — if the diff output exceeds 300 words, it is a full report, not a diff

competitor-signal-tracker

- Never rate a signal as High threat without explaining the specific roadmap item or customer segment it threatens — unjustified threat ratings lose credibility over time
- Never treat a hiring signal as definitive proof of a strategic bet — hiring signals require corroboration from product, messaging, or pricing signals before acting on them
- Never conflate a competitor's announcement with a competitor's shipped capability — press releases and blog posts often describe aspirations, not production features
- Never recommend "accelerate existing initiative" for every High signal — sometimes the right response is to differentiate harder in an adjacent area rather than race the competitor directly

competitor-teardown

- Never mark feature presence as equivalent across competitors without noting quality differences — both products may have "reporting" while one's is meaningfully better
- Never position the user's product in the most favourable quadrant without justification — a self-serving positioning map that ignores real competitive pressure provides no strategic value
- Never soften Weaknesses or Threats in the SWOT — a SWOT that only celebrates strengths is a marketing document, not a strategy tool
- Never include unverifiable claims about competitor capabilities without flagging them as assumptions — presenting rumours as facts damages analytical credibility

complaint-letter

- Never vent without asking for anything — name the specific resolution you want
- Never be abusive or sarcastic — it lets the recipient dismiss the complaint
- Never omit reference numbers and dates — they slow or stall the response
- Never make threats you won't act on — state real next steps factually
- Never bury the ask — the resolution and deadline must be impossible to miss

compliance-checklist

- Never omit the legal disclaimer — this checklist does not constitute compliance advice and must never be presented as a substitute for qualified professional review

- Never generate a generic checklist that is not tailored to the stated framework, organisation type, and maturity level — a SOC 2 checklist for a startup and an enterprise are fundamentally different documents
- Never list controls without specifying what evidence is required — a control without evidence requirements cannot be audited
- Never mark a control as "full" implementation when it is partial — overestimating readiness leads to audit failures and regulatory risk
- Never skip the "common pitfalls" section — this is where organisations most frequently fail audits for the stated framework

conference-talk-proposal

- Never write a vague abstract that could describe any talk — be specific about the payoff
- Never list topics as "takeaways" — say what the attendee will be able to do
- Never oversell a talk you can't deliver in the time — match scope to the slot
- Never ignore audience level — a mismatched talk gets rejected or bombs
- Never forget the committee's view — give them the private "why this matters now" pitch

consulting-proposal

- Never lead with "About us / our methodology" — lead with their problem; they care about themselves
- Never sell hours/day-rate as the headline — price the outcome; hours invite haggling
- Never give a single take-it-or-leave-it price — tiered options win more and at higher value
- Never leave scope fuzzy — undefined scope is how fixed-price engagements bleed
- Never pad the deliverables list — confidence comes from clarity, not volume

content-calendar

- Never fill the calendar with generic topic placeholders — every entry must have a specific, usable topic and hook
- Never stack the same pillar or format on consecutive days — variety is required
- Never produce opening hooks that start with "Did you know" or other cliché openers
- Never ignore channel norms — formats must match the platform (no long-form threads for Instagram)
- Never skip the repurposing map for High Priority posts

content-repurposer

- The same paragraph pasted into all five with different line breaks
- A LinkedIn wall of text, or a thread that's one idea split mid-sentence
- Generic hooks ("Here are some thoughts on...")

- Hashtag stuffing; CTAs that don't fit the platform

content-style-guide

- Never list abstract values ("be friendly, be clear") with no examples — examples are the guide
- Never write an exhaustive rulebook no one will read — prioritise the high-frequency decisions
- Never ignore tone-by-context — the same voice should sound different in an error vs. a celebration
- Never omit a terminology/word list — inconsistent product terms are the most visible failure
- Never skip accessibility/inclusivity — they're style rules too

context-engineering-review

- Never review the prompt template and call it a context review — the bloat lives in the dynamic parts
- Never recommend "shorten everything" — cutting the wrong 200 tokens costs more than keeping 2,000 idle ones
- Never leave contradictions in place because each section "is fine alone" — the window is read as one document
- Never treat more retrieval as more grounding — irrelevant chunks actively mislead
- Never propose structure the assembly code can't enforce — a budget without an enforcement point is a wish

context-mode

- Logging verbatim command output instead of a filtered summary (defeats the context savings)
- A resume announcement from a generic template that ignores what the log actually says
- Appending to "Last User Prompt" / "Last Action Taken" instead of replacing them (the log bloats)
- Activating silently without offering the CLAUDE.md install, so it doesn't persist across sessions
- On resume, asking the user what to do instead of continuing the in-progress task

contract-review

- Never provide legal advice or suggest the review substitutes for qualified legal counsel
- Never skip flagging unusual or one-sided clauses because they appear standard
- Never omit a plain-English summary — legal jargon alone is not useful output
- Never rate risk without explaining what specifically drives that rating

- Never ignore missing clauses — absence of key protections is itself a risk

contributor-guide

- Never assume the contributor knows the setup — spell out the exact commands
- Never leave PR expectations implicit — say what a mergeable PR includes
- Never be gatekeep-y or cold — friction and tone both lose contributors
- Never omit how to get help or who reviews — uncertainty stalls first PRs
- Never forget the Code of Conduct link — it sets the community standard

conversion-rate-optimization

- Never test trivial cosmetics (button colour) before fixing clarity, friction, and anxiety — the big levers
- Never A/B test on traffic too low to ever reach significance — use qualitative research or sequential changes instead
- Never optimise the step in isolation — a signup lift that lowers paid conversion is a loss; watch the downstream metric
- Never call a test on day two because it looks good — set the threshold and sample size before you start
- Never redesign by opinion — every change should trace to a diagnosed blocker and a hypothesis

cover-letter

- Never open with "I am writing to apply for the position of..." — it wastes the most valuable line
- Never restate the resume — the letter adds connection and context, not a duplicate list
- Never use generic flattery ("your prestigious company") — name something real and specific
- Never pad to a page — a tight 4 paragraphs beats a full page of filler
- Never sound desperate or arrogant — aim for confident and genuinely interested

creator-brand-kit

- A vague niche ("lifestyle", "tech") that positions against everyone
- Pillars that are topics-of-the-week, not durable themes
- Voice = a list of adjectives with no examples
- A clever bio that doesn't say who it helps or what they get

creator-media-kit

- A media kit that's all vanity metrics and no audience fit

- A generic "I'd love to collab!" email with no brand-specific reason
- Inventing follower/engagement numbers
- A single flat rate with no rationale or room to negotiate usage/exclusivity

cs-escalation-brief

- Never assign blame to individuals — focus on system failures and process gaps
- Never downplay ARR at risk or describe churn risk vaguely without a number
- Never leave resolution plan ownership as "TBD" or unassigned
- Never write the brief without a clear ask from the escalation owner
- Never omit the customer's own stated position — their perspective must be represented fairly

cs-health-scorecard

- Never score health dimensions on gut feel — every score needs specific supporting evidence
- Never give a Green status to accounts with unresolved P1 issues or missed milestones
- Never list risks vaguely — "low engagement" without specifics is not actionable
- Never leave recommended actions without named owners and deadlines
- Never conflate product usage frequency with product value delivery

csat-nps-analysis

- Never average NPS ratings — it's a net of percentages; averaging gives a meaningless number
- Never report the score without the why — the verbatims are where the action is
- Never ignore passives — they're the cheapest group to convert into promoters
- Never stop at the score — an analysis with no prioritised action changes nothing
- Never trust a tiny sample — flag low n; a 12-response NPS swing is noise, not a trend

customer-advisory-board

- Never turn the CAB into a product pitch or QBR
- Never stack the room with only your happiest customers
- Never let the team defend decisions instead of listening
- Never collect input and go silent — always close the loop
- Never seat direct competitors together or ignore confidentiality

customer-journey-map

- Never build the map from assumptions alone — ground at least the pain points in real customer data or research
- Never treat all journey stages as equally weighted — identify the highest-friction moments explicitly
- Never omit the emotional layer — a journey map without emotions is a process flow, not a customer map
- Never create generic touchpoints that apply to any product — each touchpoint must be specific to this product and customer
- Never leave opportunities unranked — prioritise by impact and feasibility

customer-outage-notice

- Never go quiet between updates — a "still working on it, next update by X" beats silence
- Never minimise ("minor issue") when customers are clearly blocked — it erodes trust
- Never dump engineering detail or assign blame in a live customer notice
- Never promise an ETA you're not confident in — commit to an update time instead
- Never forget the resolved message — leaving an incident "open" worries customers

customer-success-plan

- Never define success metrics that the vendor controls — metrics must reflect the customer's business outcomes
- Never set milestone dates without customer confirmation — unilateral timelines undermine joint ownership
- Never create a plan the customer hasn't agreed to — it must be mutual, not a CSM's internal plan
- Never leave ownership fields blank or assigned to "CS team" — every action needs a named owner
- Never confuse product adoption milestones with customer business outcomes — both are needed but are not the same

dashboard-brief

- Never specify metrics that the available data sources cannot actually support — always validate data availability
- Never include more than 8–10 primary metrics on a single dashboard — more creates noise, not insight
- Never skip the primary business question — a dashboard without a north-star question becomes a vanity metrics display
- Never choose chart types for aesthetic reasons — every chart type must match the data relationship it represents
- Never leave filter configurations vague — specify exact filter values, not just filter categories

data-analysis-standard

- Never present correlations as causation — always state the distinction explicitly
- Never report a metric movement without stating the time window and comparison baseline
- Never skip the "so what" — raw observations without recommended actions are incomplete analysis
- Never overstate confidence — label hypotheses clearly and note what data would be needed to confirm them
- Never ignore segment breakdowns — aggregate metrics can mask opposing trends in sub-segments

data-contract

- Never leave semantics implicit — undocumented timezone/units/null handling is the #1 silent data bug
- Never write vague SLAs — "fresh and accurate" is unenforceable; give times and thresholds
- Never allow breaking changes without notice — a deprecation window + consumer sign-off is the whole point
- Never skip ownership — an unowned dataset has no one to hold to the contract
- Never version informally — schema changes need semver so consumers know what broke

data-pipeline-spec

- Never spec a pipeline without defining SLAs — "as fast as possible" is not an acceptable freshness target
- Never omit error handling and dead-letter queue strategy — every pipeline must specify what happens to failed records
- Never design idempotent loads without documenting the deduplication key — assume reruns will happen
- Never leave data quality rules implicit — schema validation, null checks, and referential integrity must be explicit
- Never ignore schema evolution — specify how upstream schema changes are detected and handled

data-quality-audit

- Only checking for nulls and calling it done
- "Clean your data" with no specific issues or checks
- Treating all issues as equally severe regardless of the decision
- Fixing data silently with no record of what was changed

data-quality-checks

- Never block the pipeline on every check — alert fatigue makes people ignore the real failures; reserve 🚫 for critical
- Never only test the happy path — the grain key, nulls, and freshness are where the real breakage hides
- Never write checks with no failure action — a test that fails into the void changes nothing
- Never skip freshness — stale data that looks fine is the most dangerous kind
- Never check only one table in isolation — cross-table consistency (FKs, reconciliations) catches integration bugs

data-retention-policy

- Never set retention to "indefinite" or leave it blank — undefined retention is the highest-risk, least-defensible state
- Never forget backups — data deleted from production that lives on in backups is still data you hold
- Never keep data with no legal or business basis — if you can't justify it, deleting it lowers risk for free
- Never set a blanket period for all data — tax records and marketing emails have very different drivers
- Never present statutory periods as advice — flag where legal/compliance must confirm the minimums

database-migration-plan

- Never combine the expand and contract phases into a single deployment — they must be separated by a deployment cycle
- Never run DDL changes without first testing on a production-sized data clone
- Never skip the NOT VALID + VALIDATE pattern for constraint additions on large tables — it causes full table locks
- Never define a rollback as "restore from backup" — each phase must have an explicit, fast rollback procedure
- Never omit dual-write logic during the transition period — removing the old column before all writers are updated causes data loss

database-schema-design

- Never use JSONB columns as a substitute for proper relational schema design on fields that will be queried
- Never add indexes speculatively — every index must be justified by a specific access pattern
- Never omit timezone-awareness — use TIMESTAMPTZ, never plain TIMESTAMP
- Never design without documenting normalization decisions — future maintainers need the reasoning, not just the structure

- Never skip the access patterns section — schema without query patterns cannot be evaluated for correctness

dataset-datasheet

- Never describe only the happy-path contents — the gaps, skews, and noise are what cause model failures
- Never omit the consent/licensing basis — "we scraped it" is a legal and ethical liability if undocumented
- Never ignore proxy variables — removing race/gender columns doesn't remove the bias if zip code or name encodes it
- Never present label quality as perfect — state who labelled it and the agreement rate, or note it's unmeasured
- Never leave the dataset ownerless — an unmaintained dataset silently rots as the world changes

dbt-model-spec

- Never leave the grain ambiguous — an untested, unclear grain is how duplicate rows and wrong metrics happen
- Never hard-code upstream table names — use `ref()/source()` so lineage and environments work
- Never ship a model with no tests — untested models silently rot; the grain key at minimum must be tested
- Never default everything to a table — pick the materialization the use justifies (views for light, incremental for large append-only)
- Never bury business logic without comments — the next analyst must understand the rules

debt-payoff-plan

- Never recommend a method without showing the interest/time trade-off in numbers
- Never forget the minimums on non-target debts — the plan must cover all of them
- Never ignore the person's psychology — the mathematically optimal plan they quit isn't optimal
- Never assume variable-rate debt stays fixed without flagging it
- Never present this as personalized financial advice — it's an educational model to adapt

debugging-log-analyser

- A vague root cause ("something's null somewhere") instead of the specific line/frame
- A generic fix that could apply to any language, ignoring the actual stack trace
- Restating the error message instead of explaining what it means

- "Add error handling" as prevention, with no specific guardrail
- High/Low confidence with no reason behind it

decision-memo

- Never bury the recommendation at the end — the reader should know what you want in the first 30 seconds
- Never write a status update disguised as a decision memo — if there's no decision and no ask, it's not this document
- Never stack the options — strawman alternatives destroy your credibility and the decision's quality
- Never over-agonise a reversible decision — match the rigor to the cost of being wrong
- Never hide the assumptions — surfacing "what we'd need to believe" is what lets a decider pressure-test it

deck-autopsy

- Never autopsy from a deck's reputation or your memory of the company — only from the slides provided
- Never proceed without slide images — for text notes about a future deck, use board-deck-narrative or investor-pitch-deck instead
- Never treat beautiful design as evidence of a strong argument — the correlation runs the other way often enough
- Never list ten nitpicks and skip the structural weakness — one broken chain link outweighs every font choice
- Never soften findings on your own deck — the room won't

delta-briefing

- Never restate unchanged items at full length — one line in "Still watching" or nothing
- Never fabricate a delta to make a quiet cycle look productive — "nothing changed" is a valid, valuable edition
- Never diff against memory or vibes — only against the stored state record
- Never let the state record and the brief disagree — the record is written from the brief's facts
- Never track everything forever — items resolved two editions ago leave the state record

demand-letter

- Angry, insulting, or defamatory language that undermines the sender
- Vague demands ("pay what you owe") with no figure or deadline
- Threats of consequences the sender can't or wouldn't lawfully pursue

- Burying the actual demand in a wall of grievance

dependency-audit

- Never report only direct dependencies — transitive dependency vulnerabilities are often more dangerous and are the most commonly missed
- Never present raw audit tool output without interpretation — a table of 200 CVEs with no prioritisation is worse than no audit at all
- Never assign all Critical CVEs as "fix immediately" without checking whether an exploitable path exists in your usage context
- Never make license compliance decisions without legal input — flagging a GPL dependency without a recommendation is incomplete work
- Never complete the audit without including a CI/CD pipeline step — a one-time audit that leaves the door open for new vulnerabilities is not a remediation

dependency-conflict-resolver

- Never lead with `--force` / `--legacy-peer-deps` — it hides the conflict and breaks later
- Never delete the lockfile as the first move — explain what that actually does
- Never give a single fix when several are viable — rank them with trade-offs
- Never skip verifying the resolution actually installs/builds

design-critique

- Never lead with visual preference (e.g. "I don't like the colour") — every issue must reference a UX principle or user impact
- Never invent problems in the "What's Working" section — manufactured praise undermines the entire critique
- Never provide the same priority level (High/Medium/Low) to every issue — prioritisation requires genuine judgment about user impact
- Never skip the JTBD section for product screens — connecting feedback to the user's job-to-be-done is what separates UX critique from aesthetic opinion
- Never give recommendations that require a full redesign when the user is in high-fidelity — scope recommendations to the design stage

design-handoff-brief

- Never write the user goal in feature language ("design the checkout flow") — it must be written from the user's perspective with a motivation and outcome
- Never skip the "Explicitly Out of Scope" section — without it, designers will inadvertently solve problems not intended for this iteration
- Never list edge cases that are so generic they apply to any feature (e.g. "handle errors") — each edge case must be specific to this feature's failure modes

