

AUTONOMOUS AGENTS · A PRACTITIONER'S BUILD

I built an AI that runs my inbox while I sleep.

बालमुकुंद

A personal AI executive assistant — and a blueprint for agentic automation that runs unattended.

01 — THE PROBLEM

Email is where commitments are born — and where they quietly die.

The cost is rarely a dramatic miss. It's the daily tax of re-reading, re-deciding, and hand-copying a date from the third paragraph of a thread into a calendar.

THE REFRAME

“Process it without me — and only interrupt me when judgement is genuinely required.”

Everything in the design follows from that one sentence. Not a chatbot you summon. A scheduled, unattended agent — closer to plumbing than to a conversation.

02 — THREE PRINCIPLES

1

Autonomy over interactivity

Runs on a schedule, unattended. Humans are pulled in only for genuine ambiguity.

2

The model is a component, not the product

The LLM is one swappable stage — Gemini 2.5 Flash today, anything tomorrow.

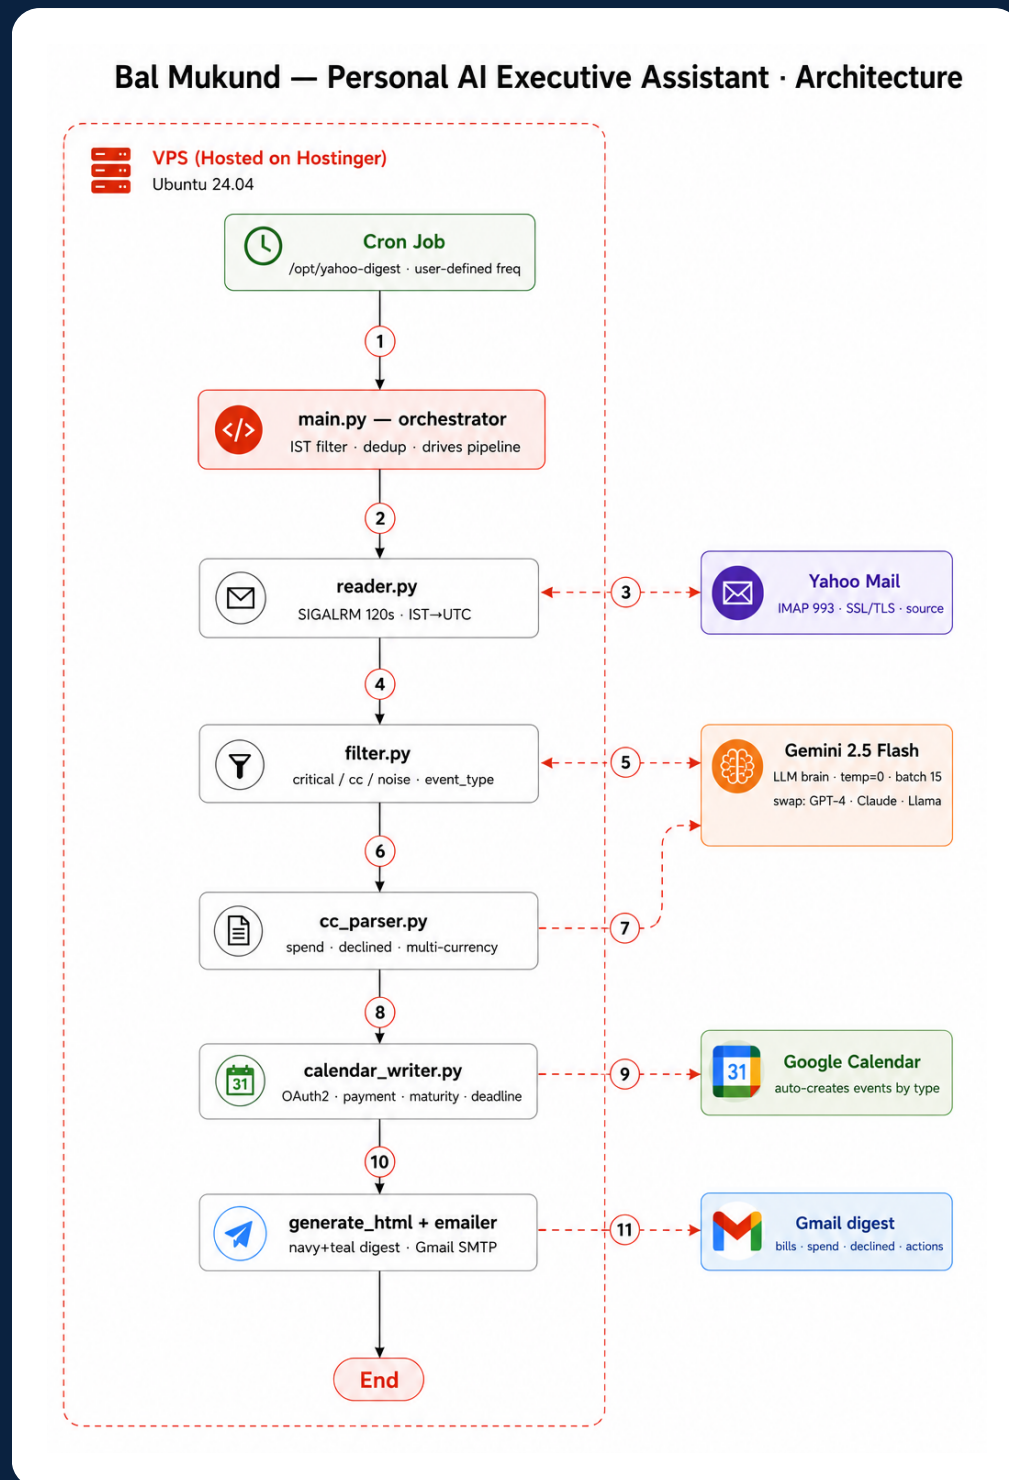
3

Idempotency is a feature

Re-runs never duplicate. State lives in local JSON files, every item fingerprinted.

