

The Morning Briefing Agent

I build AI agents that brief you before your day starts.

by AI Soft Tech Solution

*One WhatsApp message at 6 AM with everything that matters:
who is waiting on you, what money is stuck, and what today holds.
Built for your business - not rented off a waitlist.*

1. The Gap - You Start Every Day Blind

Picture the first hour of a typical business owner's day. You wake up to **40+ unread emails**. Your WhatsApp has customer messages from three time zones. Today's appointments are in one app, pending orders in another, unpaid invoices in a third, and your task list - if it exists at all - is in a fourth. Six apps before your first coffee, and the first hour of your day is burned just figuring out **what matters today**.

That hour is not free, and it is not harmless. While you are triaging, the hot lead who messaged at midnight is going cold. The 10 AM appointment you forgot to prep for is creeping up. The invoice that went unpaid yesterday quietly becomes an invoice that is 30 days late - because nobody surfaced it on day 1.

The big AI companies have noticed this pain. On **21 October 2026, Anthropic launched Claude Duet** - a proactive AI agent aimed squarely at professionals and small-business owners. It does not wait for you to ask questions. It monitors your connected apps overnight, sends you a **daily briefing** in the morning, and surfaces what needs your attention - with your approval required before it takes any sensitive action. When a company like Anthropic builds a whole product around the morning briefing, the pain is real.

But Duet is a waitlisted research preview, built for a generic professional. It does not know that in **your** business an unpaid invoice from a regular client matters more than a newsletter, or that a WhatsApp message saying "price?" from a new number is a lead worth money. A briefing agent is only as good as its wiring into **your** tools and **your** priorities.

That is the gap this guide closes. You do not need to wait for a waitlist. The pieces to build your own Morning Briefing Agent - connected to your email, your calendar, your orders, delivered to your WhatsApp at 6 AM - exist today, and most of them are free to start.

2. The Agent - Piece by Piece

The Morning Briefing Agent is a custom agent we build for business owners. It does **one job**: every morning, before your day starts, it hands you a short, ranked brief of everything that happened overnight and everything today demands - and it never acts on anything important without your say-so. Four pieces, working in sequence:

Piece 1 - The Collectors

Collectors are small connectors that pull overnight activity from the places your business already lives:

- **Email:** the **Gmail API** (developers.google.com/gmail/api) fetches everything that arrived since yesterday's brief - new leads, client replies, supplier invoices.
- **Calendar:** the **Google Calendar API** pulls today's appointments, deadlines and reminders, with start times and locations.
- **Messages & orders:** WhatsApp Business / your CRM / your store - a **Google Sheet** you already keep, or your store's API (Shopify, WooCommerce) - contributes new enquiries, pending orders and unpaid invoices.

Nothing exotic: read-only pulls, once a night, from sources you already use. The agent sees what you would see - if you had an hour to go looking.

Piece 2 - The Ranker

Raw data is not a briefing - 60 items dumped on you at 6 AM is just a faster inbox. The Ranker is an **LLM prompt** that scores every collected item by what it is worth to you, using one simple rule: **money waiting beats customers waiting, customers waiting beats routine**. It then drafts the brief in a fixed format - top decisions first, schedule next, money watch, and a short FYI tail. Here is the actual prompt template at the heart of every build:

COPY-PASTE PROMPT - the ranker:

You are the Morning Briefing Agent for {OWNER NAME}, owner of {BUSINESS NAME}.

Here is everything that happened overnight and everything scheduled today:

EMAILS: {emails}

MESSAGES / ENQUIRIES: {messages}

CALENDAR TODAY: {calendar}

ORDERS & INVOICES: {orders}

Score every item 1-10 before writing anything:

- Money waiting (unpaid invoice, pending payment, refund request): 8-10
- A customer or lead waiting for a reply: 7-9 (a brand-new lead asking price = 10)
- Today's appointments and deadlines: 6-8
- Routine updates and FYI: 1-4

Then write the brief in EXACTLY this format:

GOOD MORNING - {today's date}

NEEDS YOUR DECISION (highest score first, max 5):

- [item] - [why it matters] - reply YES and I will draft the action

TODAY'S SCHEDULE:

- [time] - [appointment] - [anything to prepare]

MONEY WATCH:

- [unpaid invoices, pending orders, payments received]

FYI (one line each, max 5):

- [routine items]

Rules: under 250 words. Never invent an item that is not in the data. If nothing needs a decision, say "Nothing urgent - clear morning."