- Never hand off the brief without confirming engineering constraints are accurate — a constraint that is wrong is worse than no constraint
- Never omit the emotional context of the user — designs without emotional grounding produce technically correct but experientially flat results

design-system-audit

- Never assess only the Figma library without checking the code implementation — Figma-code drift is one of the most common and costly design system failures
- Never score adoption without interviewing teams — audit tool metrics miss the human reasons teams build custom components instead of using the system
- Never treat all component gaps equally — prioritise gaps based on how many production screens rely on custom implementations, not alphabetically
- Never recommend adding more components without first auditing documentation quality — an undocumented component is often worse than no component
- Never schedule remediation without a named owner per initiative — design system improvements without ownership consistently stall

developer-onboarding-doc

- Never document the ideal setup — document the actual setup; real oddities and gotchas are what new engineers need most
- Never leave placeholder contacts like "ask your manager" — name specific people for each domain or the doc becomes useless when the new joiner has an urgent question
- Never write the onboarding doc without reviewing it with a recent joiner — the author is blind to what they take for granted
- Never include every piece of architectural detail — an onboarding doc that covers everything teaches nothing; link to deeper docs instead
- Never skip the "things that might surprise you" section — undocumented non-obvious patterns are the number one cause of wasted engineering time in the first week

difficult-conversation

- Never open with blame or "you always/never" — it triggers defensiveness and ends learning
- Never confuse your story with the facts — "the deadline slipped" is fact; "you don't care" is a story
- Never over-script — plan the open and the points, then stay responsive; a rigid script breaks
- Never aim to win — if the goal is to be right, the relationship loses even if you "win"
- Never avoid the actual ask — name the change or agreement you need, kindly and clearly

disaster-recovery-plan

- Never write runbook commands without testing them — an untested command in a runbook is actively dangerous during a real disaster when cognitive load is highest
- Never set RTO/RPO targets without business sign-off — technical teams often set aspirational targets that do not reflect actual business cost tolerance for downtime
- Never include only the "happy path" of each failover scenario — runbooks must explicitly cover what to do when the recovery step itself fails
- Never list Slack handles as the only escalation contact — Slack may be unavailable during a region-wide failure; phone numbers are mandatory
- Never schedule DR game days without pre-committing to fix the gaps found — a game day that produces action items no one owns is theater, not preparedness

discharge-summary

- Never invent medications, doses, diagnoses, or results to complete a section
- Never present this as medical advice — it formats clinician-provided information for handoff
- Never leave the medication list ambiguous about what changed during the stay
- Never omit pending results or follow-up ownership — that's where handoffs fail
- Never write patient instructions in clinical jargon the patient can't act on

discovery-call-prep

- Never write the call hypothesis after the call — hypotheses written post-hoc are rationalisations, not testable predictions
- Never open with a product pitch before establishing the prospect's problem — leading with pitch signals you are not there to learn, which closes discovery conversations
- Never use closed questions in the discovery phase ("Do you have this problem?") — they produce yes/no answers that confirm bias rather than reveal pain
- Never skip the "not successful" definition in success criteria — a call that ends with "send me more info" feels like progress but is not a qualified next step
- Never treat all prospect research equally — recent news (last 90 days) is more relevant to call context than static company facts from LinkedIn

discovery-interview-guide

- Never use future-tense questions ("Would you use this?") — hypothetical responses do not predict real behaviour and produce false confidence in an idea
- Never mention your product or solution before problem exploration is complete — doing so anchors the participant's responses and invalidates the discovery
- Never synthesise across fewer than 5 interviews — themes from 2–3 interviews reflect anecdote, not pattern; wait for saturation
- Never write screener questions that are too easy to pass — if participants can guess the "right" answer, you will recruit the wrong people

- Never treat participant opinions as evidence of future behaviour — what people say they will do consistently diverges from what they actually do

dispute-letter

- Never be vague about which charge/record and how much — precision is the whole game
- Never omit evidence or fail to reference it — assertions without proof stall
- Never present this as legal/financial advice or guess at statutory deadlines — flag them to confirm
- Never get emotional — a factual record is more persuasive and more useful if it escalates
- Never forget to request written confirmation of the resolution

docs-quickstart

- Never front-load concepts/architecture — get them to a working result first, explain later
- Never assume hidden setup — every step needed to run must be present
- Never show a huge "kitchen sink" example as the first one — minimal win first
- Never skip the expected output — devs need to confirm it worked
- Never leave dead-ends — always point to what's next

docx-tracked-changes

- Never paraphrase original text when creating tracked deletions — the original text must be preserved exactly, character for character, or the tracked change cannot be reviewed against source
- Never mix substantive changes with stylistic edits in the same section — reviewers need to approve substantive changes at a different threshold than copy edits
- Never write margin comments as meta-commentary about the review process ("This section needs work") — comments must be actionable instructions the author can act on
- Never flag every imperfect sentence as a change — over-redlining trains authors to accept changes without reading, which defeats the purpose of tracked review
- Never produce a redline without a summary of top-level changes — reviewers read the summary first and use it to decide which changes to scrutinise in detail

donor-update

- Never make a "thank-you" that's really just another donation ask — stewardship builds the next gift
- Never be generic ("thanks for your support") — name the specific impact their gift had
- Never present invented impact as real — mark placeholders for the org
- Never write like a corporation — warmth and a real human voice retain donors
- Never omit the story — numbers thank the head, a story thanks the heart

email-campaign

- Never include more than one primary CTA per email — competing calls to action reduce click-through by splitting attention
- Never open with "Hi, welcome to [product]" or any variation of a generic greeting — the opening line must earn attention immediately or recipients stop reading
- Never write preview text that repeats the subject line — preview text is a second chance to earn the open, not a repeat of the first chance
- Never write a sequence where each email restates the same value proposition — each email must advance the narrative or serve a distinct purpose in the buyer's journey
- Never assume all subscribers receive all emails — each email must stand alone for subscribers who missed earlier messages in the sequence

email-sequence

- Never write a newsletter — each email needs one job, not five updates
- Never ask in every email — give value first; pushing too early kills the sequence
- Never forget the exit condition — emailing converted users "buy now" erodes trust
- Never stuff multiple CTAs — one action per email or none gets taken
- Never judge by opens alone — tie each email to the step it's meant to drive

email-triage

- Never surface FYI emails in the High or Medium priority sections — burying actionable items with informational ones defeats the purpose of triage
- Never write vague "What they need" summaries ("Sarah sent an email about the report") — every summary must state the actual ask, not a description of the email
- Never apply the same tone to every reply starter — a formal email from a client requires a different opener than a casual Slack-style email from a colleague
- Never include emails outside the requested time window — time window accuracy is the core trust signal for this skill
- Never omit the filtered-out count — users need to know how much was scanned, not just what was surfaced, to trust the triage is complete

employee-engagement-survey

- Never launch a survey without committing to a communication-back date — surveys with no follow-through reduce trust and depress future response rates
- Never use only Likert scale questions — open-text responses surface specific themes that quantitative scores cannot, and are essential for action planning
- Never design a comprehensive 30+ question survey as a pulse — pulse surveys that take more than 5 minutes see sharply lower completion rates

- Never present analysis without an action planning template — raw scores without committed actions are the most common reason engagement survey data is ignored
- Never segment results below teams of 5 when anonymity is promised — small-group breakdowns allow individual identification and destroy psychological safety

empty-state-writer

- Never ship a bare "No data" / blank screen — it wastes the best activation moment
- Never treat every empty state the same — "nothing yet" is opposite to "all caught up"
- Never make a cleared/complete state look like an error
- Never offer no action on a first-use state — give the one next step
- Never over-explain — a headline, a line, and a button, not a paragraph

engagement-retro

- Never skip the money question — "the client was happy" but you lost margin means change the pricing, not repeat it
- Never write vague lessons — "communicate better" isn't actionable; "add a weekly written status" is
- Never let the engagement end without asking for the testimonial/referral — now is the easiest it'll ever be
- Never bury scope creep — name what you ate so the next SOW prevents it
- Never treat the retro as internal-only — the client-facing close-out also sets up the renewal

engineering-hiring-rubric

- Never use a single behavioral anchor description per competency — you must define what Strong Hire AND No Hire look like separately, or interviewers cannot calibrate
- Never allow "culture fit" as a standalone assessment dimension — it masks similarity bias; all judgments must use observable behavioral evidence
- Never let interviewers share scorecard feedback before the debrief — verbal pre-debrief discussion anchors everyone to the first opinion expressed
- Never set the same must-hire competency list for all engineering roles — a senior backend engineer and a frontend engineer have different non-negotiable competencies
- Never skip the calibration bias notes section — interviewers who have never been briefed on halo effect, recency bias, and credential bias will reproduce them in every loop

engineering-weekly-report

- Never fabricate metrics — if data is not available, mark the field as [data needed] rather than estimating; stakeholders making decisions on invented numbers is actively harmful

- Never write next week's priorities as activities ("work on X") — they must be outcomes ("ship X", "complete Y migration") so stakeholders can evaluate whether the team delivered
- Never bury escalations inside a risk table row — anything needing leadership attention must be called out explicitly in the Escalations section
- Never list blocked items without naming a specific owner and a concrete unblocking action — "waiting on X" is not a blocker entry, it is a placeholder
- Never write a report that exceeds two printed pages — length signals the author has not done the editorial work of deciding what matters to stakeholders

entity-relationship-diagram

- Never draw relationships without cardinality — "related" isn't a data model
- Never leave many-to-many unresolved when a join table is the right call
- Never dump every conceivable column — show the keys and the attributes that matter
- Never omit foreign keys — they're how the relationships are actually enforced
- Never break Mermaid with unquoted spaced labels

error-decoder

- Never just paraphrase the error — explain what it *means* and why it happened
- Never give a generic "try reinstalling" answer when the trace points to a specific cause
- Never invent file names or code that wasn't given — infer and label, or ask for the one missing thing only if truly blocking
- Never stop at the fix — always add the one prevention step

error-message-writer

- Never show raw/technical errors ("Error 500", "null pointer") to end users
- Never blame the user ("Invalid input") — say what to do instead
- Never write a dead-end ("Something went wrong") with no next step
- Never be jokey about serious failures (payment, data loss) — match the tone to the stakes
- Never bury the action — the recovery step should be obvious

escalation-tree

- Never leave severity fuzzy — if "critical" is subjective, everything becomes critical (or nothing does)
- Never write "escalate when needed" — time-box it so issues don't rot waiting on judgement
- Never route to named people only — use roles with primary/secondary; people leave and go on holiday
- Never forget customer comms in the tree — internal escalation without customer updates still feels like neglect

- Never over-escalate everything — tiers exist so seniors see only what truly needs them

eval-rubric-designer

- Never ship dimension names without anchors — names alone don't make scores reproducible
- Never let one quality issue load onto multiple dimensions — keep them orthogonal
- Never trust an LLM judge blind — calibrate against a handful of human labels first
- Never use a vague "overall quality 1–10" — it hides which part is broken
- Never ignore the negative case — a rubric must distinguish "wrong" from "thin", not just "great" from "okay"

evidence-lock

- Never fill source gaps with general knowledge — in this mode the provided sources are the entire universe of evidence
- Never cite documents wholesale — a lock names the passage
- Never launder inference as citation — derived conclusions are labelled as derived
- Never quietly drop claims that can't be sourced in soft mode — the register exists so the author sees what's resting on air
- Never proceed without sources "just this once" — without sources this is a normal draft, and other skills do that better

excel-model

- Never write computed results as static numbers — the whole point is that inputs recalculate
- Never hard-code an assumption inside a formula — put it on the Inputs sheet and reference it
- Never scatter inputs across sheets — one assumptions sheet, single source of truth
- Never skip formatting — an unformatted grid of numbers is hard to trust or use
- Never claim a file was produced if there was no code execution — fall back to a clear spec instead

executing-plans

- Never improvise around a broken plan — amend it visibly or stop; silent drift is unaccountable work
- Never skip verifications when steps "obviously worked" — the mysterious step-6 failure was born at step 3
- Never push through a stop condition on momentum — it was written calm precisely because you wouldn't be

- Never declare done without running the done-test — feeling-finished and being-finished diverge exactly when it matters
- Never end a session without the state note — re-derivation is the tax on every resumed task

executive-presence

- Never give generic advice ("be more confident") — coach specific behaviours for this room
- Never bury the lead — answer-first; making execs wait for the point reads as junior
- Never bluff a tough question — calm "here's what I know, I'll confirm the rest" beats a confident wrong answer
- Never equate presence with talking more — concision and comfortable silence project more authority
- Never coach a persona — it's behaviours layered on who you are, not an act that won't survive pressure

executive-summary

- Never summarise the document chronologically — an executive summary that follows the structure of the source document is not an executive summary, it is an abstract
- Never bury the recommendation at the end — executives read the first paragraph and skim the rest; the ask must be in sentence one or two
- Never use the same summary for different audiences — a CEO and a board member have different decision contexts and require different framing
- Never include background that the reader already knows — every sentence of background must earn its place by making the bottom line more actionable
- Never leave the "risks of inaction" section vague — a summary that does not quantify what happens if the reader does nothing removes the urgency needed for a decision

executive-update

- Never lead with context or background — executives read the headline first; bury the important thing below two sentences of setup and they will miss it
- Never present metrics without a comparison point — a number without context (vs. target, vs. last period) cannot be interpreted and will prompt follow-up questions
- Never soften or spin risks — executives rely on these updates to make resource and escalation decisions; sanitised risk sections destroy the update's utility
- Never present a "Decisions Needed" item without a recommendation — asking an executive to decide without your view forces them to do the analytical work the PM should have done
- Never exceed 250 words in the main body — length signals the author has not done the compression work; every word over 250 reduces the chance the update is read

expense-audit

- Never say "spend less" — every cut must name an amount and a method
- Never rank by monthly when annual reveals the real impact — annualize everything
- Never suggest cutting things the person flagged as off-limits
- Never miss the silent recurring charges — those are usually the biggest, easiest wins
- Never present this as personalized financial advice

expense-policy

- Never leave limits vague ("reasonable") with no number or guidance — that creates the disputes
- Never bury the process — people need to know exactly how to get paid back
- Never assert tax/per-diem rules as fact — flag for finance to confirm by jurisdiction
- Never omit what's *not* covered — the exclusions prevent the awkward conversations
- Never make it so strict it signals distrust, or so loose it has no teeth — aim for fair and clear

experiment-designer

- Never define success criteria after seeing preliminary results — post-hoc success definitions are HARKing (Hypothesising After Results are Known) and invalidate the experiment
- Never stop a test early because the result looks significant — early stopping dramatically inflates false positive rates; the test must run to the planned sample size
- Never treat statistical significance as the same as practical significance — a $p < 0.05$ result with a 0.1% lift is real but may not be worth shipping
- Never run the same experiment on the same population multiple times without correction — multiple testing inflates the chance of a false positive proportionally
- Never use more than one primary metric — multiple primary metrics require multiple hypothesis corrections and make the ship/kill decision ambiguous

experiment-readout

- Never call significance by eye — compute the p-value and CI; a higher number isn't a result
- Never ignore the confidence interval — a CI spanning zero (or huge) means you don't actually know the effect
- Never confuse statistical with practical significance — a tiny significant lift may not be worth shipping
- Never trust a peeked/early-stopped test — stopping when it looks good inflates false positives massively
- Never spin a null result — "no detectable difference" is honest and often the right call

exploratory-test-charter

- Never write "explore the feature" with no mission, areas, or oracles — that's aimless clicking

- Never skip prioritisation — explore the riskiest areas first
- Never turn charters into scripted step-by-step cases — exploration needs freedom within focus
- Never omit oracles — without them a tester can't tell right from wrong
- Never leave sessions open-ended — timebox them so coverage is accountable

feature-flag-guide

- Never create release flags without a cleanup date — flags without expiry dates become permanent technical debt that accumulates silently until the codebase is unmaintainable
- Never skip monitoring setup for flags between 1–99% rollout — a partially-rolled-out flag without metric comparison is a risk without a sensor
- Never nest flags inside other flags — compound flag logic makes cleanup nearly impossible and creates untestable code paths
- Never allow flag owners to leave the team without reassigning ownership — orphan flags with no owner never get cleaned up
- Never use feature flags as a permanent configuration system — flags that have been at 100% or 0% for more than 30 days must be cleaned up; using flags as permanent config couples business logic to a feature flag platform

feature-prioritisation

- Never score items against different goals — every item in a prioritisation session must be scored against the same objective
- Never omit deprioritised items — explicitly listing what was cut and why is as important as the ranked list
- Never let stakeholder politics override framework scores without documenting the override and reason
- Never mix RICE, ICE, or MoSCoW scores across frameworks in a single session — pick one framework per prioritisation exercise
- Never treat the output as final without documenting the assumptions used in scoring — assumptions change, and the list must be revisitable

feature-sunset-plan

- Never announce by blog post before dependent customers hear it from their account team
- Never use raw usage percentage as the whole case — depth and contracts decide who can veto your math
- Never promise the replacement is equivalent when it isn't — name the losses; users find them anyway
- Never grant open-ended exceptions — one immortal customer instance is the whole maintenance cost with none of the revenue

- Never declare victory at shutoff — the sunset is done when the code is gone and the retro is filed
- Never let "deprecated" become a permanent state — a deprecation without a removal date is a mood, not a plan

figma-annotation-guide

- Never annotate only the happy path — error states, loading states, and empty states must all be documented
- Never use vague spacing descriptions like "some padding" — specify exact pixel values or token names
- Never skip accessibility annotations — focus order, ARIA labels, and colour contrast ratios must be included
- Never leave interaction behaviour undescribed — every interactive element needs a documented response
- Never produce annotations without edge cases — developers need to know what happens at boundaries

figma-component-audit

- Never flag naming issues without providing a specific, consistent naming convention to adopt
- Never audit only visual consistency — also check for missing interactive states and accessibility compliance
- Never list all issues at equal priority — group by impact (Critical / Major / Minor) so the fix plan is actionable
- Never omit variant completeness — every interactive component must cover all required states
- Never leave coverage gaps without recommending specific missing components to add

figma-design-brief

- Never write a design brief that describes the solution — the brief must describe the problem and constraints, not the design answer
- Never skip the success criteria — designers need to know what "done" looks like before starting
- Never omit existing components to reuse — briefs that ignore the design system lead to inconsistent implementations
- Never leave open questions unresolved — escalate them before design work starts, not during it
- Never confuse requirements with design instructions — the brief defines what, not how

figma-design-critique-pm

- Never critique visual aesthetics — PM feedback must focus on product outcomes, user goals, and business requirements
- Never provide feedback without stating the evidence basis — distinguish between observed design facts and assumed user behaviour
- Never give vague feedback like "the flow feels confusing" — every concern must reference a specific screen state or interaction
- Never ignore what is working — balanced critique includes explicit acknowledgment of design decisions that are well-executed
- Never critique without knowing the design constraints — always ask about technical, time, or resource limitations before judging decisions

figma-design-qa

- Never produce a partial QA — every checklist category must be evaluated, not just the ones that are obviously problematic
- Never leave the handoff decision ambiguous — the output must explicitly state pass, pass with conditions, or fail
- Never skip accessibility checks — colour contrast, tap target size, and screen reader labels are required, not optional
- Never report issues without specifying which screen or component they appear on
- Never approve a design if any component is detached from the library without a documented reason

figma-design-review

- Never review a design without a list of requirements to check against — always ask for the PRD, design brief, or acceptance criteria first
- Never give a vague approval status — the decision must be explicitly "approved", "approved with conditions", or "not approved"
- Never conflate requirements gaps with UX concerns — track them separately so engineers and designers can act independently
- Never raise concerns without suggesting what information is needed to resolve them
- Never skip open questions — unresolved assumptions at review time become bugs after engineering handoff

figma-prototype-plan

- Never prototype everything — scope must be limited to the interactions that answer the specific research questions
- Never design prototype flows without also writing the test task scripts — the two must align exactly

