

6

NOT A CHATBOT · AN ACTION ENGINE · WITH A BILL

Manus AI for Professionals.

Less product tour, more value. Where it earns its credits, where it burns them, and when to reach for Claude instead.

THE BRIEF, IN THREE BEATS

- 01 Manus doesn't answer. It goes and does — that's the whole product.
- 02 The value is in *delegation*, not conversation.
- 03 Credits are the constraint. Prompt like you're paying — because you are.

Your best analyst takes a two-day research task. Manus does it while you sleep.

That's the whole pitch. Everything else in this deck is how to make it true — and how not to burn \$200 doing it wrong.

Manus doesn't compete with Claude on **answers**. It competes with your **junior analyst** on **delivered work**.

THREE COROLLARIES

- 01** The unit isn't a prompt — it's *a delegated task with a deliverable*.
- 02** You measure ROI in *analyst-hours saved*, not tokens spent.
- 03** The professionals who win are ones who already know what *good junior work* looks like.

THE VALUE CLAIM

THE HONEST TRADE

Where Manus wins vs. Claude/ChatGPT.

Wins: multi-step research, competitive landscape decks, spreadsheet builds, sentiment reports, SOPs, MVPs.

Loses: conversational depth, code refactor sessions, production apps, anything requiring your judgment mid-flow.

Rule of thumb: if the task ends with "*send me the file*", use Manus. If it ends with "*let's discuss*", use Claude.

Chatbot vs. Manus. One sentence each.

CHATBOT (CLAUDE / CHATGPT)

You ask → it answers. You keep the browser open.

- SHAPE · Turn-by-turn conversation. You drive.
- OUTPUT · Text in the window. You copy, paste, format.
- TOOLS · Only what the model was given this turn (browser, code, MCP).
- DURATION · Seconds to minutes. Session ends when you leave.
- BILL · Flat monthly subscription. Usage caps.

MANUS — ACTION ENGINE

You brief → it works. You close the tab.

- SHAPE · Multi-agent orchestration in a cloud VM. It drives.
- OUTPUT · Files — .xlsx, .pptx, .pdf, deployed URLs, dashboards.
- TOOLS · Browser, terminal, code, Gmail, Calendar, connectors — persistent.
- DURATION · Minutes to hours. Runs while you're offline.
- BILL · Credit-based. Complex task = more credits.

If you catch yourself starting a Manus task with "let's think about..." — you're using it wrong. Stop, open Claude, come back when you know the deliverable.

Six plays. Two axes: **task shape** × **frequency**.

#	PLAY	TASK SHAPE	FREQUENCY	TYPICAL DELIVERABLE
06	Research delegation	Multi-source, multi-hour, one question	Ad-hoc, weekly	Sourced brief · 8-15 pages · citations
07	Analyst-in-a-box	Data-in, chart-out — stock, KPI, survey	Weekly, monthly	.xlsx + .pptx or dashboard URL
08	Skills library	Packaged domain playbook, reusable	Build once, run forever	A Skill you invoke by name
09	Parallel fan-out	Same task × 10 targets (competitors, cities)	Ad-hoc	10 parallel runs, one merged output
10	Overnight tasks	Scheduled, low-oversight, long-running	Daily, scheduled	A morning inbox message with the file
11	Deliverable-first prompting	The prompt template that saves credits	Every task	Fewer retries, tighter output, lower bill

#01 The competitive landscape memo, delivered before your coffee.

THE BRIEF THAT WORKS

The four-line memo Manus is looking for.

TARGET · "Competitive landscape for AI coding tools, Q2 2026."

DIMENSIONS · Positioning · pricing · recent funding · notable features · churn signals.

SOURCES · Primary only — vendor docs, press releases, Crunchbase, G2/Reddit for signal.

DELIVERABLE · Google Doc, 8-12 pages, one-page exec summary on top, citations inline.

The word "brief" (not "help me think about") is the entire trick — it tells Manus the task ends with a file, not a chat.

ANALYST-HOURS SAVED

6–12 hours → 20 min

Junior analyst: 6-12h to draft. Manus: ~15-25 min wall-clock, ~800-1500 credits. Your job shifts from "gather" to "review + red-pen."

THE TRAP IN THIS PLAY

Don't ask an open question.

