user@agentic-os:~/architecting$

Architecting the Agentic OS

Mastering Claude Skills 2.0 & The Answer Engine Optimization (AEO) Shift

The Cost of AI Amnesia

The Chat Paradigm

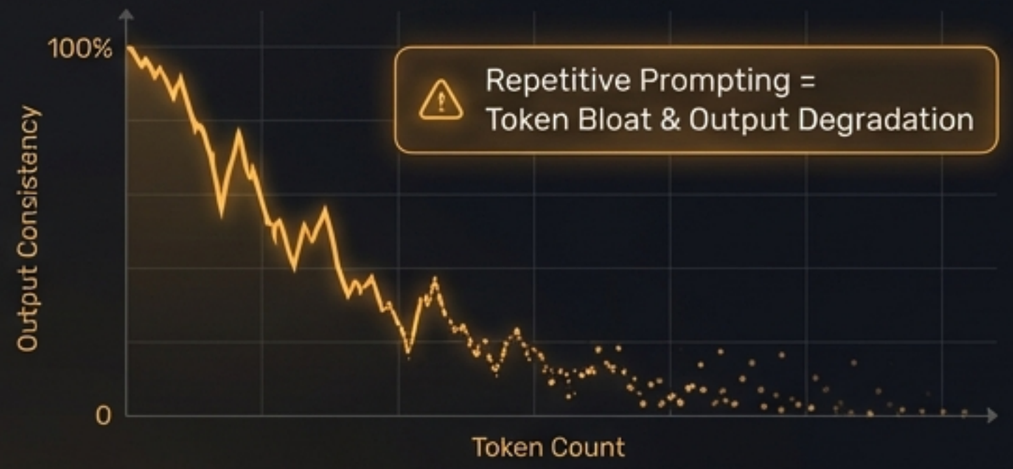

The Executable AI OS

The Context Management Matrix

	Custom Instructions	Claude Projects	Claude Skills	MCP (Model Context Protocol)
Portability	Low	High (within project)	Global	Global
Stackability	None	Low	High (Multiple skills per chat)	High
Token Efficiency	Poor	Medium (Persistent Containers)	High (Progressive Disclosure)	Low (Dumps all tool metadata upfront)
Primary Use Case	General conversational tone	Ongoing work with consistent reference docs	Executable capabilities and repeatable workflows	Accessing external databases and APIs

