

The 7 Cron Jobs That **Burned** My **Weekly Codex Budget**

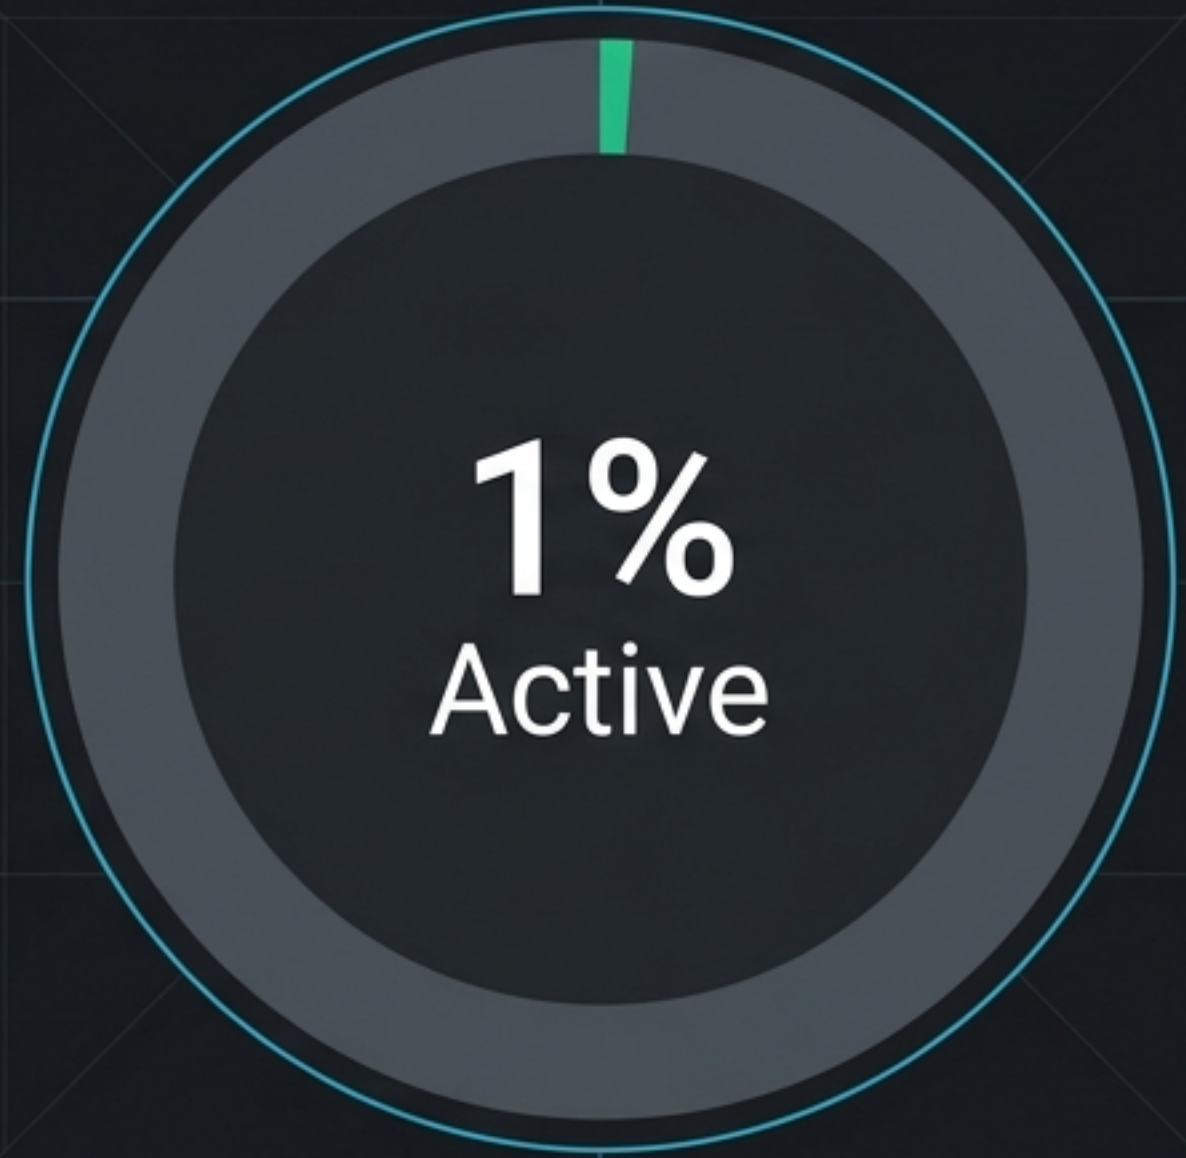
(While I Did Almost Nothing)

