ISIGHT LABORATORIES

HINDSIGHT 80/20 DAILY

One number a day. The same fatigue math that keeps wings on airplanes, running on your own record — entirely on your device. Now with the mission layer: what the spending is **for** — and, new in v2.6, whether it feeds you.

User's Guide · App v2.6.1 · From why to *i* — the wave continues.

BEFORE THE FIRST NUMBER

Put it on your home screen.

Hindsight Daily is a web app — nothing to download from an app store, nothing that phones home. Installing it just gives it an icon and a full screen. The icon on your home screen is the Hindsight mark — the triangle, the eye, the wave.

iPhone / iPad — open the app link in Safari, tap the **Share** button, then **Add to Home Screen**.

Android — open the link in Chrome, tap the ⋮ **menu**, then **Add to Home screen** (or **Install app**).

After the first visit it works offline — in a canyon, on a flight, anywhere. Entries wait on the device; nothing needs a connection except the one-time unlock.

And the app folds to fit: every section collapses to its title — tap any header to expand it. Entry sections start open; the deeper instruments start folded. Your layout is remembered, and it rides in your backups.

The whole idea, in a minute.

Wings don't fail from one terrible moment — they wear out from thousands of ordinary flexes. Engineers add those flexes up with an eighty-year-old method and read one number: how much of the structure's life has been spent.

This app points that method at a person. You give it one number a day — **how are you, 0 to 10**. Every swing between a good stretch and a hard one is one flex. The math (rainflow counting → Walker correction → S-N life curve → Miner's sum) does the rest, and the surprise it tends to find is the 80/20: **the everyday swing, not the catastrophe, spends most of the life.**

And since v2.0, the app carries the other half of the sentence: damage units measure what a life **costs** — missions name what it **buys**. New in v2.6, a third question completes the loop: does the mission **feed the airframe**?

TAB 1 · TODAY

Thirty seconds, honestly spent.

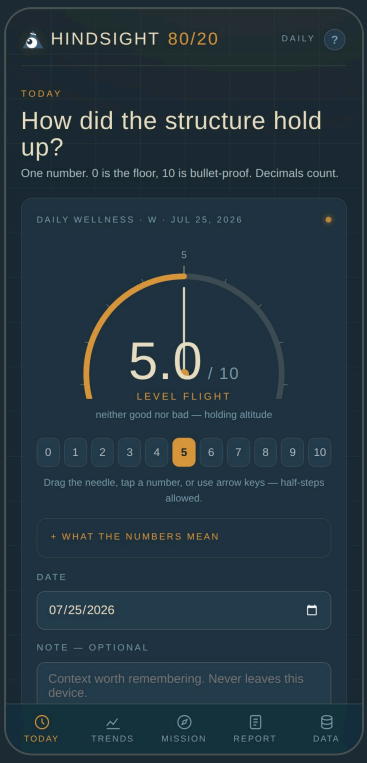

1. **Set the number.** Drag the needle, tap a chip, or use half-steps — 0 is the floor, 10 is bullet-proof. Tap **What the numbers mean** for the full scale in plain words. Don't overthink it; the method reads the pattern, not the precision of one day.

2. **Missed a day?** Change the date field and backfill. Gaps are fine — the math never counts a swing across a long gap.

3. **Note what mattered** — optional, never analyzed, never uploaded. It's for the future reader of your record: you.

4. **Save.** The record block below shows days logged, streak, and median W — and, once your record is long enough, a one-line read of your acute pace against your baseline.

Log more than once in a day — each save adds an entry; the day stores the average of all of them and keeps every distinct note. **Replace** starts the day over for corrections.

TAB 2 · TRENDS

Rate first. Level second.

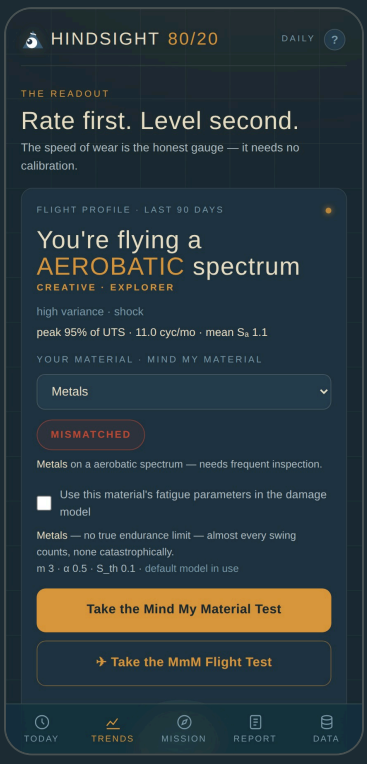

DD/DT · 30-DAY	The sprint — damage deposited in the last month. Sensitive, a little noisy.
DD/DT · 365-DAY	The marathon pace, and the honest gauge. It appears once your record spans a year.
DAMAGE INDEX D	The running total. Uncalibrated until you set your own anchor (Data tab) — the app says so under the number.
WELLNESS TRACE	Your W over any window — **30d / 90d / 180d / 1y / All** or any custom N. Remembered.
YEAR VIEW	January–December of any year — daily dots plus a trailing moving average you pick (7/14/30/90 or custom N days).
D(T) CURVE	The accumulation over the full record, live from every entry.
DAMAGE BANDS	Which size of swing is spending your life — amplitude is the lever.
STRUCTURE	New in v2.3 — a second life-model folded at the bottom of the tab; tap its title. Page 11.