- Never skip the reset process between participants — unsettled prototype state contaminates results
- Never plan a prototype without specifying which interactions are clickable vs static — ambiguity causes scope creep
- Never scope a prototype without first defining the research questions it needs to answer

figma-spacing-system

- Never create a spacing scale with arbitrary values — the scale must follow a consistent mathematical ratio (e.g. 4px base, 8-4-2 system)
- Never define spacing tokens without Figma implementation instructions — token names alone are not actionable
- Never create a spacing system that doesn't account for component-level spacing conventions — global tokens and component usage must both be documented
- Never skip grid definitions — spacing without a grid system is incomplete layout foundation documentation
- Never produce a spacing system that ignores responsive behaviour — define how spacing adapts across breakpoints

figma-user-flow-planner

- Never plan only the happy path — all error states, empty states, and edge cases must be mapped before designing starts
- Never produce a flow map that doesn't match the Figma file structure — the page structure must reflect the flow map
- Never define screens without specifying all required states — a screen without its variants is an incomplete design scope
- Never start designing before entry and exit points are fully documented — unclear boundaries cause scope creep
- Never plan user flows without tying each step back to a user goal — every screen must justify its existence

figma-variant-matrix

- Never create a variant matrix with properties that overlap or conflict — each property must be independently variable
- Never use opacity for disabled states — disabled states must use layer-level properties, not opacity
- Never enumerate every mathematical combination if many are invalid — document only valid, buildable combinations
- Never define variants without considering responsive behaviour — specify which properties change across screen sizes

- Never produce a matrix without Figma implementation guidance — variant naming conventions must follow Figma's property system

financial-due-diligence

- Never present the checklist without tailoring it to the specific transaction type and stage of diligence
- Never overlook revenue concentration risk — customer concentration above 20–30% is a material risk that must be flagged
- Never confuse EBITDA with cash — always check cash conversion and identify non-cash items
- Never skip the related-party transaction review — undisclosed related-party dealings are a common due diligence failure point
- Never produce output without noting this is a framework and qualified financial and legal advice is required

financial-model-narrative

- Never list numbers without explaining what is driving them — narrative must go beyond restating the figures
- Never mix one-off items with recurring performance without clearly distinguishing them
- Never write the same level of detail for all line items — focus depth on the items that matter most
- Never omit forward-looking commentary — a narrative without outlook is incomplete for board or investor audiences
- Never use technical accounting language without translation — the audience is executives, not accountants

financial-statement-explainer

- Never invent figures or compute ratios from numbers you weren't given
- Never drown the reader in every line — surface what matters for their question
- Never give investment/financial *advice* — explain, and point decisions to a professional
- Never assume accounting literacy — define terms as you go
- Never conflate profit and cash — they're different and the reader needs to know why

flowchart

- Never produce a linear chain when the real process has branches — capture the decisions
- Never stuff full sentences into nodes — keep labels short, move detail to notes
- Never leave a decision with only one labelled branch — show what happens on every condition

- Never use parentheses or quotes inside labels in a way that breaks Mermaid
- Never invent steps to fill gaps — flag what you assumed

foia-request

- Never write an overbroad "any and all records" request — it invites denial or endless delay
- Never omit the date range and custodian — specificity is what gets records produced
- Never forget the fee cap — an uncapped request can return a huge estimate that stalls it
- Never cite the wrong law for the jurisdiction — federal FOIA ≠ state open-records acts
- Never overstate an expedited-processing basis — it must genuinely qualify

follow-up-sequence

- Never send a generic "thank you for your time" — reference a real moment or skip it
- Never "just check in" — every touch should add something or it reads as needy
- Never follow up too fast or too often — respect the cadence; desperation repels
- Never issue ultimatums — mention a real competing timeline gracefully, never as a threat
- Never follow up forever — define when to stop and leave the relationship intact

founder-market-fit

- A résumé in prose ("10 years at BigCo, then...")
- Vague passion claims with no evidence
- Borrowed secrets (industry truisms anyone could state)
- Overclaiming — investors discount stories that don't ring true

frontend-design

- Never decorate before systematising — tokens first, UI second, or consistency is luck
- Never spend the accent everywhere — a UI where everything is highlighted has no hierarchy, just noise
- Never ship default focus rings removed with nothing in their place — that's not minimal, it's broken
- Never design only the happy state — empty/loading/error are where users actually judge the product
- Never mix density registers — a marketing hero above a data grid needs a deliberate seam, not a collision

fundraising-faq

- Dodging the hard question instead of answering it
- "No competitors" and "we only need 1% of the market"

- Over-long answers that sound rehearsed and evasive
- Pretending a real risk doesn't exist instead of framing how you'll manage it

gantt-roadmap

- Never list tasks with no dates or durations — that's a checklist, not a timeline
- Never ignore dependencies — overlapping things that can't overlap is a fake plan
- Never draw milestones as long bars — they're points in time
- Never use ambiguous date formats — stick to YYYY-MM-DD
- Never present estimated dates as commitments — flag assumptions

gdpr-compliance

- Never default every activity to "consent" — it's revocable and high-maintenance; use the basis that actually fits
- Never skip the ROPA — without the record of what you process, every other GDPR obligation is unanchored
- Never store data with no retention period — "forever" is not a lawful retention policy
- Never treat a DPIA as optional for high-risk processing — it's a legal requirement, not best practice
- Never give legal advice as settled law — flag where a DPO or counsel must confirm (esp. lawful basis and transfers)

git-troubleshooter

- Never suggest `git push --force`, `reset --hard`, or `clean -fd` without a ⚠ and a safer alternative first
- Never give commands without saying what each one does
- Never assume the remote state — ask or label it if it changes the safe path
- Never skip `git reflog` when work might be recoverable — it usually is

giving-feedback

- Never judge character ("you're disorganised") — describe behaviour ("the doc was missing the dates")
- Never use the feedback sandwich — burying the point in praise muddles both; be direct and kind
- Never be vague ("be more strategic") — if they can't picture the change, it's not feedback
- Never save it for the review — feedback works when it's timely and low-stakes, not stockpiled
- Never make praise generic ("great job!") — specific praise is what gets the behaviour repeated

glossary-builder

- Never list terms without context — "ticket" or "filter" with no definition guarantees wrong-sense translations
- Never omit the do-not-translate flags — that's how brand/product names get mangled across locales
- Never present machine translations as approved — mark them draft for native review
- Never ignore source consistency — if the source mixes "sign in/log in," the glossary should pick one
- Never forget part of speech — a term that's both noun and verb often needs two entries

go-to-market

- Never write feature descriptions instead of benefits — the GTM pack must translate features into customer value
- Never use the same messaging across all buyer personas — each role has different priorities and language
- Never create a positioning statement that could apply to any competitor — differentiation must be specific and defensible
- Never skip the "not for" section — defining who this is not for sharpens positioning and prevents misdirected sales effort
- Never list use cases without tying them to specific job titles or buyer roles

go-to-market-planner

- Never build a Tier 1 GTM plan for an incremental feature update — tier the launch appropriately before planning
- Never create activity lists without named owners and due dates — unowned tasks do not get done
- Never skip the rollback procedure for Tier 1 and 2 launches — every significant launch must have an abort plan
- Never treat marketing and engineering as separate tracks — cross-functional coordination is the whole point of a GTM plan
- Never set success metrics without a defined measurement window — "increase signups" is not a measurable target

grant-proposal

- Never write a generic proposal — every section must be tailored to the specific funder's stated priorities
- Never exceed the specified word or page limits — over-length proposals are disqualified at many funders

- Never leave the sustainability section vague — funders need to know what happens after grant funding ends
- Never use jargon the funder's reviewers won't understand — write for the panel, not the project team
- Never underspecify the budget narrative — every significant line item must be justified with method and reasoning

growth-experiment-backlog

- Never run experiments without a hypothesis and a target metric — that's just shipping changes and hoping
- Never call a test before it reaches the planned sample size — peeking and stopping early manufactures fake wins
- Never chase many tiny tests when traffic is low — you'll never reach significance; pick fewer, bigger bets
- Never ignore guardrail metrics — a conversion win that tanks refunds or retention is a loss
- Never discard losing experiments silently — the learning is the asset; record why it failed

headline-options

- Never favour clever over clear — a confusing headline isn't read twice; clarity wins
- Never write clickbait the content can't pay off — the bounce destroys trust and SEO
- Never give one-formula variations — the value is the spread to test
- Never stay vague — "Transform your workflow" says nothing; name the specific outcome
- Never ignore the medium's limit — a truncated subject line is a wasted headline

help-center-article

- Never bury the answer — front-load it; most readers want the quick fix, not your intro
- Never title by internal feature name — title by the user's task/question, or they won't find it
- Never skip troubleshooting — the "it didn't work" cases are exactly what generate the ticket you're trying to deflect
- Never use internal jargon — write the words users type
- Never cram multiple tasks into one article — one task per article = better search + clearer steps

hipaa-safeguards

- Never treat "addressable" as "optional" — you must implement it or document why an equivalent is reasonable; silence is a violation
- Never skip the risk analysis — it's explicitly required and the most-cited gap in OCR enforcement

- Never handle PHI through a vendor without a BAA — that alone is a breach
- Never present this as legal certification — flag that compliance counsel / a security assessor must validate, especially the risk analysis
- Never conflate HIPAA with SOC 2 or GDPR — overlapping controls, different legal requirements; map each separately

hiring-rubric

- Never include competencies that overlap significantly — each dimension must assess a distinct quality
- Never write behavioural questions that can be answered with a yes/no — use "Tell me about a time..." format
- Never set a scoring bar without calibration guidance — "above bar" means nothing without concrete examples at each level
- Never create a rubric with more than 6 competencies — panel interviews cannot reliably assess more
- Never omit a "must-have vs. nice-to-have" distinction in the requirements — all criteria cannot carry equal weight

hook-writer

- "Here's everything you need to know about X" (zero tension)
- Ten rewrites of the same hook
- Clickbait the body betrays (kills trust + reach long-term)
- Hooks too long for the platform (a 90-char YouTube title, a 3-line "first line")

human-in-the-loop-design

- Never gate everything — undifferentiated approval load is how the important one slips through
- Never show raw transcripts as the approval artifact — humans approve decisions, not logs
- Never let unanswered approvals auto-proceed on timeout "to keep things moving"
- Never loosen tiers on gut feel — the down-tier bar is written evidence, the up-tier trigger is any incident
- Never measure only agent mistakes — an approver who edits nothing for a month is the riskier signal

i18n-readiness-review

- Never start translating before i18n is ready — you'll translate into hard-coded strings and broken layouts

- Never concatenate sentence fragments — word order differs by language; translate whole strings with placeholders
- Never string-format dates/numbers — use Intl/ICU, or every locale shows them wrong
- Never assume text length — German/Finnish expand; fixed-width UI truncates and clips
- Never ignore RTL until late — retrofitting right-to-left into a left-to-right layout is a rebuild, not a tweak

iep-goal-support

- Vague goals ("will improve reading") with no criterion or measurement
- Inventing a diagnosis or specific data not provided
- Confusing accommodations with modifications
- Presenting drafts as final/compliant without team review

impact-report

- Never list activities as if they were impact — tie everything to outcomes
- Never present invented figures as real — mark placeholders for the org to replace
- Never hide or omit financials — transparency is what earns repeat giving
- Never drown the human story in statistics — pair numbers with one real face
- Never forget the ask — an impact report is also a fundraising moment

incident-postmortem

- Never assign blame to individuals — postmortems must focus on system and process failures
- Never write action items with vague language like "improve monitoring" — each must name a specific, ownable change
- Never skip the contributing factors — root cause alone misses the systemic issues that enable incidents
- Never omit the detection timeline — how long it took to detect matters as much as how long it took to resolve
- Never treat the postmortem as closed until all action items have named owners and due dates

incident-public-statement

- Never open with self-congratulation or context-setting — lead with the acknowledgement
- Never use evasive legalese ("issues may have impacted some users") when you can be specific
- Never speculate on cause or promise outcomes you can't guarantee
- Never state numbers/scope you haven't confirmed — bracket them for confirmation

- Never omit what the reader should actually do next

incremental-implementation

- Never slice by layer — horizontal slices defer all verification to the end, which is the big bang with extra commits
- Never "keep going" on a red state — stacking onto broken is how one bug becomes an archaeology dig
- Never skip verification on 'trivial' increments — the trivial one is statistically where it breaks
- Never delete the old path in the same increment as the last migration — cutover and removal are separate, reversible steps
- Never let increments shrink into commit-theatre (40 one-line steps) — an increment is sized by verifiable meaning, not by smallness itself

influencer-brief

- Never leave creative guidelines so restrictive that the influencer's authentic voice is lost — prescriptiveness kills performance
- Never omit the approval process — undefined approval workflows cause delays and missed publishing windows
- Never set performance metrics that the influencer cannot influence — views are a metric, algorithm reach is not
- Never skip the disclosure requirements section — FTC/ASA compliance is mandatory, not optional
- Never list deliverables without specifying format, dimensions, and platform specs

infra-as-code-review

- Never mark a finding as Low if it involves hardcoded credentials or secrets in any form — always Critical
- Never review IaC in isolation from the deployment context — networking and IAM must be evaluated together
- Never produce narrative findings without the specific resource name, file, and line number
- Never skip the "Required Actions Before Merge" summary — reviewers need a clear blocking list, not just a full report
- Never approve code where encryption at rest or in transit is missing on data stores, even if not explicitly flagged by the requester

instagram-post-downloader

- Never attempt to download private posts or content behind a login wall — this skill is for public posts only
- Never ignore 429 rate-limit responses — always implement a backoff wait before retrying

- Never save all downloads to a single flat folder when processing multiple accounts — use named subfolders per source
- Never skip PDF stitching for carousel posts — individual slides delivered without a combined PDF are incomplete output
- Never proceed if Instagram returns a login wall — surface the limitation clearly rather than returning an error silently

insurance-claim

- Never inflate or invent losses — it risks the whole claim and is fraud
- Never be vague about amounts or dates — itemise and date everything
- Never omit or fail to reference evidence — undocumented claims stall
- Never ignore the denial reason in an appeal — rebut it specifically with the policy terms
- Never present this as legal/insurance advice or guarantee an outcome — flag deadlines to confirm

interview-me

- Never fire a questionnaire — seven questions at once produces skim-answers and resentment
- Never interview when the brief is already clear — process applied without judgment is friction
- Never ask questions whose answers wouldn't change the deliverable — every question spends the requester's patience
- Never start building mid-interview "to save time" — half-brief work anchors the requester to the wrong draft
- Never skip the playback — the interview's value is captured only when the requester says "yes, that"

interview-prep

- Never produce a generic question list — prep is only useful when it's for this round at this company
- Never write fabricated achievements into STAR answers — build from the candidate's real stories
- Never over-script — answers should be structured talking points, not memorised paragraphs that sound robotic
- Never dodge the candidate's weak spots — rehearse an honest, confident response instead of hoping it won't come up
- Never ignore the round type — a behavioural prep and a case prep are different documents

interview-question-bank

- Never produce a flat question list with no evaluation guidance — that's how interviews stay inconsistent
- Never use brain-teasers or trivia that don't predict job performance
- Never let every interviewer assess the same competency — map and distribute
- Never include questions about age, family status, health, religion, or other protected areas
- Never score on "culture fit" gut feel — score on observable, job-related evidence

investing-policy-statement

- Never recommend specific stocks, funds by ticker, or "hot" assets — stay at the asset-class level
- Never set an allocation that ignores the stated time horizon (e.g. all-equities for money needed next year)
- Never omit the behavioural guardrails — they're the point of an IPS
- Never imply guaranteed returns or market-timing works
- Never present this as personalized financial advice

investor-cold-email

- Long backstory before the hook
- Generic flattery ("I love your work")
- Multiple asks or a vague one
- A follow-up that just says "bumping this" with no new information

investor-pitch-deck

- Never include a "no real competitors" slide — every company has competition and investors will discount founders who claim otherwise
- Never use a top-down TAM calculation without a bottoms-up validation — investors distrust pure top-down market sizing
- Never leave the ask vague — specify the amount, use of funds, and 18-month milestones the funding enables
- Never let traction slides show vanity metrics — focus on revenue, retention, and growth rate over downloads and signups
- Never bury the problem slide — investors must understand and feel the pain before they care about the solution

investor-update

- Never omit challenges or bad news — sanitised updates erode investor trust faster than bad results do

- Never bury the lead — use BLUF structure and put the most important news in the first paragraph
- Never send an update without a clear "Ask" section — investors who want to help need to know how
- Never use buzzwords or spin — investors see hundreds of updates and will see through vague positive language
- Never report metrics without a comparison baseline — numbers without context (vs. last period or target) are meaningless

invoice-generator

- Never invent a tax rate or tax treatment — flag it to confirm with an accountant
- Never omit the invoice number or due date — they're what makes it trackable and payable
- Never leave payment instructions vague — say exactly how to pay
- Never miscompute totals — show the math so it can be checked
- Never present this as tax/legal advice — it formats an invoice, it doesn't certify compliance

iso-27001-isms

- Never exclude a control without a written justification — silent exclusions are audit findings
- Never build the SoA before the risk assessment — applicability is *derived* from risk, not guessed
- Never treat Annex A as the whole standard — clauses 4–10 (the management system) are mandatory and where many fail
- Never mark controls "implemented" without evidence of operation — certification audits sample evidence
- Never present this as certification — only an accredited body certifies; this prepares the ISMS

jd-decoder

- Never treat every listed requirement as mandatory — most are wishes; the skill's value is telling which few aren't
- Never give a flattering fit read — a candid "stretch, here's the gap" is more useful than false confidence
- Never ignore tone and ordering — they often reveal more than the bullet list
- Never invent company facts — decode the text given; flag what needs separate research (pair with company-brief)
- Never skip the red flags — helping someone *not* apply to a bad role is a real outcome

job-application

- Never fabricate or embellish experience — only use real achievements from the provided CV
- Never use the same cover letter template for every role — every letter must reference specific details of the job description
- Never address selection criteria that aren't in the JD — match keywords the employer actually used
- Never omit ATS optimisation — ensure role-specific keywords from the JD appear naturally in the CV summary
- Never write a cover letter that re-summarises the CV — it must add context and motivation, not repeat bullet points

job-description-writer

- Never include years-of-experience requirements unless legally necessary — they exclude qualified candidates and may create legal risk
- Never list "nice to have" items in the requirements section — separate mandatory from desirable clearly
- Never use gendered or exclusionary language — run the inclusive language check before finalising
- Never write a responsibilities section with more than 8 items — prioritise the most important duties
- Never omit compensation range where legally required or culturally expected — hiding salary deters qualified candidates

job-story-mapper

- Never write job stories that describe a feature rather than a situation-motivation pair
- Never skip the social and emotional dimensions — mapping only functional jobs misses the most defensible differentiation opportunities
- Never define situations too broadly ("as a user who wants to manage their work") — the situation must be a specific moment or trigger
- Never conflate opportunity scoring with priority — a high opportunity score still requires feasibility and strategic fit assessment
- Never produce a job map without identifying current workarounds — the workaround reveals what the job is worth to the customer

kb-audit

- Never prioritise by what's easy to write — prioritise by what deflects the most tickets
- Never ignore duplicates — overlapping articles split search ranking and confuse users; merge them
- Never treat all gaps equally — a gap on a top-5 ticket driver outranks ten niche ones
- Never skip findability — a perfect article with a bad title that no one finds deflects nothing

