



TRINITY NUCLEAR • TRINITY SENTINEL

HOW IT WORKS

It reads your P6 exactly, tells you what is wrong with the schedule, drafts the fix, and never touches your database.



Reads

four formats,
one engine

Recomputes

every date,
calendar-aware

Drafts

typed proposals
with an impact preview

Stops

a person decides;
P6 is never written

The analysis is sound. The assembly is slow.

A site running eight projects typically has schedule health assessed in a spreadsheet once a month, earned value calculated somewhere else, risk analysis done rarely, and a monthly report assembled by hand from all three. By the time the report is written the data date has moved. Sentinel does the computation continuously and assembles the deliverable in one action. What it does not do is decide anything.

END TO END

- 1 It reads what P6 already exports**
 Point it at a folder. P6 XER, P6 PMXML, Microsoft Project MSPDI and generic CSV through a single engine. Nothing is converted or re-keyed, and no connection to your P6 database is made or needed. Each read is filed — headline and activity-level — so the next read can say what moved.
- 2 It recomputes every date independently**
 A calendar-aware critical path: each activity on its own P6 calendar, working week and holidays, honoring FS, SS, FF and SF logic with lags in working time. A forward pass from the schedule's data date holds the actuals and forecasts the remaining work; a backward pass from the mandatory finish gives late dates and signed float. It then compares its answer against what P6 stored and reports the agreement on screen.
- 3 It assesses the schedule**
 DCMA 14-point, grouped by the question each check answers rather than as a flat list, negative float included. Failures first, and every activity named in a finding is a link that opens the grid filtered to it. Alongside it: EIA-748 earned value to the work-package level, Monte Carlo risk under a spread stated per project, variance against thresholds the project owns, and named baselines. Every finish is measured at the end milestone a PM names, with its driver.
- 4 It drafts, and then it stops**
 “Shorten PU-150 by five days” becomes a typed, validated proposal with an impact preview — what it does to the finish date, to float, to health. The monthly status update arrives the same way, as one card. Nothing is applied until a person decides; the person who proposed a change cannot approve it, and the refusal is written permanently to the change log.
- 5 It hands you a file, not a connection**
 Approved changes become a scoped import workbook containing only the fields that were approved, for P6's own Import Wizard. Sentinel holds no credentials for your P6, opens no connection to it, and cannot modify it.
- 6 It produces the deliverable, and seals it**
 One action produces a DOE-format project report — a cover carrying the headline and the six figures a manager reads first, performance at the data date, what changed since last period and what caused it, cost and schedule variance with drafted narratives and the status of every VAR, the budget log reconciled against the contract figures, slipping activities, staffing, baseline history and change record, risk, corrective actions, signature block. HTML, PDF or Word. Every issued period is archived, and can be sealed into an evidence package anyone can check — held under legal hold, or disposed of under a recorded authority, when the records schedule says so.

The gap between this tool and your system of record is a human being clicking approve. That is a deliberate design constraint, not a missing feature.

Why the agreement figure is the foundation.

Every other figure in the application is derived from the computed dates. If Sentinel could not reproduce P6's own arithmetic, nothing else it said would be worth reading — so it measures that continuously and shows you the result rather than asking you to trust it.

100%

P6 date agreement, validation library

706

of 706 activity finishes, exactly — and 37 of 37 on a real site export

4

file formats through one engine

5×8 4×10

9/80

working patterns resolved per activity

Computed against stored, activity by activity. The percentage is measured on every read, not asserted once.

THE DATA DATE, AND WHERE IT COMES FROM

A fact about the schedule, not the installation

The date everything is measured “as at” is read from the file when it declares one — XER last_recalc_date, PMXML DataDate, MSPDI StatusDate. When the file is silent, the latest actual it carries is offered as *inferred* — *confirm it*. A person can set it, under a permission of its own, and the change is audited.

The provenance is always printed

Hero band, report header, Word and PDF all say where the date came from. If a file's actuals sit *after* its data date, the engine says so rather than forecasting from a date the schedule has already passed.

THE ASSISTANT, AND ITS LIMITS

It runs on your machine

The model is local. The code refuses any non-local endpoint — pointing it at a commercial AI service is not a setting an administrator can enable.

It may phrase, never calculate

Every figure in a model-written sentence is checked against the computed data; anything invented is discarded and logged. A sentence claiming work it did not do is corrected before it is shown.

