SPERT® Scheduler

Quick Reference Guide

Probabilistic project scheduling with SPERT® three-point estimation and Monte Carlo simulation

WHAT IS SPERT® SCHEDULER?

SPERT® Scheduler is a lightweight, probabilistic scheduling tool for project managers working on traditionally-scheduled projects. It combines Statistical PERT three-point estimation with Monte Carlo simulation to answer: *"How confident are we that we finish by a given date?"* — and to calculate the schedule buffer needed. It is independent of SPERT® Forecaster (agile tool).

Three-point estimates	Enter Optimistic (min), Most Likely, and Pessimistic (max) durations in working days per activity. Enable Heuristic Mode in Settings → Preferences to auto-calculate Min and Max as configurable percentages of Most Likely (e.g., Min = 75%, Max = 200%). A zero-uncertainty estimate — where Min, Most Likely, and Max are all the same value — is valid in every distribution: use it for genuinely fixed-duration activities (a booked training week, a contractual notice period), or leave new activities at the 1/1/1 default while you lay out the full activity list and dependencies before estimating anything.
Confidence level	**Applies only to T-Normal, LogNormal, and Beta-PERT.** Maps your subjective confidence (Near certainty → Guesstimate) to a spread: for T-Normal and LogNormal, a standard deviation via the SPERT Ratio Scale Modifier (RSM); for Beta-PERT, the spread on the Statistical PERT® Beta Edition scale. Has no effect on Triangular or Uniform, where the grid shows a dash instead of a level.
Distribution types	T-Normal, LogNormal, **Beta-PERT**, Triangular (default), Uniform. Beta-PERT stays between Min and Max, peaks exactly at Most Likely, and takes its spread from Confidence (Medium: SD = range ÷ 6) — it is *not* the fixed (O + 4M + P) ÷ 6 formula. At Medium it schedules shorter than Triangular (10 / 12 / 40 days: 16 vs. 20). An amber **suggestion dot** beside Distribution flags a better fit judged from the three numbers alone; it never suggests Uniform.
Deterministic schedule	Each activity's scheduled duration = inverseCDF(activity probability target, P30–P95, default P50). P30/P40 build the schedule on optimistic durations and leave more uncertainty to the buffer (Project target: P50–P99). The schedule is calendar-aware, respecting your configured work week and all holiday layers.
Monte Carlo simulation	Default 10,000 trials per run; up to 100,000 trials available for audit-grade precision where reduced variance is needed. Each trial samples a duration for every activity from its distribution. Parkinson's Law clamping (configurable per scenario): no simulated duration falls below the deterministic scheduled duration, which is set by your Activity target — not always P50. For an In Progress activity, sampling is conditioned on the elapsed time already observed — durations shorter than what has already elapsed are excluded — so the buffer widens honestly as a long-running activity slips, instead of optimistically assuming it is nearly done.
Schedule buffer	Buffer = the Monte Carlo percentile at your project target minus the deterministic schedule span (project start through constraint-adjusted finish). Measuring against the span rather than the sum of activity durations means **Finish w/Buffer always equals the target percentile row's date in the Percentile Summary**. Where a scenario waits on hard date constraints, the idle working days are disclosed separately as **Constraint delay**, so the arithmetic reads cleanly: Duration + Constraint delay + Schedule Buffer = Duration w/Buffer. Set a **Finish Target date** to see RAG (red/amber/green) schedule health in the summary card and on the Gantt chart. Thresholds configurable in Settings (defaults: Green = P80, Amber = P50).

■ **Tip:** Use T-Normal, LogNormal, or Beta-PERT when you want the Confidence level to influence the spread. Use Triangular or Uniform when you prefer a simpler, assumption-free distribution.

STEP-BY-STEP WORKFLOW		
#	**Where**	**What to do**
1	Projects page	Create a new project (tile-based list) or open an existing one — click the tile or its name, or Tab to the name and press Enter/Space. Grab anywhere on a tile to drag and reorder it. Search or archive projects for tidy housekeeping. Assign a tile color (Slate, Sage, Sky, Lavender, Rose, Amber, Teal, Clay) to color-code projects by program — shown as a soft full-tile color wash plus a 4px accent strip. Hover a tile to reveal its action icons: **Export**, **Archive**, and **Clone** in the bottom-right corner (Clone produces a fully-detached copy named '{original} (Copy)', useful for scrubbing sensitive data before sharing); **Share** in the bottom-left (cloud mode, owned projects only) opens the sharing dialog directly from the dashboard. The trash icon in the top-right is always visible; it asks first, and after a delete you stay on the Dashboard. Beside the Projects heading, **Export All** and **Import** handle bulk backup and restore (Import opens the file window straight away), **Load Sample** creates a complete worked example in one click (see the Sample Project section below), and **New Project** asks for a name and creates the project. A theme toggle (Light / Dark / System) in the app header lets you switch color modes globally.
2	Calendar & holidays	Configure your work week (any combination of days), add country public holidays via the Holiday Loader (100+ countries), and enter manual holidays. All scheduling respects the active calendar.
3	Add activities	Type activity names directly into the grid, or import in bulk: Settings → Import Activities from Spreadsheet accepts a CSV file or clipboard paste with a downloadable template. Hover between any two rows to reveal an insert strip — click to add an activity at that position with the name auto-focused (it hides in a locked scenario and during a drag). Use the **+ Section** button (beside + Add Activity) to insert section header rows — each with an editable name and optional color — to visually group activities in the grid and Gantt chart.
4	Estimate & configure	Enter Min / Most Likely / Max per activity (or let Heuristic Mode auto-fill). Choose a Distribution, then a Confidence level (T-Normal / LogNormal / Beta-PERT only). Use Settings → Estimation Heuristics Suggester for evidence-based min/max percentage defaults across 23 industries and 73 activity subdomains.
5	Dependencies & targets	Toggle Dependency Mode to add Finish-to-Start / Start-to-Start / Finish-to-Finish relationships with optional lag, apply scheduling constraints (MSO, MFO, SNET, SNLT, FNET, FNLT), and view CPM critical path and float analysis. Set a Finish Target date in the summary card (each scenario keeps its own) for RAG schedule health tracking.
6	Simulate, report & export	Click Run Simulation. Review the histogram, CDF chart, and Percentile Summary (P5–P99, each with its projected calendar finish date). The RAG indicator reflects schedule health vs. your Finish Target. Copy the Gantt as an image for PowerPoint or Outlook — the copy always holds the whole chart; enable **Fit to Window** in the Gantt appearance panel to make a multi-year program compact. Compare scenarios side-by-side, run Sensitivity Analysis, and export the schedule as XLSX or CSV.