- Never audit without the ticket data if it exists — it's the map of what actually matters

landing-page-copy

- Never offer competing CTAs — multiple asks split attention and lower conversion; one goal per page
- Never open with "Welcome to [company]" — lead with the visitor's outcome
- Never list features without benefits — visitors buy outcomes, not specs
- Never hide the price/effort/objections — unanswered doubt is a silent exit
- Never write "Submit"/"Learn more" buttons — say what happens and the value

last-30-days-research

- Never treat SEO blog posts or vendor-authored content as community signal — only count independent sources
- Never report findings without applying the date filter — prior knowledge mixed with recent search results produces stale, unverifiable claims
- Never fabricate or guess at URLs — every link in the Sources section must have been retrieved during the research session
- Never report a single mention as a "finding" — a finding requires corroboration from at least two independent sources
- Never rate Signal Confidence as High when fewer than 5 credible sources were found — this misleads the reader about how much to rely on the output

launch-post

- Never use marketing hype ("revolutionary", "game-changing") — devs downvote it
- Never hide limitations — naming them earns trust and pre-empts the top comment
- Never bury the what-it-does under backstory — lead with substance
- Never make claims without proof — show the code/benchmark/demo
- Never write a generic post — tune tone and format to the actual channel

launch-readiness

- Never mark a function as "Ready" without evidence — green status must be backed by a completed checklist item, not an assumption
- Never issue a Conditional Go without specifying exactly what conditions must be met and by when — vague conditions are not conditions
- Never treat the rollback plan as complete unless it has been tested in staging, not just documented
- Never assign blockers to "the team" — every blocker must have a single named owner or it will not be resolved before launch

- Never skip the analytics verification step — unverified tracking events mean the launch will be invisible and cannot be evaluated

launch-tiering-framework

- Never launch everything at T1 — attention and effort are finite
- Never treat a strategic launch as T3 because the team is busy — flag the gap
- Never skip enablement on a tier that sales needs to sell
- Never measure a T3 with T1 metrics (or vice versa)
- Never ignore existing house tier definitions if provided

layoff-communication

- Never hide behind euphemism ("rightsizing", "graduating talent") — name it plainly and humanely
- Never let affected people learn via the all-hands, press, or rumour — sequence individuals first
- Never use passive voice to dodge accountability — leadership owns the decision
- Never over-promise or give false hope about reversal or rehire
- Never treat this as legal advice — flag jurisdiction-specific obligations for counsel

legal-brief

- Never present uncertain legal positions with confident language — areas of legal ambiguity must be flagged explicitly, not smoothed over
- Never omit the disclaimer — every legal brief output must include the professional review caveat before the user treats it as advice
- Never structure the brief chronologically — IRAC format (Issue, Rule, Application, Conclusion) must be used regardless of how the user framed the request
- Never cite cases or statutes from memory without flagging them as [REQUIRES VERIFICATION] — hallucinated citations are worse than no citations
- Never conflate jurisdiction — legal positions in England & Wales, US, and EU can differ materially; always confirm jurisdiction before stating the rule

lesson-plan

- Objectives that can't be observed or measured ("students will appreciate...")
- A flow that's all teacher talk with no student practice
- No formative checks until a final test
- One-size-fits-all with no differentiation

lifecycle-crm-plan

- Never batch-and-blast — untriggered, irrelevant sends train users to ignore and unsubscribe
- Never measure success by opens/clicks alone — tie journeys to the lifecycle outcome (activation, retention, revenue) with a holdout
- Never forget exit conditions — a user who already activated should not keep getting "activate now" emails
- Never ignore frequency capping — overlapping journeys are how you fatigue and burn a list
- Never skip deliverability guardrails — a great journey in the spam folder reaches no one

linkedin-profile

- Never make the headline just your title — it's prime keyword + value real estate
- Never bury the hook — the opening lines are all most viewers see; don't waste them on "passionate professional"
- Never write About in third person — LinkedIn is personal; "I" converts better
- Never ignore keywords — recruiters filter by them; a profile without them is invisible to search
- Never copy the resume verbatim — LinkedIn is warmer and slightly more narrative

literature-review

- Never summarise papers one by one — evidence must be synthesised thematically across multiple studies, not presented as a sequence of abstracts
- Never omit methodological critique — a literature review that only reports findings without assessing study quality is not a critical review
- Never organise by chronology when thematic organisation is possible — chronological reviews bury the conceptual structure of the field
- Never present contested findings as settled consensus — where evidence is mixed, name both sides and why the evidence diverges
- Never skip the gap analysis — identifying what the field still needs is a core deliverable, not an optional addition

llm-cost-latency-budget

- Never budget on average latency — users feel the p95, and the tail is where AI features feel broken
- Never default every call to the most capable model — most requests don't need it; tiering often cuts cost by more than half
- Never forget output tokens cost more than input — verbose responses are often the hidden cost driver
- Never ship without a spend cap and alert — an unbounded LLM feature is an unbounded bill
- Never optimise cost before measuring it — itemise the real token usage first, then pull the biggest lever

llm-guardrails-spec

- Never rely on the system prompt alone — prompt-only guardrails are bypassable; defend in layers
- Never trust retrieved or tool-returned content as instructions — that's the injection vector
- Never grant the model broad tool/action access "for flexibility" — least privilege, allowlist
- Never ship without a red-team set — untested guardrails are decoration
- Never log raw prompts/outputs with PII or secrets in the name of debugging

load-testing-plan

- Never set thresholds without grounding them in actual SLO targets or production baselines — arbitrary numbers produce meaningless pass/fail results
- Never run the load generator on the same host as the service under test — this contaminates both the test results and the service metrics
- Never use production user data in load test seeding — all test data must be synthetic, tagged, and cleaned up after each run
- Never skip the soak test on first deployment — only a soak test reveals slow memory leaks and connection pool exhaustion that short tests miss
- Never treat a passing baseline test as evidence the service handles spikes — baseline, stress, spike, and soak scenarios test fundamentally different failure modes

local-dev-setup

- Never write setup steps from memory without testing them on a clean machine — steps that skip implicit knowledge break for new engineers
- Never leave environment variables undocumented — every variable in .env.example must appear in the Variables table with a description and source
- Never write troubleshooting entries for theoretical issues — only include problems that have actually occurred during real onboarding sessions
- Never assume Docker Desktop is configured correctly — memory limits and platform (M1/M2) compatibility must be explicitly called out
- Never omit expected output for key commands — without "expected output", engineers cannot tell whether a step succeeded or silently failed

localization-brief

- Never equate localization with translation — payments, legal, formats, and imagery decide whether it feels native
- Never ignore region — fr-FR ≠ fr-CA, es-ES ≠ es-MX; the variant changes copy, formats, and norms
- Never localize SEO by translating keywords — research how locals actually search

- Never skip local payment methods — the best-localized UI converts nothing if they can't pay how they pay
- Never launch without in-market native review — machine/relay translation misses the embarrassing stuff

managing-up

- Never bring a problem with no recommendation — "what should I do?" offloads your job; bring options + a pick
- Never communicate in your preferred style — match theirs (a detail-lover and a headline-skimmer need different messages)
- Never surprise your manager — surfacing a risk late is the fastest way to lose trust
- Never escalate everything (looks like you can't decide) or nothing (looks like you hide things) — calibrate
- Never frame the ask around what *you* want — connect it to what *they're* measured on

marketing-funnel-plan

- Never pour budget into the top of the funnel when the leak is mid-funnel — more visitors through a leaky funnel just wastes more money
- Never list every channel — focus beats breadth; name the 1–2 that fit the motion
- Never set tactics without a metric and target per stage — unmeasured tactics can't be cut
- Never treat attribution as truth — state its limits and lean on leading indicators
- Never ignore retention/referral — acquisition-only funnels buy growth they can't keep

marketing-psychology

- Never invent fake scarcity, countdowns, or fake social proof — it's a dark pattern and it backfires
- Never list every principle — pick the few that fit the specific friction
- Never stay theoretical — every principle needs a concrete application to the asset
- Never use confirm-shaming, forced continuity, or hidden costs — short-term lift, long-term trust loss
- Never ignore friction — sometimes the fix is removing a step, not adding persuasion

marketplace-listing-optimizer

- Never stuff the title with every keyword — relevance + readability beat a keyword soup the algorithm discounts
- Never repeat title keywords in the backend field — it wastes indexable space
- Never optimize keywords while ignoring conversion (images, reviews, price) — ranking without conversion decays

- Never invent search volume or claims — flag them to verify with the marketplace's tools
- Never give a flat list — rank fixes so the seller knows what to do first

mcp-server-spec

- Never mirror the REST API — 40 endpoint-tools is the #1 way MCP servers fail
- Never write descriptions for developers ("wraps the /v2/items endpoint") — write them for a model choosing a tool
- Never return full API payloads — context windows are the scarce resource
- Never expose destructive actions ungated because "the client will be careful"
- Never skip the never-exposed list — an MCP server without one hasn't been threat-modelled
- Never ship without running the agent test plan — schema-valid and agent-usable are different properties

media-pitch

- Never write a pitch that leads with the company's history or description — the story angle must come first, not who the company is
- Never use vague data points ("significant growth", "thousands of users") — every statistic must be specific and verifiable
- Never send the same pitch to multiple journalists in a BCC — pitches must be individually tailored to each journalist's beat and recent work
- Never offer an exclusive without setting a response deadline — an open-ended exclusive invitation is ignored or used to delay indefinitely
- Never follow up with "just checking in" — a follow-up must contain new information or a fresh angle, otherwise it is noise

meeting-notes

- Never assign action items to "the team" or "everyone" — every action item must have exactly one named owner or it will not be completed
- Never capture verbatim transcript content — meeting notes record decisions and commitments, not the full conversational path to get there
- Never omit the context for decisions — a decision without its rationale is useless when someone asks "why did we do that?" six months later
- Never leave open questions without an owner and deadline — an unanswered question with no follow-up assigned is a blocked decision
- Never delay sending notes beyond 2 hours after the meeting — notes sent the next day miss the window when action item owners can act on commitments while fresh

messaging-framework

- Never lead with features — buyers care about the outcome; features are proof, not the message
- Never make claims without proof — an unbacked superlative ("the best", "revolutionary") reads as noise
- Never try to speak to everyone — messaging for all audiences resonates with none; pick the segment
- Never use internal jargon the customer wouldn't say — if they can't repeat it, it won't spread
- Never confuse this with positioning — decide the category/competitive frame first (see product-positioning-doc), then write the words

metric-semantic-layer

- Never leave the definition fuzzy — "active users" without the exact rule is how three dashboards disagree
- Never omit default filters — if one tool counts test accounts and another doesn't, the metric is broken
- Never ignore additivity — summing a non-additive metric (like a distinct count) across days gives a wrong number
- Never define metrics in BI tools instead of the semantic layer — that's how definitions fork
- Never skip timezone/null/dedup edge cases — they cause the subtle, hard-to-find discrepancies

metric-tree-builder

- A "tree" that's just a flat list of unrelated KPIs
- Stopping at output metrics no one can directly move
- Ignoring how drivers combine (treating everything as additive)
- No view on which lever actually matters most

metrics-framework

- Never set a North Star metric that measures business activity (revenue, pageviews) rather than customer value delivered — this creates incentives misaligned with product quality
- Never define metrics without specifying the formula or data source — an ambiguous metric will be measured differently by different people
- Never skip counter-metrics — optimising any single metric without a guard rail will eventually produce perverse incentives
- Never include more than 4–5 metrics in a daily team view — a dashboard with 20 metrics is a dashboard nobody looks at
- Never classify all metrics as "leading" — be honest about which are lagging outcome metrics and which genuinely predict future outcomes

microcopy-writer

- Never use vague labels ("OK", "Submit", "Click here") when a specific verb communicates the outcome
- Never write clever copy that obscures what the button does — clarity beats personality at decision points
- Never ignore character limits or the existing voice — microcopy must fit the UI and the brand
- Never hide consequences behind a generic confirmation — name what will happen
- Never invent product terms — reuse the established vocabulary for consistency

microservices-decomposition

- Never define service boundaries before completing the domain analysis — services derived without bounded context mapping will split the wrong things and couple the wrong things
- Never assign multiple teams as co-owners of a single service — shared ownership is no ownership; every service needs exactly one team accountable for it
- Never default to synchronous REST calls for all inter-service communication — using sync calls where async events would decouple services creates cascading failure modes
- Never propose more than one service per bounded context without a clear justification — over-decomposition (nanoservices) creates operational overhead that exceeds the decomposition benefit
- Never begin migration without deploying distributed tracing first — migrating without observability means flying blind when the first extraction causes a production incident

mind-map

- Never produce a flat list dressed up as a map — there must be real hierarchy
- Never make one branch huge and the rest empty — balance the structure
- Never use long sentences as nodes — keep them to a few words
- Never break indentation — Mermaid mindmaps derive structure from it
- Never silently drop ideas from the source — place or park them

model-card

- Never report a single aggregate metric and call evaluation done — averages mask the slices where a model fails worst
- Never leave "intended use" open-ended — an undefined boundary is an invitation to misuse
- Never omit known biases because they're uncomfortable — an undocumented risk is a worse liability than a documented one
- Never present accuracy without the class balance / base rate — 95% accuracy on a 95/5 split is meaningless

- Never ship without a monitoring plan — a model card without a rollback trigger is a snapshot, not a contract

model-migration-plan

- Never skip shadow because offline evals passed — real traffic finds what golden sets miss
- Never migrate the feature and the prompt redesign in one change — you won't know which moved the metrics
- Never compare models with an unpinned judge, or a judge that is the target model grading itself
- Never leave the old model path in code indefinitely "just in case" — set the removal date in the plan
- Never treat a cheaper model as free savings without re-checking quality at the tails, not just the mean

model-selection-advisor

- Never default to the biggest model "to be safe" — pay only for the capability the task needs
- Never pick on price alone — a cheap model that fails the bar costs more in rework and trust
- Never recommend without an eval to confirm the quality bar is actually met
- Never hardcode a single model name as the answer — reason by tier and let the eval pick the current best in it
- Never ignore the long tail — design for the hard cases via escalation, not by oversizing everything

monitoring-setup-guide

- Never create alerts without a specific on-call action — an alert that just says "investigate" trains engineers to ignore it
- Never set alert thresholds from a template without calibrating against production baselines — uncalibrated thresholds cause either alert fatigue or missed incidents
- Never log PII, tokens, or secrets — a logging standard is incomplete without an explicit list of what must never be logged
- Never measure only the four golden signals without adding at least one business metric alert — infrastructure health can be green while the business-critical path is silently failing
- Never deploy distributed tracing without verifying that trace IDs propagate across all service boundaries — partial tracing is worse than no tracing because it produces misleading incomplete traces

morning-intelligence

- Never skip the interview and write a generic master prompt — a brief that is not tailored to the user's specific role and topics will be ignored after the first day

- Never proceed to write the master prompt without confirming the "What I Heard" summary — errors in the summary will silently propagate into a prompt that produces the wrong briefing every morning
- Never use broad topic labels in the master prompt (e.g. "AI", "tech news") — every topic must have a specific angle or focus to produce signal-to-noise ratio worth reading
- Never omit the NEVER INCLUDE section — without explicit exclusions, the briefing will fill with noise that the user said they wanted filtered out
- Never ask all 15 questions at once — the interview must run one question or small group at a time to produce specific, considered answers
- --

multi-source-signal-synthesiser

- Never echo surface-level feature requests as insights — translate every request to the underlying need before including it as a finding
- Never assign High confidence to insights supported by only one source type — confidence requires corroboration across at least two distinct source types
- Never treat all sources as equally weighted — a single interview quote and a pattern across 200 support tickets are not comparable signals
- Never collapse divergent signals into a single finding — where user segments have genuinely different needs, name the segments explicitly rather than averaging them away
- Never omit the research gap section when key decisions rest on thin data — acting on low-confidence findings without flagging the gaps misleads product teams

nda-analyser

- Never present the analysis as legal advice — the disclaimer must appear prominently and the output must recommend qualified legal review before any signing decision
- Never skip the residuals clause check — residuals clauses allow the receiving party to use disclosed information from memory, which is one of the highest-risk provisions in any NDA
- Never evaluate only the clauses explicitly flagged by the user — a complete analysis must cover all standard clause types even if the user only asked about one
- Never assess breadth of the confidentiality definition without checking for oral disclosure coverage — oral disclosures with no written confirmation requirement are a common enforcement gap
- Never omit the plain English verdict — a clause-by-clause analysis without a summary conclusion leaves the user unable to act on the findings

net-worth-statement

- Never inflate asset values — use realistic current/market values, not purchase prices
- Never omit liabilities or net them silently — show both sides