The Anatomy of a Claude Skill

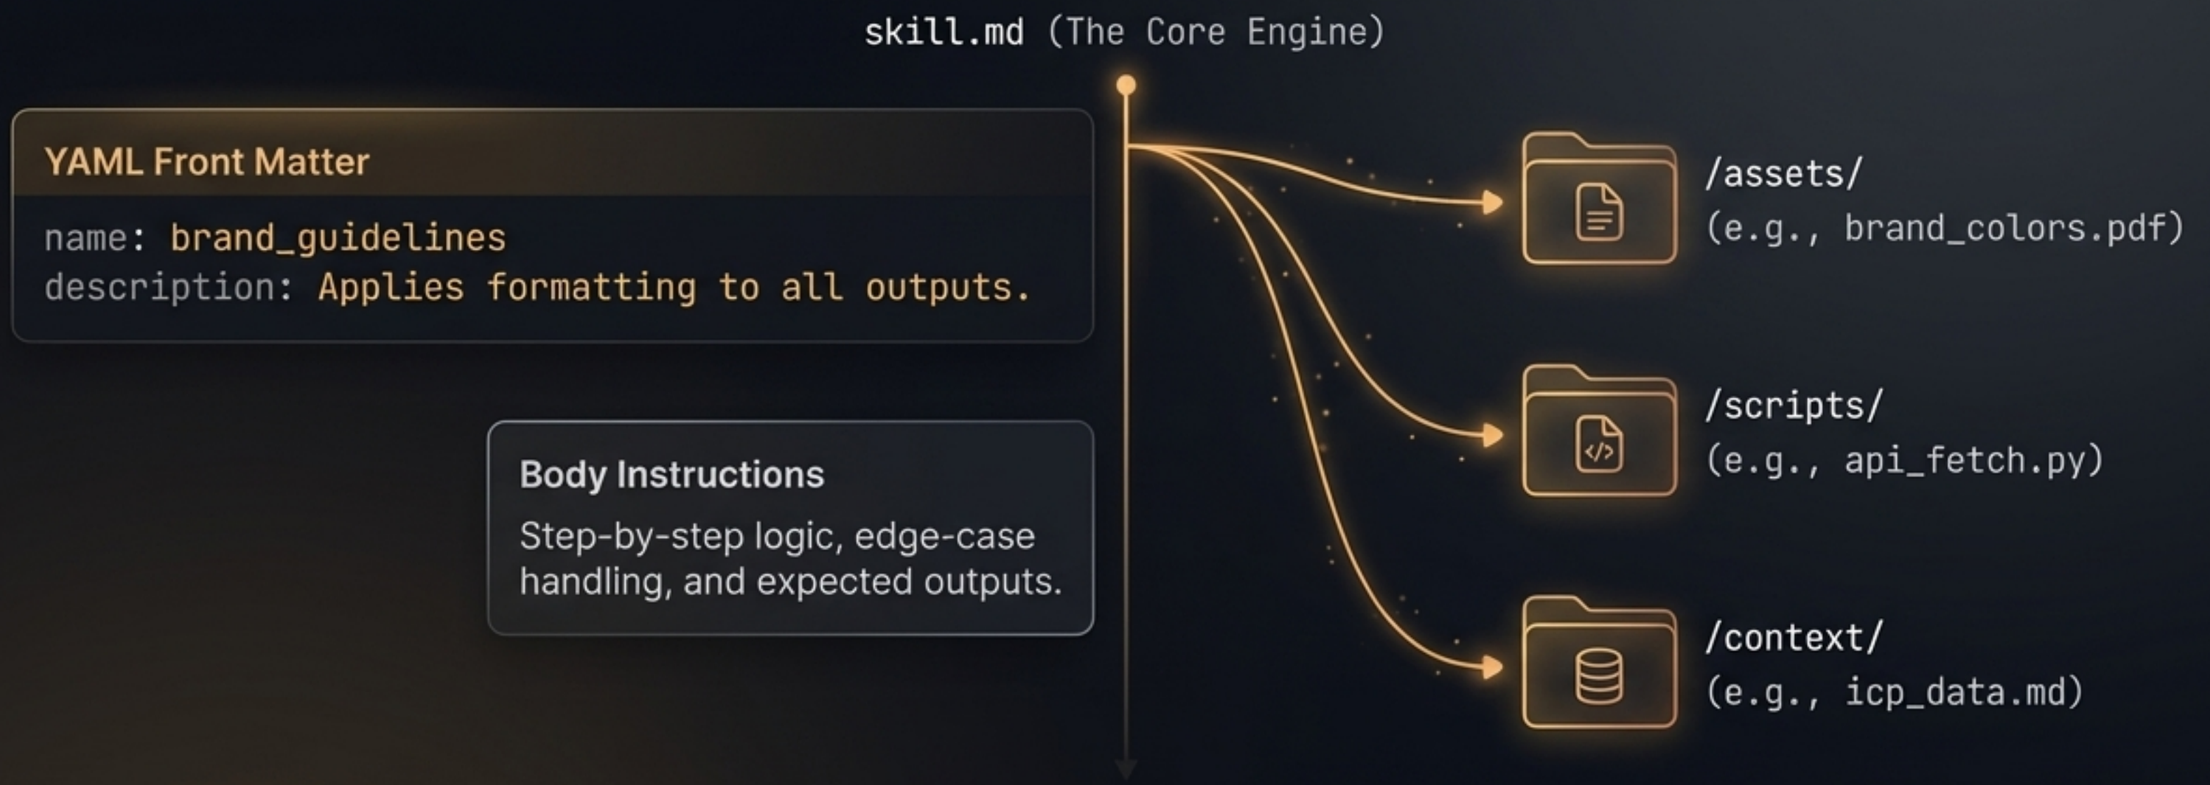

Only add context that Claude doesn't already possess. Lean architecture equals faster execution.