Note: Editing estimates, distributions, or the work calendar clears the scenario's saved simulation results — the summary card then prompts you to run again. Click **Run Simulation** to refresh everything, or turn on **Auto-run simulation** in Settings → Preferences (off by default) to re-run automatically about half a second after each change. Once a scenario records a start, work that has not started is scheduled no earlier than today (the next working day on a weekend or holiday), so results simulated on an earlier day can go out of date: they then show as not run — on the page, in print, in Compare and the schedule export, and to Connect AI — with a sentence saying why, until you run again. With Auto-run on, the open scenario runs again by itself.

ACTIVITY GRID — COLUMN GUIDE	
Activities bar	A bar above the grid reads **Activities (N)** with **Hide** / **Show**, so you can skip past every row to the charts. While collapsed it shows an amber '· N flagged' count, and clicking an activity in the validation summary reopens the grid at its row. The grid, Milestones panel, and Dependencies panel each remember their collapsed state per project, in this browser only.
Name	Activity name (up to 200 characters). Click to rename inline; a cut-off name shows in full on hover. A blank name is refused; a nameless activity reads (unnamed). Hover any row to reveal a pencil icon that opens the Activity Edit modal directly. The # column shows sequential numbers when Show Activity IDs is enabled (Gantt toolbar — persists per project). Status icon prefix: circle = Planned, filled square = In Progress, checkmark = Complete.
Min / ML / Max	Three-point estimates in working days, editable in the grid. The three cells save *together* when you leave them, so Tab between them freely: **Enter** saves and moves on, **Escape** puts all three back. An out-of-order estimate is saved and flagged; a cleared or negative one is not saved and stays red. Heuristic Mode fills Min and Max from Most Likely but never overwrites a value you typed.
Distribution	T-Normal, LogNormal ('LogNorm' in the grid), Beta-PERT, Triangular, or Uniform. **Hover for a chart** of the distribution the simulation samples from this row's own Min / ML / Max (and Confidence) — the estimate as entered, before Parkinson's floor or any in-progress conditioning. A **green** dashed line marks a Planned activity's scheduled Duration, so switching distributions shows why that number moves; a **blue** one marks Most Likely. The curve meets the baseline where a distribution ends at Min or Max, and is cut off where T-Normal or LogNormal run past them. While you are typing, it shows the numbers last saved. The amber suggestion dot — click to apply, hover for the reason — appears when your numbers fit T-Normal or LogNormal better, or contradict the chosen curve; it never shows on Beta-PERT or Uniform rows. The text turns grey when Min = ML = Max (the distribution cannot change that duration). Changing distribution with heuristics on prompts you to recalculate Min/Max.
Confidence	Subjective confidence label (Near certainty → Guesstimate), abbreviated in the grid (full wording on hover). **Applies to T-Normal, LogNormal, and Beta-PERT.** Shows a dash (—) wherever it cannot change the activity — Triangular and Uniform rows, or Min = Max — and the saved level returns if the activity changes. Hovering a level shows the same chart as Distribution; hovering a dash says why. Blank for Triangular and Uniform in exports.
Scheduled	Calculated start and end dates for this activity based on the deterministic schedule (the Start and End columns). Respects work week, holidays, converted work days, and constraints. A recorded Actual Start or Actual Finish that the schedule uses as entered shows in **green**, as the columns' Sched. / Actual subheading says.
Duration	Deterministic scheduled duration = inverseCDF(activity target, e.g., P50). Shown in working days.
Status / Actual	Set activity status: Planned, In Progress, or Complete. Setting a Planned activity **In Progress** records today as its **Actual Start** (on a weekend or holiday, the last working day before) and, when the schedule can be calculated, keeps its **Planned Start** and **Planned Finish**; setting it **Complete** records its **Actual Finish** the same way — or, if it is still Planned and its planned finish has come, its planned dates as the actual ones. The **Actual** cell holds the working days done: for In Progress, counted when you set the status — it does not count on by itself, so update it (or clear it to recount); for Complete, the total. Change the actual ones in the Activity Edit modal → General. The simulation conditions each in-progress activity's sampled duration on the elapsed time already observed, so the buffer reflects an overrunning activity honestly rather than assuming it is nearly done. A bulk status change gives each activity the dates it would get on its own.

Progress notes	Once a scenario records a start, work that has not started is scheduled no earlier than today, so while such work drives the finish, the forecast moves on with the calendar; started work is never held back for work before it that has not started. A grey **Progress notes** box above the summary card explains such moves — for example, which activities were moved to today, or that one started before its predecessor finished. They are notes, not errors, and they print with the report.
Task progress bar	Thin bar beneath the activity name (blue = tasks in progress, green = all tasks complete). Click to open the Activity Edit modal → Tasks section. Up to 50 tasks per activity.
Deliverables bar	Second thin bar (indigo = deliverables in progress, green = all delivered). Shown alongside the task bar when both are present. Click to open the Activity Edit modal → Deliverables section. Up to 50 deliverables per activity.
Notes indicator	A small violet bar beneath the activity name (alongside task and deliverable bars) indicates the activity has notes. Click to open the Activity Edit modal → Notes section.
Constraint	Clickable badge showing the constraint type (MSO, SNET, etc.) and Hard/Soft mode. Opens the Scheduling Constraint section of the Activity Edit modal. Always shown in Dependency Mode. In sequential mode the column appears as soon as at least one activity carries a constraint — the sequential scheduler honors constraints too, so a constrained date is no longer explained only in the Gantt chart or the edit modal. Once shown, every row gets a badge or a clickable — affordance to add one.
Section headers	Bold divider rows inserted via the **+ Section** button (next to + Add Activity). Each section has an editable name, an optional color (applied as a faint tint to the section header row itself), and a delete button. Up to 50 per scenario: once an unlocked scenario holds 50, **+ Section** is greyed out and a line under it says so. Locked scenarios prevent changes. Carry no scheduling logic and do not affect simulation results.
Validation summary	An amber list above the grid names each flagged activity — '#4 Activity name: Min is above Most Likely' — and clicking a name jumps to its row. Flags cover out-of-order, cleared, and negative estimates, and a LogNormal activity at 0 / 0 / 0. While anything is flagged, every Run control refuses and names what to fix (by activity number when the project shows numbers); if the schedule cannot be calculated, a red panel links to the activity and shows your three numbers.