- Never present a one-time number as the goal — emphasize the trend
- Never ignore liquidity — high net worth that's all illiquid is a real risk worth flagging
- Never present this as personalized financial advice

newsletter-writer

- A subject line that's a label ("Newsletter #42")
- Burying the value under a long personal preamble
- Three competing CTAs, or none
- A wall of text with no sub-heads or callout — unskimmable

notebooklm-connector

- Never proceed with any browser action before the full request has been parsed into discrete steps — ambiguous source references must be clarified before navigating
- Never guess at source URLs if the user says "add my research sources" without specifying them — ask for the explicit list before starting
- Never batch actions speculatively — each action must be confirmed complete before the next one begins to avoid compounding failures
- Never wait for audio overview rendering to complete — audio overviews take 5–15 minutes server-side; report the trigger and move on rather than blocking the session
- Never attempt this skill if the Claude Chrome extension is not active — report the missing prerequisite immediately rather than attempting browser steps that will fail

notes-humanizer

- Never remove all em dashes — only the ones functioning as parenthetical substitutes should be removed; genuine dramatic pauses are valid
- Never invent problems to justify changes when the original is already clean — report what was found honestly, even if the answer is "this text is mostly fine"
- Never add the aside or opinion generically — the injected human signals must connect to an actual claim or argument in the text, not float in as decoration
- Never list changes by category in the change log — every individual change must be listed separately with the specific reason for that specific instance
- Never apply humanisation changes that alter the factual claims or intended meaning of the original text — the skill rewrites style, not substance
- --

offer-letter

- Never present this as legally vetted — flag terms for HR/legal and don't assert jurisdiction-specific law

- Never make it cold and purely transactional — the offer is also a recruiting moment
- Never omit contingencies or the expiry date — ambiguity causes problems later
- Never invent benefits, equity terms, or PTO numbers — mark them to confirm
- Never send the written offer with no verbal first — surprises lose candidates

okr-builder

- Never accept output-based key results — any KR phrased as "launch X" or "complete Y" must be rewritten as an outcome with a baseline and target
- Never write OKRs without asking for the company or product North Star — OKRs disconnected from the strategic context are just a goal-setting exercise
- Never write more than 4 KRs per objective — too many KRs dilute focus and make scoring ambiguous at quarter end
- Never use binary KRs (ship/don't ship) — every KR must be scorable on a 0.0–1.0 scale based on degree of achievement
- Never skip the health check section on baselines — OKRs without current baselines cannot be scored objectively at quarter end

onboarding-copy

- Never tour every feature — guide to the first win; the rest can be discovered
- Never write blocking, un-skippable walls of modals — let users get to the product
- Never explain what's obvious ("This is the menu") — spend words where there's real friction
- Never forget the success moment — activation should feel rewarded
- Never be generically chirpy — encouragement should be specific to what they just did

onboarding-plan

- Never produce a generic plan that could apply to any role — the plan must reference the specific role, team, tools, and priorities provided, not use placeholder text
- Never skip the Before Day 1 manager checklist — IT access and system provisioning failures on day 1 destroy first impressions and waste the new hire's first week
- Never set milestones without distinguishing between the orient, learn, contribute, and lead phases — collapsing phases produces plans where new hires are expected to lead before they understand the product
- Never omit the 90-day review questions — the review is the accountability mechanism for the entire plan, and skipping it makes the milestones meaningless
- Never treat the plan as a task list — each phase should have a clear theme and a milestone that describes an observable capability, not just a set of completed activities

oncall-runbook

- Never write alert runbooks with vague diagnostic steps like "check the logs" — every step must specify the exact command, dashboard link, or query to run
- Never include an alert in the runbook that has no specific on-call action — an alert that pages someone with no defined response path creates panic, not resolution
- Never leave the rollback command undocumented or untested — a rollback procedure that has never been run will fail when needed most
- Never list escalation contacts without phone numbers and Slack handles — email-only escalation paths are useless during a 3am incident
- Never write the runbook once and treat it as permanent — runbooks go stale after incidents; every incident must trigger a review of the relevant runbook entries

one-on-one-prep

- Never turn the 1:1 into a status report — status belongs in writing; use the live time for decisions, feedback, and growth
- Never avoid the hard topic — name it, framed constructively; the 1:1 is the safest place to raise it
- Never arrive without an ask — "anything you need?" wastes the leverage
- Never let career/growth fall off when things are busy — it's the first thing dropped and the most costly
- Never over-pack — 3 real topics beat 10 skimmed

one-pager

- Never overflow the page — a "one-pager" that's two pages has failed its only constraint; cut content, not font size
- Never bury the ask — the reader must finish knowing exactly what to do next
- Never write an abstract problem ("inefficiencies in the market") — name the concrete pain and who feels it
- Never list features instead of the outcome — lead with what it does for the user
- Never make claims without proof — one real metric beats three adjectives

open-house-plan

- Never treat it as just unlocking the door — it's a promoted, lead-generating event
- Never skip lead capture — foot traffic with no contacts is a wasted Saturday
- Never forget follow-up — leads go cold within a day
- Never under-promote — most attendance comes from the days-before push
- Never ignore agent safety and securing valuables during the open house

org-chart

- Never invent reporting lines that weren't given — chart only what's known, flag gaps
- Never mix solid and dotted lines arbitrarily — solid = direct, dotted = indirect
- Never flatten a real hierarchy into a list — show the levels
- Never break Mermaid with special characters in names/titles
- Never editorialize on individuals — structural observations only

outcome-tracker

- Never reinterpret a claim after the fact so it scores as a hit — the original wording is the contract
- Never record point estimates when the author thinks in ranges — bands are honest, points are theatre
- Never let a framework take credit for hits and blame "execution" for misses — score the prediction as made
- Never compute calibration on fewer than ~10 resolved predictions per framework — report "insufficient history" instead
- Never skip recording because the decision feels obvious — obvious bets that miss are the most valuable calibration data

outreach-message

- Never write a long message — every extra sentence lowers the reply rate
- Never make it about you — lead with why *them*, then a tight credibility line
- Never use a generic hook ("I came across your profile") — it signals a mass blast
- Never stack multiple asks — one easy request, or none will be answered
- Never be pushy in the follow-up — one graceful nudge, then stop

paid-acquisition-plan

- Never set budgets before deriving the CAC ceiling from unit economics — spending you can't recoup is just buying revenue at a loss
- Never trust platform-reported conversions as truth — every channel over-claims; verify with incrementality
- Never under-invest in creative testing — in modern paid, creative beats targeting as the primary lever
- Never scale a winner or kill a loser inside the learning phase — let it gather signal first
- Never spread a small budget across many channels — concentrate until a channel proves out

parent-communication

- Labelling the child instead of describing the behaviour
- Jargon or edu-speak parents won't parse

- A concern with no path forward or offer of support
- Over-long; burying the point under throat-clearing

partnership-proposal

- Never write the value proposition from your own perspective — the "For Partner" section must be written from the partner's point of view, in the language of their goals and their customers
- Never leave commercial terms as structure without numbers — a proposal that says "revenue share" without stating the percentage is not a proposal, it is a conversation opener
- Never omit the "What we're not proposing" section — leaving unstated assumptions creates misaligned expectations that derail negotiations later
- Never set success metrics unilaterally — metrics that only your company controls or cares about will not earn partner commitment
- Never write a go-to-market plan with "TBD" owners — every activity must have a named owner on at least one side before the proposal goes out

patient-communication

- Never use medical jargon without a plain-English explanation — write for the patient, not the clinician
- Never omit a clear "next steps" section — patients must know exactly what to do after reading
- Never produce final content without flagging that clinical review is required before sending
- Never write above a Grade 8 reading level without a compelling reason — accessibility is the default
- Never include Latin abbreviations (e.g. "p.r.n.", "b.d.") without spelling them out — they are not universally understood

paywall-optimization

- Never gate the core aha — users who never feel value never pay
- Never hit users with the wall on first open, before any value — it just bounces them
- Never use dark patterns (hidden cancel, forced continuity, fake urgency) — short lift, long-term churn
- Never optimize conversion while ignoring churn/refund guardrails
- Never present plans without a clear recommended/anchor option — choice overload kills conversion

pentest-report

- Never omit the authorization/scope — an unbounded, unauthorized-looking report is unusable and unsafe

- Never give a severity without impact and remediation — clients fix what they understand and can prioritize
- Never write findings only engineers can read (or only execs) — serve both audiences in their sections
- Never leave evidence unredacted — protect the very data you're helping secure
- Never produce this for testing that wasn't authorized in writing

performance-budget

- Never set budget thresholds without measuring a current baseline first — targets must be anchored to reality
- Never define global averages only — critical user journeys need individual budgets as they may diverge significantly
- Never omit CI enforcement — a performance budget that is not enforced in the build pipeline will not be respected
- Never leave the breach response process without named owners and escalation channels
- Never set budgets that apply only to one environment — production and staging targets should be documented separately if they differ

performance-review

- Never inflate positive language to avoid difficult feedback — growth areas must be clearly stated, not buried
- Never include feedback that isn't supported by specific examples — every development point needs evidence
- Never write a review that only covers what happened in the last month — the full review period must be considered
- Never omit development goals — a review without forward-looking guidance is incomplete
- Never use language that could be read as discriminatory — avoid references to personality traits unrelated to work performance

personal-bio

- Never open with empty adjectives — "an experienced, passionate professional" says nothing; lead with the proof
- Never make the three versions inconsistent — they should be the same story at different resolutions
- Never stuff every accomplishment into the short bio — pick the strongest; that's what "short" means
- Never use buzzword filler ("synergy", "thought leader") — specifics earn credibility, labels don't

- Never forget the audience — a conference bio and a startup about-page emphasise different things

persuasion-brief

- Never lead with the reason that persuades *you* — lead with what moves *them*
- Never rely on logic alone — people decide on emotion and justify with logic; address both
- Never ignore the unspoken objection — the stated reason ("no budget") often hides the real one (risk/ego)
- Never ask for the big commitment first — a reversible pilot is far easier to say yes to
- Never manipulate — use true reasons; a win built on a distortion costs you the next ask

pm-weekly-review

- Never report metrics without comparing to target or the prior week — absolute numbers without context are not useful
- Never list blockers without a named owner and proposed resolution — unowned blockers stay blocked
- Never write a weekly review that is longer than one page — it must be scannable in under 2 minutes
- Never include more than 3 priorities for next week — a list of 8 "top priorities" means nothing is prioritised
- Never skip the insights section — observations that inform future decisions are a PM's key value add

policy-memo

- Never bury the recommendation at the end — decision-makers read the top
- Never present a fake menu (one real option + straw men) — options must be genuine
- Never dump all the research — include only what's needed to decide; annex the rest
- Never hedge into non-recommendation — name a choice and own the trade-off
- Never ignore feasibility/cost/politics — an un-implementable recommendation is useless

portfolio-page

- Never list artifacts without the story — a screenshot with no context proves nothing
- Never blur your contribution into the team's — "we shipped" leaves the viewer unsure what you did
- Never omit outcomes — "redesigned the flow" without a result is a task, not a case study
- Never pad with weak projects — each extra mediocre one dilutes the strong ones
- Never leak confidential data — anonymise to ranges instead of dropping the impact entirely

pptx-slide-auditor

- Never flag stylistic preferences as issues — only report genuine layout problems, overflow, and consistency errors
- Never produce a flat list of issues — group by severity (Critical / Major / Minor) so fixes can be prioritised
- Never skip slides without commenting — every slide must have an explicit pass or issue status
- Never suggest redesigning content — the audit scope is layout, consistency, and readability, not messaging
- Never report the same issue type repeatedly across slides without summarising the pattern — consolidate repeated issues

pr-crisis-response

- Never go silent or delay — issue a holding statement, then update; absence writes the story for you
- Never speculate, guess at cause, or admit unverified fault — acknowledge and commit to updates instead
- Never let channels drift off-message — one core message, tailored, not contradictory versions
- Never forget employees — they're your first responders and they'll hear it anyway
- Never over-spin — minimising or blaming others erodes the trust you're trying to keep

pr-description

- Never just paste the commit list — explain intent the diff can't convey
- Never say "tested" without saying how — give the reviewer something to trust
- Never hide risk or migrations — surface them so they're reviewed deliberately
- Never write a novel for a one-line change — match effort to size
- Never omit the "what to focus on" — undirected review is slow review

pr-description-writer

- Never write a description that only restates what changed — explain why the change was made
- Never skip the testing steps — reviewers need to know how to verify the change works
- Never omit the reviewer notes for high-risk PRs — flag deliberate trade-offs and areas needing careful review
- Never describe implementation details that are obvious from the diff — add context that the diff cannot convey

- Never produce a single paragraph — structure with headers so reviewers can navigate to what they need

prd-template

- Never write requirements from the company's perspective — every requirement must trace back to a user need
- Never include vague requirements like "the system should be fast" — every requirement must be testable
- Never conflate MVP with future phases — be explicit about what is and is not in scope for the first release
- Never leave success metrics as percentages without baselines — specify the current state and the target
- Never skip open questions — unresolved assumptions are risks; surfacing them is the PM's job

press-release

- Never bury the news — the most important information must appear in the first paragraph (inverted pyramid)
- Never use promotional language or superlatives — press releases must read as news, not advertising copy
- Never omit the boilerplate — every press release needs the standard "About [Company]" paragraph at the end
- Never forget the embargo date and media contact — journalists need both to use the release
- Never write a headline longer than 12 words — it must be scannable and specific

pricing-calculator

- Never model a price rise assuming volume holds — always state an elasticity, even a conservative one
- Never compute margin on revenue — use contribution (price – variable cost)
- Never present one elasticity as fact — show the break-even elasticity so the reader judges the risk
- Never ignore fixed costs in break-even — contribution must cover them before profit
- Never confuse this with strategy — the number doesn't decide the model/packaging; pair with pricing-strategy

pricing-page-copy

- Never list raw features with no grouping or value framing
- Never hide the value metric or make overages ambiguous

- Never over-anchor with fake "most popular" if it isn't
- Never promise guarantees or terms you weren't given
- Never use the same CTA on a self-serve and an enterprise plan

pricing-sensitivity-model

- Never fabricate or extend survey responses — with no data, deliver the survey design and stop
- Never read OPP as "the right price" — it is the least-resisted price, and premium strategies ignore it on purpose
- Never hide the dropped respondents — non-monotone answers are evidence about the survey, not noise to delete
- Never report a single point without the range — the range is the finding; the point is a summary of it
- Never pool segments that obviously differ (SMB with enterprise) — the pooled curves cross somewhere nobody actually is

pricing-strategy

- Never base pricing solely on cost-plus — pricing must reflect value delivered to the customer
- Never design tiers where the middle tier is clearly worse value — it undermines trust and pushes customers to extremes
- Never change pricing without a migration plan for existing customers — surprise price changes cause churn
- Never set enterprise pricing as "contact us" without a floor — it deters self-serve evaluation and qualification
- Never skip competitive positioning — pricing in isolation from the market is incomplete strategy

prior-authorization-letter

- Never invent diagnoses, codes, dates, prior treatments, or evidence to strengthen the case
- Never be vague about the request — name the exact service/drug and codes
- Never ignore the stated denial reason in an appeal — address it head-on
- Never present this as medical advice or submit without clinician review and signature
- Never pad with generic boilerplate that buries the patient-specific justification

privacy-policy-drafter

- Generic boilerplate that doesn't match what the product does
- Claiming GDPR/CCPA compliance as a fact rather than reflecting practices

- Vague "we may share with third parties" with no categories or purpose
- Overpromising security ("your data is 100% safe")

process-documentation

- Never write steps without specifying who is responsible for each — ownership must be explicit throughout
- Never omit the escalation path — every process must say what happens when something goes wrong
- Never document the ideal process if the real process differs — document reality, then note improvements separately
- Never skip edge cases and exceptions — they are where most process failures actually occur
- Never produce documentation without a review date — undated process docs quickly become incorrect

product-description

- Never list features without their benefit — "5000mAh battery" means nothing without "2 days without a charge"
- Never keyword-stuff — it reads as spam and channels penalise it
- Never invent specs or make unverifiable claims (waterproof, organic, FDA-approved) — flag to confirm
- Never bury the value in a long paragraph — shoppers skim, lead with the hook and bullets
- Never ignore the channel's limits (Amazon title/bullet lengths, etc.)

product-health-analysis

- Never report a single aggregate metric without segment breakdowns — averages hide opposing trends
- Never flag a metric as healthy just because it is above the target — check if the target itself is meaningful
- Never list metric movements without root cause hypotheses — observations without explanations are not analysis
- Never mix product health metrics with business KPIs without explaining the relationship between them
- Never omit recommended actions — a health report that only describes problems without prioritised next steps is incomplete

product-launch-checklist

- Never apply a Tier 1 checklist to an incremental update — tier the launch appropriately before generating the checklist

- Never launch on a Friday without confirmed weekend engineering coverage
- Never leave the Go/No-Go decision owner as "the team" — it must be a named individual
- Never skip the rollback plan for Tier 1 and 2 launches — know the revert time before going live
- Never close the launch without scheduling the post-launch retrospective — it must be booked at launch time, not after

product-naming

- Never hand back a flat list with no evaluation or recommendation
- Never claim a name is "available" — you can't verify trademarks/domains; list the checks to run
- Never over-index on clever/coined names for features that just need to be findable
- Never ignore pronounceability/spelling — a name people can't say or type costs you word-of-mouth
- Never skip cross-language/meaning checks for names going to multiple markets

product-positioning-doc

- Never write positioning that could describe any competitor — differentiation must be specific, provable, and hard to copy
- Never mix category design with category entry — know whether you are creating a new category or competing in an existing one
- Never create persona messaging that uses the same headline for all personas — each persona has different priorities
- Never include proof points that are claims without evidence — every proof point needs a supporting data point or reference
- Never skip the "not for" section — defining who this is not for sharpens targeting and prevents off-persona deals

professional-brain

- Never paraphrase a source into the durable layer without keeping the original in source/ — the audit trail is the point
- Never drop provenance tags when reusing a fact — an untagged claim is an unfalsifiable one
- Never answer a recall from general knowledge and present it as something the brain "knows" — say when memory is empty
- Never overwrite a decision when it changes — append a new dated entry so the history survives
- Never build a vector database or hide memory behind embeddings — the brain stays plain, grep-able markdown a human can read and correct

professional-translator

- Never translate word-for-word — convey meaning and tone the way a native would phrase it
- Never ignore register — the wrong formality (over-familiar or stiff) can offend or undermine
- Never translate idioms literally — render the equivalent expression or the plain meaning
- Never translate brand/product/proper names or code — keep them verbatim
- Never silently resolve ambiguity — flag it; the author may mean something specific

programmatic-seo

- Never generate near-duplicate pages that differ only by a swapped keyword — that's doorway spam
- Never publish without real, maintained data behind each page
- Never dump thousands of pages at once — phase it and watch indexation
- Never ignore search intent — a page that doesn't answer the query won't rank or convert
- Never skip the kill criterion — unmaintained thin pages become a sitewide quality drag

project-status-report

- Never rate project health as Green while listing unresolved critical blockers
- Never report milestone progress as a percentage — milestones are binary: complete or not complete
- Never bury risks at the bottom — if something is high risk, it belongs in the executive summary
- Never leave decisions required without specifying who must decide and by when
- Never write an executive summary that requires reading the full report to understand — it must stand alone

promotion-packet

- Never argue from tenure or effort ("I've been here 3 years", "I work hard") — committees reward demonstrated scope and impact
- Never leave a target competency unevidenced — one unbacked competency is the gap reviewers latch onto
- Never frame it as potential — "could do the next level" loses to "is already doing it"
- Never pad with low-level wins — they signal you're operating *below* the target level
- Never submit with known gaps to "see what happens" — a failed packet is costly; close gaps first

promotion-plan

- Never discount without the margin math — a deep cut can lose money on every order

