

The Agent-Run Inbox

The exact gameplan your AI agent runs to set up, warm, and protect your cold email infrastructure. You approve. It operates.

The product is `INSTALL.md`. This PDF is the human-readable manual: the reasoning, safety controls, proof gates, and failures behind the system.

Operating model

You approve. The agent operates.

V1 scope

Google Workspace only.

The manual explains. INSTALL.md operates.

You bought a lightweight operating system for an AI agent, plus this manual. The executable product is a single self-contained file called `INSTALL.md`. The manual exists so you can understand and audit what the agent is doing without becoming the operator.

1**Paste `INSTALL.md`**

Copy the entire file into Muse, ChatGPT, Claude, or another capable agent.

2**Answer five questions**

The agent determines where you are and finds the earliest incomplete mission.

3**Approve human steps**

The agent does what its access allows, gives you the shortest manual path when needed, and verifies afterward.

1**Diagnose****2****Pick domain****3****DNS****4****Identity****5****Warm****6****Go cold****7****Operate****REQUIRED Proof before progress**

Authentication, separation, downstream delivery, gradual ramping, thresholds, monitoring, and automatic pauses do not change.

YOUR CALL Your operating preferences

Approval style, who drafts, report cadence, send timing, and exact volume inside the ramp are configurable.

V1 operates with Google Workspace only. No dedicated warmup software is required. Where the agent cannot act directly, it gives you the shortest manual path and verifies afterward.

Read for understanding, not execution

The executable missions live in INSTALL.md. This manual explains why the system works, what proof looks like, and where shortcuts fail.

How to install	2	Missions 6 and 7	10
Mission 1: Diagnose	4	The standing instruction	11
Mission 2: Pick the domain	5	Under the hood	12
Mission 3: DNS	6	The graveyard	13-14
DNS reference	7	Required vs Your call	15
Mission 4: Identity	8	Connections and credentials	16
Mission 5: Warmup	9		

THE STRICT DEFAULT

"Never send any email without my explicit approval. Never treat Sent folder presence as proof of delivery; confirm downstream. Pause all sending on any bounce, complaint, or reputation dip, and tell me before resuming. Report your own misses before I ask."

Relax only the lines marked Your call. Required controls are what keep the system safe.

Diagnose before changing anything

The system starts with a public, pass-or-fail baseline for every domain you own. That prevents an agent from treating a guess, an old configuration, or an unverified status as safe infrastructure.

What this phase establishes

The agent checks MX, SPF, DKIM, DMARC, and SURBL for active brands, newsletters, dormant properties, and spare domains.

Why it comes first

Domain selection and repair depend on knowing what is already configured, what is clean, and what must never carry cold-email risk.

Proof gate

REQUIRED Every domain has a clear green or red result for all five checks. Unknown is not green.

Domain	MX	SPF	DKIM	DMARC	SURBL
yourdomain.com	Pass / fail	Pass / fail	Pass / fail	Pass / fail	Clean / listed

Interpret Spamhaus correctly. Code 127 . 255 . 255 . 254 from a public Domain Name System resolver means the resolver is blocked, not that the domain is listed. Recheck through another verification source such as MXToolbox.

No account access is needed. This phase uses public DNS and blacklist data. The owner supplies the domain list; the agent produces and records the evidence.

Separate the reputation boundary

Cold email belongs on a clean, neutral sending domain whose reputation is isolated from the website, newsletter, customer support, and primary inbox.

What this phase establishes

The agent ranks candidates by ownership age, name neutrality, and blacklist status, then removes unsafe options.

Why it comes second

Authentication and mailbox work should happen only after the permanent sending boundary is approved.

Prefer

- Aged ownership history
- Neutral startup or studio name
- No blacklist hit

Reject

- Main brand or newsletter domain
- Topic-specific names
- Personal-name domains

Proof gate

REQUIRED The chosen domain is neutral-sounding, clean on SURBL, and separate from every brand-critical identity. Age is preferred, not mandatory.

Connected or manual

Connected: With registrar access, the agent can register the approved domain after the owner approves the irreversible purchase.

Manual: The owner registers the approved domain. The agent verifies ownership and configuration afterward.

Never force a bad choice. If every owned domain is unsuitable, stop and register a clean, neutral domain before continuing.

Authentication is a four-part proof

The agent generates, applies, and verifies the mail records. The owner does not need to learn every token, but the system does not send until every record is observable and correct.

The four records

- MX routes incoming mail to Google.
- SPF authorizes Google's sending servers.
- DKIM signs mail with a 2048-bit key.
- DMARC begins at p=none with aggregate reporting.

The two paths

Connected: a one-zone DNS token lets the agent add and verify records.

Manual: the agent supplies exact values; the owner pastes them; the agent verifies.

Security: use a least-privilege token scoped to DNS edits on this one zone. Revoke it when setup is done.

Proof gate

REQUIRED MX, SPF, DKIM, and DMARC all verify green. Red or unknown means fix and recheck. Nothing sends early.