Chris Johnson, CryptoFlex LLC

Weekly Subscription Budget

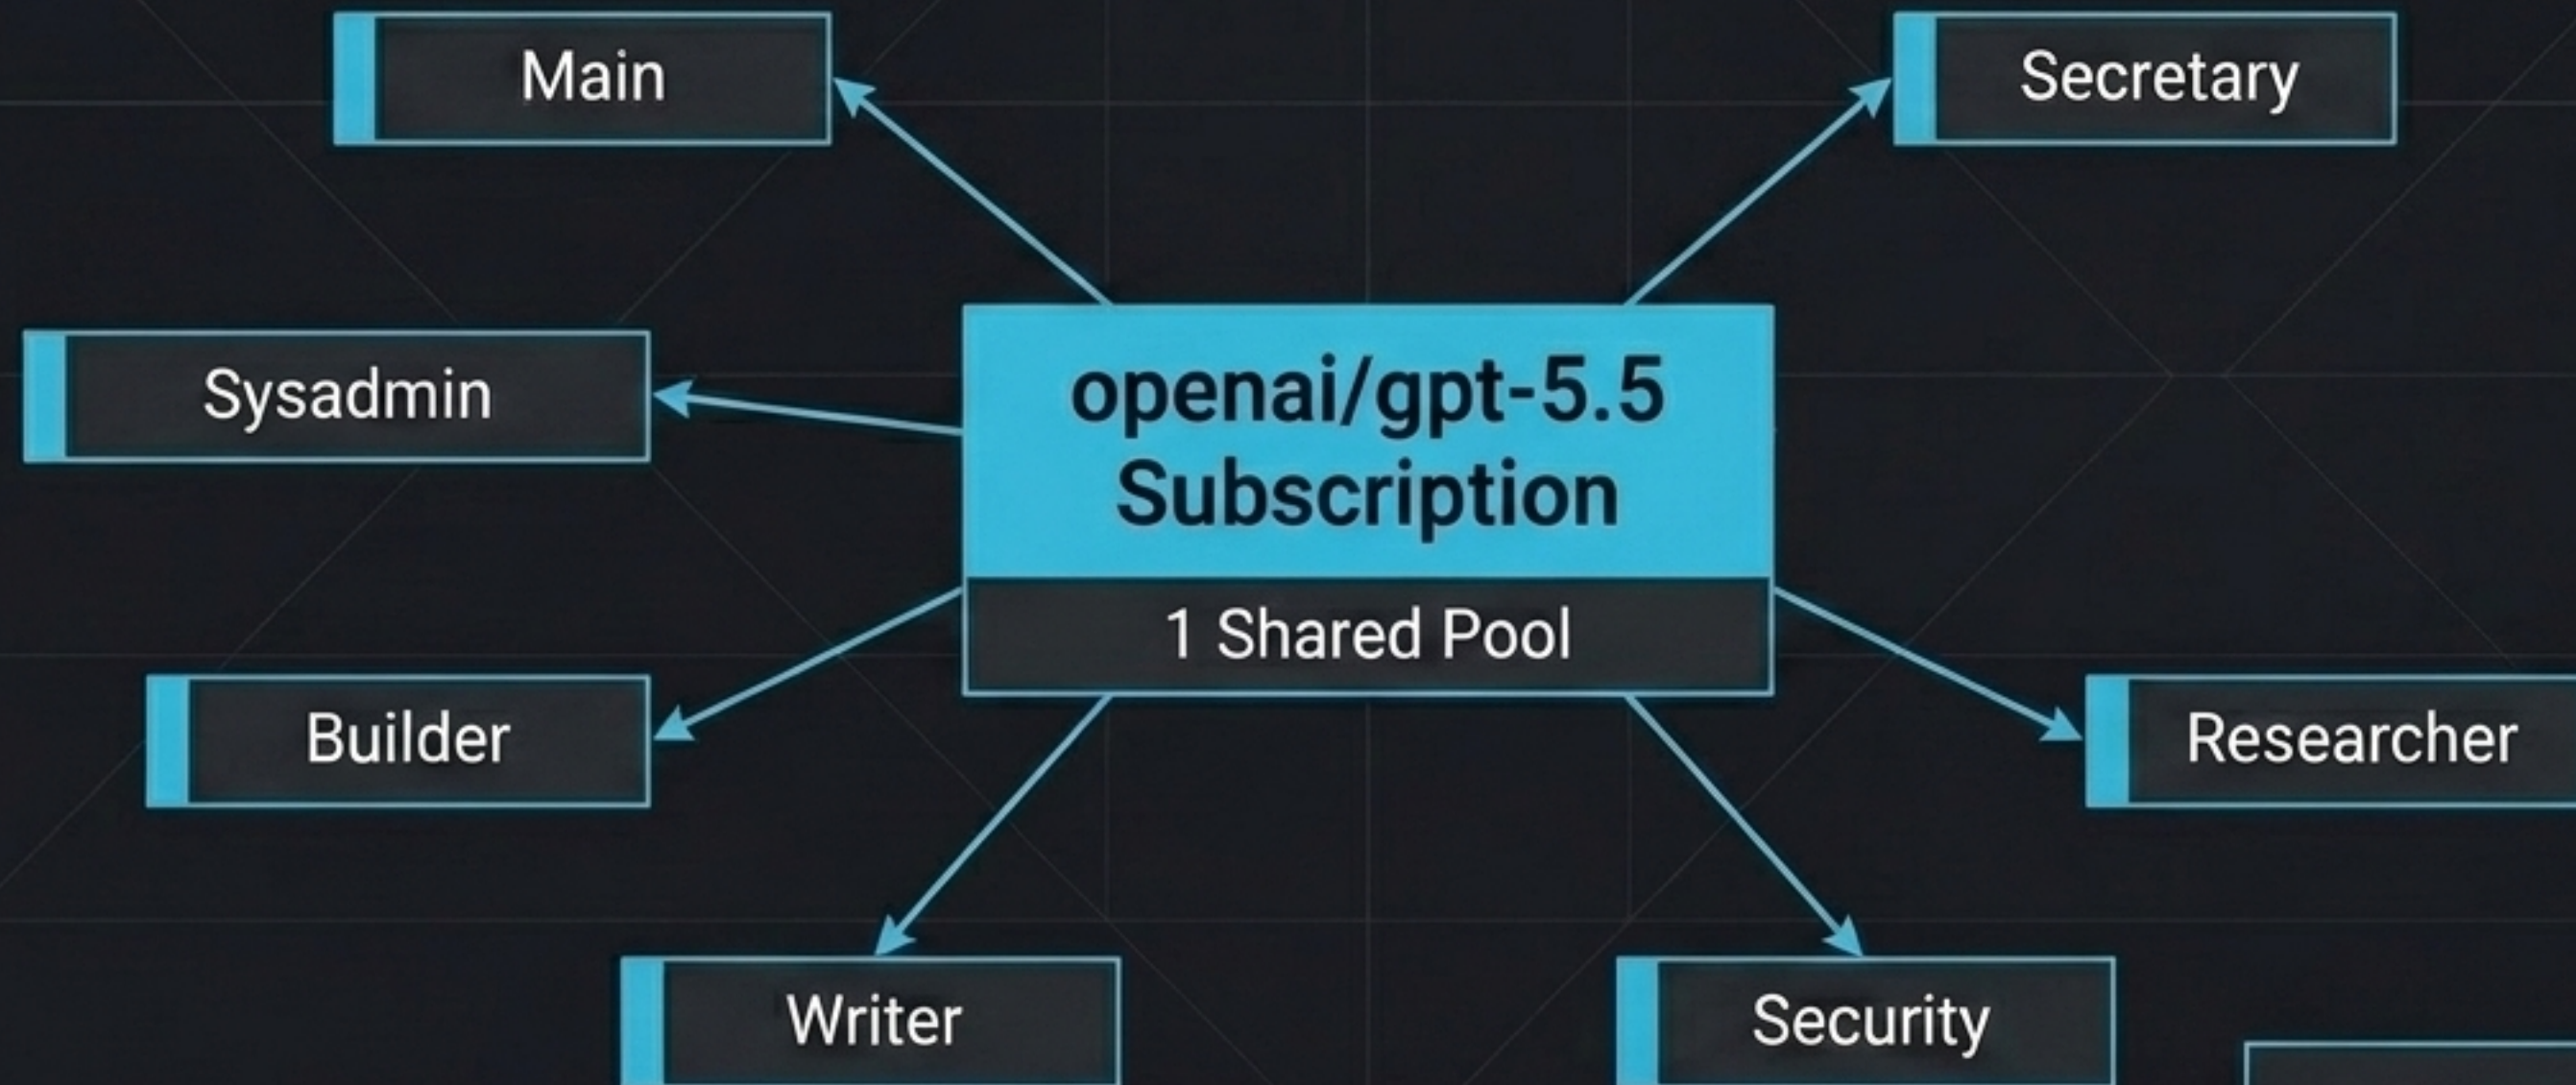


Rolling Five-Hour Window



The drain occurred steadily in the background over preceding days, not during current interactive sessions. The