- Never pick a percentage by reflex — match the mechanic to the goal (AOV vs. acquisition vs. clearance)
- Never blast everyone — discounting full-price buyers is margin you didn't need to give away
- Never fake urgency/scarcity — countdowns that reset and "last chance" that isn't erode trust
- Never run a promo with no success metric — you won't know whether to repeat it

prompt-optimizer

- Never return a critique without the rewritten prompt — the rewrite is the deliverable
- Never pile on every technique at once — apply the levers that match the diagnosed failure, and say why
- Never add examples that contradict the instructions — the model copies the example over the rule
- Never make the prompt longer when the fix is to make instructions clearer and earlier
- Never claim a fix works without a way to test it — ship the test set

prompt-regression-suite

- Never test only happy paths — the suite exists for the inputs that hurt you
- Never let anyone update golden expectations in the same PR that broke them, without review
- Never use an unpinned judge model — a judge that upgrades itself moves your baseline silently
- Never treat pass-rate-vs-baseline as the only gate — one dead canary matters more than 2% aggregate drift
- Never grow the set unboundedly — a suite too slow to run on every change protects nothing

property-investment-analysis

- Never omit vacancy, maintenance, management, and capex — that's how a "good" deal becomes a money pit
- Never fold debt service into operating expenses — it breaks NOI and cap rate
- Never invent operating costs as fact — use the user's figures or labelled placeholders
- Never present one optimistic scenario — show the downside sensitivity
- Never give investment advice or guarantee returns — analyse and point to a professional

property-listing

- Never use buyer-preference or steering language ("perfect for a young family", "great for...") — describe the home
- Never invent or inflate facts (square footage, upgrades, permits, schools) — flag to confirm
- Never pile on clichés and exclamation marks — specifics sell, hype doesn't

- Never bury the best feature — lead with it
- Never present this as legal/compliance certification — flag for MLS/fair-housing review

property-offer-letter

- Never include anything about race, religion, family/children, disability, or national origin — it's a fair-housing risk and can sink the offer
- Never write generic flattery — name the specific features that won the buyers over
- Never invent or restate legal offer terms — this is the cover letter, not the contract
- Never be pushy or guilt-trippy — warmth and respect, not pressure
- Never present this as legal advice or assume a letter is allowed — flag to confirm with the agent

proposal-writer

- Never lead with the solution before establishing that the problem is understood — the proposal must demonstrate problem comprehension first
- Never use vague investment language like "competitive pricing" — every proposal must state a specific price or range
- Never omit a "not included" section — undefined scope leads to disputes after the proposal is accepted
- Never forget a "valid until" date — proposals without expiry create awkward situations and stale pricing
- Never list next steps without naming who is responsible for each and what the expected timeline is

public-comment

- Never submit a bare "I support/oppose" — agencies weigh substance, not vote counts
- Never argue in generalities — tie every point to the proposal's actual text
- Never just object — propose the specific alternative you want instead
- Never omit evidence — unsupported assertions are easy to dismiss on the record
- Never go off-topic — comments outside the proposal's scope carry no weight

qa-release-signoff

- Never give a thumbs-up with no evidence — sign-off is a record, not a vibe
- Never hide untested areas or thin coverage — that's the risk the reader needs
- Never conflate severity and priority, or omit blocker status on open bugs
- Never recommend "go" while ignoring a known critical defect without naming the risk/decision
- Never ship without a rollback/mitigation note for when it goes wrong

qbr-deck

- Never fill the QBR with product activity metrics — lead with business outcomes the customer cares about
- Never present a roadmap without linking each item to a customer goal — vendor priorities are not a QBR agenda
- Never run a QBR as a one-sided presentation — it must include structured time for the customer to speak
- Never close a QBR without documented mutual commitments with named owners on both sides
- Never skip the "what's not working" slide — suppressing problems erodes trust and misses renewal risks

quiz-generator

- All recall, no application or analysis
- Obvious throwaway distractors
- Trick questions that test reading, not the subject
- Answer key with answers but no explanations

quote-card

- Never fabricate or embellish a quote — only use words that are in (or faithfully trimmed from) the source
- Never change the meaning to make it punchier — fidelity over flash
- Never pick a generic line ("It's great!") when a specific, vivid one exists
- Never hide your edits — the editing note must reflect every change
- Never invent attribution — use only what's given, flag what's missing

raci-matrix

- Never assign more than one Accountable per task — shared accountability means no accountability
- Never create a RACI with more than 5–6 roles — it becomes unreadable and unenforceable
- Never include tasks so broad that the RACI cannot be acted upon — break down to decision-level granularity
- Never skip the conflict resolution process — RACI matrices without a process for disputes are unused after the first disagreement
- Never confuse Responsible with Accountable — document the distinction clearly for each role

rag-architecture-review

- Never recommend fine-tuning the model when the failure is in retrieval — fix what's retrieved first
- Never review only the generation prompt — most RAG quality is won or lost before the LLM sees anything
- Never present findings without severity and priority — a flat list doesn't tell the team what to do Monday
- Never assume the corpus is fine — stale or badly-structured source data caps every downstream stage
- Never skip the eval gap — without separated metrics, every fix is a guess

rag-design-doc

- Never jump to "fine-tune the model" when retrieval is the problem — fix what's retrieved first
- Never evaluate only the final answer — a good answer from luck and a bad answer from bad retrieval look different and need different fixes
- Never force an answer when nothing relevant was retrieved — an honest "I don't know" beats a confident hallucination
- Never ignore metadata filtering — semantic similarity will happily return the right-sounding chunk from the wrong document or wrong tenant
- Never pick a chunk size by default — it's the single biggest lever on retrieval quality

rate-card

- Never pick a rate by gut — compute the floor from income/costs/capacity, or you'll quietly run at a loss
- Never assume full billable capacity — ~30–40% goes to sales/admin/holidays; pricing on 100% underprices badly
- Never default to hourly — it caps income and penalises you for being fast; package and value-price where possible
- Never justify price by effort/cost — clients pay for ROI; anchor there
- Never present one rate — tiers convert better and lift the average deal

readme-writer

- Never bury what-it-does under a wall of badges or backstory — pitch first
- Never write a quickstart with missing steps — it must actually run
- Never inline the entire documentation — summarize and link
- Never over-promise; reflect the real project status (alpha/beta/stable)
- Never skip the license — it determines whether anyone can legally use it

recruiter-outreach

- Never write a generic "I came across your profile and was impressed" blast — it reads as spam
- Never dump the whole job description — sell the fit and the next step
- Never fabricate a personal connection — flag placeholders for true details
- Never over-follow-up or guilt-trip — respect a non-reply and end gracefully
- Never over-promise comp, level, or scope to get a reply — it backfires later

red-team-review

- Never produce vague, generic objections ("it might be risky") — name the specific failure mode and trigger
- Never only criticise — every review must end with concrete, prioritised ways to strengthen the plan
- Never give all personas the same critique reworded — each lens must find something the others miss
- Never soften the most dangerous risk to be polite — surface it first and plainly
- Never invent facts about the plan — infer plausibly and label assumptions as *(assumed)*

redundancy-consultation

- Never proceed without a prominent disclaimer that qualified HR and legal advice is required before taking any action
- Never use template letters without customising them for the specific individual and situation
- Never omit the genuine exploration of alternatives — redundancy consultation must consider alternatives before confirming decisions
- Never leave out statutory redundancy pay guidance — employees have legal entitlements that must be referenced
- Never conduct a redundancy process without documenting the selection criteria and scoring — undocumented decisions create legal risk

refactoring-plan

- Never refactor and add features in the same commit — separate them
- Never start without tests — add characterization tests first if coverage is thin
- Never plan a big-bang rewrite — sequence small, reversible steps
- Never change behavior and call it refactoring — behavior must stay identical
- Never skip running tests between steps — that's the whole safety mechanism

reference-letter

- Never rely on generic adjectives ("hardworking, dedicated") with no evidence — they signal nothing

- Never present invented examples as real — mark them for replacement
- Never write a one-size-fits-all letter — tailor the evidence to the decision
- Never overpraise to the point of incredibility — calibrated specifics are more persuasive
- Never bury the recommendation — make your endorsement explicit and early

referral-program

- Never set a reward without checking it against CAC/LTV — that's just burning money
- Never reward signups/clicks alone — reward the conversion you actually want
- Never ask before the user has felt value — timing is half the program
- Never ignore fraud — reward farming can quietly eat the whole budget
- Never make sharing clunky — every extra step kills participation

referral-program-design

- Never pay for signups instead of activations — you'll fund fraud and low-quality users
- Never claim virality from a k-factor below 1 — be honest that it's lowering CAC, which is still worth doing
- Never bolt the ask onto onboarding before the user has felt value — nobody refers a product they haven't experienced
- Never ignore the sharer's social risk — give them a reason that makes them look generous/smart, not spammy
- Never skip fraud guardrails — an ungated incentive is an arbitrage opportunity, not a growth loop

regex-builder

- Never give a regex with no test cases — always prove it
- Never ignore the flavor — `\d`, lookbehind, and named groups differ across engines
- Never produce an unreadable one-liner when a commented/verbose version or a non-regex approach is clearer
- Never silently assume anchoring — state whether it matches the whole string or a substring

regression-test-plan

- Never "re-run everything" by default — it's slow and trains teams to skip it
- Never test only the changed file — cover its dependencies and shared components
- Never silently drop coverage — when you de-scope, state the risk
- Never automate flaky or rarely-run cases first — start with stable, high-value ones
- Never let the suite grow unbounded — prune and tier it as the product changes

regulatory-impact-analysis

- Never assume regulation is the answer — establish the problem and test the baseline first
- Never present only the preferred option — compare real alternatives
- Never fabricate precise numbers — use ranges and label assumptions where data is thin
- Never ignore who bears the cost — distributional/small-business impact matters
- Never omit proportionality — the benefit must justify the burden imposed

renewal-playbook

- Never start renewal conversations less than 90 days before the renewal date for accounts over \$50K ARR
- Never build a renewal strategy without first honestly assessing account health — wishful thinking leads to last-minute churn
- Never treat all renewal objections as negotiating tactics — some objections signal genuine dissatisfaction that requires resolution first
- Never offer discounts as the first response to price objections — explore value gaps before reducing price
- Never close the renewal without confirming the expansion opportunity — every renewal is also an expansion conversation

rental-application

- Never send a bare "I'd like to apply" — give the reliability signals that win competitive listings
- Never overshare exact salary/bank details unsolicited — reassure without exposing yourself
- Never hide a likely concern — address it positively before the landlord wonders
- Never sound desperate or over-familiar — confident and professional wins
- Never invent references or history — bracket real details to provide

research-protocol

- Never write an analysis plan as "to be determined" — the analysis approach must be pre-specified before data collection
- Never skip the ethical considerations section — all research involving human participants requires ethical review
- Never define research questions so broadly that the study cannot answer them within scope and budget
- Never conflate the research question with the hypothesis — state them separately and clearly
- Never omit sample size justification — an underpowered study wastes resources and produces inconclusive results

resume

- Never list job duties — "managed a team" is a responsibility; "grew the team 4→11 and cut attrition 30%" is an achievement
- Never use multi-column layouts, tables, headers/footers, or icons — they scramble in ATS parsers
- Never write a generic resume — tailor the summary, skills, and emphasis to the target role
- Never pad with soft-skill filler ("hard-working team player") — show it through results
- Never invent or inflate metrics — use real numbers, or a defensible estimate clearly framed

retention-analysis

- Never recommend "improve onboarding" without specifying what specific step to change and why
- Never analyse retention without segmenting by cohort — aggregate retention curves hide cohort-specific patterns
- Never treat DAU/MAU below 5% as a retention problem — at that level, it is a product-market fit problem
- Never skip qualitative research — churned user interviews reveal reasons that quantitative data cannot
- Never set a monitoring alert without specifying the threshold that triggers it

retention-loop-design

- Never optimise retention before the curve flattens for some segment — if it decays to zero there's no PMF to retain, fix that first
- Never bolt on streaks/badges without a real reward — gamification on a product with no core value just annoys
- Never spam notifications to force engagement — manufactured frequency drives uninstalls and erodes trust
- Never ignore the investment phase — without stored value/data, there's nothing raising the cost of leaving
- Never report only average retention — cohorts and the best-retaining segment tell you where to aim

retro-analysis

- Never assign blame to individuals in the retrospective brief — observations must describe patterns, not people
- Never produce Start/Stop/Continue prompts that are vague categories — each must name a specific behaviour
- Never recommend an experiment that cannot be completed within one sprint — small, testable experiments only

- Never treat carry-over tickets as a velocity problem without first identifying the root cause category
- Never run the same retrospective format every sprint — vary the format to prevent engagement fatigue

return-refund-policy

- Never hide unfavourable terms in dense legalese — clarity builds trust and cuts tickets
- Never omit the faulty/wrong-item path — that's the case that most needs a clear, no-cost route
- Never state statutory rights as fact across regions — flag for jurisdiction review
- Never leave window/shipping/refund-timing vague — those are exactly what shoppers check
- Never contradict the marketplace's own policy if selling there — note where it overrides

review-response

- Never get defensive or argue facts in public — you're writing for the next shopper, not to win
- Never paste an identical canned reply on every review — personalise to the specific point
- Never expose order numbers, emails, or other private details in a public reply
- Never over-apologise or grovel on a minor issue, or under-respond on a serious one — match the severity
- Never bribe for removal or incentivise changing the review in ways the platform forbids

rfc-writer

- Never write the RFC as a persuasion document — its purpose is to expose trade-offs, not sell a decision
- Never list alternatives without explicit rejection reasons — "we preferred the proposed solution" is not a reason
- Never leave the security implications section blank or write "N/A" without a reasoned explanation
- Never write open questions without assigning a named owner and a resolution deadline
- Never skip the "impact of not solving this" section — without it, reviewers cannot assess urgency

rfp-response

- Never skip any mandatory requirement — one missed "shall" can disqualify the whole bid
- Never answer the criteria you wish they'd asked — answer their actual scoring rubric
- Never lead with a company brochure — lead with the buyer's problem and outcomes
- Never make unsupported claims — evaluators score evidence, not enthusiasm
- Never ignore format/page rules — non-compliant submissions get rejected unread

rfp-writer

- Never leave evaluation criteria unstated — undefined scoring invites bias and disputes
- Never write open-ended questions that produce incomparable marketing answers
- Never blur must-haves and nice-to-haves — vendors (and evaluators) can't prioritise
- Never omit out-of-scope — scope creep starts in a vague RFP
- Never set the weights after seeing the bids — decide what matters up front

rice-impact-matrix

- Never treat the combined score as a definitive ranking — use it to structure a conversation, not replace one
- Never rate strategic alignment as "high" because an initiative feels important without mapping it to a specific OKR
- Never place all initiatives in the "Now" quadrant — a matrix with no "Drop" recommendations is not credible
- Never ignore the conflict flag when RICE rank and strategic alignment sharply diverge
- Never accept 100% confidence on estimates that have not been validated with data

rice-prioritisation

- Never default to 100% confidence on estimates that lack supporting data — this inflates scores and misleads planning
- Never treat RICE scores as a final decision — a ranking that surprises the team must be investigated before it is accepted
- Never omit effort estimates from engineering — PM-only effort estimates are frequently optimistic and skew results
- Never forget to note dependencies that would change the sequencing even if RICE scores suggest otherwise
- Never score every initiative at the same impact level — if everything is "high impact," the framework produces no useful signal

risk-register

- Never assign risks to "the team" or "TBD" — every risk must have a named individual owner
- Never write mitigations as "monitor and review" — mitigations must describe what is actively being done to reduce likelihood or impact
- Never delete closed risks — they provide an audit trail; archive them instead
- Never confuse risks with issues — a risk is something that might happen; an issue is something that has already happened
- Never leave Critical or High risks without a contingency plan — what happens if the mitigation fails must be documented

roadmap-narrative

- Never produce a list of features with dates and call it a narrative — every initiative must connect to a strategic theme
- Never omit the "what's not on the roadmap" section — without it, the narrative lacks strategic discipline
- Never write progression as a chronological list — show causal links between quarters (Q1 enables Q2 because...)
- Never write the executive summary last and treat it as a summary — write it as the version stakeholders will repeat
- Never let orphaned initiatives appear without a theme — either create a theme or flag the gap explicitly

roadmap-presentation

- Never put specific dates on NEXT or LATER items — use quarters or halves to signal appropriate confidence levels
- Never show the same level of detail to executives and engineers — calibrate depth to audience or you lose both
- Never omit the "What We're NOT Building" section — a roadmap without explicit deprioritisation becomes a wish list
- Never present LATER items as commitments — frame everything outside NOW as directional, not promised
- Never skip the success metrics section — without it, stakeholders cannot evaluate whether the roadmap is working

roi-estimator

- Never ignore ongoing costs — a low upfront, high-recurring option can lose to a pricier one-time buy
- Never present a single benefit number as fact — it's the softest input; give a range and flag it
- Never skip discounting for multi-year cases — \$1 in year 3 isn't \$1 today
- Never bury the assumptions — a business case is only as credible as its stated inputs
- Never omit payback — a great 5-year ROI with a 4-year payback may still be too slow to fund

role-redesign-for-ai

- Never redesign the role without the people in it — a charter discovered in a reorg deck creates the resistance it deserved
- Never keep old throughput metrics "for continuity" — they now measure the vendor, not the human