The Progressive Disclosure Funnel

YAML Front Matter

Always Active.
`~100 Tokens`

Claude reads the name and description to decide if the skill is needed.

skill.md Body

Loads on Trigger.
`~1,000 to 5,000 Tokens`

Loads only when a user prompt matches the YAML description.

Reference Files & Scripts

Loads on Demand. Unlimited Capacity

Loaded into context only when a specific step in the skill.md requires it.

Claude Skills 2.0: The Evaluation Engine

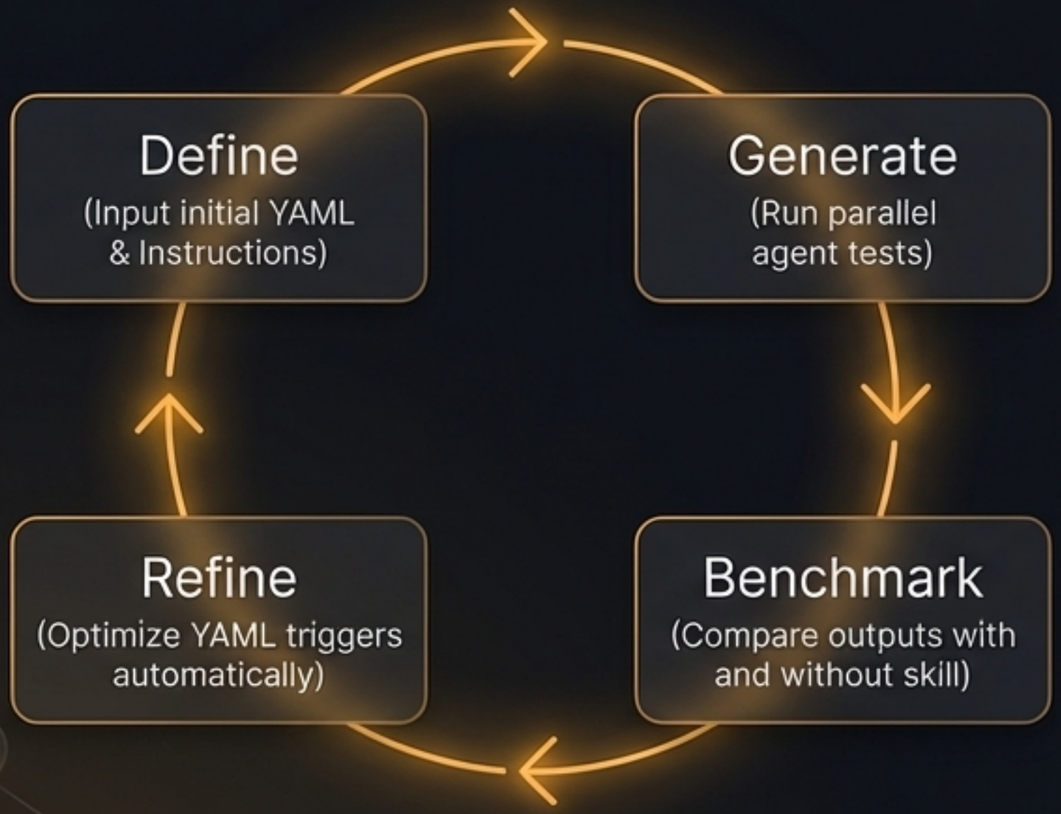

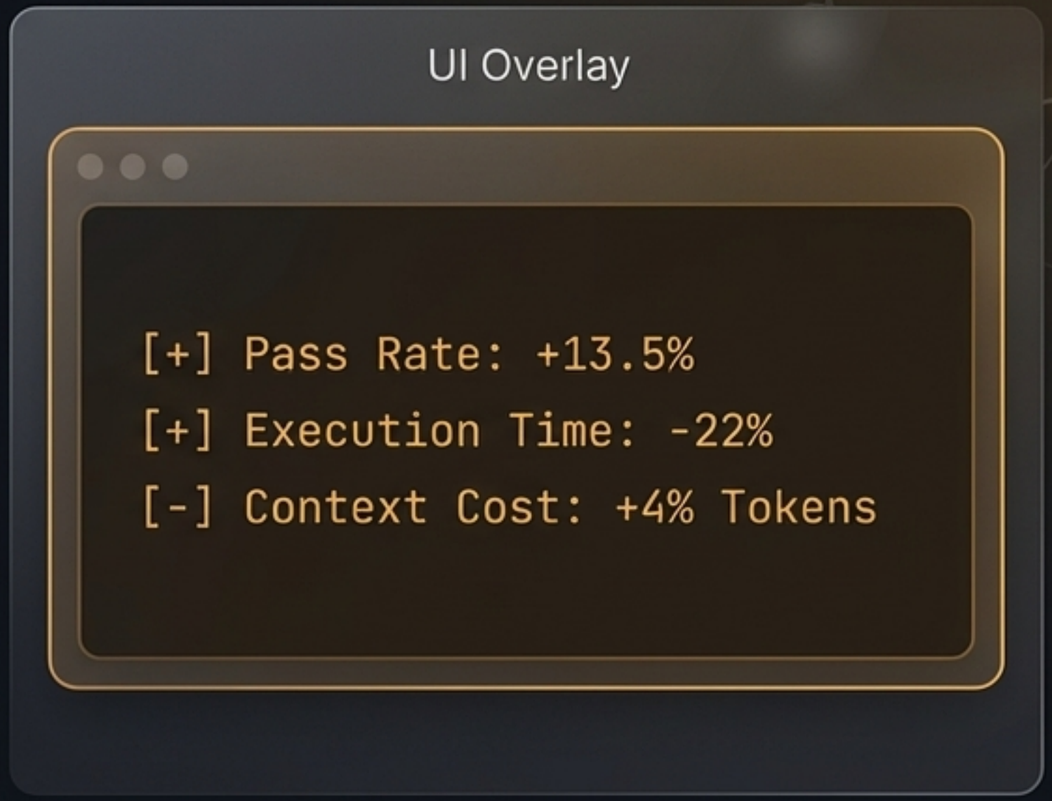

You are no longer building skills on vibes.
You are compiling tested, deterministic AI software.

Skill Typology: What to Build

Capability Uplift

Definition: Fills a gap in the base model's current intelligence (e.g., Front-end UI design, complex PDF extraction).

Lifecycle: Temporary. Has an expiration date. When Opus 5.0 drops, this skill may become obsolete or restrictive.

Evaluation Focus: Pass rate vs. Base Model.

Encoded Preference

Definition: Encodes your unique business logic, brand voice, or multi-step SOPs (e.g., NDA review checklists, Brand Theme generation).

Lifecycle: Durable. Survives model updates because it represents your specific way of doing things.

