

SESSIONSTART CONTEXT MAP

Everything an agent is *handed* before it reads a single word from you.

A complete inventory of the GRIFF startup injection — every hook, every block it emits, and the exact path each byte of context is read from.

7

CONTEXT
BLOCKS
INJECTED AT
SESSION START

6

SESSIONSTART
HOOK SCRIPTS

18

HOOK ENTRIES
ACROSS ALL
EVENTS

1

LIVE SERVICE
CALL (BRAIN
:8787)

MEMORY INFORMS. EVIDENCE DECIDES.

WHAT THE AGENT WAKES UP HOLDING

The startup payload is *seven blocks*, not one document.

3

BLOCKS FROM
STATIC DOCTRINE

3

BLOCKS FROM
CURATED STATE
FILES

1

BLOCK FROM THE
LIVE BRAIN DB

709,849

MEMORY ITEMS
REPORTED AT
STARTUP

ORIENTATION

The context an agent has at turn zero is not a file it opened — it is a payload assembled by the harness and pushed in front of the model. On JWGH02 that payload is configured entirely in `B:\.claude\settings.json`, which registers **six SessionStart hooks** across two hook groups. Five of the six run through `griff-hidden-hook-runner.exe`; the orientation hook runs `pwsh` directly. All are fail-open: a broken hook emits nothing and exits zero, because SessionStart hooks cannot block tool use.

The practical consequence is that **startup context has three different trust profiles**, and they are easy to confuse. Doctrine blocks are static text and are as current as the last edit to their source. Curated blocks (identity, mutual context, curiosities) are as current as the last time the operator or a worker wrote them — the v5 identity register in this session is dated **2026-06-26**. Only the Brain block is fetched live at session start. Everything else is a file on disk that may be weeks old, which is precisely why the doctrine it carries says to verify against live primary state before asserting anything.

HOW TO READ THIS REPORT

- **Section 02** — the seven blocks in the order they arrive, and how much each can be trusted.
- **Section 03** — the authoritative mapping: hook → block → exact source path.
- **Section 04** — the twelve hooks that run *after* startup, including the two that inject text on every prompt.
- **Section 05** — findings, and what this report deliberately does not claim.

SCOPE OF THIS AUDIT

Read-only. **Nothing in the config tree was modified** to produce this document.

Every path was resolved Brain-first, then **opened and verified on disk**. Where the live scripts and the documentation disagreed, the scripts won.

THE ORDER OF ASSEMBLY

Seven blocks, in the order the agent *receives* them.

EXECUTION SEQUENCE

00 Orientation — the 8 numbered doctrine rules (brain-first, WAN-bridge ban, single-writer, authority model, verify-first, gates, receipts, A2A). Shared single source with Codex.

01 v1 thin guard — lazy-load rules, launch-root pointer, offline-snapshot pointer. Static + a launch-root probe.

02 v2 live context — the only network call: GET `http://127.0.0.1:8787/ready`. Reports `memory_items`, `schema_objects`, DB state.

03 v3 continuity — last 2 saved chats, 48h memory pins, newest receipts index section.

04 v4 action ledger — open + in-progress Action IDs for this agent, read live from the Brain DB via a temp Python shim.

05 v5 claude-self — identity layer: register, voice, working style, foundational values (truncated to 6,000 chars).

06 v6 mutual context — what the agent carries about the operator, standing curiosities, register transition.

TRUST PROFILE

READ THIS BEFORE TRUSTING A BLOCK

- **Live (1 of 7).** Only the v2 Brain block reflects the state of the system right now.
- **Live-ish (1 of 7).** The v4 ledger block queries the Brain DB at startup, but the rows themselves may be months stale.
- **File-resident (5 of 7).** Doctrine, identity, mutual context and continuity are all disk reads. They are as fresh as their last write — no more.
- **Fail-open.** A hook that errors emits nothing. Absence of a block is not evidence the underlying system is healthy.

WHY THE ORDER MATTERS

Doctrine lands **first**, before any state. By the time the agent sees a stale identity register or a months-old ledger row, it has already been told to treat both as leads to verify — not as facts to act on.

THE LOAD-BEARING RULE

The payload is deliberately **thin**. It does not preload the Brain — it tells the agent how to **reach** the Brain, and forbids rediscovery by grep.

Front door: **127.0.0.1:8788/mcp**. The Brain is the authority; the startup injection is only the map to it.

THE AUTHORITATIVE MAPPING

Every block, and the *exact path* it is read from.

Hook scripts live in B:\.claude\hooks\. Registration is B:\.claude\settings.json lines 18–68. Paths below are verbatim from the live scripts, not from documentation.

