



TRINITY NUCLEAR

TRINITY SENTINEL

Every date, traceable to the logic that produced it.



Complements P6

reads XER, PMXML,
MSPDI and CSV

Runs offline

no egress,
enforced in code

AI assistant

answers from the engine,
drafts for the gate

Nothing executes

every change is a
proposal at the gate

Your master schedule is assembled by hand.

A site runs thousands of schedules. Someone maintains a summary project in P6 and re-types the dates into it every cycle. The integration logic lives in that person's head, the roll-up is stale the day after it is built, and nobody can answer “which project is actually driving the program finish” without a week of work.

Sentinel reads every schedule where it already lives, reproduces P6's own dates with a calendar-aware CPM, and rolls the estate up in milliseconds — then stops, and hands a person a decision.

The summary project itself is **built, not re-typed**: choose the schedules and Sentinel assembles one roll-up from them in a single action, carrying every activity, its logic and its progress through unchanged — and recording the recipe, so next cycle is a rebuild rather than a rewrite.

WHAT IT IS NOT


Sentinel is **not a P6 replacement** and never writes to P6's database. P6 stays the system of record. Sentinel emits a scoped import file that a scheduler reviews and imports — or does not.

THE GOVERNED CYCLE — EVERY ACTION, EVERY TIME



Anything that would move a date becomes a proposal at the approval gate. Refusals are written to the hash-chained change log exactly like approvals — the record shows what was declined, by whom, and when.

ONE SCHEDULE, ANSWERED IN ONE SCREEN



Health, cost, schedule and forecast side by side — each figure computed by the engine, none of it re-keyed.

One engine, from a single schedule to the whole estate.

Reads what you already have

P6 XER and PMXML, Microsoft Project MSPDI, and CSV. No conversion step, no data entry, no asking the customer to change how they work.

Reproduces P6's dates — then forecasts

A calendar-aware CPM that schedules each activity on **its own** P6 calendar — workweek plus holidays — so computed finishes match P6's exactly — then a statused pass from the data date, every finish measured at the end milestone the project names (p. 8).

DCMA 14-point, with the fix

All fourteen checks, grouped into the five questions they answer, a plain verdict at the top — and every failing check carries its resolution. Unambiguous fixes draft themselves to the approval gate; a person decides.

EVMS & variance, to EIA-748

CPI, SPI, CV, SV, EAC, VAC, TCPI and BEI to the work-package level, the per-period series behind them, variance narratives with a lifecycle, a budget log reconciled to the dollar, and an evidence map for all 27 EIA-748-E guidelines (p. 9).

Risk, in time and in money

Monte Carlo P10/P50/P80/P90 says *when*; a DOE G 413.3-7A-style risk register says *why* — 5×5 matrix, owners, handling plans, candidates the simulation drafts and a person accepts. Priced against the schedule's burn rate it also says what the exposure *costs*, with a contingency position beside it, and it imports and exports so a program can share one register.

A DOE-format report, in one action

The report a monthly review expects — a cover that tells the story, what changed since last period and what caused it — assembled, archived, and sealable into one evidence package.

The roll-up, built not re-typed

Pick the schedules; Sentinel combines them into one — activities, logic, calendars and progress carried through exactly, each source a level-1 band. It invents no logic *between* projects: those are declared by a person. The program pages (p. 5) sit over it.

A build area for the controller

Rename bands, drag activities between them, add one, save and come back. Restructuring the WBS moves no dates; anything that *would* become a proposal at the gate, visible in place until it is decided.

THE ASSISTANT

Ask in a sentence; it answers, or drafts the click.

Findings, variances, the end driver, what changed since last month — answered from the engine. Asked to do something a page does, it drafts that page's own form for a person to confirm (p. 7). It runs **on your machine**, restates only numbers the engine computed, and can be turned off.

100%

P6 date agreement
706 of 706 finishes
and 37 of 37 on a real
site export

THE CLAIM EVERYTHING ELSE RESTS ON

Every activity finish in the library reproduces P6 exactly.

706 comparable finishes across every schedule in the validation library. Each activity is scheduled on its own P6 calendar, so 5×8, 4×10 and 9/80 patterns resolve independently within one file. And it is proven against the real thing: a genuine P6 site export, validated on the machine that holds it — only aggregate metrics ever leave — scored **37 of 37, 100%**. A tool that disagrees with P6 about a date is not usable at a DOE site, so agreement is measured on every read and shown on the page.

Measured on the shape of a real site, not a demo.

