



TRINITY NUCLEAR • TRINITY SENTINEL

SECURITY & AI ASSURANCE

Most tools answer the AI question with policy. This one answers it with architecture.



No egress

enforced in code,
verifiable by inspection

Local model

may phrase,
never calculate

Attributable

hash-chained log;
refusals recorded

Honest

every open gap
listed in full

Summary for a reviewer.

This is a first-party assessment. It has not been independently reviewed or penetration-tested, and it says so wherever that matters. Every open gap is listed later in this document rather than left for a reviewer to find.

QUESTION	ANSWER
Does it send data outside the boundary?	YES No. Enforced in code, not by configuration.
Does the AI send data anywhere?	YES No. The model runs locally; the adapter refuses a non-local address.
Can the assistant act, or say it has?	YES No. A change becomes a proposal, any other click a draft a person confirms; purges, roles, the log, a Critical Decision get a logged <i>Won't do</i> (v5.15).
Does it collect usage, or certify?	YES No. A scrubbed packet of engine counts can be shown to an administrator; a person carries it out (v5.17). No surface calls an EVMS certified (v5.16).
Does it connect to P6?	YES No. It reads exported files. It holds no database credentials.
Can it modify the system of record?	YES No. It emits a file a person imports.
Is activity attributable?	YES Hash-chained log with actor, role and identity mode; every refusal is an entry.
Is identity authenticated on a shared server?	YES Only behind an authenticating front end; a request without its identity header is refused (v4.72).
Cross-site request forgery?	YES Closed v4.61: token, origin check, cookie scope.
Encrypted in transit?	YES TLS terminates at the front end on the server install. Desktop is loopback only.
Are records protected from deletion?	YES Every version of every governed record is kept; an evidence package is disposed of only by a records officer, under a legal-hold check, leaving a record (v4.89). Only an administrator's Purge removes a schedule's history — passphrase, logged.
Is access reviewed?	YES Every account confirmed by name on a stated interval; stale ones flagged; logged (v4.91).
Is the ownership position stated?	PARTIAL Copyright on every file; LICENSE (licensed, not sold; FAR 12.212 / DFARS 227.7202) and NOTICE shipped (v4.97). EULA with counsel.
Third-party code?	YES Three runtime libraries plus the window or server, pinned exactly; a CycloneDX SBOM with a SHA-256 per wheel.
Documented verification?	PARTIAL One command runs every check and writes a V&V record for that exact code, each requirement traced to its check (v4.73). Not yet independent.
Encrypted at rest?	NOT YET The store is plaintext; full-disk encryption is the expected control (S-6).
Threat model and System Security Plan?	PARTIAL Both drafted, first-party, to NIST SP 800-171; nine of twelve findings closed in code the same week.
Independently assessed? Code-signed?	NOT YET A-1, P-4 — open, and listed on page 5. Release hashes are published and the installer verifies them (v4.74).

1 — No egress, and why that claim is unusually strong.

Every outbound network call in the codebase has been traced. There are four, and all four are bound to loopback. The read-aloud worker opens none.

COMPONENT	PURPOSE	CONSTRAINT
Model adapter	Local language model	Constructor raises unless the host is localhost, 127.0.0.1, ::1 or 0.0.0.0
Alternate extractor	Not wired into the application	Same guard; the running app imports only the rule-based path, which has no network code
Startup port check	Detects an already-running instance	127.0.0.1 only
Single-instance check	Window management	127.0.0.1 only
Read-aloud worker	Optional local speech (Kokoro)	A child process that opens no socket; every file it needs is on disk

There is no configuration flag, environment variable or settings field that permits a non-local endpoint. Directing Sentinel at a commercial AI service is not something an administrator can enable by mistake or a user can enable deliberately — it would require modifying and redeploying source code. That is the strongest form this control can take short of an egress-filtered host, and it is verifiable by inspection. On the server install the process is additionally confined to loopback by its systemd unit.

There is no telemetry, no update check, no license call-home, no crash reporting and no analytics. The application has no reason to open an outbound socket and does not contain code to do so.

2 — THE AI BOUNDARY

Locality

The model executes on the same machine, served by a local runtime. No prompt, no schedule content and no computed figure crosses the boundary, because there is no path across it.