Those ranking rules are yours to tune. A salon owner ranks "no-show risk" high; a freelancer ranks "client waiting on a deliverable" high. The template above is the working default from our builds.

Piece 3 - The Delivery

A briefing you have to go and fetch is a briefing you will stop reading. The agent delivers to where you already are: **one WhatsApp message at a fixed time** - 6:00 AM, or 30 minutes before your first appointment, your choice. Delivery runs through the **WhatsApp Cloud API** (developers.facebook.com) or **Twilio** (twilio.com), both of which can send a scheduled template message for a fraction of a cent. One message, on your lock screen, readable in under five minutes with your coffee.

Piece 4 - The Approval Lock

This is the piece that makes the whole thing safe - and it is the same design principle the big launches are converging on. **The agent never sends, replies, books, cancels or pays on its own. Ever.** Items that need action are listed under "NEEDS YOUR DECISION" with a drafted action attached. You reply **YES** (or "yes 2" for item 2) on WhatsApp, and only then does the agent execute - send the follow-up, draft the invoice reminder, confirm the appointment. No reply means no action. You get an assistant that prepares everything and a business that never wakes up to a message it did not approve.

What the agent is NOT - honest limits: it does not think for you (it ranks by rules you set), it does not handle angry customers or negotiations (those land in your decision list, where they belong), and it is only as current as the sources you connect - which is why the checklist at the end starts with connecting them properly.

3. Why It Matters - The Business Impact

The briefing agent pays for itself in the most boring, reliable currency there is: **the owner's morning hour**. Run your own numbers:

Morning triage, done by hand	Time it eats
Scanning email + WhatsApp + apps	45-60 min / day
Across a working week	~5 hours / week
Across a working year	~260 hours / year
Owner time at \$25/hour	~\$6,500 / year
With the agent: read the 6 AM brief	5 minutes / day

That is the floor of the return - roughly **\$6,000+ a year of owner time** handed back, for a build that runs on free or near-free tiers. But the bigger wins are the disasters that quietly stop happening:

- **The lead that did not go cold.** 78% of buyers buy from the business that responds first. A lead that lands at 11 PM is in your 6 AM brief with a drafted reply - answered before your competitors have opened their inbox.
- **The invoice chased on day 1, not day 30.** The Money Watch section surfaces an unpaid invoice the morning after it falls due, while the client still remembers what it was for. Chasing at day 30 is awkward and slower to pay; a polite nudge at day 1 is routine.
- **The appointment you actually prepped for.** Today's schedule, with prep notes, in the same message - no more 10 AM surprises at 9:55.

Add it up honestly: one saved hot lead or one invoice collected two weeks earlier typically pays for the entire build. Everything after that - 260 hours a year of it - is profit.

4. The Build Roadmap - Step by Step

Two routes build the same agent. The **no-code route** (Make or n8n) gets a working brief in a weekend with zero programming. The **code route** (a small Node or Python script) is cheaper to run at scale and easier to customise. Both use the same four pieces from Section 2.

Step 1 - Connect your sources

Create a Google Cloud project (console.cloud.google.com) and enable the **Gmail API** and **Google Calendar API** - free. For orders/invoices, start with the **Google Sheet** you already track them in (or your store's built-in API keys). Grant **read-only** access only: the agent reads overnight activity; it never needs write access to your mailbox.

Step 2 - Set up the backbone

No-code: create a free account at **Make** (make.com) or self-host **n8n** (n8n.io). Create a scenario/workflow scheduled for **5:45 AM** daily with three starting modules: Gmail (watch emails), Google Calendar (get today's events), and your orders source. **Code:** a short script using the official Google client libraries does the same three pulls and writes the results to one JSON file.