weekly cap hit 100% before the interactive workload even began.

One pool, shared by everything.



- 0 Cheap Tiers
- 0 Per-Token Fallbacks

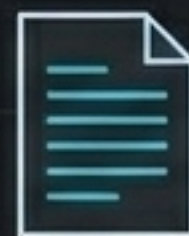
Input Payload

7 Agent Transcripts
+ Full Memory Context

400,000 to
600,000
tokens per day

Output Value

1 Summary Note
(Read approximately twice a week)

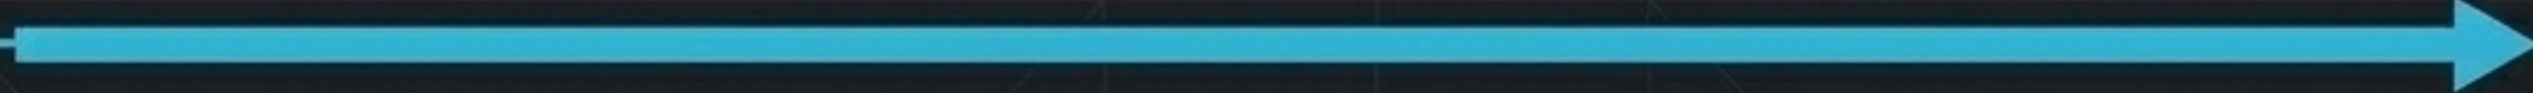


One daily roll-call job was burning more budget than the entire interactive workload combined.