Do not preconfigure future tools. Add a sending tool's SPF include and DKIM only when that tool is actually used.

Abort path: if DNS cannot be changed, stop. The shortest safe next action is to grant one-zone access or paste the records manually.

The four Google Workspace records

Use this page to audit what the agent generates. Keep the values exact.

MX at @

```
1 aspmx.l.google.com.  
5 alt1.aspmx.l.google.com.  
5 alt2.aspmx.l.google.com.  
10 alt3.aspmx.l.google.com.  
10 alt4.aspmx.l.google.com.
```

TXT at @

```
v=spf1 include:_spf.google.com ~all
```

TXT at google._domainkey

```
Generate in Google Admin: Apps > Google Workspace > Gmail > Authenticate email  
Key length: 2048-bit    Selector: google
```

TXT at _dmarc

```
v=DMARC1; p=none; rua=mailto:you@yourdomain.com
```

REQUIRED The only acceptable pre-send result: four green checks. Connected means the agent writes and verifies. Manual means the owner pastes and the agent verifies.

Create the identity, then prove delivery

A sender identity is not ready because a setting exists. It is ready only when a message arrives downstream with the intended From header and no bounce-back.

Connected or manual

Connected: with Google Workspace Admin access, the agent can add the secondary domain, alias, alternate email, and send-as identity.

Manual: the owner makes the same changes in Admin; the agent gives the shortest click path and verifies afterward.

What actually failed

The alias existed in Gmail, but not on the Workspace user

Gmail accepted the message and placed it in Sent. Delivery then rejected it because the sending address was missing as an alternate email on the Workspace account.

The fix

1. Add the alias as an alternate email in Admin.
2. Re-add the Gmail send-as entry.
3. Complete verification.
4. Retest end to end.

The proof

- Correct received From header
- No bounce-back
- No silent identity substitution
- No reliance on Sent

Sent proves that Gmail accepted an instruction. It does not prove that the recipient's mail system accepted the message.

Proof gate

REQUIRED Correct downstream From header, zero bounce-backs, and no identity substitution.

Let evidence control the warmup

The agent owns the queue, delivery checks, logs, pauses, and reports. The owner controls recipients, voice, risk tolerance, and whether each send requires approval.

Week	Daily volume	Operating rule
1	10-20/day	Warm, friendly contacts. Space sends across the workday.
2	30-50/day	Increase only if Week 1 stayed clean.
3	75-100/day	Require healthy reputation and bounces under 2%.
4	150+/day	Use only after three clean weeks. Freeze on any dip.

<2%

Required bounce ceiling

<0.1%

Required complaint ceiling

Weekly

Required Postmaster review

Proof gate

Every send has a downstream delivery check and log entry, and the full history stays within the required thresholds.

YOUR CALL **Approval, timing, and drafting:** choose per-send approval or auto-send, nightly or exceptions-only reporting, and human or agent-written copy. Warmup mail should still read like ordinary human mail, with no fake question or manufactured reason to reply.

Real teardown

Why two paid warmup services were removed

One pinned the From header to the primary Workspace identity. The other could not authenticate a send-as alias over Gmail Simple Mail Transfer Protocol. Neither warmed the intended identity, so both were removed with zero emails sent.

Earn the cutover, then keep operating

Going cold is an evidence-backed operating decision, not the end of the system. After launch, the agent continues monitoring and protecting the infrastructure.

Mission 6

Go cold only after proof

Require 3 to 4 clean warmup weeks, bounces under 2%, complaints under 0.1%, and healthy Postmaster reputation. The owner approves the cutover.

After approval, cold mail stays on the warmed domain and DMARC tightens from p=none to p=quarantine.

Mission 7

Operate after launch

On every cycle, the agent checks limits, bounces, complaints, authentication, queue state, and reputation. It updates state and reports only what matters.

Capable agents can schedule the cycle. Pure chat agents resume when the owner says, "Run today's Agent-Run Inbox cycle."

Cutover proof gate

REQUIRED Clean history documented, thresholds met, healthy reputation confirmed, and explicit owner approval recorded.

A calendar date is never proof. Another day passing does not permit the system to scale through a bounce, complaint, authentication failure, or reputation dip.

Status	Meaning
CONTINUE	All required signals are healthy; run the next allowed action.
HOLD	Do not scale or advance; evidence is incomplete or a signal needs review.
PAUSE	Stop sending, record the event, diagnose, tell the owner, and resume only after resolution.

Install the human gate once

The strict version below protects the owner's control and the system's proof gates. Preferences can be relaxed later. Required controls cannot.

THE AUTHOR'S VERSION (STRICT)

"Never send any email without my explicit approval. Never treat Sent folder presence as proof of delivery; confirm by absence of bounce-backs. Pause all sending on any bounce, complaint, or reputation dip, and tell me before resuming. Report your own misses before I ask."

Relax any line marked Your call. The lines marked Required are not preferences; they are what keep the system out of spam.