- Never treat verification as slack time — reviewing machine output at quality is skilled work with hours
- Never write "focus on higher-value work" without naming the work — that phrase is where redesigns go to die
- Never skip the intent question — a redesign that won't say whether headcount changes will be read as concealing it, correctly

rubric-builder

- Descriptors that just add adjectives ("good" → "very good" → "excellent")
- Overlapping levels a grader can't tell apart
- Criteria that measure effort/length instead of the learning goal
- A rubric only the teacher can read

runbook-writer

- Never write steps as vague actions like "run the deploy script" — every step must include the exact command
- Never leave the rollback section as a placeholder — a runbook without a tested rollback procedure is incomplete and dangerous
- Never omit expected output for each step — without it, the on-call engineer cannot tell if the step succeeded
- Never write escalation contacts as "[Team name]" — every escalation row must have a real contact or an explicit flag to fill in
- Never assume the reader knows the system — write for someone who has never touched it before

runway-calculator

- Never report runway off gross spend — net burn (after revenue) is the real number
- Never assume flat burn silently — if headcount/spend is rising, say the runway is optimistic
- Never plan to zero cash — a raise takes 3–6 months; runway should be measured to "must-raise-by," not "broke"
- Never ignore growth — a fast-growing company can be default alive even while burning
- Never present one scenario — show the cut-vs-raise-vs-grow trade-off

runway-monte-carlo

- Never report only the median — the median is the number that feels fine right up until the P10 path happens to you
- Never silently invent volatility — a made-up σ changes the answer more than the burn does; label defaults

- Never model the hoped-for fundraise inside the simulation — runway exists to time the raise, not assume it
- Never extend the horizon to make survival look better — report the horizon with the number
- Never present 56.8% survival as "about half" in one place and "likely fine" in another — one number, one interpretation, used consistently

runway-planner

- Flat cash ÷ burn when burn is clearly growing
- Ignoring that raising takes months (planning to start at 2 months left)
- Vague advice ("extend runway") instead of a quantified lever and date
- Treating gross burn as net (ignoring revenue)

saas-metrics

- Never include new customers in NRR — that's a different (and misleadingly flattering) number
- Never mix monthly and annual figures without converting — label MRR vs ARR clearly
- Never report a metric without its definition — "120% retention" is meaningless without the formula
- Never vanity-pick metrics — show churn and contraction alongside the growth numbers
- Never present computed values to false precision — round sensibly and flag assumptions

salary-negotiation

- Never compare base to base — total comp is the real number, and equity/bonus often change which offer wins
- Never negotiate without a walk-away decided in advance — it's the source of your leverage
- Never bluff a competing offer you don't have — if it's called, you lose all credibility
- Never anchor low or accept the first number — the first offer almost always has room
- Never fixate only on base — sign-on, equity, level, and start date often move when base is capped

sales-battlecard

- Never minimise or ignore genuine competitor strengths — sales reps who encounter them unprepared lose credibility
- Never write differentiators without proof points — a claim without evidence is marketing, not a battlecard
- Never make the battlecard exhaustive — it is a one-page cheat sheet, not a full competitive analysis

- Never include a "When we lose" section that is dishonestly optimistic — honest loss scenarios build rep trust
- Never skip the review date — an outdated battlecard with wrong information is worse than no battlecard

sales-demo-script

- Never do a feature tour or "let me show you the settings"
- Never start clicking before confirming the pain
- Never demo on empty or unrealistic data
- Never rush the "aha" — that's the moment that sells
- Never end with "any questions?" — end with a next step

sales-enablement-kit

- Never list every feature; reps sell outcomes, not spec sheets
- Never invent metrics, logos, or case studies — mark them [to confirm]
- Never write objection answers that dodge the real concern
- Never make the demo a full product tour; find the one "aha"
- Never bury the CTA — reps need to know exactly what to ask for next

sales-forecasting-model

- Never present a single forecast number without scenario analysis — a forecast without upside and downside cases hides risk
- Never use 100% confidence on conversion rates that are not backed by historical data — flag them as assumptions
- Never skip the activity sanity check — a forecast number that requires unreachable activity levels is not credible
- Never use top-down quota as the only forecast method when pipeline data exists — bottom-up is more accurate and defensible
- Never omit the coverage ratio — without it, stakeholders cannot assess whether the pipeline is sufficient to hit target

sales-page

- Never use fake scarcity or fake countdowns — it works once and destroys trust; use real deadlines
- Never over-hype beyond what the proof supports — believable beats biggest
- Never bury the offer or the price — clarity converts; confusion kills
- Never agitate into manufactured fear — name real costs, don't invent dread
- Never switch the ask — one offer, one CTA, repeated

savings-goal-plan

- Never give a monthly number without checking it's realistic against their means
- Never ignore money already saved or any starting balance
- Never assume an interest rate without saying so — be conservative
- Never present a single rigid plan when the date is too tight — offer the trade-off levers
- Never present this as personalized financial advice

saying-no

- Never give a false maybe — "let me see" to avoid the moment creates a worse letdown later
- Never over-justify — a long list of reasons sounds defensive and invites picking each apart
- Never say a flat "no" to a boss when a trade-off works better — make the cost visible, let them choose
- Never apologise excessively — "I can't take this on" is fine; grovelling undermines the boundary
- Never cave on first pushback — decide the line beforehand and have a response ready

schedule-recipe

- Never pick a runner the user doesn't have — a perfect n8n flow is useless without n8n
- Never write a run prompt that says "as usual" or relies on the agent remembering prior runs without a stored previous edition
- Never schedule without failure alerting — silence and success must not look identical
- Never default to hourly/daily to "be safe" — match the cadence to how often the inputs change
- Never put secrets inline in the setup block — reference the runner's secret store

schema-markup

- Never mark up content that isn't visible on the page — Google treats that as spam
- Never fabricate ratings/reviews or self-apply AggregateRating without real reviews
- Never omit required fields — the rich result simply won't trigger
- Never use the wrong type (e.g. Product for an article) — it won't validate
- Never ship without validating in the Rich Results Test

screenshot-teardown

- Never analyse a product from training-data memory when screenshots are provided — the pixels are the source of truth, and the product has probably changed
- Never proceed without images — that's competitor-teardown's job

- Never present inferences as facts — "they're struggling with churn" is a reading, not a screenshot
- Never sneer — "cluttered" is not analysis; name what the clutter costs and whom it serves
- Never extrapolate a whole strategy from one screen — say when the evidence is thin

security-incident-response

- Never wipe/rebuild before preserving forensic evidence — you lose the ability to understand the breach
- Never skip credential rotation — attackers persist via stolen keys/tokens
- Never go quiet on comms — silence with customers/regulators creates legal and trust damage
- Never blame individuals in the review — blameless analysis surfaces the real systemic causes
- Never declare "recovered" without monitoring for re-compromise
- Never act on systems you don't own or aren't authorized to defend

security-review

- Never produce a generic checklist — tie each finding to this code/design and its exploit path
- Never rank everything the same — separate critical from hardening nits
- Never report an issue without a fix — give the concrete remediation
- Never miss authorization (IDOR/privilege) — it's the most common real-world web flaw
- Never review code you don't own or aren't authorized to assess

security-threat-model

- Never restrict STRIDE analysis to only the API layer — threats exist at every component including the database and internal services
- Never leave mitigations as vague directives like "improve security" — every mitigation must be specific enough to become a ticket
- Never accept risks without a named owner and a review date — unowned accepted risks are not managed risks
- Never write a threat model that covers only theoretical threats — prioritise by likelihood and impact using the risk register
- Never omit the asset register — without knowing what is being protected, the STRIDE analysis has no anchor

self-review

- Never list activities — map every accomplishment to an outcome and a competency

- Never disguise a strength as a weakness ("too detail-oriented") — it reads as evasive; name a real growth area
- Never claim team wins as solo, or undersell your role out of modesty — be precise about your contribution
- Never inflate the self-rating beyond what the evidence supports — it costs credibility in calibration
- Never write in vague superlatives — "drove significant impact" means nothing without the number

seo-content-brief

- Never write an outline that answers a different question than the actual search intent — the brief must match what the searcher wants, not what the brand wants to say
- Never set keyword density targets so high that they produce unnatural writing — 3–5 natural mentions is guidance, not a quota
- Never skip the competitor gap analysis — without it, the brief produces content that duplicates what already ranks
- Never leave the FAQ section without real "People Also Ask" questions — fabricated questions miss search volume opportunities
- Never write a title tag longer than 60 characters — it will be truncated in search results and undermine ranking

sequence-diagram

- Never show only the happy path when a failure path matters — note the 401/timeout/retry
- Never blur requests and returns — use `->>` vs `-->>`
- Never reorder messages for neatness — sequence order is the whole point
- Never put colons inside message text — it breaks parsing
- Never invent participants — model only the systems actually involved

service-catalog-entry

- Never write aspirational SLO targets — targets must be agreed with stakeholders and based on historical data, not copied from a template
- Never leave runbook links as TODO placeholders — broken or missing links make the catalog entry worse than useless during an incident
- Never omit the "Known Limitations" section to make the service look better — undisclosed limitations cause incorrect integrations and downstream incidents
- Never list API error codes without testing them — aspirational error documentation misleads consumers
- Never write the "What It Does" section with jargon — a new engineer from another team must understand it in under 2 minutes

session-handoff

- Never write a vague status ("worked on the feature") — state exactly what's done and what's not
- Never omit dead ends — repeating failed attempts is the most common handoff waste
- Never bury the next step — it should be obvious and immediately actionable
- Never assume shared memory — the reader may have zero prior context
- Never pad it — a handoff nobody reads is worthless; keep it tight and scannable

short-form-script

- A slow intro ("Hey guys, so today I wanted to talk about...")
- Long-form structure crammed into 30s
- No on-screen text or visual cues (it's a *video* script, not an essay)
- Multiple competing CTAs

skill-security-auditor

- Never pass a skill as safe without reading its scripts — prose can look clean while a script exfiltrates data
- Never treat every mention of "API key" or "curl" as malicious; weigh intent and context
- Never give a vague verdict — always land on install / caution / do-not-install with reasons
- Never ignore zero-width or invisible characters; they are a classic way to hide instructions
- Never assume a high star count or popular author means a skill is safe — audit the content itself

slide-deck

- Never write topic-label titles ("Metrics") — use the takeaway ("Retention drove 80% of growth")
- Never cram multiple ideas onto one slide — split them; one point per slide
- Never paste paragraphs as bullets — phrases on the slide, detail in speaker notes
- Never vary styling slide to slide — consistency is what makes a deck look professional
- Never claim a file exists without code execution — fall back to the outline / the playground's PPTX export

slo-error-budget

- Never set SLO targets at 100% — this discourages feature development and does not reflect how users experience reliability
- Never measure internal system metrics as SLIs — SLIs must reflect what users directly experience, not internal CPU or memory

- Never write an error budget policy with vague triggers — "discuss as a team" is not an actionable policy; triggers must be specific percentages
- Never base targets on aspirational round numbers — always derive from historical baseline data
- Never configure only one burn-rate alert window — a single window misses both fast burns and slow burns that exhaust the budget quietly

soap-note

- Never invent vitals, labs, exam findings, or results to fill a section — mark them "not documented"
- Never present this as diagnosis or medical advice — it formats clinician-provided information
- Never blur assessment and plan into one block — they serve different readers and purposes
- Never drop pertinent negatives the clinician noted — they're part of the reasoning
- Never reorganise so heavily that the clinician's original meaning changes

soc2-readiness

- Never confuse a readiness assessment with a passed audit — readiness is self-assessed; the report comes from a licensed CPA firm
- Never claim a control is "met" without the evidence to prove it — auditors test operating effectiveness, not intentions
- Never over-scope criteria — every criterion you add is more controls to evidence; include only what's true and needed
- Never leave gaps unowned or undated — an unowned gap is a gap that's still open at audit time
- Never try to backfill Type II evidence — controls must demonstrably operate across the whole period

social-ad-campaign

- Never combine multiple campaign objectives in one campaign — pick one measurable goal or the algorithm cannot optimise correctly
- Never skip retargeting suppression — existing customers receiving acquisition ads wastes budget and damages brand perception
- Never launch without completing the tracking setup checklist — campaigns without verified pixel firing cannot be optimised or attributed
- Never run A/B tests changing more than one variable at a time — multi-variable tests produce uninterpretable results
- Never allocate equal budget across TOFU, MOFU, and BOFU — BOFU audiences convert at higher rates and deserve proportionally more budget per conversion

social-media-audit

- Never score platforms against guesswork — every score must be based on actual metrics provided or observable data
- Never write recommendations as "post more content" or "improve engagement" — every action must be specific and measurable
- Never use competitor benchmarks that are not based on real data — fabricated benchmarks invalidate the competitive gap analysis
- Never audit content mix based on what should have been posted — analyse what was actually posted during the audit period
- Never sequence the action plan by effort alone — sequence by impact × effort so the highest-value actions come first

social-media-strategy

- Never recommend every platform — justify each choice with where the target audience actually spends time
- Never define content pillars that serve only the brand — each pillar must deliver specific value to the audience or it will not earn attention
- Never set a posting cadence that exceeds the team's realistic capacity — an unsustainable strategy fails faster than a modest one
- Never use vanity metrics (likes, followers in isolation) as primary KPIs — tie KPIs to the stated business goal
- Never skip the tone of voice section — without it, multiple contributors produce inconsistent content that erodes brand identity

sop-writer

- Never write steps that contain more than one action — each step must be a single, auditable action in imperative form
- Never omit a role from any step — every action must be assigned to a specific role or the SOP cannot be enforced
- Never skip the non-conformance section — an SOP without a deviation process cannot meet audit or regulatory requirements
- Never produce an SOP without a review date and version history — undated documents cannot be relied upon for compliance
- Never use passive voice in procedure steps — write "Open the system" not "The system should be opened"

sourcing-strategy

- Never rely on a job post and inbound for a hard role — lead with proactive sourcing
- Never define the profile so narrowly that no one qualifies — include adjacent talent

- Never invent exact funnel numbers — use ratios and a worked example
- Never list channels without prioritisation — say where to spend effort first
- Never skip the weekly cadence — strategy without activity targets doesn't fill the role

sprint-brief

- Never write a sprint goal as a task list — the goal must be a single outcome-focused statement that can be scored pass/fail
- Never leave the critical path unnamed — "the important tickets" is not a critical path
- Never list risks without a mitigation or owner — a risk without a response is just a worry list
- Never ignore carry-over items' impact on this sprint's capacity and goal
- Never write a Definition of Done that mixes task completion with outcome criteria — they must be observable and agreed before the sprint starts

sprint-planning

- Never write sprint goals as task lists — goals must be outcome-focused and scoreable pass/fail at sprint end
- Never commit to 100% of available capacity — always recommend 80% to preserve slack for unplanned work
- Never carry stories with no acceptance criteria into the sprint — flag them as blockers before committing
- Never allow stories estimated at 8+ points into the sprint without splitting them first
- Never ignore carry-over items when calculating capacity — they consume capacity and must be accounted for before new work is pulled in

sprint-velocity-analysis

- Never generate the velocity chart from placeholder data — it must reflect the actual sprint data provided
- Never diagnose trend direction without computing trailing vs leading averages — "it looks like it's declining" is not a diagnosis
- Never list carry-over as a generic observation — identify root cause categories with counts for the analysis to be actionable
- Never produce recommendations without a named owner, a start date, and a measurable target
- Never score health dimensions without citing evidence in the Evidence column — unsupported Red/Yellow/Green scores are not credible

sql-optimizer

- Never say "add indexes" generically — name the columns and explain which predicate/join they serve
- Never ignore the engine — Postgres index tuning and BigQuery partition pruning are different games
- Never optimize a query that should be a model — repeated heavy logic belongs in a materialized/dbt model
- Never wrap indexed columns in functions in the WHERE clause — it kills index usage (non-sargable)
- Never recommend changes without an expected impact or a way to measure it

sql-query-explainer

- Restating the SQL in pseudo-code instead of explaining what it *does* and *returns*
- Optimisation advice with no before/after query, or no reason the new one is faster
- Ignoring the dialect (writing Postgres-only syntax for a MySQL user)
- "Looks fine" with no read on correctness, performance, or row grain
- Rewriting the query from scratch instead of explaining/optimising the user's

stakeholder-influence-mapper

- Never approach high-influence blockers before aligning their sponsors — approach order determines outcome
- Never create talking points that lead with your agenda — always lead with the stakeholder's stated concern
- Never treat every stakeholder as equally important — focus depth on the decision-makers and key influencers
- Never omit the "do not approach until X is aligned" flags — sequencing mistakes can permanently close doors
- Never build the map based only on org chart position — influence often lives outside formal authority

stakeholder-update

- Never bury the status assessment at the bottom — BLUF means the most important information comes first
- Never report metrics without a target or prior-period comparison — raw numbers without context are not useful
- Never list risks without mitigation actions and clear flags for stakeholder help needed
- Never write decisions needed as questions without providing a clear recommendation — executives need options, not open-ended questions
- Never allow the update to exceed one page — if it requires more, the message needs editing, not expanding

startup-idea-validator

- Cheerleading (validating because the founder wants a yes)
- Generic critique that applies to any startup
- Recommending a 6-month build as the "test"
- A verdict with no path forward

statement-of-work

- Never leave scope open-ended — without exclusions, clients reasonably assume everything is included
- Never omit acceptance criteria — "deliver a website" with no bar means endless revisions
- Never skip change control — it's the clause that turns scope creep into billable change requests
- Never ignore client dependencies — if their delay silently becomes your problem, you eat the cost
- Never present this as final legal advice — recommend counsel review for significant contracts

strategic-narrative-generator

- Never produce a narrative that lists initiatives chronologically without showing causal progression — the story must show why each phase enables the next
- Never use abstract strategic language that cannot be repeated by a non-technical listener — test whether someone could explain it back without the document
- Never omit the "what's not on the roadmap" section — what you are choosing not to do is as important as what you are doing
- Never set themes without measurable metrics — a theme without a metric cannot be tracked or held to account
- Never skip the hard questions section — preparing for objections in advance is the purpose of the narrative exercise

strategy-memo

- Never write goals and call it strategy — "grow revenue, delight customers" is a wish list, not a bet
- Never skip the non-goals — a strategy that sacrifices nothing commits to nothing
- Never build on a flattering diagnosis — naming the uncomfortable truth is the hardest and most important part
- Never list every initiative — coherent actions reinforce one bet; a long list dilutes it
- Never leave the bet unfalsifiable — if no evidence could prove it wrong, it can't be pressure-tested

student-feedback

- "Good job!" / "Needs work" with nothing concrete
- Marking every single error so the student can't see what matters
- Criticism with no model of the better version
- A tone that discourages instead of pointing forward