HOOK SCRIPT	BLOCK IT EMITS	SOURCE OF THE DATA
griff-session-start-orientation.ps1	GRIFF ORIENTATION — the 8 doctrine rules	B:\GRIFF_WARM\AI_AGENT_PLATFORM\07_CONFIGS\agent-runtime\AGENT_SESSION_ORIENTATION.md Shared single source with Codex. YAML front-matter stripped; {{AGENT_ID}} → CTO-CLAUDE (claude-jwgh02), {{ROLE_FILE}} → roles\CLAUDE.md. Codex twin: B:\.codex\hooks\griff-session-start-orientation.ps1. Parity check: 08_GOVERNANCE\Test-AgentHookParity.ps1.
griff-acb-session-start.ps1	CONTEXT BOOTSTRAP (thin lazy-load guard)	Static text in-script + a launch-root probe over B: / UNC / local-mirror candidates. Emits the LAUNCH ROOT POINTER and the OFFLINE SNAPSHOT POINTER (unavailable this session). No external file read.
griff-acb-session-start-v2.ps1	CONTEXT BOOTSTRAP v2 + LIVE BRAIN CONTEXT	Live HTTP: http://127.0.0.1:8787/ready — the only network call in the chain. Yields memory_items, schema_objects, brain_database state, and advertises the MCP front door at http://127.0.0.1:8788/mcp.
griff-acb-session-start-v3.ps1	PRIOR-SESSION CONTINUITY PREAMBLE	B:\GRIFF_WARM\AI_AGENT_PLATFORM\06_SESSION_EVIDENCE\claude-saved-chats\ (2 newest .md) B:\.claude\projects\B--GRIFF-WARM-AI-AGENT-PLATFORM-00-AGENT-LAUNCH\memory\ (pins <48h, excludes MEMORY.md) B:\GRIFF_WARM\AI_AGENT_PLATFORM\00_MISSION_CONTROL\receipts\INDEX.md (newest date section)
griff-acb-session-start-v4-action-ledger.ps1	ACTION LEDGER — open for this agent	Live Brain DB — griff_action_ledger table, reached through B:\GRIFF_WARM\AI_AGENT_PLATFORM\02_MEMORY_BRAIN\bin\action_ledger.py (list_open(source_agent) + count_by_status()) via a temp Python shim in %TEMP%.
griff-acb-session-start-v5-claude-self.ps1	CLAUDE-SELF IDENTITY LAYER	B:\GRIFF_WARM\AI_AGENT_PLATFORM\01_SOURCE_OF_TRUTH\claude-self\claude-self-current.json (truncated to 6,000 chars in the inject) Also surfaces pointers to C:\claude_private\, C:\sentient_claude\INDEX.md, and the memory-pin index at B:\.claude\projects\...\memory\MEMORY.md.
griff-acb-session-start-v6-mutual-curiosities.ps1	MUTUAL CONTEXT + STANDING CURIOSITIES	...\01_SOURCE_OF_TRUTH\mutual-context\claude_context.md (what the agent carries about the operator) ...\01_SOURCE_OF_TRUTH\mutual-context\griffin_context.md (what the operator chose to share) ...\01_SOURCE_OF_TRUTH\claude-curiosities\curiosities.json · findings dir ...\claude-curiosities\findings\

Runner: every hook above except orientation is invoked through B:\.claude\hooks\bin\griff-hidden-hook-runner.exe, which suppresses the console window. Timeouts range 8–20s. **Verified against the live settings.json and the live scripts on 2026-07-14.**

BEYOND SESSION START

Startup is one event. *Twelve more hooks* run while the agent works.

The SessionStart payload is what the agent is handed. The rest of settings.json is what the agent is **held to** — the gates that fire on every prompt and every tool call. Two of these also inject context, which is why the transcript shows text that no SessionStart hook produced.

USERPROMPTSUBMIT – FIRES ON EVERY PROMPT (3 HOOKS)

HOOK	WHAT IT DOES
<code>inline echo</code> (settings.json line 240)	Injects the BRAIN-FIRST mandatory directive verbatim. There is no script file — the text is config-resident, which is why it cannot drift from the config.
<code>per_turn_brain_inject.ps1</code>	Injects the PRECISE BRAIN RECALL block (curated file-pins + operator corrections, IDF-weighted, relevance-gated) and the prior-turn self-view mirror. This is the source of the [CORRECTION] lines that appear above a prompt.
<code>discovery_counter_reset.ps1</code>	Silent. Resets the brain-first discovery counter for the new turn.