ACTIVITY EDIT MODAL

Open by clicking a Gantt bar, the task or deliverables progress bar, the constraint badge in the grid, or by hovering any activity row and clicking the pencil icon. The modal has up to eight collapsible sections. Estimates section is collapsed by default. Blue dot on section headers (Scheduling Constraint, Dependencies) and violet dot on the Notes section header indicate content is present. **Save** / **Cancel** stay pinned while only the sections scroll. **Cancel** never saves: with unsaved edits it asks **Keep editing** or **Discard**. **Escape** or a click outside offers **Keep editing**, **Discard**, or **Save**; the keyboard starts on Keep editing, so Enter never discards or saves. When Save is greyed out, a line beside it says why — “Save needs a name for this activity.”, “…every estimate to be 0 or more.”, “…a date for the scheduling constraint.”, or “…the actual dates corrected, as the line under them says.” — one reason at a time, and the discard prompt gives the same reason, so a blocked field never silently discards your other edits. If the activity is changed somewhere else while the window is open — by a collaborator or by Connect AI — an amber line under the title says what changed, and **Save** writes only the parts you changed.

1 · General	Activity name and status (Planned / In Progress / Complete), then a full-width **Description** (up to 2,000 characters, with a live counter): a plain-language summary of what the activity entails, distinct from Notes. Description is the shareable “what this activity is” — it appears in the CSV and Excel schedule exports and in the printed report. Editing it invalidates simulation results (Notes, tasks, and deliverables deliberately do not). Below it, a line that follows the Status — **Planned:** Sched. Start, Sched. Finish, Sched. Dur. (Actual Dur. greyed out); **In Progress:** Planned Start, Actual Start, Sched. Finish, Sched. Dur., Actual Dur.; **Complete:** Actual Start, Planned Finish, Actual Finish, Actual Dur. Only the Actual fields can be edited. An actual date cannot be after today, and one entered on a non-working day moves to a working day (a start forward, a finish back); changing it recounts Actual Dur. on the project's work calendar, and on a Complete activity with an Actual Start a count you type moves Actual Finish. **Planned Start** and **Planned Finish** are the dates the plan gave the activity when it left Planned (before work not yet started was moved to today), or those it was imported with; the schedule never moves them. With none kept (an activity started before v0.81.0, say), In Progress shows a dash and says why, and Complete shows its Sched. Finish. The activity's progress notes appear under the line.
2 · Estimates	Min, Most Likely, Max, Distribution type, Confidence level, and Activity probability target (P30–P95). Heuristic Mode auto-recalculates Min/Max when ML changes. A line under the estimates warns of an out-of-order value before you save (it saves as typed and the grid flags it); a negative estimate disables Save and says why; an emptied field is flagged, so a cleared estimate never looks applied. Confidence shows a dash as soon as you pick Triangular or Uniform. The section opens by itself when the saved estimates are flagged.
3 · Scheduling Constraint	Select constraint type (MSO, MFO, SNET, SNLT, FNET, FNLT), date (non-working days snap forward), Hard or Soft mode, and an optional note (≤ 500 chars). Live conflict preview shown. A Type with no date says “Choose a date, or click Clear constraint to remove it.” under the Date field. Blue dot on section header when a constraint is set.
4 · Dependencies	Dependency Mode only. Add predecessors/successors directly in the modal. Choose relationship type (Finish-to-Start / Start-to-Start / Finish-to-Finish) and optional lag days. Negative lag (lead time) is supported — enter a value like -5 to overlap activities. Lag values are clamped to ±365 days on save. Add / Edit opens the dependency window on top of the activity, so your unsaved edits survive. Blue dot when dependencies exist.
5 · Tasks	Add, toggle, reorder (drag), and delete task checklist items. Up to 50 tasks. Progress count shown in section header. A task typed in the box but not yet added is added when you click Save (a deliverable too).
6 · Deliverables	Track deliverables (documents, sign-offs, artifacts) with the same checklist-style UI as Tasks. Up to 50 deliverables. Progress count shown in header. Deliverables appear in schedule export and print report.
7 · Notes	Free-text activity notes field (up to 2,000 characters), opening at 5 lines with a drag handle in the corner to resize it taller. Violet dot on section header when notes are present. Notes are for reference only and do not affect scheduling or simulation.
8 · Schedule Analysis	Dependency Mode only. Read-only. Shows this activity's Duration, Total Float and Free Float, and no dates (they are on the General line). An activity on the critical path reads “Critical path — 0 days float”.

SCHEDULING CONSTRAINTS (DEPENDENCY MODE)

Scheduling constraints impose date boundaries on an activity's start or finish. Apply them via the Activity Edit modal → Scheduling Constraint section, or by clicking the constraint badge in the activity grid. Hard constraints adjust the schedule and are enforced in every Monte Carlo trial. Soft constraints flag violations without forcing the schedule.

Type	Full Name	Description
MSO	Must Start On	Activity must start on exactly the specified date.
MFO	Must Finish On	Activity must finish on exactly the specified date.
SNET	Start No Earlier Than	Activity cannot start before the specified date.
SNLT	Start No Later Than	Activity cannot start after the specified date.
FNET	Finish No Earlier Than	Activity cannot finish before the specified date.
FNLT	Finish No Later Than	Activity cannot finish after the specified date.

Hard mode	The scheduler adjusts start/finish dates to satisfy the constraint. Each Monte Carlo trial also respects the constraint — buffer calculations reflect the constrained schedule. A hard **Must Finish On** or **Finish No Earlier Than** constraint simulates the activity finishing *on* the constraint date, matching the deterministic schedule and every calendar-date display exactly.
Constraint delay	Hard date constraints (and milestone start floors in Dependency Mode) can park a schedule on idle working days that are not work content. Those days are reported as their own figure in the Scenario Summary Card, the printed report, and the schedule export — so Duration + Constraint delay + Schedule Buffer = Duration w/Buffer. The figure appears only when there is idle time to explain; unconstrained scenarios look exactly as before.
Soft mode	The constraint date is advisory only. Violations are flagged in Scheduling Warnings with an amber indicator, but the schedule is not adjusted.
Scheduling Warnings	Collapsible panel above the scenario's summary card listing all constraint conflicts (hard — red) and soft violations (amber) with plain-language explanations.