Evaluation Focus: Fidelity to the established workflow.

Replacing the Subscription Stack

The Micro-SaaS Graveyard

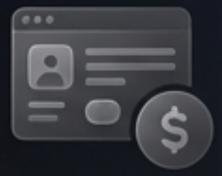

Outbound CRM Tool ($99/mo)

`lead_scorer.md`

Parses CSV, enriches via API, outputs prioritized pipeline.

Survey Analysis Tool ($49/mo)

`survey_analyzer.md`

Ingests hundreds of rows, generates Excel + 10-slide summary deck.

Content Repurposer ($29/mo)

`youtube_to_leadmagnet.md`

Scrapes transcript, formats to Notion, applies custom brand guidelines.

A well-architected skill directory replaces isolated UI dashboards with a unified, context-aware command center.

Blueprint: The Automated Analyst

1. Trigger

Natural Language: Run a pulse check on my projects.

2. Skill Activation

pulse_check.md loads instructions.

3. MCP Integration

Claude delegates to a specialized sub-agent. Searches ClickUp via MCP using Hardcoded List IDs.

4. Data Synthesis

Ingests raw JSON, cross-references with business priorities (Reference File).

5. Output

Generates executive summary flagging manual follow-ups.

Pro-Tip: Hardcode static IDs in the skill.md to prevent expensive API search queries.