1,000

schedules in the portfolio roll-up

2,024_k

activities across the estate

25–55_{ms}


to assemble the roll-up page

531/531

routes and renders, no failures, both themes

PORTFOLIO ROLL-UP • TIME TO ASSEMBLE THE PAGE



1,000 schedules • 2,024,000 activities • 291 MB on disk. Log scale.



Finding out *what is there* is a directory listing. *Reading what changed* is the expensive part, and it is skipped entirely for any file whose size and modification time are unchanged. Nothing is silently omitted – a schedule not yet read is reported as pending on the page.

ENGINE THROUGHPUT • FULL CPM AND DCMA PASS

A real DOE site schedule is 10,000–50,000 activities. Every page after the first is under half a second.



The estate, rolled up – seventeen projects here, one thousand on the customer's shape.

One room. Every project. Run the way the review runs.

A program manager walks into the monthly review with one question per project: **what needs attention first?** Sentinel reads every schedule in the estate and answers it on one page — then walks the agenda the room actually walks: CPI and SPI, cumulative and for the period; the variances by control account; what changed in the schedule and the scope; what finished. Where projects share people, money or a handoff, it says so — with the coverage stated on every figure. The program pages sit in their own group in the navigation, one click from anywhere.

PROGRAM REVIEW • ONE ROW PER PROJECT, ATTENTION FIRST



Forecast, slip, estate float and pushes, EVMS, data date and one priority-ordered sentence each — then a review block per project. Printable pack and workbook.

Program Gantt

Every project on one axis with its declared links, the driving chain in red, and every push one project puts on another. A slip scenario shows what a late anchor does to the rest — and writes nothing. Unacknowledged dependencies are filed as *Start On or After* proposals at the successor's gate.

Resource overlap

Same-named resources stacked across projects by month from the assignments; activity-code concurrency as the fallback where a schedule is not resource-loaded. Declare a capacity and the months over it are marked; declare none and the page says the limit is unknown.

Program budget

The funding profile by fiscal year, plan against forecast, so a slip reads as the funding-profile change it is. EAC against budget ranked by project with TCPI and the CPI-based estimate; the standing reserve set against the time-only cost of P80 schedule risk.


Cost risk, priced

Three-point cost impacts and drivers per risk, expected value, and a simulated contingency position for each project and across the program — beside the time-only figure, never silently replacing it. The register imports and exports so a program shares one.

WHAT THE PROGRAM PAGES REFUSE TO DO

They read; they never move a date. If eleven of eighteen projects carry a cost baseline, every program figure says it was computed over those eleven and names the seven it could not see — nothing reads as 1.00 or zero by default.


RESOURCE OVERLAP • THE DEMO ESTATE: EIGHTEEN CONNECTED SCHEDULES, ONE CRAFT OVER ITS DECLARED CAPACITY



The number is the easy part. The record around it is the review.

An EVMS surveillance does not ask whether CPI is 0.94. It asks who owns the variance and by when, under what authority the budget moved, which thresholds judged the score, and whether the evidence can be produced. Sentinel keeps each as a dated act by a named person — and prints it in the deliverable, so the answer is on the page before the question is asked.

THE COVER TELLS THE STORY • PAGE 1 OF THE PROJECT REPORT



The headline verdict and the six figures a manager reads first, on a full-bleed cover — by the same formatters and thresholds as the body. The test reads page 1 of the PDF and proves it cannot disagree with page 2.

The issued report carries its own typefaces, embedded in the HTML and the PDF, so it prints the same on any machine with nothing fetched. Every chart states its axis, its reference line and its threshold bands; a VAR prints as one card that never splits across a page.

The footer attributes the tool and not the record: produced with Trinity Sentinel; the report's content belongs to the issuing project.

A variance record with a life

Breached → assigned to a CAM with a date → narrative submitted → approved → corrective action open → closed and verified. Every step a person's dated act; out-of-order acts refused; overdue computed. The report carries the status of every VAR.

The budget log, reconciled

The opening position and every change to CBB, PMB, MR and UB, each under its authority. CBB = PMB + MR holds at every entry; the report replays the log and says, line by line, whether the contract figures and the cost file agree — or differ by how much.

Thresholds and assumptions, recorded

DCMA thresholds are a named set per project and every stored score records which set judged it. The Monte Carlo duration spread is set per project and printed beside every confidence figure it produced. Nothing is judged under a rule the page does not state.

Evidence, held and disposed of on the record

A legal hold suspends disposition. A disposition is permitted only after retention has run, under a stated authority, and leaves an append-only record carrying every fingerprint the package had — *destroyed under authority*, never merely *missing*. Access is recertified on an interval; the store is re-hashed at every start.