SCENARIO FEATURES

Baseline scenario	Every project starts with a Baseline scenario representing your original plan. Clone it to model alternatives. Any scenario can be deleted — the last remaining scenario in a project is the only one protected from deletion. A project holds up to **50 scenarios**, which comfortably supports adding one per week as a rolling snapshot of an active project for more than a year.
Scenario tabs	Clone any scenario to create alternatives (e.g., optimistic, pessimistic). Scenario tabs scroll horizontally when more scenarios exist than fit in the view — the Compare button stays pinned on the right edge. Drag the grip handle on any scenario tab to reorder tabs; double-click a tab to rename it. Each tab is reachable by keyboard in the normal tab order, and the active scenario is announced to screen readers. Lock/unlock a scenario to prevent accidental edits.

Cloning & reforecasting	A clone sits immediately to the right of the scenario you cloned (the + button's Add Scenario window, which also copies a scenario, does the same). A copy keeps its source's random seed, so an unchanged clone reproduces its source's results exactly when both are run — any difference then comes from your change (**New** beside a seed still draws a fresh one). The Clone window offers a new name each time, such as 'Baseline (Copy 2)'. Its **Drop completed activities (reforecast)** box copies only the work left: Complete activities are left out, each In Progress one keeps its estimates less the days in its Actual cell, and the copy starts on the **Reforecast from** date (the next working day by default). A Section stays above its activities in the copy, even when the first of them was Complete; a Section whose activities were all Complete stays, empty.
Finish Target date	Set a target finish date for each scenario (e.g., the committed date from your project charter) directly in the Scenario Summary Card. A scenario keeps its own date: a clone starts with its source's, and a scenario that Import Activities adds starts with none. The date saves when you leave the field — a half-typed date is not saved and the field returns to what it held; a locked scenario's field is greyed out. A dashed vertical line marks the active scenario's target on the Gantt chart (toggle via **Finish Target** in the Gantt toolbar — one choice for the whole project; the chart timeline extends automatically if the target date falls past the project buffer end).
RAG schedule health	After simulation, the Finish Target date and Gantt line are color-coded: Green = simulation finishes by the green percentile threshold (default P80), Amber = within the amber threshold (default P50), Red = at risk. Thresholds are configurable in Settings → Finish Target — Schedule Health.
Scenario Notes	Free-text notes field per scenario (up to 2,000 characters). Access via the memo icon in the Scenario Summary Card. A blue dot on the icon indicates notes are present.
Parkinson's Law toggle	Per-scenario toggle (default: on) in the Scenario Summary Card. When on, no simulated duration falls below the deterministic duration set by your Activity target (P30–P95 — not always P50). Disable to model optimistic scenarios where activities may finish early — except an In Progress activity, which can never sample shorter than its elapsed time even with the toggle off. An amber note appears when disabled. The default is configurable in Settings → Preferences.
Dependency Mode	Enable to add Finish-to-Start / Start-to-Start / Finish-to-Finish relationships with optional lag, apply scheduling constraints (Hard/Soft), compute CPM critical path (highlighted in Gantt), and surface Total Float and Free Float per activity. Float appears in the Activity Edit modal → Schedule Analysis, in Gantt hover tooltips, and in schedule exports.
Milestones	Dependency Mode only. Named milestones (a name is required) with optional target dates. Assign activities to milestones; the simulation computes buffer days per milestone. Health reads the same on the summary card, Milestones panel, and printed report: **On Track** (5+ working days of slack), **At Risk** (0–4), **Late** (negative). A grey dash means no result yet — hover for what to do. Milestone markers appear on the Gantt chart.
Scenario comparison	Compare 2–3 scenarios side-by-side: dates (Finish Target right after Start Date), duration, buffer, and a CDF overlay chart. Tick a scenario's box to add it (ticking never switches the scenario on screen; once three are ticked the rest grey out). An unrun scenario shows dashes and takes no part in Duration w/Buffer or Mean; a scenario saved with validation errors reads '1 flagged' and is never marked best. Copy-image buttons on the metrics table and CDF chart let you paste a clean PNG directly into slides or documents (Firefox-aware: disabled with tooltip on unsupported browsers). A **Run** button under any compared scenario that has no results and can be run runs it in place (its own trials and seed; the tab on screen does not change). A scenario whose sample data was not stored draws no curve until it is run again, and a note under the table says so. Access from the scenario toolbar.

Buffer used rows	When a later scenario records more progress than the first — more activities In Progress or Complete, or more working days in their Actual cells — three rows after Duration w/Buffer — Buffer used (days), Buffer left (days) and Buffer used (%) — show how much of the *first* scenario's buffer that status update has used: the working days its End Date (no buffer) has moved past the first's, below 0 meaning days were recovered. A note under the table names the scenario they are measured against or says why a cell is blank; drag a tab to the left to make it the reference. The Finish Target and buffer rows print too.