It cannot supply a number

Output passes a numeric-fidelity check against the block of computed facts it was given. Any sentence containing a figure not present in that block is discarded, the rejection is logged, and deterministic wording is substituted.

It cannot take an action

Actions are named tools with declared permissions. Routing is deterministic — Sentinel matches the request to a tool by pattern, and the model never selects what executes. No tool applies a schedule change. A hostile schedule whose activity names are instructions cannot move it beyond the asker's own role — tested every build (v4.73). Nor may it say it acted: a sentence claiming work not done is removed before it is shown (v5.11).

It is optional — and its voice is data, not weights

With no model installed, every calculation, page and deliverable functions unchanged. What makes the assistant sound like a project-controls analyst ships as a versioned **knowledge pack** — prompts, house style, number-free exemplars — verified by hash before use; a changed file means the built-in prompts stand, and nothing a site adds to it ever leaves the site (v4.87).

Stated precisely: **the model may choose words; it may never supply a value.**

3 — Data handling, identity, and the record.

At rest	One SQLite file, plaintext — see gap S-6. Filed documents and sealed evidence packages beside it, each checkable by SHA-256.
In transit	Desktop: loopback only. Server: TLS terminates at the front end; Sentinel is plain HTTP on loopback behind it, and its session cookie is Secure. There is no plain LAN mode: a network interface is served only behind the front end (v4.72).
Identity	Desktop: the Windows account. Server: whatever the front end authenticated, carried in a header the front end sets and the client cannot — Sentinel refuses a request without it. Every log entry records the mode.
Source files	Read-only. Sentinel never modifies a schedule file on disk.
Outputs	Written to a local output folder, or a configured share; on the server, one of exactly three writable paths.
Classification	A configurable marking is printed on every deliverable and in the status bar.

THE CHANGE LOG

Hash-chained: every event links to the one before it, so rows cannot be removed without the application reporting its own log as tampered. Deletion and purge deliberately never touch it. A wrong passphrase, a refused self-approval, a request without an identity and a refused cross-site request are all entries. Purge is gated three ways: the Administrator role, the admin passphrase typed at the moment of use, and the schedule must already be deleted. The daily server backup verifies the chain before it snapshots and copies the filed documents and evidence packages beside it. Since v4.72 the chain head is anchored where the service account cannot reach: every sealed entry writes it to the system journal, every issued report prints the head at issue in its footer, and a failed append is itself journaled — a rewritten store no longer agrees with copies its writer could not touch. Every entry carries its UTC offset.

DEVELOPMENT DISCIPLINE

Real client schedules are never used in development. A synthetic library stands in for them, and validation against a genuine export runs `tools/validate_real_schedule.py` on the machine holding the file, printing aggregate metrics only. No activity names, IDs, dates or WBS text leave that machine. That pattern is permanent.

ISSUES FOUND AND FIXED DURING ASSESSMENT

Stored XSS — a user-supplied folder name rendered unescaped; escaped, the download route contained, and a sweep added as the regression gate (v4.24–v4.29). **Hard-coded hours-per-day on the write path** — a 12-day activity written as 9.6 days on a 4×10 calendar (v4.27). **Reject recorded as an approval**, permanently, and **deleting a baseline destroyed it** — both found by the August QA pass and fixed (v4.52). **One person switching schedules switched everyone** on a shared server — the per-person workspace (v4.60). **No CSRF protection** — closed (v4.61). **Two writers could break the change log** under load — found by the concurrency test, now one transaction (v4.62). **Clearing a record deleted it, and Revert deleted the approval record** — every version now kept (v4.68). **Backups never carried the filed documents** — they do now (v4.69). **The report footer and status bar printed a fixed marking** whatever the administrator had set — marking-driven (v4.72). **Two routes had slipped the permission discipline** — combining schedules into a new roll-up, and installation-wide colours — found by the new route inventory, which now fails the build on any such route (v4.73).

Every gap, stated in full — closed ones included.

Each is tracked and prioritized. A list a reviewer has to assemble themselves is a list they will not trust — so the closed ones stay on it, with the build that closed each.