THE BUDGET LOG • AGREES, OR DIFFERS BY HOW MUCH

- **AGREES** Contract budget base: header \$3,150,000 agrees with the budget log
- **AGREES** Performance measurement baseline: header \$3,030,000 agrees with the budget log
- **AGREES** Management reserve: header \$120,000 agrees with the budget log
- **AGREES** Undistributed budget: header \$0 agrees with the budget log
- **DIFFERS** Control accounts: the cost file's BAC total \$2,850,000 vs PMB - UB per the log \$3,030,000 (difference -\$180,000) — distributed budget and the control accounts disagree

Scope distributed by BCR but not yet in the cost file: the log says so, in dollars, rather than letting two figures sit apart on two pages.

A colleague who knows every number — and stops at the gate.

Ask it what a project manager asks on an ordinary day. It answers from the same engine as every page, in plain words, with the next step offered. Ask it to change something and it drafts the change for a person to decide. It runs on your machine, it works with no model installed, and it never says it did something it did not.

“WHAT SHOULD I WORRY ABOUT?”

DR Dana Reyes

What should I worry about this week?

Sentinel

ANSWER · from the computed data

Here's where I'd put your attention.

4 things I'd look at first, in order:

- PU-150 Instrumentation and controls is +20d on the driving path** — every day there moves the finish — a recovery on it buys the most.
- Milestone PU-210 Operational readiness review is forecast +19d late (2027-03-16)** — the date people outside the project are watching.
- 2 control accounts outside threshold, worst CA-3130 (S. Okafor, CV -\$52K)** — each needs a VAR with a cause and a corrective action this period.
- TCPI 1.08 against CPI 0.94** — the remaining work has to run better than it has so far to hold the budget.

Want me to start on the first one?

Copy Read

3 steps · 0.5s

“SHORTEN PU-150 BY 10 DAYS”

Shorten PU-150 by 10 days

Sentinel

PROPOSAL · waiting at the Approval Gate

On it.

Drafted 1 change for the Approval Gate on the live schedule — nothing has changed yet.

- PU-150** Original duration: 38 → **28** working days

If approved:

- Project finish: 2027-03-29 → **2027-03-15** — 14 days earlier on the calendar (10 days of working time)
- Schedule health: 64% — unchanged
- 5 activities move:
 - PU-150 finish 2027-01-26 → 2027-01-11
 - PU-160 finish 2027-02-22 → 2027-02-05
 - PU-170 finish 2027-03-15 → 2027-03-01
 - PU-180 finish 2027-03-29 → 2027-03-15
 - PU-210 finish 2027-03-16 → 2027-03-02

A person approves it on the gate — not me, and not the person who asked for it.

Want to see what it does to the finish first, or adjust it?

STOP · 1 PROPOSAL(S) AWAITING APPROVAL

Nothing has changed yet. Approve or reject it on the gate.

Left: a ranked list from the computed data — the driving-path slip, the late milestone, the account outside threshold — with an offer to start on the first. Right: the change as a proposal, the finish in calendar and working days, every moved activity listed, and nothing applied.

Asks like a manager; answers from the engine

What to worry about, why it is late or over budget, any activity or control account by its ID, what changed since last period, whether it is fit to baseline, a status note to paste for your manager. On all 25 demo schedules her figures are checked against the dashboard and the report on every build.

Drafts any click — a person confirms

Add a risk, close R-004, mark the CD-0 statement done, set the *commitment date*: the reply carries that page's own form and a Confirm button — same role check, validation and log. *Approve* it is refused.

A new schedule, from a sentence

Pour a pad 1 week, place the box 2 days, wire it 2 weeks, energize 1 day, start 10-30-26 — the engine lays it out on the working calendar. Submitted, it becomes a request on the gate, routed to a named approver; only their approval adds it to the library.

Honest, and fenced

A local model may choose words, never a value; a claim of work not done is removed. Purges, roles, the change log, a Critical Decision, a legal hold: a plain *Won't do*, logged. Nothing leaves the machine.

THE REQUEST ON THE GATE · DATES BY THE ENGINE, DECISION BY A PERSON

NEW SCHEDULE NS-001 — Electrical box replacement 15KV

ROUTED TO any approver

Requested by Dana Reyes · 2026-10-15 20:07 · drafted in a conversation, dates by the engine · 2026-10-30 → 2026-12-03 · the draft .xer

ACTIVITY	NAME	WORKING DAYS	START	FINISH
EBRK-1000	Pour a pad	5	2026-10-30	2026-11-05

One date the project is managed to — and what is driving it now.