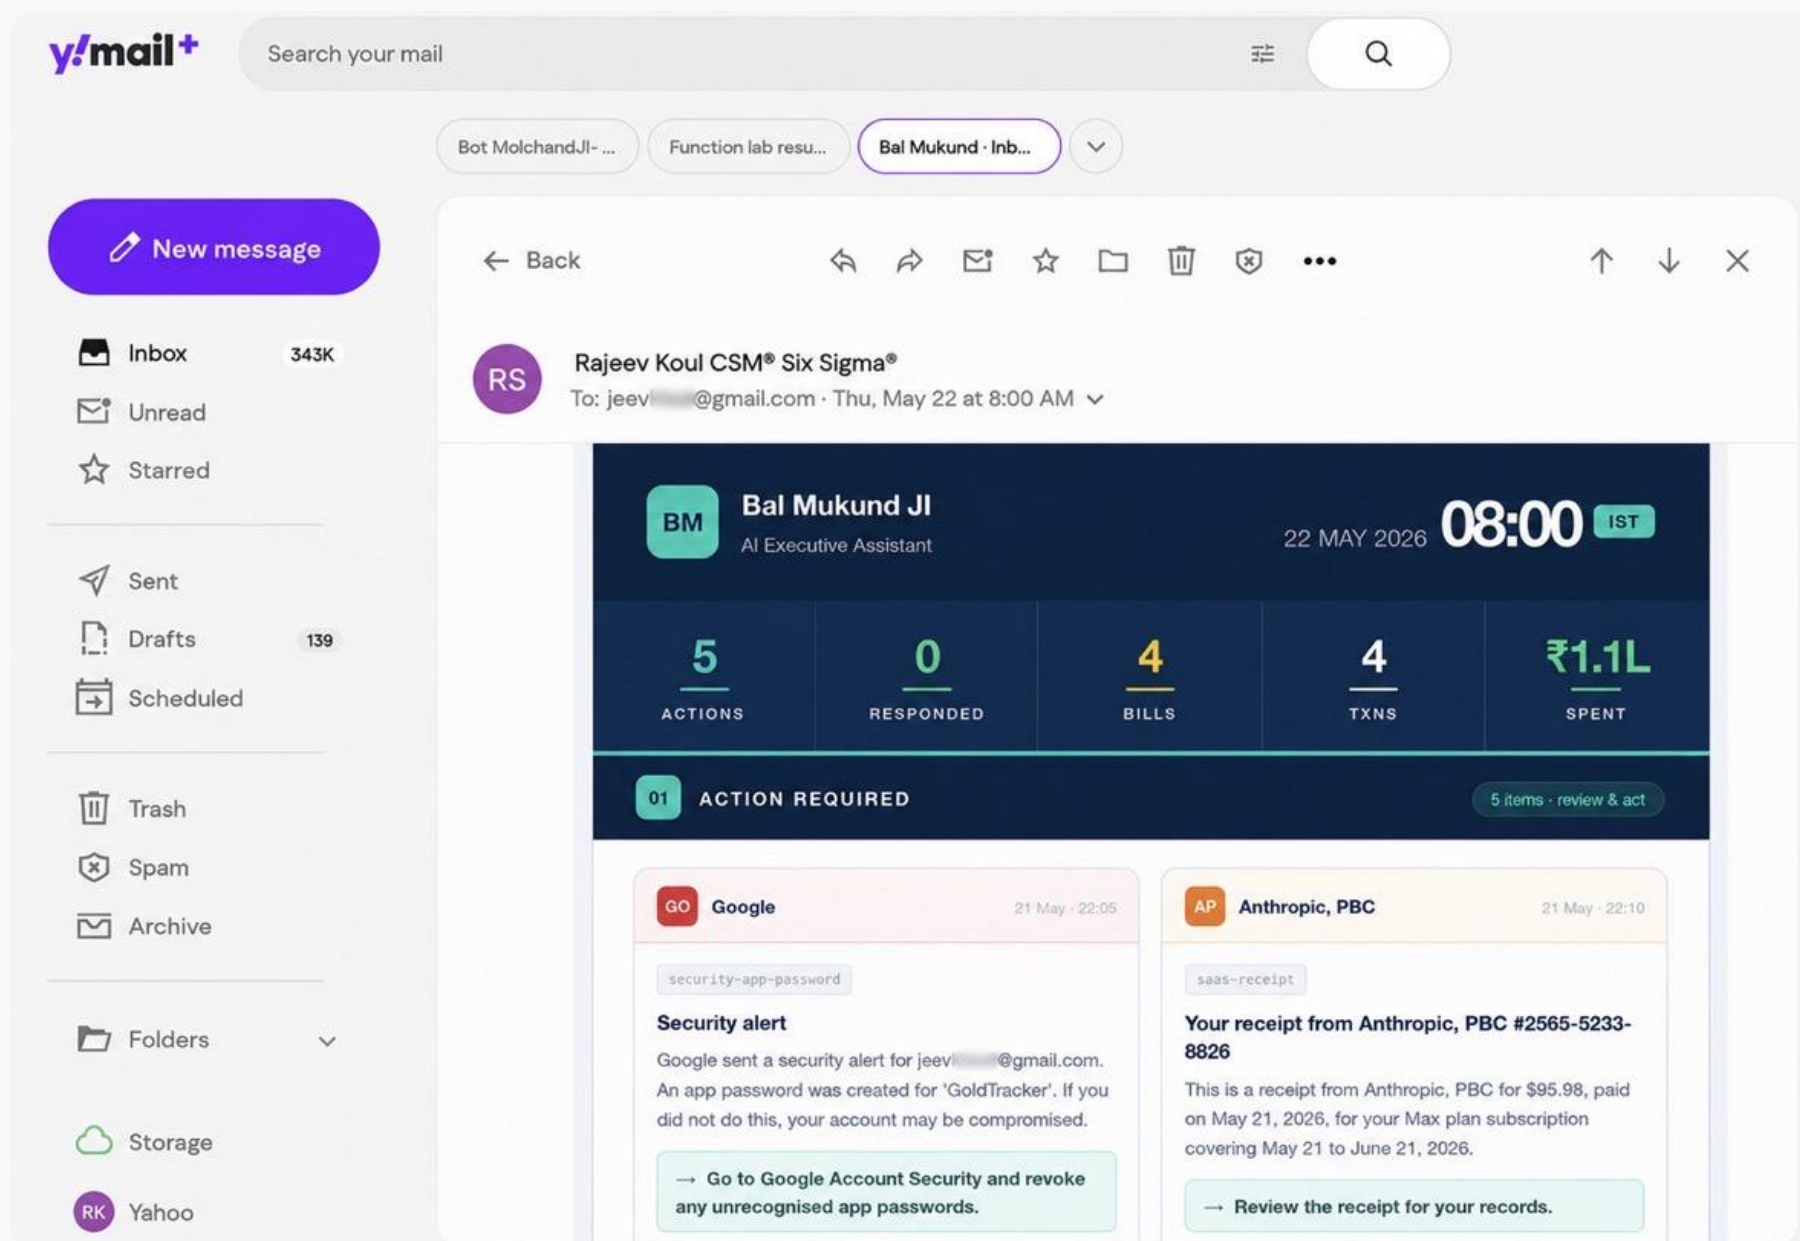
03 — THE ARCHITECTURE

One VPS. Seven modules. One swappable model.



THE PRODUCT


The only thing I see each morning.



One email by 8 PM IST — actions to take, bills due, spend, and declined transactions.

PERSONAL → ENTERPRISE

The same pipeline scales up.



Personal

One executive's inbox & calendar on a low-cost VPS.



Small Team

Many inboxes unified behind one shared digest.







Enterprise

Private model, RBAC, auditable calendar writes.

THE THROUGH-LINE

Push intelligence to the edges. Keep the spine deterministic.

-  **Keep the model swappable**
-  **Make the boring parts deterministic**
-  **Design for unattended failure**
-  **The digest is the product**

WHAT RELIABILITY LOOKS LIKE

1x/day

UNATTENDED, ZERO
CLICKS

~\$/mo

TOTAL RUN COST

Safe

RE-RUNS NEVER
DUPLICATE

Reliability comes from boring guarantees — not clever ones. When a run fails, it stops cleanly and the next day reconciles state.

ABOUT THE AUTHOR

Rajeev Koul

AI Transformation Advisor

Rajeev Koul is an AI Transformation Advisor and program leader with 20+ years of experience delivering complex business and technology transformations across organizations including Samsung, Barclays, and Infosys.

A PMP, Certified ScrumMaster, and Six Sigma professional, he combines enterprise-scale delivery expertise with applied AI to build intelligent systems, autonomous agents, and workflow automations that drive measurable business outcomes.

His work also extends into evidence-based health and performance research — spanning cardiology, sports medicine, athlete performance, nutrition, metabolic health, and longevity — integrating clinical evidence, wearable analytics, and AI-driven insights into practical solutions.

Balmukund is one of several production-grade AI systems developed by Rajeev, combining advanced AI, structured knowledge, and automation to transform information into actionable intelligence.

Follow for more · Connect on LinkedIn