GANTT CHART

Views	Two views: **Deterministic** (scheduled dates) and **With Uncertainty** (bars extended to a Monte Carlo percentile). Switch via the toolbar.
Appearance panel	Click the palette icon (Appearance settings) at the right end of the Gantt toolbar: the Appearance Panel opens right under the toolbar. Controls: **Name Column** (Narrow / Normal / Wide), **Font Size** (Small / Normal / Large / XL), **Row Size** (Compact / Normal / Comfortable), **Bar Label** (Duration / Dates / None — auto-hidden when bar is too narrow; dates read 'Oct 21', no year), **Timeline Labels** (Sparse / Normal / Dense), **Fit to Window** checkbox, **Non-work day shading** (gray bands for weekends and holidays), and **Row guide lines** (faint horizontal lines every 3 rows — on by default). All settings are saved with the project and survive export/import round-trips. Print Gantt mirrors all appearance settings.
Color presets	Ten built-in themes: Classic (default blue/orange), Professional, Colorful, Grayscale, Contrast, Forest, Ocean, Sunset, Lavender, and Earth — each with light and dark mode variants. Override Planned, In Progress, or Completed bar colors with any custom color via the swatch picker or native color input. Custom colors clear when switching presets.
Timeline Labels (multi-year)	For projects spanning more than 18 months (>540 days), controls tick density: **Dense** — monthly ticks (Apr '26, May…) · **Normal** — quarterly ticks (Q1 '26, Q2, Q3, Q4, Q1 '27…) · **Sparse** — semi-annual ticks (H1 '26, H2, H1 '27…). Year labels appear on the first tick and at every year boundary (shown in bold). For projects ≤18 months, ticks auto-select daily / weekly / biweekly / monthly based on span. Milestone names, dates, and the timeline sit in separate rows; crowded names alternate between two.
Interactive actions	Click a Gantt bar → Activity Edit modal (even where an arrow overlaps it; very short bars get a wider target). Click a dependency arrow in open space (Dependency Mode) → Dependency Edit modal (change type, lag, or delete). Click an activity name → opens that activity after a brief pause; **double-click** it to rename inline (Escape cancels). Hover arrows for visual highlighting. A **dashed** arrow is a link an actual date overtook (the activity started or finished earlier than the link allowed): it is never drawn as critical, and the legend, on screen and in print, adds “Link overtaken by an actual date”. Terminal markers (right-edge stripe) identify activities with no successor. In sequential mode (Dependency Mode off), the Gantt auto-draws non-interactive finish-to-start arrows to make the implicit ordering visually explicit.
Sections	Each Section's row is shaded in its color (grey if it has none), with its name in bold. Once it holds activities, a bar on that row runs from their first start to their last finish: hover it for those dates and how many activities it holds. The printed chart shows the shading and the bar too. Click a Section's name to rename it inline (not in a locked scenario).

Toolbar toggles	Today line · Project name · Show Activity IDs (#1, #2… — persisted per project) · **Finish Target** (dashed RAG line — always visible when toggled on, extending the chart timeline if the target date falls past the project buffer end; greyed out when the active scenario has no target date — the tooltip says to set that scenario's date in the summary card, or to unlock the scenario first if it is locked; clearing a date switches the toggle off only when no scenario has a date left) · Dependency arrows show/hide (available in both modes — hides auto-drawn arrows in sequential mode) · Critical path highlight · Palette (appearance panel).
Copy & print	Copy the Gantt as a PNG to clipboard (disabled with tooltip in Firefox). The copy always includes the whole chart, however wide your window or wherever it is scrolled, and in dark mode it is the light chart. Print the full project report — Gantt mirrors all appearance settings including layout dimensions, bar colors, font scaling, weekend shading, and bar labels. Printed activity and section names use the full width of the name column and scale with your Name Column and Font Size settings, so only genuinely over-long names are shortened. A 'Date prepared' label appears in the bottom-right corner.

■ **Tip:** For a presentation-ready multi-year Gantt: enable Fit to Window for a compact one-screen chart, set Timeline Labels to Normal (quarterly) or Sparse (semi-annual) depending on project length, then click Copy as PNG. The complete chart — all activities across the full timeline — copies to clipboard in one shot. The Appearance Panel settings are stored on the project, not globally — each project can have its own layout and theme.

CALENDAR, WORK WEEK & HOLIDAY LOADER

Work week	Click interactive day pills to toggle work days on/off — any combination including non-contiguous schedules (e.g., Mon/Wed/Fri). A warning appears if the configuration produces no valid work days.
Country holiday loader	Load public holidays from 100+ countries via the Nager.Date API. Multi-country support: holidays from multiple countries are merged additively. Built-in US holidays serve as an offline fallback.
Manual holidays	Enter custom non-work dates (e.g., company shutdowns, regional observances) with an optional locale label. Smart validation warns if a date is already a work day or falls on an API holiday.
Converted work days	One override editor per project handles both cases: mark a weekend as a working day, or make a company holiday a working day for this project only. Adding a company-holiday date asks for confirmation first; if the holiday spans several days, a **Convert all N days** button converts the whole range in one click (weekend days inside the range are never silently converted). Project-specific holidays (added on this project's own holiday list) cannot be converted — you're guided to remove them from that list instead. Each override chip shows live status: it grays out if a holiday later lands on that date, with a one-click action to restore it as a forced override when possible. An amber warning appears if the list exceeds 50 entries.
Project vs. global	Calendar settings (work week, converted days) are per-project. Holiday lists are global across all projects. All scheduling, Gantt charts, and exports respect the active calendar configuration. Editing a project's own calendar clears the saved simulation results of *every* scenario in it; editing the global Company Holidays does not.

DATA, STORAGE & EXPORTS

Local storage	All project data is stored in your browser's localStorage by default — nothing leaves your device. No analytics, no telemetry, no third-party data collection. An amber caution banner appears on every page while data is stored locally, reminding you to export before closing. Dismiss with **Got it** for the session, or suppress it permanently via Settings → Notifications. Clearing browser site data permanently deletes local-only projects.