Every project has an end milestone — by default the finish milestone with the latest planned finish; a PM or CAM can name another, and it is logged. The forecast, the slip, the P50/P80 and every roll-up measure it. Turnover and closeout after it are listed apart and move none of them.



THE DRIVING CHAIN • END DRIVER IN GOLD



WHEN IT FINISHES • AND THE COMMITMENT



Left: the chain into the end milestone walked on the statused forecast; the unfinished work at its head is the end driver. Right: the simulated finish at the end milestone, each P-level a date, the commitment placed on the curve — met in about 65% of runs on the pilot.

The end driver, named

What controls the date right now — and when it changes from one filed period to the next, the page says so in plain words rather than leaving it to be noticed.

Measured against two dates

The baseline, and a commitment date a person enters — the approved CD-4, a contract finish, a return to service. The margin, the float on the end milestone and the P80 against it are stated.

Risk at the date that matters

The Monte Carlo is run at the end milestone, criticality counted only on work leading to it, so a training course after the review is no longer schedule risk to the finish.

A record, period by period

Each filed period keeps the end milestone's forecast, slip, float, P80 and end driver: the monthly review reads a trend, not one number. The assistant answers *what's driving the finish?* from the same figures.

One engine. A profile for each kind of work.

Line-item capital projects, reportable earned value, operations and maintenance, small task orders. Each project carries a profile, chosen by a person and logged, and the pages it needs follow from it — on the dashboard, in the report, on the program review.

Capital — DOE O

413.3C

The Critical Decisions as a ruleset: CD-0 to CD-4 and closeout, deliverables tailored to cost, hazard category and first-of-a-kind; the approval authority for the band.

Reportable EVM

The full EVMS engine and variance narratives; work authorization, the RAM, the WBS dictionary and the period close; the EIA-748-E evidence map — built so an EVMS run on it is certifiable.

Operations

Schedule adherence week by week, the backlog by age, recurring work found by name, the driving work to an outage's finish, spend against the fiscal-year plan.

Task order

Schedule health, the status update and the simple report — nothing a small scope does not need.


THE CRITICAL DECISIONS • CD-0 RECORDED, CD-1 IN PROGRESS



EIA-748-E • ALL 27 GUIDELINES, WHAT THE DATA BEARS OUT

ALL 27 GUIDELINES AT A GLANCE

ORGANIZATION




PLANNING, SCHEDULING AND BUDGETING



PROGRESS ASSESSMENT AND DATA COLLECTION



ANALYSIS AND MANAGEMENT REPORTS



REVISIONS AND DATA MAINTENANCE



✓ shown × needs attention ... needs data ○ attestation attestation on file

Each guideline marked from the schedule and cost data — WBS coverage, control accounts and their managers, the network, the baseline, work packages, earned value techniques, variances and their analyses, EAC credibility, work authorization, the WBS dictionary, the period close and every retroactive change — shown, needing attention, needing data, or the contractor's attestation, with the system-description section beside it.

IT RECORDS DECISIONS; IT NEVER MAKES THEM

A Critical Decision is the approving official's act: Sentinel records the date and the name a person enters, and will not record one from a conversation. Sentinel is built so an EVMS run on it is **certifiable** against EIA-748-E; it is not itself certified — that belongs to the contractor's system and DOE's surveillance. No surface says an EVMS is certified, compliant or validated, and the battery checks for the words.

A variance is only as good as the baseline it is measured against.

Large projects re-baseline — often a dozen times, as estimates firm up. Each one resets what every figure afterwards is compared to. Sentinel treats that as a record to be kept, not an event to be survived.

Every capture records why — and what it moved

A baseline is taken with a stated reason, and the movement against whichever baseline was active at that moment is computed **at capture** and stored with it: how many activity dates moved, later versus earlier, net days, the largest mover named. A comparison recalculated later is an opinion; this is a record.

The change that authorized it, on the same page

Change number, approving board, approval authority, date, change level, cost impact and funding action — the governance half, beside the measured half. Save it with gaps and it names them back to you, because a record with holes in it is what an audit finds first.

The paper behind the numbers, filed beside them

Signed change forms, directives, the spreadsheet a register arrived on, review packs — filed against the schedule, numbered, and attached to the baseline or risk they belong to. Kept on your machine, next to the data. Nothing is uploaded anywhere, and only document types are accepted, checked by content rather than by name.