Step 3 - Add the Ranker

Add an AI step - **Claude API** (console.anthropic.com) or **OpenAI API** (platform.openai.com) - and paste in the Ranker prompt template from Section 2, mapping the four {fields} to your three pulls plus calendar. Set the output format exactly as the template specifies. Test it on yesterday's real data and read the brief yourself: if the ranking feels wrong, adjust the scoring rules in plain English and re-run.

Step 4 - Wire the Delivery

Create a **WhatsApp Cloud API** number (developers.facebook.com - the test number is free) or a **Twilio** account (twilio.com). Add a final step that sends the Ranker's output as a WhatsApp message to **your own number** at **6:00 AM**. You are the only recipient - this is a private operations tool, not a marketing channel.

Step 5 - Add the Approval Lock

Create a second, small workflow that listens for your WhatsApp replies. When you reply **YES** to a decision item, it executes the drafted action (send the drafted email via the Gmail API's send scope - granted separately, only for this step) and confirms back to you: "Done - follow-up sent to {name}." Any other reply - or no reply - and nothing happens. This lock is not optional; it is what makes the agent trustworthy.

Step 6 - Deploy and run the 7-day test

No-code: turn the scenario ON - Make/n8n runs it in the cloud. **Code:** put the script on a schedule with a free cron - a **Vercel** cron job (vercel.com) or a **Railway** worker (railway.app). Then run it live for **7 days** without changing anything, and score each morning's brief: did it catch what mattered? did it miss anything? Tune the ranking rules once, at the end of the week, with real evidence.

ADVANCE - when the basics work

- **Evening preview.** Add a 9 PM mini-brief: tomorrow's first appointment and anything to prepare tonight.
- **Weekly roll-up.** Every Monday, the agent summarises last week: money collected, leads answered, invoices still open - the brief, zoomed out.
- **More sources, same rules.** Stock alerts, website form submissions, review notifications - each new collector joins the same ranking pass. Add one at a time, and only when the current brief is boringly reliable.

5. Troubleshooting - 4 Common Problems

1. The brief is too long

The Ranker is surfacing routine items. Tighten the prompt: cap NEEDS YOUR DECISION at 5 and FYI at 5 (already in the template), and add "If in doubt, leave it out - I can open the app myself." A brief you finish beats a brief you skim.

2. Something important was missed

Almost always a source problem, not a ranking problem: the item never reached the Collectors. Check that the source is connected and that the overnight window (brief time minus 24h) covers when it arrived. Fix the collector first; only then tune the scores.

3. The brief arrives late

Move the backbone schedule earlier (5:30 AM collection for a 6:00 AM brief) and make sure your host is not on a free tier that sleeps - Railway and Vercel cron both wake reliably; a laptop that is off does not. The agent must run in the cloud, not on a machine that sleeps.

4. YES replies do nothing

The Approval Lock workflow is separate from the briefing workflow - check it is enabled and listening on the same WhatsApp number. Test it with a dummy decision item before trusting it with a real client follow-up.

6. Build Checklist

Keep this open while you build. Tick each box as you go:

- ☐ **Sources connected:** Gmail + Calendar read access on; orders/invoices source (Sheet or store API) pulling last night's activity
- ☐ **Ranking rules written:** my scoring order set in the Ranker prompt - money waiting, customers waiting, then routine - in my words, for my business
- ☐ **Delivery time set:** one WhatsApp message to my own number at a fixed morning time (6:00 AM), arriving on schedule
- ☐ **Approval Lock ON:** the agent drafts actions but executes nothing until I reply YES - tested with a dummy item
- ☐ **7-day test done:** a full week of briefs scored - catches what matters, misses nothing important, ranking tuned once with real evidence

7. What's Next

You now know how the morning briefing works end to end: collect overnight, rank by what things are worth, deliver on WhatsApp at 6 AM, and lock every action behind your YES. A weekend to build the no-code version, five minutes a morning to run it - against roughly 260 hours a year of triage, and the leads, invoices and appointments that stop slipping while you are still finding your coffee.

Want it built for you instead? This is what we do at AI Soft Tech Solution: I build custom AI agents, apps and websites for businesses - briefing agents, support agents that finish tasks, ordering and booking systems - wired into the tools you already run, with your rules and your approval lock. The free guide above is the full recipe either way: build it yourself, or hand the wiring to someone who does it every week.