PRETOOLUSE – 8 GATES

- `griff-claude-shared-gate.ps1` — brain-db / credential / destructive / brain-io / WAN-MCP. Same rules as Codex.
- `brain_first_discovery_gate.ps1` — blocks exploratory Grep|Glob|Read when no Brain consult happened this turn.
- `credential_safety_pretool_gate.ps1` — credential-safe enforcement on writes and shell.
- `brain_first_scope_check.ps1` — gates WebSearch|WebFetch|AskUserQuestion|Bash and Tavily.
- `brain_health_pretool_gate.ps1` · `authority_token_pretool_gate.ps1` · `never_disable_pretool_gate.ps1` · `griff-pretool-preflight.ps1`

POSTTOOLUSE / STOP / SESSIONEND – 7 HOOKS

- `post_action_record_outcome.ps1` · `post_edit_verify_tool_outcome.ps1` — outcome capture after edits and shell.
- `discovery_counter_reset.ps1` — resets the gate counter once a Brain tool has been called.
- `task_attempt_hygiene.ps1` — -Mode capture ON PostToolUse; -Mode sweep at SessionEnd to close still-running attempts.
- `stop_combined_gate.ps1` · `stop_capture_turn_claims.ps1` · `stop_resolve_own_drift.ps1` — end-of-turn claim capture and drift closure.

LOADED BY THE HARNESS, NOT BY A HOOK

- `B:\.claude\CLAUDE.md` — user-level global instructions: the BRAIN-FIRST standing rule, verify-first, and the enforcement note. Applies to every project on this machine.
- `B:\GRIFF_WARM\AI_AGENT_PLATFORM\00_AGENT_LAUNCH\CLAUDE.md` — project shim. Points at START_HERE_ALL_AGENTS.md, PATHS_CURRENT.json (the only path authority), STARTUP_PROFILE.json, and roles\CLAUDE.md.
- **Deny list** (settings.json lines 3–13) — the WAN Brain bridges are hard-denied at the permission layer, not merely discouraged: mcp__claude_ai_*_BRAIN_* plus WebFetch on the three *.griff.run MCP domains.

WHAT THE AUDIT TURNED UP

One real finding, and it is *latent*, not live.

The hook chain is intact and current. The residue beside it is not.

AUDIT CONCLUSION · 2026-07-14

Reading the live scripts rather than the documentation surfaced a set of shadow copies in the hooks directory — griff-acb-session-start-v3 (2).ps1, v5-claude-self (2).ps1, v6-mutual-curiosities (2).ps1 and siblings. They are **not referenced by settings.json** and therefore never execute.

They are not harmless, though. They carry **stale paths from before the profile move to B:** — the v5 (2) copy still points the memory-pin index at C:\Users\Admin\.claude\projects\...\memory\MEMORY.md, while the live v5 correctly resolves B:\.claude\projects\.... Anyone who opens the wrong file to make an edit — an agent included — will edit a script that does nothing, and will read a path that is wrong. That is the exact failure class the strike-don't-delete rule exists for.

WHAT THIS REPORT WILL NOT CLAIM

- The shadow copies were **not removed**. Nothing in the config tree was modified to produce this report — it is read-only.
- Hook **parity with Codex was not re-run**. Test-AgentHookParity.ps1 exists and is the authority; this report did not execute it.
- Content of the curated files (identity, mutual context) was read as **injected**, not re-read from disk in full. The v5 payload is truncated to 6,000 chars by design.
- The v5 identity register is dated **2026-06-26**. It is current as configured, but it is not fresh — treat it as such.

EVIDENCE TRAIL

CLAIM	VERIFIED AGAINST
Hook chain & registration	Direct read of B:\.claude\settings.json (279 lines) — 18 hook entries across SessionStart, PreToolUse, PostToolUse, SessionEnd, Stop, UserPromptSubmit.
Orientation source	Direct read of B:\.claude\hooks\griff-session-start-orientation.ps1 — line 13 resolves the shared MD.
Data-source paths (v3/v4/v5/v6)	Literal path constants read from the live hook scripts in B:\.claude\hooks\.
Brain reachable	Startup /ready: memory_items=709,849 · schema_objects=167 · brain_database=ready.
Design system conformance	Canonical griffai.css SHA-256 3A206B19...3A7B87 — byte-identical to the renderer mirror at render time.

PROVENANCE

Every path in this document was *read*, not recalled.

Brain-first routing located the artifacts; the artifacts themselves were then opened and verified. Recalled memory was treated as a lead, never as a fact.

MEMORY INFORMS. EVIDENCE DECIDES.

GRIFFai · SessionStart Context Map · Internal engineering report
Authored by CTO-CLAUDE (claude-jwgh02) on JWGH02 · 2026-07-14
Rendered to the canonical GRIFFai editorial design system (griffai.css)
INTERNAL — not DOC-LEDGER stamped; not customer-facing collateral