Cloud storage	Enable by clicking the auth chip in the app header to open the **Storage & Sign In** modal, or navigate to Settings → Cloud Storage. Sign in with Google (personal or work accounts) or Microsoft (work or school accounts only; personal Microsoft accounts are refused). Cloud mode syncs projects to Firestore across all your devices. Share projects with teammates from Settings → Sharing: add collaborators by email individually, or paste multiple email addresses at once for bulk invitations — existing SPERT Scheduler users are added immediately; new users receive an invitation email and claim the project on sign-in. Owner, editor, and viewer roles are supported. Requires acceptance of the current Terms of Service and Privacy Policy (the October 26, 2026 editions — you are asked to accept them at your next sign-in; using Scheduler without signing in is unaffected). If you are signed out for not having accepted them, the projects and settings this browser keeps for use without signing in are not erased.
Loading & signing out	While your projects load, the Dashboard reads “Loading your projects from cloud storage…” — it no longer shows this browser's signed-out projects meanwhile — and New Project, Load Sample, a tile's Clone, Import Activities, Import, and Choose File in Settings → Import Projects are greyed out with a note saying why. If Scheduler cannot finish loading, the Dashboard says some projects may be missing and to reload the page. Signing out asks you to confirm first — it explains that this device's locally-cached copy of your cloud projects will be removed (your cloud data is unaffected and reappears on next sign-in); this browser's own projects and settings then come back straight away. If a shared project was last saved from a newer app version than the one you're running, it's set aside and shown in a "Newer version" notice instead of being opened with data missing — update the app to restore access; nothing about the project is harmed. After any update, reload every tab that was open before it: a tab on an older version cannot open projects a newer one has saved, and without cloud storage the tab that saves last decides what is kept.
Export / Import (JSON)	Export any single project as a JSON file for backup or transfer — hover a project tile on the dashboard and click its download icon for a one-click export, no trip to Settings required. **Export All** (beside the Projects heading) exports every active project in one click. **Import**, beside it, opens the file window straight away and restores one or more projects from JSON — including all scenarios, activities, constraints, dependencies, tasks, deliverables, notes, section headers, and Gantt appearance settings; the file's preview, with Confirm Import and Cancel, or what is wrong with it, shows below the buttons (Settings → Import Projects works the same way). Every exported filename includes a date and time stamp, so multiple exports from the same day are easy to tell apart.
Import checks	Import detects both ID and name conflicts, each with per-conflict resolution (skip / replace / copy — ID conflicts default to skip, name conflicts to copy). Choosing to keep both gives each copy a distinct name ('{original} (Copy)', then '(Copy 2)'), even when two same-named projects arrive in the same file. An opt-in preferences import toggle (default: off) applies imported preference settings. The import preview also checks each file for a **broken dependency graph** — a circular chain, a link to an activity not in the file, or an activity depending on itself — and names the affected projects and scenarios before anything lands, with the count repeated in the summary afterwards. Nothing in your file is altered: the projects import exactly as they are, and skipping one remains your choice. Affected scenarios cannot compute a schedule until the problems are fixed. Projects holding an unnamed item or an out-of-order estimate import and open normally, flagged. A project the Projects page *could not load* can still be exported: its stored data goes into an export file exactly as stored, which Import Projects reads and checks like any other and accepts once what is wrong is fixed. Refusals name the project and the field at fault (e.g., scenarios.0.activities.0.name).

Import from spreadsheet	Settings → Import Activities from Spreadsheet. Accepts a CSV file or clipboard paste, with a live preview that names each row with a problem. Download the template for format guidance; the **Import guide (PDF)** beside it covers the columns, date forms and messages, and the **P6 & MS Project Import Guide** on the About page shows how to bring in a Primavera P6 or Microsoft Project schedule. Beyond the template's columns, it reads Actual Start, Actual Finish, Planned Start and Planned Finish, so a project already underway comes across with its recorded progress. A Description column is ignored. Import as a new project or add as a new scenario to an existing project. The new scenario takes your heuristic preferences, has no Finish Target date, and starts today — or on the earliest Actual Start it imports.
Schedule export (XLSX/CSV)	Download the activity schedule from the Scenario Summary Card buffer row when the scenario has current simulation results. XLSX: bold headers, frozen row, auto-width columns, metadata summary block (Constraint Delay alongside duration and buffer, and the Date Format of its dates). After End Date come **Planned Start** and **Planned Finish** (for an activity not yet started, its dates in the plan before work was moved to today), then **Actual Start** and **Actual Finish**. Dependency Mode adds Predecessor/Successor columns (e.g., 1FS+2d, 2SS, 3FF) and float columns (Total Float, Free Float); Tasks and Deliverables columns (summary + details) are always included. Constraint columns (type, date, mode, note) are included in Dependency Mode, and in sequential mode once any activity carries a constraint — matching the grid. **Description** is the final prose column, immediately before Type, and **Distribution** precedes **Confidence**. Section headers appear as labeled rows. The export is for reading and sharing: it cannot be imported back as it is, so use JSON export to move a project.
Print report	The printer icon in the project header prints the full report. Validation errors or an uncalculable schedule are explained at the top of page 1. With Compare showing two or three scenarios, a **Scenario Comparison** section (table, notes, CDF) comes first. Progress notes print in their own section after Constraints, and out-of-date results print as not run, with the sentence saying why. On-screen notices never print, and dark mode prints clean white.
Simulation CSV export	Export raw Monte Carlo trial results from the Simulation Results panel. One row per trial — useful for external analysis or audit.
Preferences	Settings → Preferences: default distribution (any of the five, Beta-PERT included), confidence level, activity (P30–P95) and project (P50–P99) probability targets, trial count, heuristic percentages, Parkinson's Law default, **Auto-run simulation** (off by default), date format, and RAG schedule health thresholds. Settings → Export Schedule: configure schedule export options. Settings → Notifications: suppress the localStorage caution banner permanently.

Note: Saved results are cleared whenever an edit could change them. After bulk-editing, changing calendar settings, or importing a project, click Run Simulation in the Scenario Summary Card to refresh them (or enable Auto-run simulation in Settings → Preferences).

SIMULATION RESULTS & CHARTS

Histogram	Project duration distribution across all Monte Carlo trials. Buffer shading highlights the range between the deterministic finish and your project target percentile. Parkinson's Law clamping is reflected in the distribution shape.
CDF chart	Cumulative distribution function. Read directly: 'There is an X% probability of finishing within Y days.' Hovering over the curve shows the projected calendar finish date alongside duration and probability. Enter a target finish date in the field above the chart to see the probability of finishing by that date, with a RAG-colored reference line (green/amber/red based on your Schedule Health thresholds). The project target percentile is marked. Scenario comparison CDF tooltips also include projected finish dates.

Percentile Summary	P5 through P99 in working days, each with its projected calendar **Finish date** in a second column — same working-day, holiday-aware math as the rest of the schedule, following your date format preference. The target percentile row's date matches **Finish w/Buffer** on the summary card. (The date column appears whenever the scenario has a start date.) A copy-as-image button in the panel header drops the whole table into a status report or email in one click. Use the table to select a project target that matches your risk tolerance.
Bootstrap CI	Optional 95% confidence interval on each percentile estimate. Enable via the toggle in the Simulation Results panel. Shows simulation precision.
Sensitivity analysis	Ranks each activity on its own estimates alone — sort by Impact, Variance Contribution, or Relative Uncertainty — using the mean and SD of its own distribution (a Triangular 10 / 15 / 30 shows 18.3d and 4.2d at every Confidence level). It ignores dependencies, Parkinson's Law, and status; in a dependency-driven scenario a note warns that an activity with slack can rank high without moving the finish. Out-of-order activities (and an all-zero LogNormal) are left out, and a line counts them. Available for projects with 2+ activities.
Scenario comparison	Compare 2–3 scenarios side-by-side (dates, duration, buffer) plus a CDF overlay chart for visual comparison; the dashed line marks the Project target. Ticking rules: see Scenario Features.