You retain

- Which credentials to connect
- Manual-mode clicks
- Recipient judgment
- Your voice
- **YOUR CALL** Every-send approval
- The cold cutover decision

The agent owns

- Diagnostics
- Record generation or edits
- Verification
- Connected sends or scheduling
- Logging
- **YOUR CALL** Report cadence

The operating model: decide what to connect and which preferences to relax. The agent removes repetitive work, but Required proof never changes.

Under the hood

INSTALL.md can run the system without this page. Keep it for auditing the agent's work and troubleshooting failures.

Authentication

- MX routes mail to Google.
- SPF authorizes Google's senders.
- DKIM signs messages with a 2048-bit key.
- DMARC begins at p=none with aggregate reporting.

Tightening path

- Start DMARC at p=none.
- Move to p=quarantine after proof.
- Start SPF at ~all.
- Consider -all later, after every legitimate sender is authorized.

Control	Threshold	Agent response
REQUIRED Bounces	Under 2%	Pause on any spike and investigate.
REQUIRED Spam complaints	Under 0.1%	Pause on any complaint or rise.
REQUIRED Postmaster reputation	Healthy, checked weekly	Do not scale through a downgrade.
REQUIRED Warmup	10-20, 30-50, 75-100, 150+ per day	Promote only on clean evidence.

Reference rule: the technical values exist to audit the agent, not to turn the owner into the operator.

Three shortcuts that failed

The system was earned through failures that looked fine until checked end to end.

Failure 01

Trusting Sent

A message appeared in Gmail's Sent folder and then bounced. **Lesson:** Sent proves acceptance of an instruction, not delivery.

Failure 02

The alias bug

The send-as entry existed in Gmail, but the address did not exist as an alternate email on the Workspace user. **Fix:** add it in Admin, re-add send-as, verify, and retest.

Failure 03

Warmup services that could not warm the identity

One service pinned From to the primary identity. Another could not authenticate the alias over Gmail Simple Mail Transfer Protocol. **Lesson:** verify that a tool warms the exact identity before paying or connecting it.

Rule produced by the failures: verify the entire path, not the interface that launched the action.

Contain the blast radius

The last two failures are avoided by architecture and operating discipline, not clever copy.

Failure 04

Burning the brand domain

Cold outreach shares a reputation boundary with the website, founder inbox, customer mail, or newsletter. One bad list can damage all of them.

REQUIRED **System rule:** cold email always lives on a separate, neutral domain.

Failure 05

Scaling through a dip

A volume schedule creates false confidence. Continuing through a reputation decline turns a warning into a larger problem.

REQUIRED **System rule:** pause first, diagnose, correct, and resume only after the problem is resolved.

Observed signal	Agent response
All guardrails healthy	Hold or increase within the approved band.
One signal deteriorates	Hold volume and investigate.
Any bounce, complaint, or reputation dip	Pause and report before resuming.

The system should fail small. Separate domain, human approval, downstream verification, hard thresholds, and automatic pause behavior.

Required versus Your call

The strict setup is a safe default, not a claim that every preference protects deliverability.

Control	Class	Why
Separate sending domain	REQUIRED	Contains reputation damage.
All four DNS records green	REQUIRED	Authenticates the sender before any mail leaves.
Three-to-four-week gradual ramp	REQUIRED	Builds reputation without a sudden volume spike.
Bounces under 2%; complaints under 0.1%	REQUIRED	Defines the safety ceiling.
Postmaster monitoring and automatic pause	REQUIRED	Stops a warning from becoming a reputation event.
Downstream delivery verification	REQUIRED	Sent folder presence can lie.
Per-send approval and human-written copy	YOUR CALL	Use strict approval or trust the agent to send.
Nightly reports, 9am start, hourly spacing	YOUR CALL	Choose a cadence that fits inside the ramp bands.

Optional does not mean risk-free. Auto-send can be enabled, but the owner remains responsible for recipient judgment, copy, and account risk. Required controls stay in force in every mode.

Connect narrowly, then step back

Automation is a dial. Connect only what the current mission needs. Use the manual path whenever access is unavailable or undesired.

Connection	What it needs	What it unlocks	Afterward
Registrar	Registration access.	Register the approved sending domain.	Revoke unless managing renewals.
DNS	Token scoped to one zone.	Add authentication records; tighten DMARC later.	Revoke after setup; reconnect for cutover.
Google Workspace	Minimum Admin domain and alias scope.	Add the secondary domain, alias, and alternate email.	Revoke after identity testing.
Gmail	Send scope for the warmed identity.	Schedule and send from the exact identity.	Keep only while auto-send is enabled.

Connected path

The agent performs the action and verifies it. Setup credentials are temporary and narrow.

Manual path

The owner performs the shortest necessary action. The agent waits, then verifies the observable result.

The credential lifecycle

1

Scope. Grant only what this mission needs.

2

Act. Complete the named action only.

3

Verify. Demand observable evidence.

4

Revoke. Remove temporary access.

Security rule: never grant a broad credential when a zone-scoped or task-scoped token exists.