REF	GAP	BEARING
S-1-3	CLOSED HTML escaping (v4.29, xss_sweep the gate); CSRF — token, Origin/Referer, refusals logged (v4.61); cookie flags, Secure behind the front end (v4.72).	—
S-4	CLOSED Shared-server identity — only from the authenticating front end, one header; self-declared mode retired, a network interface is refused without the front end (v4.72).	—
S-5	CLOSED TLS in shared mode — at nginx in the service install; Sentinel is loopback-only behind it.	—
S-6	OPEN Store is not encrypted at rest.	The site's full-disk encryption (draft SSP); confirm at install
S-7	PARTIAL Threat model and draft System Security Plan written (v4.71, first-party, NIST SP 800-171); nine of twelve findings closed in code (v4.72–74).	Site ISSO review, then independent
S-8-11	CLOSED The threat model's findings — one identity header, hostile-file guards, record hygiene, chain head anchored off-box — closed v4.72.	—
A-1	PARTIAL Not code-signed. Release hashes published with every build; the Linux installer verifies them (v4.74).	Needs a certificate; blocks a Windows site install
A-2	CLOSED Pinned dependencies and SBOM — closed v4.71; hash-checked server install (v4.72); every pin queried against OSV each build (v4.74).	—
A-3	PARTIAL One-command battery and V&V record per build (v4.71); every requirement traced to its check in each record (v4.73). Independent verifier open.	NQA-1 / DOE O 414.1D
R-1	CLOSED Records — every version kept; evidence packages with retention, legal hold and a disposition record (v4.69–89); access recertified (v4.91). Not an EDMS, by design.	Records schedule
A-4	PARTIAL Control-responsibility matrix — NIST SP 800-171 in the draft SSP (v4.71); NQA-1 mapping open. The August 2026 QA review's 23 findings are closed (v4.91).	NIST 800-171, NQA-1
P-6	PARTIAL Cross-project critical path — declared-link pass and Combine shipped; inferred logic and import emission not.	The master-schedule capability
P-1	OPEN Scoped import not yet fed to a real P6 Import Wizard.	Biggest unknown; gates the write side
L-1	PARTIAL Licensing — LICENSE, NOTICE and notices shipped (v4.97); EULA, beta licence, legends, registration and trademark search with counsel.	Contract terms at award

RECOMMENDED POSTURE TODAY

Single-user desktop mode (identity is the Windows account), or the Linux service behind the site's own sign-in for a team — there is no third mode. The service install is verified in a build environment; the customer's own hardware is the next step.

The two rules that do not bend.

RULE 1

No outside AI touches your data

Sentinel runs entirely on your machine or your server. There is no egress, and it is enforced in code rather than promised in a policy. The AI assistant runs on the same machine against a locally held model, and can be switched off entirely.

RULE 2

Nothing that changes a schedule executes

Every action the assistant can take is a named tool with a declared permission. Anything that would alter a date becomes a proposal at the approval gate, where a person decides — and the gate refuses the person who proposed the change.

WHAT THIS ASSESSMENT DOES NOT CLAIM

It has not been independently assessed or penetration-tested. The threat model and System Security Plan are first-party drafts.

It is not accredited or certified under any framework; it is architected for accreditation, and the customer's site owns the Authority to Operate.

Shared mode has been exercised by automated tests — up to 50 simulated people at once — not yet by real people on the production box.

The scoped P6 import has not yet been run through a real P6 Import Wizard.

Fidelity is measured and displayed, not asserted. The Technical Architecture sheet describes the engine and the verification tools; the User Manual describes every screen. All three are at trinitynuclear.ai/sentinel, and the source of this assessment is available to a sponsoring customer's ISSO on request.



Trinity Nuclear Management Group, Inc.

Project controls for DOE and defense nuclear work — Primavera P6, EVMS, DCMA schedule assurance and risk.

CONTACT

Jeffery S. Gammon, President & CEO
jgammon@trinitynuclear.ai
 505-205-1307

ONLINE

TrinityNuclear.ai
trinitynuclear.ai/sentinel

REGISTRATION

UEI JXLHYSWM1AR9
 CAGE 22FY1 · SAM.gov Active
 NAICS 541330 (primary)

CLEARANCE & BUILD

DOE Q cleared
 Trinity Sentinel v5.25