■ **Tip:** Press **?** anywhere to open the shortcuts panel: Ctrl/Cmd+Z undo, Ctrl/Cmd+Shift+Z or Ctrl/Cmd+Y redo, Tab / Shift+Tab through the grid, Enter / Escape to confirm or cancel a name or estimate edit. In a text box, Undo acts on your typing (except in Scenario Notes), and while a window such as the Activity Edit modal or the Connect AI panel is open, Undo and Redo never change the page behind it. Undoing or redoing a change to the project clears the simulation results in all its scenarios, so run the simulation again.

SAMPLE PROJECT

Load Sample — beside the Projects heading, and offered as **Load Sample Project** when you have no projects yet — builds a complete worked example in one click. It is an ordinary project in every respect: open it, edit it, clone it, export it, or delete it.

What you get	**Cloud ERP Solution** — a vendor-neutral cloud ERP implementation led by a systems integrator, starting after contract signature and running roughly fourteen months. It ships 40 activities across eight color-coded sections, 68 dependencies, and 4 milestones. Every activity carries a description, two deliverables, three checklist tasks, and notes covering completion advice and common pitfalls. A US holiday calendar spanning the schedule is attached and fully editable.
Why it is worth opening first	The sample is built to exercise the product rather than merely populate it. Dependency Mode is enabled, and the Financials, Order Management, and CRM tracks run concurrently before re-converging at system integration testing — so the Gantt chart, critical path, and schedule buffer all have non-trivial behavior to show. Four of the five distributions appear (no Beta-PERT), matched to the character of the work: LogNormal across data migration, integration, and test, where real programs develop long tails, and Uniform on the two externally-gated activities.
Always current, always safe	The scenario starts on the first Monday from today that is not one of its holidays, and milestone targets are derived as offsets from that start, so the sample always reads as a current plan rather than a frozen fixture. Loading it more than once is safe — each copy is completely independent, with fresh identifiers throughout.

CONNECT AI

Connect AI lets an external AI assistant read and modify the project you have open, through the shared SPERT® Suite MCP (Model Context Protocol) server. SPERT® Scheduler never calls an AI API itself: the assistant writes operations into a per-browser-session document, and your browser applies each one *in order* through the same validation path as a manual edit. Both local-only and cloud-synced projects are supported. Start from the **Connect AI** button in the project header (it becomes **AI** with a status dot while a session is live).

Connecting	A consent screen opens first: neither box is ticked for you, and **Connect** waits until you tick Write Mode. The session panel then shows a session code in word-and-number form (for example, NEITHER · 0384216791). **Copy Code** copies the code alone; **Copy Prompt** copies a ready-made instruction block that also tells your assistant which tools exist and how to use them. Paste it into your AI chatbot — it first asks what the project is, whether activities can run in parallel, and tells you where to switch on Dependencies and Read. Codes are valid 15 minutes and refresh automatically. The panel shows **Waiting for AI** until your assistant pairs, then **AI connected**.
Two-tier consent	**Write** is required for a session: the AI can create and edit activities, estimates, milestones, section headers and notes in this project, and, with Read, update each activity's progress. **Read** is optional and can be turned on or off at any time from the session panel — it lets the AI read the full schedule, *including section headers, activity notes, descriptions, checklist and deliverable text, and computed, recorded and planned dates*, so it can make context-aware edits. Read is required for dependency edits, reordering activities and updating progress. Scheduler asks for consent again after a material change to what the AI may do, as it did at v0.84.0.
What the AI can do	Create activities (optionally seeding a description and an initial note); rename them; change three-point estimates, confidence level, and distribution type (if unnamed, the app picks T-Normal, LogNormal, or Triangular from the numbers — Uniform and Beta-PERT only on request); set or clear the description; append notes; add and toggle checklist and deliverable items; create milestones, edit their name and target date, and assign or unassign activities; create, update, and remove dependency edges; reorder a scenario's activities; add, rename, recolor, move, and delete **Sections**; and, with Read, record each activity's **Status**, **Actual Start**, and **Actual Finish**. A date left out is the day the request was sent, or the working day before it if that is not one; Planned to Complete needs both dates; and the AI reads today from your computer, never its own clock. Bulk equivalents handle up to 100 items per call (500 dependencies or assignments, 50 Sections), and a single import call can build an entire schedule — activities, milestones, assignments, and dependencies together.
What the AI cannot do	By design it cannot delete activities, milestones, scenarios, or projects — removing a dependency edge and removing a Section header are the only deletion-shaped exceptions. Progress is add-only: it cannot set an activity back to Planned, reopen a Complete one, change a recorded date, or set a Planned date. It cannot create or clone scenarios, edit scheduling constraints, write scenario settings (dependency mode and lock state are read-only), or trigger simulation runs. Those remain yours.
Guardrails	Every AI edit passes the same validation as a manual one. Dependency work requires Read Mode *and* a scenario with Dependency Mode already enabled by you — re-checked in your browser at the moment of applying, so an import carrying dependencies is declined whole rather than half-built. Reordering is all-or-nothing and is rejected outright if the list no longer matches your scenario. Within a bulk call, items apply independently: anything skipped is reported with its reason. A flagged activity's name, description and estimates accept only a repair.
Undoing AI work	Close the Connect AI panel first: Ctrl+Z does nothing while it is open, and the Undo button sits behind it. One Ctrl+Z then takes back everything the AI did in that stretch, which ends at an edit of your own or a pause of 5 minutes or more. Like every Undo, it clears the simulation results in every scenario, so run the simulation again: Auto-run re-runs at most the scenario on screen, and not at all after an Undo of Sections alone.

Watching it work	**Recent AI activity** in the session panel lists each operation as it lands, with per-item outcomes and a summary by part for imports — so a chain of skipped items reads as one root cause rather than many failures. A Section delete's row names what went, for example “Deleted from Baseline: Phase 1, Testing”. **Disconnect** asks first (“End this AI session?” — the AI can no longer change the project and the feed is cleared), then ends the session. The × in the panel's title row, Esc, or a click outside only closes the panel; the session keeps running. Signing out ends it.