Integrity you can prove, not assume

Each document is stored under a cryptographic fingerprint of its own contents, recorded in the hash-chained log the day it is filed. Any time you ask, Sentinel re-reads every document and re-checks it — and refuses to hand back one that no longer matches. Each period can be sealed the same way: one package, every file fingerprinted.

THE BOX THAT USUALLY GETS TYPED — MEASURED INSTEAD

SCHEDULE IMPACT — MEASURED, NOT STATED

These figures were computed when the baseline was captured, against whichever baseline was active at that moment. They are not editable here and never were: on a paper change form, this is the box that gets optimistic.

MEASURED AGAINST	PMB Rev 1
ACTIVITIES WHOSE DATES MOVED	5 (0 later, 5 earlier)
NET MOVEMENT	-38 day(s)
ACTIVITIES ADDED / GONE	0 added, 0 gone
LARGEST SINGLE MOVE	PU-160 moved -10 day(s)

On a paper change form this box is filled in by the person requesting the change. Here it is read-only, computed when the baseline was taken.

AND IT REACHES THE DELIVERABLE

The project report no longer just states a variance. Its header names the baseline that variance was measured against — and, where a record exists, the change that authorized that baseline: “**measured vs PMB Rev 3 · authorized by BCP-2026-014 (Project Change Control Board, 2026-07-02)**”. Naming the comparator says what the number means; naming its authorisation says why that comparator is the legitimate one.

THE PROJECT DOCUMENT REGISTER

NO.	DOCUMENT	KIND	THEIR NUMBER	DATED	FILED AGAINST	FILED BY	FILE
1	...	...	...	...	...	...	...

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. Three runtime libraries plus the desktop window or the server, all auditable; every shipped file checksum-listed. The AI assistant runs **on the same machine**, against a model held locally — nothing is sent to a vendor, and it 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. The approval gate refuses the person who proposed the change.

AND THE ASSURANCE BEHIND BOTH

The AI calculates. It does not learn.

Sentinel reasons with two instruments: a deterministic engine — same input, same answer, every time — and an optional local language model for phrasing. Neither trains on your data. Nothing you load is absorbed into a model, nothing leaves the machine, and every figure traces to a calculation you can audit. A system that quietly learned from controlled schedules would walk out the door as a liability; Sentinel is built so that cannot happen. What persists is a **record** — assessments, comments, the hash-chained log — never training data.

Hash-chained change log

Every event links to the one before it. Rows cannot be removed without the application reporting its own log as tampered — so deletion and purge deliberately never touch it.

Records-safe lifecycle

Nothing a person wrote is destroyed by a button: every version of every record is kept, an evidence package leaves only under a legal-hold check and a recorded authority, and the schedule file on your disk is only ever read.

The assistant may phrase, never invent

A numeric-fidelity guard means it can only restate figures the engine has already computed. It cannot introduce a number of its own.

Separation of duties — across people

Roles, an admin passphrase re-typed at the moment of use, and attribution written to the log with every decision. On a shared server identity comes only from the site's own sign-in, everyone sees who else is on a schedule, every account is recertified on an interval, and nobody is ever locked out.

THE PROPOSAL, BEFORE ANYONE DECIDES


Sentinel
 TRINITY NUCLEAR

- Dashboard
- Schedules
- Gantt
- Programs
- Portfolio

[< Back](#)

Approval Gate

The governed edit layer — describe a change, see exactly what it does to the schedule, then decide. Read → Propose → Validate → **STOP** → Apply → Scoped import.

REVIEWS ROUTE TO — not set —

DESCRIBE THE CHANGE

Proposing against Pu Facility Glovebox Line — Pilot (basel)

P6 stays the system of record.

Sentinel sits beside it and does the work that 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

Nothing else — ever, to P6 directly

DEPLOYMENT

Runs on a single workstation, or as one server for a team behind the site's own sign-in. No cloud tenancy, no external service, no account. Installation is a folder. On a network that never reaches the internet, it behaves identically — because it never needed to.

THE NAME

A sentinel is the watchman on the wall.

From the Latin *sentire* — to perceive. The sentinel is posted where watching matters most: first to see trouble coming, raising the alarm — and never leaving the post, never acting alone. That is this product's job: it stands watch over your schedules, senses drift the day it starts, reports what it sees — and never moves a date on its own. **Trinity** is the heritage: the name the nuclear era began with, in the New Mexico desert — carried today by a New Mexico company built for the sites where that history still works. **Born of nuclear history, standing watch over its future.**



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

REGISTRATION

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

CLEARANCE & BUILD

DOE Q cleared
Trinity Sentinel v5.25