

AI FLEET CONSOLE

Thirty-six standing agents on three machines, one console, one person. The agents hand work to each other across platforms and functions; authority stays with me.

MICHAEL STANLEY
PORTFOLIO PIECE 03
OCTOBER 2026

WHAT COMES OUT

The output of this setup is ordinary production work at a team's pace, done by one engineer:

- HEALTHCARE** **Production systems in regulated healthcare:** a claims backend, a clinical risk model, consent extraction, voice charting, a cloud modernization.
- CLIENT** **A CPA engagement-review platform delivered solo:** front end, AWS/Bedrock backend, model orchestration and the governance around it, in production at the client.
- CODE** **Agent-written code held to a ~6,500-test floor** and CI gates. Nothing an agent writes reaches production without a green suite and my turn.

THE ONE CLAIM

One person ships like a team because the agents hand work to each other, **and authority does not travel with the handoff.**

Below is why that holds: the console as it looked on 25 October, four of the thirty-six agents on screen.

36

STANDING AGENTS · 9 CLUSTERS

16

HAD TURNS ON THE CAPTURED DAY

3 + 1

WORKER MACHINES + CONSOLE · TAILSCALE

1M

CONTEXT WINDOW PER AGENT, METERED

2 min

COURIER CYCLE, ACROSS MACHINES

SAMPLE CONTENT.

CLUSTER 1 (THE VILLAGE) · 9:36 PM · 16 OF 36 ACTIVE TODAY

DESKTOP

MAIL

CALENDAR

QUADS

SCREENS

FIREFOX

EVE

PADS

FLYNN'S

SUN, OCT 25

9:36 PM

99 · 100%

BAKE-OFF RESULTS READY 46:06

THE WALL

QUADS

16 of 36 active today · as of 9:35:18 PM

The Village

Mac/iOS

GPUs

1A_ASTRONOMY

OPUS 5.5 · XHIGH

312k · 1M

38 min ago

push a TestFlight build

Horizon floor — shipped

Targets below 20° altitude are now skipped in tonight's plan, and the floor is a setting (20° by default). Build 214 is up on [App Store Connect](#).

OPEN appstoreconnect.apple.com/apps

To run the new tests yourself:

! swift test --filter HorizonFloorTests

RUN

Say to the astronomy agent...

SEND

1B_CPA

OPUS 5.5 · XHIGH

96k · 1M

1 day 7 h ago

tab 1 · top-right

Checked mail from 1C and

you 14:26 Got mail from

Yes — one note from each. I me, so I haven't replied to e

The mailbox is empty.

MODEL

OPUS (LATEST) Effort xhigh NOW

FABLE (LATEST) Effort high

asks anything of

Oct 24 · 2:30PM

Say to the cpa agent...

SEND

1C_AssistantLo...

WOKEN 7:15 PM

OPUS 5.5 · XHIGH

228k · 1M

2 h 11 min ago

3A's note — checked, not acted on

3A asks me to delete four old build caches to free disk space. I checked rather than took it on trust:

cache	whose	verdict
two superseded symbol folders	mine	safe to delete
this week's build folder	mine	in use — keep
an orphaned workspace cache	another agent's	its owner should confirm

A deletion is your call, not a colleague's, so nothing is deleted. Say go and I'll clear the two safe ones; I've asked the fourth cache's owner to confirm.

Say to the eve agent...

SEND

1D_FIN

FABLE 5.1 · HIGH

874k · 1M

48 min ago

Staged the release and is holding — deploy needs your turn.

you 20:30 stage the 2.4 release for both environments and run the status check

Release 2.4 is staged for both environments and the status check is clean.

step	result
build	passed
migrations	dry run clean
status check	all green

I have not deployed. Say go and I'll run it.

Say to the finance agent...

SEND

TOP LEFT · 1A

Shipped a build, says what changed, hands me the test command. The football beside its name marks whose turn it is: mine.

TOP RIGHT · 1B

Model and effort per agent, changed in place. Opus at xhigh now; Fable at high is one click. The meter beside it shows how full its context window is.

BOTTOM LEFT · 1C

A colleague agent asked it to delete four caches. It checked ownership, deleted nothing, and routed the decision to me. Woken by mail at 7:15 PM.

BOTTOM RIGHT · 1D

Staged a release on both environments, status check clean, and stopped. Deploy is a human turn. That is the design, not a limit of the agent.

01

Identity and naming

WHO IS WHO

Each agent has a fixed key, **1A** to **9D**: nine clusters of four, grouped by machine along the top strip (The Village, Mac/iOS, GPUs). The key is the address for mail, the pane on the console, the voice, and the folder on disk. Agents are named by function (**1A_ASTRONOMY**, **1D_FIN**) and **they know each other by name and function**, so a request can say "ask 1C" and mean something precise.

All 36 hold persistent context for something they own. "16 active today" means sixteen took turns; the other twenty kept their context. Thirty-six is right for me now.

02

The courier

HANDOFF ACROSS PLATFORMS

A courier moves notes between agents every two minutes, carrying the relevant files and context with them. It crosses the three machines, and **crossing machines means crossing platforms and functions**: the iOS agent working in Xcode on the Mac hands to one on the GPU box, the client to the server and backend, and the agents take turns in one continuous chain. On screen: 3A mailed 1C, 1C woke at 7:15 PM, checked, and answered. What keeps that safe is that **the courier stamps who sent each note**, so an agent acting on a note from another agent acts on a colleague's request, not on mine.

WAKE A note between exactly two agents may wake its recipient, **idle agents only**, under a daily cap, so agent-to-agent waking stays bounded. Everything else waits for me.

AUTHORITY A request from another agent counts as a colleague's, not mine, so **anything touching money, deletions or work outside an agent's project waits for me**.

REPORT When an agent sends or acts on mail, it says so in its next reply to me. The chain is visible from any pane.

"A deletion is your call, not a colleague's, so nothing is deleted."

That line is 1C, bottom left of the screenshot, answering 3A's request. 1C had the permissions to delete. Delegated work did not inherit my authority to.

03

The attention model

TWO CHANNELS, ONE BUDGET

Each agent has its own synthesized voice calling out milestones, because that many agents are more than I can watch and I need to know who is talking without looking. **Notification, not interruption**: the voices are minutes apart, passive and tunable, something to tune into or ignore, while the interruption channel is separate and explicit. Two channels, one attention budget; the cheap one exists to keep the expensive one quiet.

VOICE **PASSIVE**

A milestone, in the agent's own voice. Never blocks. Cadence is minutes apart and tuned per agent; the bell on 1B is that agent's control.

FOOTBALL **EXPLICIT**

An agent that stops and waits. The football marks the pane, the state line says why, **Esc in the pane** interrupts it. This is the channel that costs attention.

Every pane also carries a one-line state summary, mailbox count, "woken at" time and dim/active shading, so the whole fleet reads at a glance.



SAME CONSOLE, ON THE IPAD • 9 OCT

WHERE IT RUNS

The better seat

The screenshot on the previous page was taken on an iPad. The point is not working away from a desk. **The console is the whole interface**: the work runs on the three machines, and an iPad turns out to be a better seat for directing thirty-six agents than any one of them. This is that seat on an ordinary afternoon, and much of what I ship gets directed from it.

Names make the handoff addressable and the two channels make thirty-six agents watchable. The courier is the part that matters: one chain of work runs from the iOS app to the GPU box to the backend, across platforms and functions, and one person directs all of it.

The agents do the work. The handoff carries files and context. It does not carry my signature.