style-fingerprint

- Never fingerprint from fewer than 3 samples — you'd be fingerprinting one mood
- Never describe voice with adjectives ("professional yet approachable") — extract mechanics
- Never merge conflicting registers into one mushy card — name variants ("exec", "team") instead
- Never include the user's confidential content in the card — rules and short quoted phrases only
- Never overwrite an existing style card silently — diff against it and show what changed

subagent-orchestration

- Never parallelise for the feeling of speed — two colliding agents are slower than one sequential pass
- Never write briefs that reference your context ("as discussed", "the usual way") — subagents weren't in the room
- Never average contradictory results — a contradiction is a defect to resolve, with a cause
- Never merge everything then verify once — verify at each join while causes are still traceable
- Never delegate the judgment-bearing core (the decision, the synthesis, the taste) — delegate the legwork around it

substack-notes-scraper

- Never proceed without a valid Substack handle or profile URL — scraping without a specific target cannot be completed
- Never ignore rate-limit responses from Substack — implement backoff and reduce request frequency before retrying
- Never export data without conditional formatting and summary stats — raw data without visualisation is not the expected output
- Never attempt to access private or subscriber-only notes — this skill is for public Notes content only
- Never produce output without a clear date range filter — undated exports make trend analysis impossible

subtitle-caption

- Never exceed reading speed — a perfectly accurate caption no one can read in time has failed
- Never translate verbatim for subtitles — the target overruns the slot; condense to the gist
- Never break lines mid-phrase — split at clause/phrase boundaries for readability
- Never exceed 2 lines per cue — split into multiple cues instead
- Never omit sound cues in SDH — they're the point of accessible captions

support-macro

- Never write "Dear valued customer" / robotic openers — acknowledge the actual person and problem
- Never make it un-personalisable — a macro with no slots gets sent cold and feels like spam
- Never bury the answer in apology — empathise briefly, then solve
- Never over-promise on escalations — give an honest, realistic timeframe
- Never write one macro for a scenario with multiple outcomes — give the resolved/more-info/escalating variants

support-runbook

- Never write a tip list instead of an ordered path — agents need "do this, then this," with branches
- Never leave escalation vague — "escalate if needed" means everyone escalates differently; time-box and specify the target + attachments
- Never omit the customer comms — resolution + silence still feels like bad support
- Never ignore severity — treating a full outage like a how-to question loses trust fast
- Never let a high-frequency issue stay a runbook forever — flag the root-cause fix to product

sycophancy-challenger

- Never open with a softening phrase or acknowledgment before the challenge — the first sentence must be the critique
- Never retreat from a position when the user pushes back without providing new evidence — update only when genuinely persuaded
- Never invent flaws — every criticism must connect to something real in what the user described
- Never provide a list of weak objections — identify the single strongest case against the idea
- Never end the session because the user seems satisfied — end only when something genuinely changed or was defended

synthetic-user-research

- Never run synthetic "validation" for launch or investment decisions — that's laundering a model's agreeableness into evidence
- Never build personas from vibes or market-report archetypes — stereotypes in, stereotypes out
- Never ask personas how they *feel* or what they'd *pay* — the fluent answer is the false one
- Never report synthetic findings in the same register as real research — a stakeholder who can't tell the difference wasn't told loudly enough
- Never let a synthetic pass replace the discovery interview it was supposed to prepare — the lane is *before* human research, never instead of it

system-design-interview

- Never jump to solutions before clarifying requirements — always establish functional and non-functional requirements first
- Never present a design without discussing trade-offs — every architecture decision has costs and benefits that must be acknowledged
- Never use vague capacity estimates — show the actual calculation (QPS, storage bytes, bandwidth) not just "this handles scale"
- Never design for unlimited scale by default — match the design to the requirements stated
- Never skip the data model — a system design without entity definitions and data flow is incomplete

tax-planning-checklist

- Never provide specific tax advice — always recommend qualified tax advice and note this prominently
- Never present threshold figures as definitive without noting they require verification for the current tax year
- Never produce a generic checklist without tailoring it to the entity type (individual, sole trader, limited company)
- Never omit timing-critical items — some reliefs require action before year-end and deadlines must be called out
- Never conflate UK and non-UK tax rules — clarify jurisdiction before generating any checklist

tdd-workflow

- Never write the implementation first and the test after — that's not TDD, it's rationalization
- Never write five tests then all the code — one red→green→refactor cycle at a time
- Never over-build in green — only enough to pass the current test
- Never test implementation details — test observable behavior so refactors don't break tests
- Never skip the refactor step when there's obvious duplication or a bad name

teaching-lesson-plan

- Never design a lesson plan without explicitly stating the learning objectives — activities must trace back to outcomes
- Never allocate timing that does not add up to the total session length — the plan must be time-feasible
- Never create activities with no assessment component — learning must be measurable, not just delivered
- Never ignore differentiation — a plan with no accommodation for different learning levels or abilities is incomplete
- Never front-load all content delivery without interactive breaks — passive listening degrades retention after 15–20 minutes

team-health-check

- Never run a health check without first establishing psychological safety — without it, scores reflect fear, not reality
- Never treat a single health check as a trend — one data point cannot show improvement or regression
- Never keep results with management without sharing them with the team — transparency is a prerequisite for trust
- Never generate action items that are vague commitments like "improve communication" — every action must be specific and verifiable
- Never assign actions to "the team" — each improvement action needs a single named owner

team-offsite-planner

- Never fill the entire agenda with structured sessions — unstructured social time is essential for team bonding and must be built in
- Never schedule more than 90 minutes of intensive working sessions without a break
- Never design an offsite without clearly linking each session to the stated goals — purpose must be explicit
- Never neglect logistics — venue, travel, dietary requirements, and accessibility must be confirmed before the agenda is finalised
- Never plan without an energy management arc — high-energy collaboration sessions should not appear directly after lunch

tech-radar

- Never place a technology in Adopt without evidence it is proven at the team's scale — aspirational placements mislead engineers
- Never add a blip without a written rationale paragraph — table rows without context are unusable

- Never create a Hold entry without specifying a concrete migration path or target technology
- Never skip the maintenance process — a radar with no process for updates becomes stale within two quarters
- Never omit ring definitions — engineers need to know what they should do in response to each ring, not just what the ring means

technical-debt-register

- Never score debt items arbitrarily — priority scores must be calculated using the documented formula
- Never conflate technical debt (deliberate shortcuts) with bugs (unintended defects) — they require different remediation strategies
- Never underrate security and dependency items because they feel abstract — score based on actual business impact
- Never create "permanently deferred" items — every accepted item must have a review date and named owner
- Never include resolution plans that are vague descriptions — each plan must have specific, ticketable steps

technical-spec-template

- Never include solution language in the problem statement — the problem must be described independently of the proposed solution
- Never omit alternatives considered — a spec that considers only one approach has not been properly evaluated
- Never leave open questions as "TBD" without a named owner and due date — unresolved questions are blockers
- Never skip security and privacy sections for any feature that touches user data
- Never write a non-goals section that is empty — always list at least two things that might be assumed in scope

tenant-screening-guide

- Never screen on or mention protected characteristics (or proxies for them) — it's illegal and unfair
- Never apply criteria inconsistently between applicants — inconsistency is where discrimination claims live
- Never run credit/background checks without consent or skip adverse-action notice — FCRA requires them
- Never present this as legal advice or jurisdiction-specific compliance — flag for an attorney
- Never use blanket criminal-history bans where the law restricts them — flag to confirm locally

test-case-writer

- Never write only happy-path cases — the bugs live in the edges and negatives
- Never write vague steps ("test the login") — give the exact actions and data
- Never use unverifiable expected results ("it should work") — state the observable outcome
- Never combine many checks into one bloated case — keep cases atomic and traceable
- Never skip preconditions/test data — they're why a case is reproducible

test-strategy-doc

- Never write a test strategy without a risk table that drives test priority — generic coverage targets are not a strategy
- Never leave the "out of scope" section blank — every test strategy must explicitly name what is not being tested and why
- Never specify test types without naming a concrete tool for each — "some testing framework" is not actionable
- Never define a Definition of Done that is not measurable — "QA is happy" is not a completion criterion
- Never create P0 risk areas without corresponding P0 test cases — risk rating must map to test coverage

threat-model

- Never list generic threats — tie each to a specific boundary/data-flow in this system
- Never skip categories silently — at least consider each STRIDE class
- Never rate everything "high" — prioritize by realistic likelihood × impact
- Never propose vague mitigations ("add security") — name the specific control and where it lives
- Never model an attack on a system you don't own or aren't authorized to assess

thumbnail-creator

- Never generate thumbnails without incorporating brand colours and style specs when provided — off-brand outputs must be regenerated
- Never skip the evaluation step — all candidates must be scored before being presented to the user
- Never present only one thumbnail candidate — always generate multiple options for comparison
- Never include the full image generation prompts in a separate document — they must be included in the evaluation report for iteration reference
- Never claim a thumbnail is final without offering an iteration round

transcreation

- Never translate literally — transcreation recreates the feeling; identical words that lose the punch is failure
- Never give one option — creative work needs choices; offer distinct routes
- Never omit the back-translation — clients need to know what the new copy literally says
- Never ignore cultural connotation — a fine word in one market can be odd or offensive in another; flag it
- Never bust the character limit — an ad headline that truncates is unusable

unit-economics

- Never compute LTV on revenue instead of gross margin — it inflates LTV and hides an unviable model
- Never ignore payback — a great LTV:CAC with a 30-month payback can still starve a business of cash
- Never treat blended CAC as paid CAC — separate organic from paid or the model lies
- Never present assumptions as facts — label estimated churn/CAC and validate them
- Never optimise the smallest lever — model which input actually moves the outcome

user-interview-synthesis

- Never mix single-source signals into main themes — insights cited by only one participant must be flagged separately
- Never write implications that are observations restated rather than product decisions enabled
- Never include themes that only support the project hypothesis — contradictory findings must be surfaced, not omitted
- Never present findings without quotes — every theme requires verbatim evidence from at least 3 participants
- Never leave research questions unanswered — each question from the study brief must be explicitly addressed, even if the answer is inconclusive

user-journey-map

- Never score everything positively — the map's value is exposing the painful steps
- Never list features instead of the user's actions — stay on the user's side
- Never skip the "why" behind low scores — a score without a reason isn't actionable
- Never put colons inside task names — it breaks the Mermaid journey syntax
- Never invent research — label inferred sentiment as an assumption

user-research-synthesis

- Never list every individual comment — synthesis must identify patterns across participants
- Never make interpretive leaps without supporting evidence from the data
- Never focus on feature requests before understanding the underlying problem — always identify the job-to-be-done first
- Never ignore contradictory data — conflicting findings must be surfaced and noted
- Never present results without quantifying prevalence — state how many participants held each view

user-story-writer

- Never write user stories from a technical perspective — every story must be from the user's point of view and state their goal
- Never write acceptance criteria that are untestable — every criterion must have a clear pass/fail condition
- Never create stories that are too large to complete in a single sprint — break epics into estimable, independently deliverable stories
- Never omit edge cases — unhappy paths and error states are required, not optional
- Never skip the Definition of Done — without it, "done" means different things to different people

ux-research-plan

- Never write a research plan without clearly stated research objectives — every methodology choice must flow from the objectives
- Never design a plan that mixes generative and evaluative research without clearly separating them
- Never omit screener criteria — recruiting unqualified participants invalidates the research
- Never write discussion guide questions that are leading — questions must be neutral and open-ended
- Never skip the incentive recommendation — uncompensated research has lower participant quality and completion rates

value-proposition

- Never list features — a value prop is the outcome, features are the proof later
- Never write for everyone — "for modern teams" resonates with no one; pick the segment
- Never use empty superlatives ("revolutionary", "seamless", "all-in-one") — they're noise
- Never skip the alternative — value is relative; "better than what?" must be answered
- Never make an unbacked claim the headline — if the proof isn't there, soften or earn it first

vendor-contract-checklist

- Never skim only the order form — the MSA/terms is where the risk lives
- Never ignore auto-renewal and notice windows — they quietly lock you in
- Never accept an SLA without checking the remedy (credits ≠ reliability)
- Never present this as legal advice — flag material/legal items for counsel
- Never produce a flat list — prioritise what's actually worth negotiating

vendor-evaluation

- Never weight all evaluation criteria equally — the scorecard must reflect the relative importance of each criterion
- Never evaluate vendors only on features — security, support, contract terms, and financial stability matter too
- Never produce a recommendation without explaining why the runner-up lost — this enables future vendor conversations
- Never skip contract terms to negotiate — identifying leverage points is part of the procurement decision
- Never recommend a vendor without stating the conditions under which the recommendation would change

vendor-security-review

- Never size diligence by the vendor's brand — a small vendor with privileged access to PII outranks a famous one with none
- Never accept a SOC 2 Type I as equivalent to Type II — Type I is a point-in-time design check, not operating effectiveness
- Never skip the sub-processor question — your data may flow to fourth parties you never assessed
- Never approve high-risk vendors on a promise — require evidence and bind it in the contract (DPA, breach notice SLA)
- Never treat the review as one-and-done — set a re-review cadence tied to the tier

verification-before-completion

- Never verify against your memory of the ask — memory is where the third requirement went to die
- Never treat a clean-looking output as evidence — polish and correctness are uncorrelated at exactly the worst moments
- Never skip the pass under time pressure — the pass is minutes; the rework it prevents is hours
- Never produce a zero-findings record on complex work — that's theatre; look harder or say what you couldn't check

- Never hide residuals to seem finished — an honest "untested under X" builds more trust than the failure it predicts

viral-content-framework

- Never create a single generic framework — hook formulas and content structures must be platform-specific
- Never confuse reach with virality — high reach alone is not viral; content must drive sharing, saves, or resharing
- Never produce hook formulas without testing guidance — frameworks without a testing system produce one-off results
- Never ignore the shareability trigger — all content must have a clear reason why someone would send it to another person
- Never design hooks that work only once — the framework must be repeatable, not a collection of one-time tactics

voice-agent-design

- Never port the chatbot script to voice — text tolerates paragraphs and menus; ears don't
- Never hide the human escape hatch to protect containment metrics — callers find the exit anyway, angrier
- Never let the agent bluff on regulated topics (medical, legal, financial advice) — pass or read the approved statement
- Never re-ask a failed question unchanged — the caller heard you; the strategy failed, not their ears
- Never launch without the mid-flow hang-up metric — it's where voice agents quietly hemorrhage trust

voice-of-customer-program

- Never collect feedback with no one accountable to act on it
- Never build a taxonomy so complex no one tags consistently
- Never rank purely by volume — a few high-value accounts matter
- Never skip the customer follow-up; silent VoC erodes trust
- Never treat VoC as a survey; it's every channel, continuously

vuln-triage

- Never treat the scanner's rating as the answer — adjust for reachability and real impact
- Never ignore exploit status — a known-exploited (KEV) bug jumps the queue regardless of score
- Never give only "patch it" with no interim mitigation when patching will take time

- Never assign a generic SLA — tie urgency to the contextual severity
- Never triage assets you don't own or aren't authorized to assess

which-skill

- Never recommend more than two skills — a router that returns a list has not routed
- Never route on topic keywords ("competitor" ≠ always competitive-analysis); route on the deliverable
- Never ask a chain of clarifying questions — one at most, and only if it changes the pick
- Never invent skill names — if nothing in the catalog fits, say so and suggest `SKILL_REQUEST.md`
- Never recommend a general skill when a specific one exists for the exact artifact

whiteboard-to-spec

- Never proceed without an image — this skill reads boards, it doesn't imagine them
- Never "clean up" the room's thinking into what it should have decided — transcribe what it did decide
- Never guess illegible words silently — a confident wrong guess is worse than a flagged gap
- Never ignore spatial grouping — merging two separate clusters into one list destroys the meaning
- Never drop the marginalia — initials, dates, and edge notes are often owners and deadlines

win-loss-analysis

- Never treat "price" as a root cause without checking whether it's really value perception
- Never average away segment differences — a 60% overall win rate can hide a 20% enterprise rate
- Never fabricate buyer quotes or inflate sample size; state the n
- Never list 15 actions — rank ruthlessly and name the top few
- Never blame sales or product reflexively; let the data assign the theme

word-document

- Never fake headings with bold text — use heading styles, or the document's structure breaks
- Never dump unstructured text — apply the section hierarchy the doc type needs
- Never hand-format what a style should do — consistent styles beat per-paragraph fiddling
- Never invent contract/legal terms silently — mark drafted clauses and recommend review for anything legal
- Never claim a file was produced without code execution — fall back to the markdown export instead

workshop-facilitation-guide

- Never design a workshop without explicitly linking every activity to a session goal — purposeless activities waste participant time
- Never schedule more than 90 minutes of continuous structured activity without a break
- Never close a workshop without capturing decisions and actions before the room empties — post-session follow-up is too late
- Never plan a workshop without considering psychological safety for sensitive topics — establish ground rules at the start
- Never underestimate timing — add 20% buffer to all activity estimates, especially for groups over 8 people

writing-great-skills

- A vague description with no trigger phrases — the skill never gets picked
- An Output Format that *describes* the artifact instead of *templating* it
- Quality Checks that aren't observable ("output should be good")
- Leaving `<!-- TODO -->` or `[bracketed]` placeholders in the final file
- Overlapping so heavily with an existing skill that the model can't choose between them

writing-plans

- Never plan happy-path-first when a hard unknown exists — sequence to kill the plan early if it's killable
- Never write steps without verifications — unverifiable steps are where sprawl enters
- Never bury discoveries — when execution reveals the plan is wrong, revise the PLAN visibly (see executing-plans), don't improvise around it
- Never gold-plate a checklist task into a project plan — ceremony must earn its cost
- Never treat the plan as the deliverable — a beautiful plan for the wrong goal fails the interview-me test; brief first, plan second

youtube-script-writer

- Never write paragraphs of dialogue without accompanying visual cues. YouTube is a visual-first medium; every paragraph of speech needs visual transitions.
- Never pitch sponsors, channel subscriptions, or external links during the hook (first 60 seconds).
- Never create a single generic hook; always provide 3 distinct hook variations (Contrarian, Story, Pattern Interrupt) to give the creator flexibility.
- Never use a generic outro that triggers the "viewer exit ramp" (e.g., "That's all for today's video, hope you enjoyed, see you next time!"). Suggest another video to keep viewers on the platform.