Run standard shell scripts



Trigger

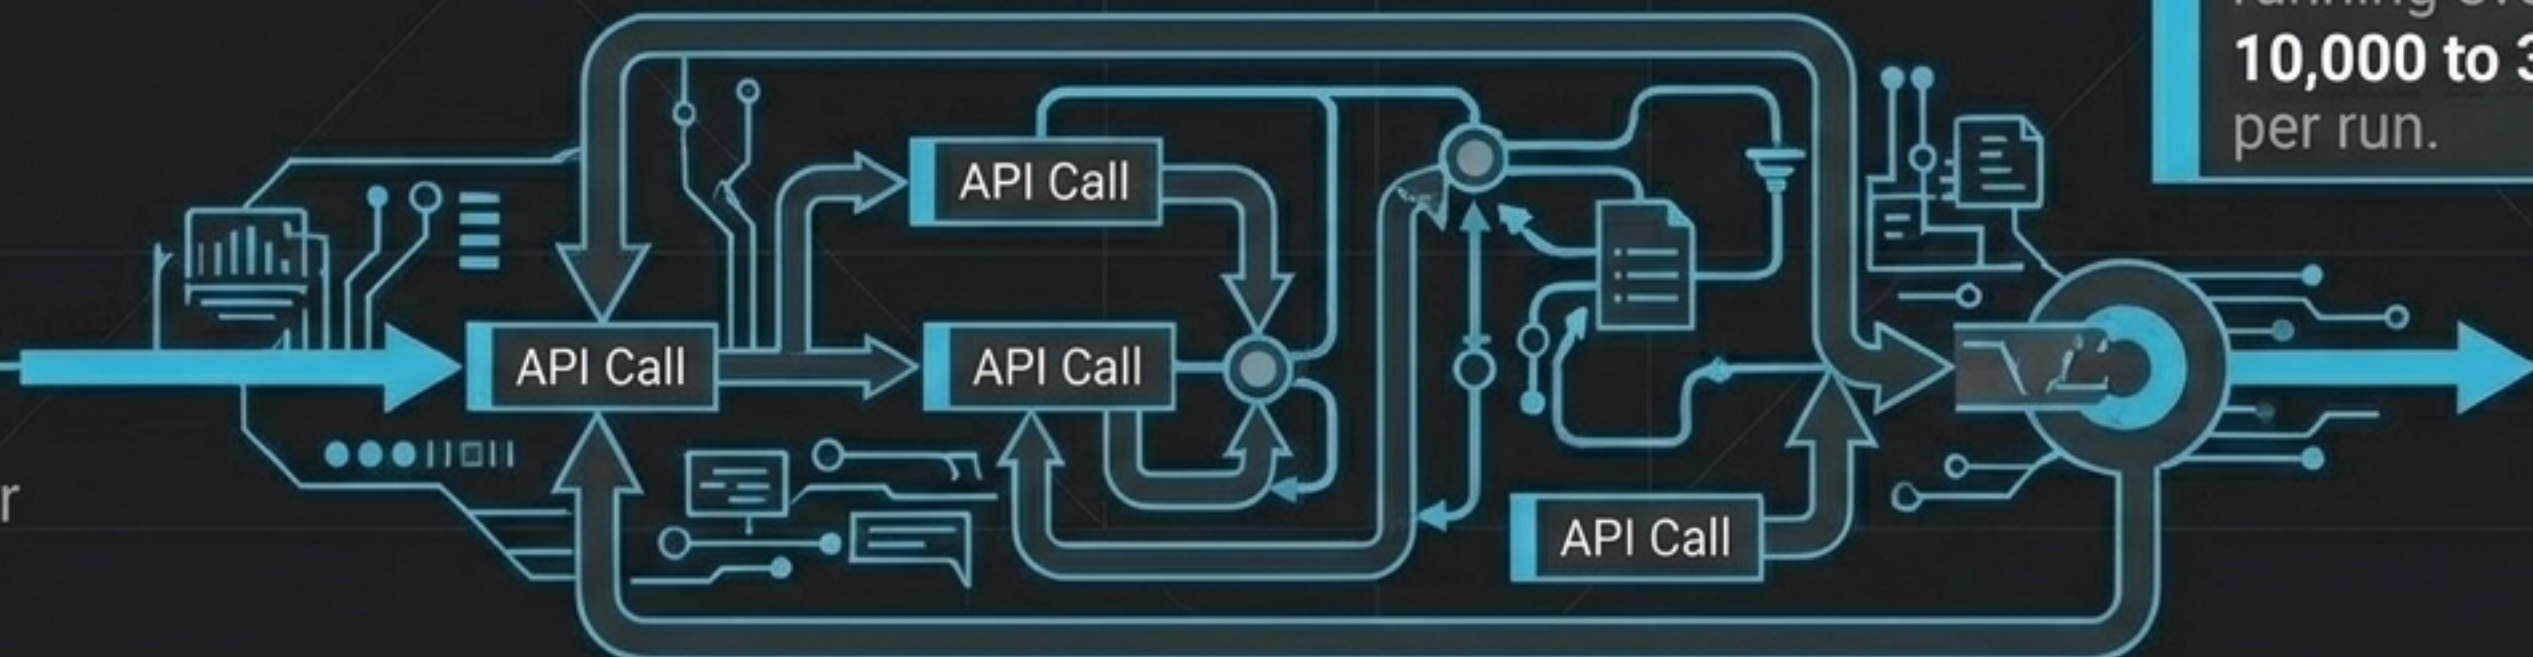


“Launching rockets to check the weather.”

agentTurn Payload



Trigger



3 Healthcheck Jobs running every 8 hours.
10,000 to 30,000 tokens per run.



Rate-limit failures triggered the exact anomalies the system was designed to detect.