"What should we know about the AI coding market?" burns 3000+ credits and returns a Wikipedia-shaped answer. Name the file you want — and Manus builds THAT file.

#02Give it the raw file. Get the deck.

THREE REAL TASKS THAT WORK

01 · TICKER TO DECK

"TSLA last 5 years — fetch, chart price + volume, tag 6 news events, output pptx." Manus pulls Yahoo Finance, builds a 6-slide deck.

02 · CSV TO DASHBOARD

Upload your team's 12-month KPI export. "Build a Streamlit dashboard, deploy it, send me the URL." Get a working URL in 8-12 min.

03 · SURVEY TO SENTIMENT SLIDES

1,200 Zendesk tickets. "Classify by sentiment + top 5 pain themes, output 8-slide pptx with quotes." One .pptx, cited quotes.

ANALYST-HOURS SAVED

4–8 hours → 15 min

Chart-work + slide-formatting is where junior time evaporates. Manus does the pandas + matplotlib + pptx dance in one shot.

THE TRAP IN THIS PLAY

Verify the numbers. Manus is confident, not necessarily right.

Financial figures, chart axes, sentiment counts — spot-check 3 randomly before the deck goes to a VP. This is judgment work you don't outsource.

#03 Turn every good task into a Skill. Then invoke it by name.

WHAT A SKILL IS, IN ONE PARAGRAPH

A Skill is an operating manual you hand Manus — procedures, standards, examples of good output.

CONTAINS · System instructions + few-shot examples + tool preferences + output format spec.

EXAMPLES · "Legal contract review · v3", "Weekly board update · v2", "Customer churn diagnosis · v1".

INVOCATION · "Run *Weekly board update* for last week's data" — one line, not a 200-word brief.

Rule: any task you'll do 3+ times becomes a Skill. Anything below that stays a one-off brief.

COMPOUNDING MATH

1 Skill × 50 uses/year

1 hour to build the Skill · 50 runs at 5 min setup each → you save ~4 person-days per Skill per year. Ten Skills = 40 days.

THE TRAP IN THIS PLAY

Don't Skill-ify too early.

Ship a task 3-5 times as one-offs first. Watch what varies vs. what stays constant — that's what goes in the Skill. Premature Skills are brittle.

#04 Same task × 10 targets. One merged deliverable.

THE PATTERN THAT BEATS SERIAL

- 01 · Take one recipe — e.g. "One-pager on company X: positioning, pricing, funding, hires."
- 02 · Fan it out across N targets in *separate* Manus tasks (each runs in its own VM).
- 03 · When all N finish, run a final *synth* task: "Merge these 10 one-pagers into a comparison table + ranked shortlist."
- 04 · Real use: 10 competitors · 20 cities for market entry · 15 candidate suppliers · 12 acquisition targets.

The clever move: give each fan-out task the *same* output schema so the synth task just concatenates.

WALL-CLOCK LEVERAGE

Serial 3.5h → Parallel 25 min

10 competitor briefs, serial: ~3.5 hours. Parallel with 5 concurrent tasks + synth: ~25 min. Same credit total; radically different wall clock.

THE TRAP IN THIS PLAY

One inconsistent schema ruins the synth.

If fan-out #7 returns freeform prose while the others return a JSON table, the merge step burns 500+ credits patching the odd one. Lock the schema in the recipe.

#05 The 7am inbox message with the file already in it.

TASKS WORTH PUTTING ON A SCHEDULE

- DAILY · 6am** · "Overnight news on my 12 watchlist companies. One-liner per co, quotes for anything material. Email me the digest."
- WEEKLY · Mon 7am** · "Pull last week's KPI export from Sheets. Regenerate the board dashboard. Update the exec summary section."
- DAILY · 8pm** · "Yesterday's calendar → note who I met, tag by account, draft 2-sentence follow-ups in Gmail drafts."
- MONTHLY · 1st** · "Freshen the competitive landscape doc. Diff vs. last month. Highlight what changed."

THE CONNECTORS THAT UNLOCK IT

- Gmail** · read source threads, send digests, park drafts.
- Google Calendar** · meeting notes, follow-up triggers.
- Google Sheets / Drive** · pull KPI exports, save outputs.
- Slack (via MCP)** · post the morning digest to a channel.

THE TRAP IN THIS PLAY

Set-and-forget quietly burns your monthly credit budget.