The application is complete without it. With no model installed the assistant still answers — what to worry about, any activity or account by its ID, a status note for your manager — from the deterministic path. Ask it to change something, or to lay out a new schedule, and it drafts a proposal or a request for another person to approve; it cannot act.

Scale — what is proven, and what is not.

Stated plainly, because it is the question a multi-project site asks first.

WITHIN A SINGLE SCHEDULE — PROVEN

Measured to 20,000 activities, which covers essentially any single site schedule. A full assessment — calendar-aware critical path, DCMA 14-point, float and driving path — completes in under nine seconds, and every page after the first is served in well under a second.

ACTIVITIES IN ONE SCHEDULE	FULL ASSESSMENT	SUBSEQUENT PAGES
184	0.06 s	instant
10,120	4.2 s	under 0.5 s
20,240	8.7 s	under 0.5 s

ACROSS MANY SCHEDULES — PROVEN

The portfolio reads a schedule once and keeps the result. A file unchanged since it was last read is not read again, so opening the roll-up costs a directory listing rather than a re-analysis. Measured on 1,000 schedules of 2,024 activities — two million activities, 291 MB on disk:

WHAT YOU DO	WHAT IT COSTS
Open the portfolio or a program roll-up	25–55 ms
First read of a brand-new estate	6.8 min, once, in the background
Come back after one schedule was updated	0.7 s
Come back after a normal week of updates	19 s

Every figure the roll-up reports is identical to what a full re-read of the same folder produces. Reading once is a saving in time, not in rigor — and a schedule not yet read is reported as pending on the page rather than quietly left out of the totals.

ACROSS PROJECTS — DECLARED, NEVER INFERRED

The master schedule, built not re-typed — and the program, reviewed

Combine assemblies one roll-up from the sources, recipe recorded; the Estate page computes the driving path from the links people declared. Over them, the program pages (v4.75–80): a Program Review with one row per project and the review block the meeting walks; a Program Gantt with a slip scenario that writes nothing; resource overlap by month; the funding profile by fiscal year.

What it does not yet do

It does not infer logic between projects — a wrong automatic guess between two nuclear schedules is a real problem — it does not level or reassign a resource, and the import of a combined schedule back into P6 waits on the write-side test against a real Import Wizard. All are defined work, not configuration.


[← Back](#)

Estate

The critical path across projects, declared links only, nothing here changes any schedule.

Estate — the driving path across projects, from the links a person declared; nothing here changes any schedule.

What it runs on.

Deployment	Desktop application on Windows — installation is a folder with its own runtime. Or a Linux service on a shared box behind the site's authenticating front end, serving a team from one store with each person's workspace their own.
Network	None required. No cloud service, no external API, no telemetry. The service binds loopback; only the front end reaches it.
Identity	Desktop: your Windows account. Shared server: whatever the front end authenticated — Sentinel refuses a request that arrives without it, and a browser cannot name itself.
Data	Stays on the machine. One database file and the filed documents beside it, backed up daily on the server with the change-log chain verified first. Every version of every record is kept.
P6	Read-only, via the exports P6 already produces.
Dependencies	Flask, openpyxl, python-docx — plus pywebview for the desktop window, or waitress as the shared server.
Scale	20,000 activities in one schedule; 1,000 schedules in one estate; 50 simulated people at once on one server in the lab.


WHAT YOU WOULD NOTICE FIRST

The monthly report becomes an action rather than an exercise.

Schedule health is a live number, not a monthly opinion.

Every change to a schedule has a name, a reason, a decision and a time attached to it — permanently, and in a form that can be shown to an auditor.

Nothing leaves the site.



The Approval Gate — describe the change in plain language; the impact preview is computed against a throwaway copy before anyone decides.

P6 stays the system of record.

Sentinel sits beside it and does the work P6 leaves to a person: reading the whole estate, checking it, explaining it, and drafting the change — then stopping so a scheduler can decide.

WHAT SENTINEL READS

Primavera P6 · XER

Primavera P6 · PMXML

Microsoft Project · MSPDI

CSV · mapped on import

WHAT SENTINEL WRITES

A scoped P6 import file

A DOE-format project report

A program review pack and workbook

An archived, reproducible period

A sealed evidence package — and its disposition record

Nothing else — ever, to P6 directly

GETTING STARTED

The Quick Start covers the first thirty minutes; the User Manual covers every screen. Two one-click builders put a synthetic demo estate and six complete demo projects on your machine, so every page can be walked before a real file is loaded. The documents, with the licence and notices, are at trinitynuclear.ai/sentinel.



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