Note: Do not switch projects while the AI is building — the session is bound to the project you had open. Note-appending is not idempotent: running the same append again adds the notes a second time. Reordering a **sequential** scenario changes start and finish dates; in a dependency-driven scenario it changes display order only. Either way, simulation results are cleared, and a Section stays with the activity it starts at, so a reorder can change what it groups. Sections are display only: adding, changing or deleting one moves no date and keeps simulation results. Recording progress clears that scenario's results.

■ **Tip:** Always use **Copy Prompt** rather than pasting the code alone — it gives your assistant the full tool list and the batching rules, which is the difference between a schedule built in a handful of calls and one built activity by activity. For a printable walkthrough, see the **Connect AI Guide** on the About page; the suite-wide **AI Privacy Notice** is linked in the footer.

TIPS & BEST PRACTICES

Activity status workflow	Mark an activity In Progress the day work begins and Complete the day it ends, so its recorded dates are right (correct them in the Activity Edit Modal → General section). At each status update, bring each In Progress activity's Actual count up to date and run the simulation again. The simulation uses actuals for completed activities and conditions on elapsed time for in-progress ones, so a long-running activity's buffer widens honestly the longer it runs instead of staying optimistically "almost done." If an activity has run far past its original estimate, treat that as a cue to re-estimate it or split it into smaller pieces.
Parkinson's Law toggle	Keep Parkinson's Law on for baseline planning (work expands to fill available time). Turn it off to model an optimistic scenario where activities may finish ahead of schedule. Compare the two scenarios to understand the planning assumption's impact.
Constraint best practices	Use Hard constraints only for genuine external date commitments (regulatory deadlines, vendor delivery dates). Use Soft constraints for aspirational targets. Review Scheduling Warnings after adding constraints — hard conflicts must be resolved before the schedule is valid.
Choosing a distribution	Triangular is the simple default. With Beta-PERT, lowering Confidence moves the middle toward the far side of the range (10 / 12 / 40 days: 13 days at Near certainty, 23 at Guesstimate), and at Medium a whole project comes out shorter than as Triangular — compare the two in a cloned scenario. Hover the Distribution cell to see the shapes side by side with the green Duration line. Treat the amber suggestion dot as a prompt, not a verdict.
Description vs. Notes	Use **Description** for the shareable “what this activity is” — the sentence or two you would want to read back months later, or hand to someone picking the work up. It exports to CSV and Excel and appears in the printed report. Keep **Notes** for the private catch-all: risks, jottings, review comments. Tasks and Deliverables still cover structured breakdowns. One practical difference: editing a Description invalidates simulation results, because it is schedule-shareable content — editing Notes, tasks, or deliverables does not.
Activity Notes & Scenario Notes	Record assumptions, risks, or review comments in activity Notes (up to 2,000 chars) via the Activity Edit modal. Use Scenario Notes (memo icon in summary card) for scenario-level context — e.g., 'This scenario assumes vendor delivers by April 15.' A violet indicator beneath the activity name and a violet dot on the Notes section header show when activity notes are present. A blue dot on the Scenario Notes icon indicates scenario notes are present.

Bulk activity actions	Select activities via checkboxes to reveal the Bulk toolbar. Stage a Status (Planned / In Progress / Complete), Distribution type, and/or Confidence level change, then click Apply to apply all at once. Staging Triangular or Uniform greys out Confidence and drops any level you had staged; with no distribution staged, Confidence stays available. A bulk Status change gives each activity the dates it would get on its own, and marking several Complete at once is one step to undo. Delete and reset confirmations start on Cancel, so pressing Enter dismisses them.
Finish Target & RAG	Set a Finish Target date in the summary card — each scenario keeps its own, so a clone can test a different commitment — to activate RAG health tracking. Green = simulation P80 (or your configured threshold) finishes on or before the target. Amber = P50 meets the target. Red = at risk. Adjust the thresholds in Settings → Finish Target — Schedule Health to match your organization's risk appetite.
Scenario workflow	Start with the Baseline, then clone it to explore alternatives. Use Scenario Comparison to present the trade-offs. Drag scenario tabs to put them in a logical order. Lock scenarios you've approved to prevent accidental edits. Run each scenario before comparing them (Compare's Run button does it without leaving the tab). Because a copy keeps its source's seed, a clone differs from its source only by what you changed.
Sensitivity analysis	Run Sensitivity Analysis after building your full schedule to identify high-variance activities. Prioritize those for tighter estimation, additional risk mitigation, or closer tracking during execution. It ranks activities one at a time, so in a dependency-driven scenario confirm a candidate by changing it in a cloned scenario and comparing the two.
Section headers	Use the + Section button to group related activities visually. Assign a color for a tinted section header row in the grid; section headers appear in the Gantt as a shaded row with a bar across their activities' dates, in printed reports, and in CSV/Excel exports (as labeled rows; the schedule export does not import back). Export to JSON to preserve section headers across round-trips. An assistant on Connect AI can read and change Sections too; deleting one leaves its activities in place.
Gantt tooltips	Hovering a Gantt bar, the Finish Target line, or a dependency arrow shows a tooltip only after the cursor rests there for about a second and a half — sweeping across the chart won't trigger a flicker of tooltips. Pause deliberately on the element you want details about.
Reading the buffer honestly	When a scenario carries hard date constraints, read three figures together rather than the buffer alone: Duration is the work content, **Constraint delay** is the idle time you are waiting on external dates, and Schedule Buffer is your protection against variability. Only the third is risk cover — a large Constraint delay is a scheduling fact to renegotiate or design around, not contingency you can spend.
Working with Connect AI	Give the assistant a target before it starts — the activity list, the shape of the dependency network, roughly how long things take — and let it build in bulk rather than one item at a time. Turn on Read Mode when you want context-aware edits, dependency work, or recorded progress, and remember it shares activity notes and descriptions. Review **Recent AI activity** as operations land, keep the project you connected open, and export before a large AI-driven restructure: one Ctrl+Z takes back the AI's whole stretch of work (with the panel closed) and clears every scenario's simulation results, so a file on disk is the surer guarantee. Ask for Uniform or Beta-PERT by name — the app never picks either one from the numbers.

■ **Tip:** The Schedule XLSX export is the most complete artifact from SPERT® Scheduler — it includes activity dates, float values, tasks, deliverables, constraint details, and dependency relationships. Share it with stakeholders who do not have access to the app.

© 2026 William W. Davis, MSPM, PMP · GNU GPL v3. https://scheduler.spertsuite.com