A daily 800-credit job = ~24,000/month — that's a Pro plan. Add a monthly credit budget alert & a task-count cap on the schedule.

#06 Name the file first. Everything else follows.

VAGUE VS. DELIVERABLE-FIRST — SAME TASK

✗ VAGUE — burns ~2500 credits

"Look into the AI coding tools market and help me understand what's going on."

✓ DELIVERABLE-FIRST — ~900 credits

"Build me a Google Doc: title 'AI Coding Tools · Q2 2026 Landscape'. Structure: (1) 1-page exec summary, (2) 8 competitor one-pagers with positioning/pricing/funding/notable feature, (3) comparison table, (4) 3 emerging signals worth watching. Sources: primary only, citations inline. Length ~10 pages."

FIVE SLOTS THE BRIEF MUST FILL

Format · Google Doc · .xlsx · .pptx · dashboard URL

Structure · Sections in order, with headings

Sources · Primary only / include vendor docs / etc.

Length · ~10 pages · 8 slides · 500 words

Done-when · "Ready to send to CFO" · "Peer review draft"

THE TRAP IN THIS PLAY

Don't over-specify the *process*.

Name the destination + structure + sources. Do NOT prescribe "first do X, then Y, then Z" — that fights the multi-agent planner and wastes credits on your bad plan.

API You're chatting with the intern. Stop.

FIVE HABITS FROM CHATGPT TO UNLEARN

- 01 · "Actually, can you also..." mid-run — interrupts the plan, credits spent on backtracking.
- 02 · "First let's discuss..." — you're paying VM time to think out loud. Think in Claude, execute in Manus.
- 03 · Watching the live log for 20 min — Manus is designed for you to leave. Close the tab.
- 04 · Follow-up messages after done — the VM is fresh; you're paying re-context. Start a new task, or use a Skill.
- 05 · "Can you make it more comprehensive?" — the most expensive four words. Specify length up front instead.

COST OF THE HABIT

2–4× credit inflation

Every mid-run interruption forces the planner to re-scope. A 900-credit task turns into 2,500 with three "one more thing" edits.

THE EXIT

Draft the brief in Claude first. Paste it into Manus once.

Two-minute investment in Claude = 60% fewer mid-task interruptions in Manus. The two tools are complements, not substitutes.

AP2 "It built the app!" is not the same as "it's in production."

WHAT MANUS SHIPS WELL – AND WHAT IT DOESN'T

SHIPS WELL · MVP dashboards, internal tools, one-off scripts, demoware, prototypes, PoCs, static reports. The "get a working URL by lunch" tier.

SHIPS BADLY · Multi-tenant apps, auth flows, database migrations, anything with SLAs. The output *compiles* but rarely survives real load.

SHIPS WORST OF ALL · "Fix the bug you introduced" — the debug loop with Manus burns 5× the credits vs. Claude Code, and often doesn't converge.

THE RULE

Manus builds v0. Humans (or Claude Code) ship v1.

Treat Manus output as the fastest possible scaffold — not the codebase.

THE EXIT

Ship the MVP with Manus. Hand off the repo the moment users appear.

Have Manus push to GitHub, then move debug/hardening to Claude Code or a human. Don't let Manus own the second dev cycle.

AP3Credit-based pricing punishes lazy briefs.

#	THE LEAK	WHY IT DRAINS	THE PLUG
01	Open-ended briefs. "Research the market."	Planner keeps expanding scope until it thinks it's done. 3-5× the credits of a scoped brief.	
02	Mid-task edits. "Also add..."	Every interruption re-plans. VM stays hot, meter runs.	Batch changes into a single follow-up task after done.
03	The debug spiral. "Fix the bug. Now the new bug. Now..."	Multi-step debugging inside Manus loses context. Same bug re-appears.	After 2 fix attempts, export repo, move to Claude Code.
04	Unbounded scheduled tasks. Daily digest running forever.	800 credits/day × 30 = 24,000/month. That's a Pro plan just for one job.	Cap per-task credit; alert at 70% of monthly budget.

AP4 Manus is an individual power-tool, not a team platform.