The Ultimate Arena for Agentic Leverage

We build **Agentic Skills** to scale our output. But output but is useless without discovery. The highest-ROI application for your Claude OS today is automating the transition from traditional SEO to **Answer Engine Optimization (AEO).**

.claude/skills/

76% of all AI Overview citations come from the top 10 traditional search results. The foundation hasn't changed; the extraction has.

The Search Paradigm Shift

↘ The Threat

Informational keywords (How-tos, What-is) are losing up to 35% of clicks to zero-click AI Overviews.

✨ The Opportunity

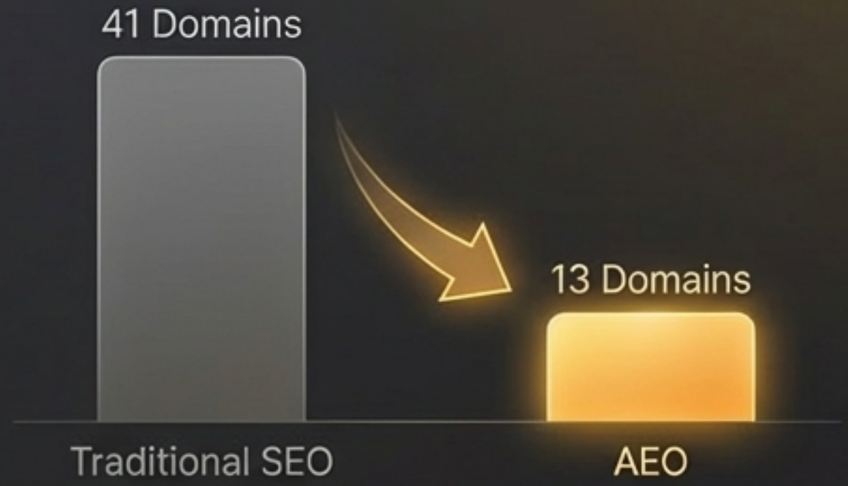

Traditional Page 1 SEO requires ~41 referring domains. Ranking in AI Answer Engines requires an average of just ~13. The barrier to entry has plummeted.

AI is better than humans at explaining surface-level facts. If a query requires no nuance, the click is dead. You must pivot to intent.