ANOMALY TIMELINE AND CASCADING FAILURE GRAPH









CASCADING FAILURE CHAIN



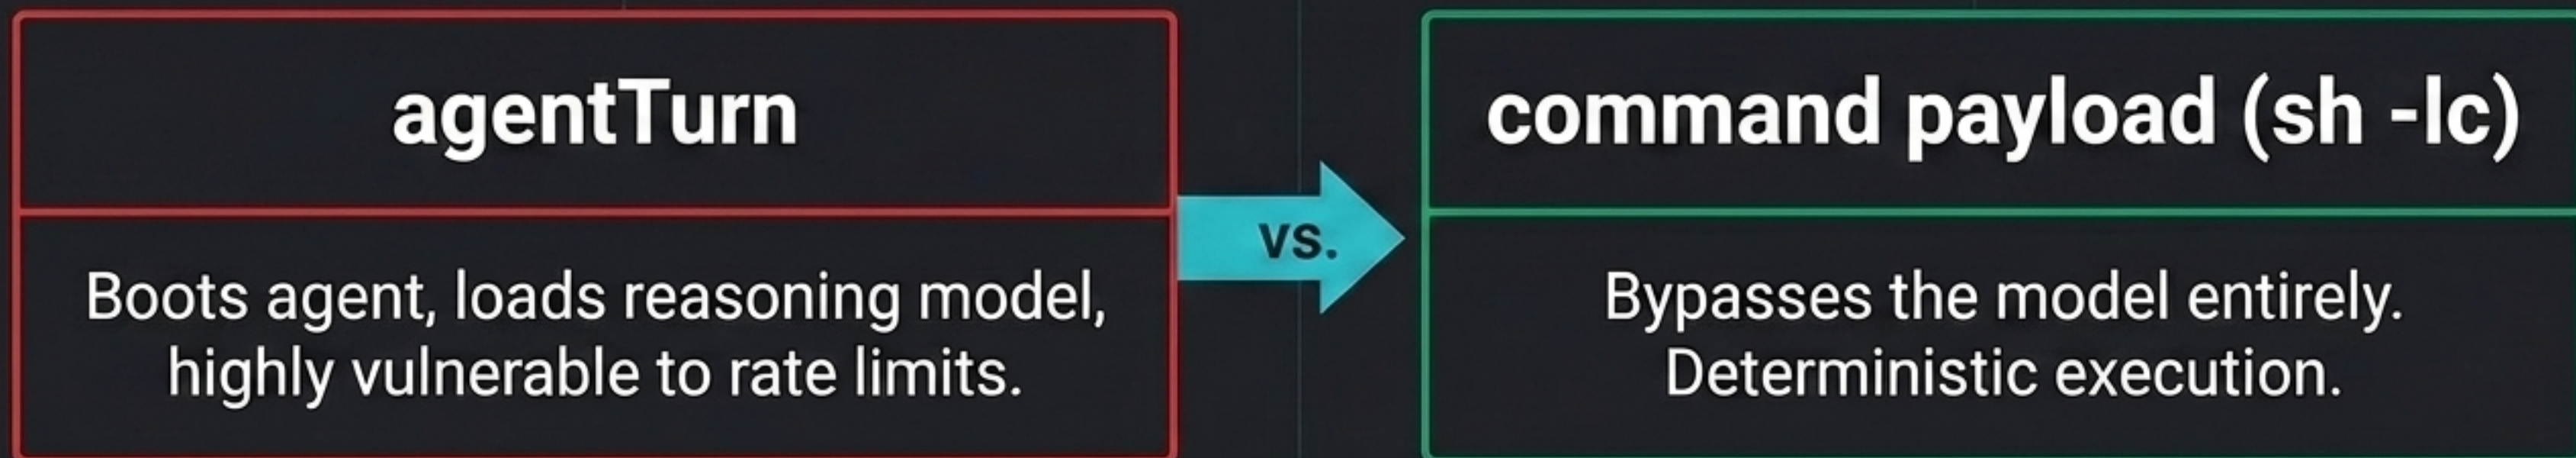
CONTEXT NOTE: Synced configuration files preserved outdated modification timestamps. The system was running stale logic while masking the available fix.