WHAT Q2 2026 MANUS DOESN'T DO

- 01 · No true shared workspace — each seat runs its own tasks in isolation.
- 02 · No role-based access control — can't restrict who runs which Skill.
- 03 · No cross-user Skill governance — a Skill shared with the team is a copy, not a synced source.
- 04 · No audit trail teammates can inspect — the "reasoning replay" is per-task, per-user.
- 05 · No SOC-2 workflow enforcement, DLP policy, or centralized billing granularity (Q2 2026).

THE CONSEQUENCE

One person's insight stays one person's.

Every teammate rediscovers the same "good brief" independently. No compounding across the team.

THE EXIT

Keep the shared IP in Git, not in Manus.

Store Skills / briefs / prompt templates in a team repo (Notion, GitHub, Confluence). Manus is the runtime; the source-of-truth is somewhere your teammates can PR.

Six slots. Fill them in this order. Ship better briefs.

MANUS BRIEF · V1

GOAL

What business question does this answer? One sentence.

DELIVERABLE

Format: [Google Doc / .xlsx / .pptx / deployed URL].

Filename: [exact title].

Length: [pages / slides / rows].

STRUCTURE

Section 1, section 2, section 3 – each with sub-bullets.

SOURCES

Primary only / vendor docs / avoid Reddit / cite inline.

DONE-WHEN

"Ready to send to [CFO / peer / customer]."

BUDGET

"Stop if credits exceed [N]."

WHY THIS WORKS

Goal stops research from drifting.

Deliverable tells the planner when to stop.

Structure lets it plan in one pass, not five.

Budget is the safety net you'll forget until you get the bill.

PRO TIP · SKILL IT

After 3 briefs using this template for the same job (e.g. "competitor one-pager"), lift the template into an Agent Skill.

Then invoke it with just: *"Run competitor-one-pager on {company}."*

Three tiers. One question: how many tasks per week?

PLAN	PRICE	MONTHLY CREDITS	~TASKS PER MONTH	WHO IT'S FOR
Free	\$0	300 daily (expire 24h)	~1-2 light tasks/day	Trying it out. One-off questions. Not for work.
Starter	\$39/mo	3,900 + 300 daily refresh	~4-8 medium tasks/week	Individual professional. Weekly research + analysis.
Pro	\$199/mo	19,000 + 300 daily + concurrent	~15-30 tasks/week	Power user. Daily delegation + parallel fan-out + scheduled jobs.

CREDIT-PER-TASK RULES OF THUMB

Light (single-source lookup, simple format): ~200-400 credits.

Medium (multi-source research, chart/deck output): ~800-1,500 credits.

Heavy (parallel fan-out, code + deploy, long-running): ~2,000-4,000 credits.

ROI MATH

\$199/mo ÷ 20 medium tasks = \$10/task.


Each task saves ~3 analyst-hours. At \$50/hr blended cost, that's \$150 saved per task. ROI: 15×. The math works unless you burn credits on vague briefs.

One question: does the task end with "send me the file"?

THE TASK SHAPE	REACH FOR	WHY THIS TOOL WINS HERE
"Research X, deliver an 8-page doc."	Manus	Multi-step browsing + file assembly + close-the-tab autonomy.
"Help me think through strategy."	Claude	Conversational depth, nuance, back-and-forth reasoning.
"Refactor this module + add tests."	Claude Code	Codebase context, test loop, diff-only editing, persistent repo.
"Browse + book + compare flights."	ChatGPT (Operator)	Consumer-grade web automation, plugin ecosystem.
"CSV → dashboard → deployed URL, by lunch."	Manus	End-to-end: code, build, deploy — in one autonomous run.

The professionals who get the most from Manus use all three tools. Claude for thinking. Manus for doing. ChatGPT for consumer tasks.

Find your rung. Pick the next one.



Which task on your list ends with "send me the file"?

(a) · Research brief

One question. One doc. By morning.

Use the template from slide 16. Fill the six slots. Let Manus run overnight. Judge the output in the morning.

(b) · Analyst task

Data in. Charts + deck out.

Upload the CSV/spreadsheet. Ask for the specific chart + slide format. Spot-check 3 numbers before forwarding.

(c) · Competitive fan-out

Same recipe × 5 targets.

Write one competitor one-pager recipe. Fan it out across 5 companies in parallel. Merge with a synth task.

My bet is (b). Pick yours — but the best first task is one you already know what "good" looks like for.