

eve.

The full comparison. *One that actually scales.*

Most teams can produce one great post. The hard part is doing it every week, across multiple execs, without the quality dropping.

	DIY ChatGPT / Claude Projects	BUILD Custom in-house	BUY Eve
WHAT IT COSTS			
Software cost	\$20–200/mo	API fees — typically \$200–500/mo at scale	\$1,500/mo (3 execs, 10 seats, no per-post fees)
People cost	10–15 hrs/week across the team — a few hours from the exec for input, the rest from whoever's prompting, editing, and posting	\$100K+/yr in engineering time to build, plus a dedicated PM to maintain	5–30 min/week exec time — team handles the rest
Hidden cost	Every hour your exec spends prompting AI is an hour not spent running the business	Engineering cycles pulled from core product — the pipeline is never their top priority	None — maintenance, model updates, and prompt evolution are included
WHAT IT TAKES			
Time to first draft	Weeks to months of setup — pulling writing samples, crafting system prompts, configuring your project — before you produce anything usable	8–12 weeks before the pipeline produces anything	Same day — voice model built during 30 min onboarding session
Total time per post	What used to take ~5 hours becomes ~1 — AI gets you ● 75% there, but prompting, editing, and polishing still take real time	● Minutes if well-calibrated, but calibration takes months of engineering	Under 15 min total — 5 min exec input, 5–10 min team ● review. Multi-agent pipeline gets closer to final on the first pass
How execs contribute	Starts with getting time on the exec's calendar for ● input, then a series of prompts and edits to get the output right	● Depends on the interface you build — could be easy or could require training	● Send a voice memo or email from your commute — draft is waiting when you sit down
Time to publish-quality voice	Weeks to months — you manually refine prompts and samples, adjusting by feel after each post	3–6 months of calibration after the build	First draft — execs consistently say "this sounds like me"

	DIY ChatGPT / Claude Projects	BUILD Custom in-house	BUY Eve
WHAT YOU GET			
Voice fidelity	<ul style="list-style-type: none"> Recognizable tone with samples loaded, but often sounds like AI on first pass — nuance and personality require heavy editing 	<ul style="list-style-type: none"> Only as good as the prompts and data you feed it — varies widely based on engineering investment 	<ul style="list-style-type: none"> Learns how you think, not just how you write
Editorial coaching	<ul style="list-style-type: none"> None — it generates text, you decide if it's good enough 	<ul style="list-style-type: none"> Only if you build a feedback layer — most teams don't 	<ul style="list-style-type: none"> Flags weak angles, suggests stronger framing, and adapts coaching by channel — users say the notes are as valuable as the drafts
Channels covered	<ul style="list-style-type: none"> Can request multiple, but you provide the guidance per channel or it defaults to generic 	<ul style="list-style-type: none"> Each channel is a separate build 	<ul style="list-style-type: none"> LinkedIn, blog, newsletter, podcast, board memos, and more
Team collaboration	<ul style="list-style-type: none"> Solo tool — no approvals, no versioning 	<ul style="list-style-type: none"> If you build it 	<ul style="list-style-type: none"> Draft status, collaboration notes, versioning, brand rules — approvals and publishing calendar coming soon
Multiple executives	Separate project per person, managed manually	Multiplies complexity with each voice	Each exec gets their own isolated model — 3 included
RISK			
Data & privacy	Writing samples and outputs live in OpenAI/Anthropic's consumer platform — subject to their data policies	You own the architecture, but you're responsible for data handling and compliance across every API you touch	Voice data fully isolated per client — never used to train other models. Compliance layer built in.
Knowledge retention	Tribal knowledge lives in one person's head — when they leave, the prompts, samples, and context walk out with them	Builder leaves, system decays — new hires face a steep ramp to understand the pipeline	Voice models, brand rules, and editing history live with Eve — your organization owns the IP, not any one person
Does it improve over time	Incremental at best — even after years of use, the tool still asks basic clarifying questions because it doesn't retain context	Possible with dedicated engineering cycles, but improvements are manual and don't transfer across voices	Compounds automatically — every edit, every post, every interaction teaches Eve how you think and how your business is evolving. Less editing by month 3, minimal by month 6
Ongoing maintenance	Prompt upkeep, sample curation	Models update, prompts break — doable if you staff it, but it's ongoing work	Fully managed — and as foundation models improve, so does your output
Bottom line	~1 hour of back-and-forth per post, even after months of setup	3–6 months before voice quality is usable	Day one publish-ready from first draft