TELEMETRY CONSOLE EXECUTIVE MONITORING DASHBOARDS

Before	After
 Total Scheduled Jobs: 7	 Total Scheduled Jobs: 3
 Agent Runs per day: 11	 Agent Runs per day: 2 (1 weekly synthesis, 1 daily backup)
 Primary Drain: Roll-call and heavy Healthchecks	 Addition: 1 zero-token digest

11 daily agent runs reduced to 2.

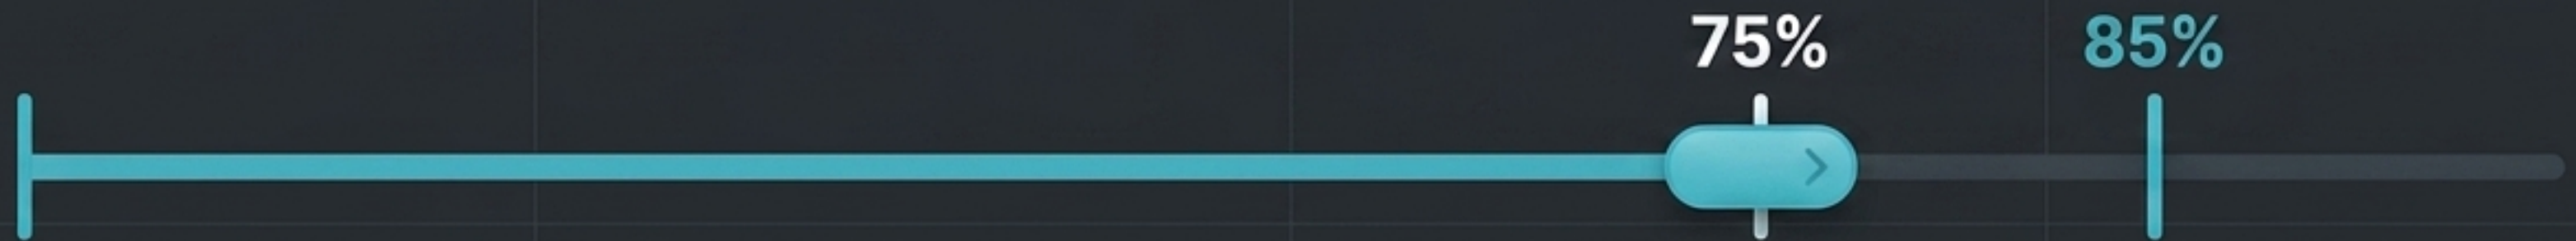
COMMAND PAYLOAD ARCHITECTURE



Local Python script assembling CLI outputs into a Telegram message.

0 tokens consumed. 2 seconds total execution time.

LCM contextThreshold raised from 0.75 to 0.85.



Reduces the frequency of model-driven compaction calls for lightly used agents.

CORE INSIGHT

Memory compaction runs on a model too. The default threshold is a configuration choice, not a static rule. Audit it before assuming it fits the workload.



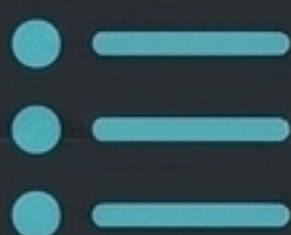
Audit the Scheduler.

Background jobs create a 100% budget floor before any interactive use occurs.



Match Tool to Task.

Do not pay a reasoning model to execute deterministic shell commands.



Prefer Deterministic Digests.

Structured CLI outputs cost nothing compared to full-transcript ingestions.



Watch the Weekly Window.

A maxed weekly cap paired with low hourly usage is a direct pointer to background drain.