The B.I.D. Keyword Vetting Method

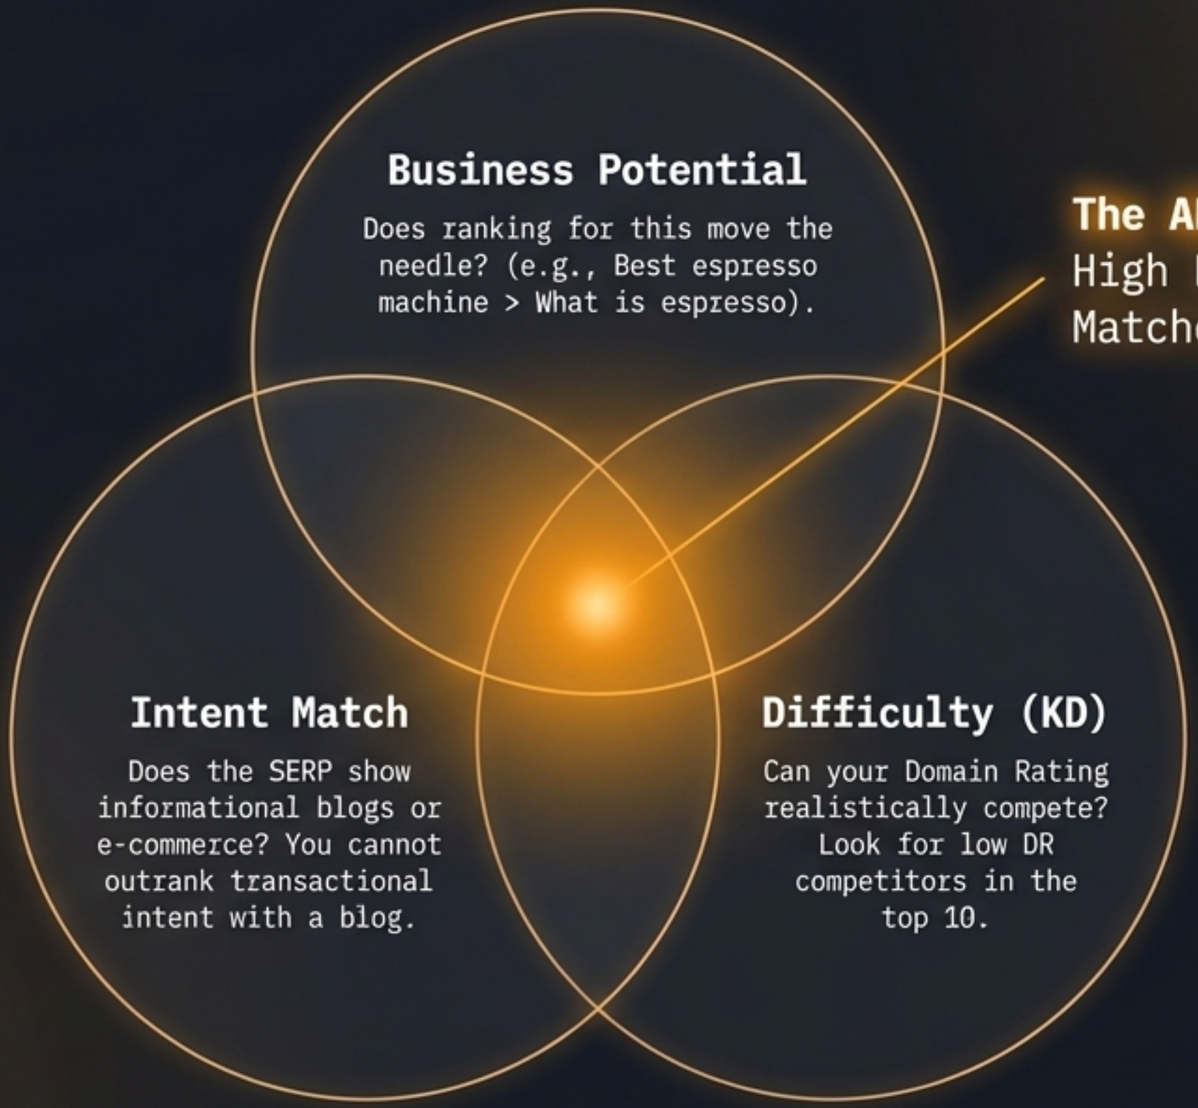

The AEO Sweet Spot:
High Business Value, Matched Intent, KD < 30.

The AEO Triple Crown

AI Overviews trigger an average of 3 other SERP features simultaneously.

Zone 1: AI Overviews

Synthesize intent-based answers.
Citation-worthy depth.

Zone 2: Featured Snippets

Over 50% overlap with AIOs.
Prioritize exact-match phrasing.

Zone 3: People Also Ask (PAA)

Appearing in 60%+ of AEO results. Use question-based H2 headers that match PAA queries exactly.

Target the **9,000+ keywords** that trigger all three features. This is how you dominate the entire first screen.

Brand Mentions as Training Data

The Old Game

Backlinks from high DR sites.

The New Game

Co-occurrence in LLM training data. Every Reddit thread, YouTube comment, and brand partnership feeds the AI.

Case Study: The Electric Bike

A strategic YouTube sponsorship led to massive spikes in Reddit discussions. This user-generated content was ingested by OpenAI and Google, resulting in the brand dominating AI queries without traditional SEO.

Stop optimizing just for Google's crawler. Start optimizing to be the most cited entity in the LLM's latent space.

Synthesis: The Agentic AEO Pipeline

Input | Target B.I.D. Keyword + Raw Draft

Skill Activation | `aeo_content_optimizer.md`

Execution Steps

1. Encapsulate answers in the first 2-3 sentences under H2s (for Featured Snippets).
2. Format H2s as exact-match "People Also Ask" questions.
3. Inject strategic CTAs to guide the user off the AI Overview and onto your site.

Output | A fully formatted, AEO-compliant markdown file ready for publishing.

Combine the precision of Encoded Preference Skills with the visibility of Answer Engine Optimization.
This is the blueprint for the next era of the web.