THE PACE ALERT

When the last month breaks pattern, the app says so.

Every time you open Trends, the app compares your **acute rate** (damage in the last 30 days) against your **baseline pace** — the 365-day rate once you have a year of record, or the whole-record pace before that. Three outcomes:

Running hot — acute rate 2.5× your chronic pace. The last 30 days are spending life faster than your baseline. In airframe terms, the inspection interval just shortened. The lever is swing amplitude: open this month's Report to see what drove the big swings, protect the recovery days that break the pattern — and if this stays hot a few more weeks, say it out loud to someone you trust.

Easing — acute rate 0.4× your chronic pace. The last 30 days deposited less than your baseline. Whatever changed is working; worth a note in tomorrow's entry so future-you knows what it was.

Between roughly half-pace and double-pace, the app stays quiet — steady flight needs no callout. The same read appears as a one-line row on the Today tab, so a hot month is visible without visiting Trends.

The thresholds (2× hot, 0.5× easing) are engineering ones, not clinical ones — the alert is a gauge talking, not a doctor. It never sees a crisis; only a person can.

THE FLIGHT PROFILE

Every life flies a mission profile.

At the top of Trends, the app reads your last 90 days of cycles the way a loads engineer reads a usage spectrum — how big the swings, how often, how tired the baseline, how spiky — and names the profile you've been flying:

FIGHTER	CEO · Business — high amp · high rate · hot. Big swings, arriving often.
SAILPLANE	Monk · Lifestyle — low amp · low cycles · benign. The gentle spectrum.
TRANSPORT	Parent · Service — high mean · relentless. Moderate swings on a tired baseline, without pause.
AEROBATIC	Creative · Explorer — high variance · shock. Calm air broken by hard pulls.
Ferry flight	Fewer than six cycles in 90 days — too quiet to call. A reposition between missions; keep logging.

The material check. Pick your Mind My Material family (or take the free test — linked in the panel) and the app reads the pairing off the profile-vs-material matrix: **Thrives** — optimized to the load path; **Serviceable** — flies it on a maintenance schedule; **Mismatched** — needs frequent inspection.

Go further: fly as your material. An opt-in toggle feeds your material's own fatigue constants into the damage model — a Ceramic shrugs small loads but pays dearly for big swings; an Elastomer barely notices a tired baseline. Every material's scale is pinned so 1.0 means the same thing, and your own anchor recalibrates automatically when you switch. Off by default — the standard model stays the reference.

A playful crossover between two iSight instruments — not a verdict. The gauges stay the instrument; the profile is the story your last three months were telling.

TAB 3 · MISSION — NEW IN V2.0, REFINED THROUGH V2.6

What the spending buys.

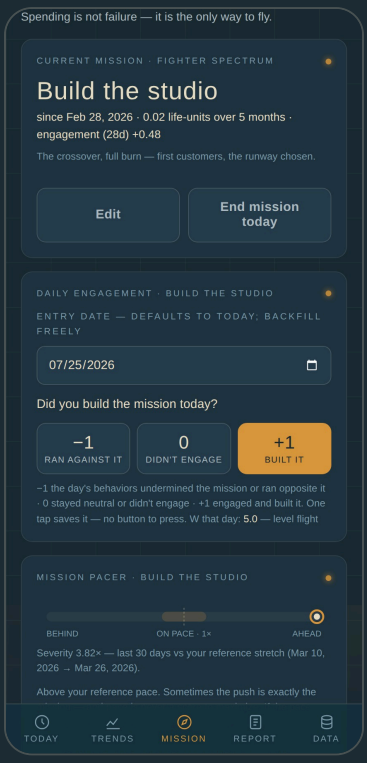

D measures spending; the Mission tab names what the spending is **for** — the ASIP Task I move, deciding what the airframe is built to do.

1. **Start a mission in three taps.** Pick a spectrum — Fighter, Sailplane, Transport, Aerobatic, Trainer, Rescue, Test bed, Depot (✶ marks the ones your MmM material typically flies well) — name it in your own words, **Start**.

2. **Dates live under an Advanced fold.** Start defaults to today; an end date is only for logging past missions into the history — so a whole life's record can be mapped from journals, retroactively.

3. **Then, each day: one chip.** "Did you build the mission today?" — **−1** ran against it · **0** didn't engage · **+1** built it. One tap saves it; there is no button. The chips have their own date field, so backfilling engagement works like backfilling W.

One mission at a time; gaps between missions are allowed and never counted against you — ground time is preparation for engagement. If the log sits quiet more than a week, a flat flag says so.

The service log, and the alignment map.

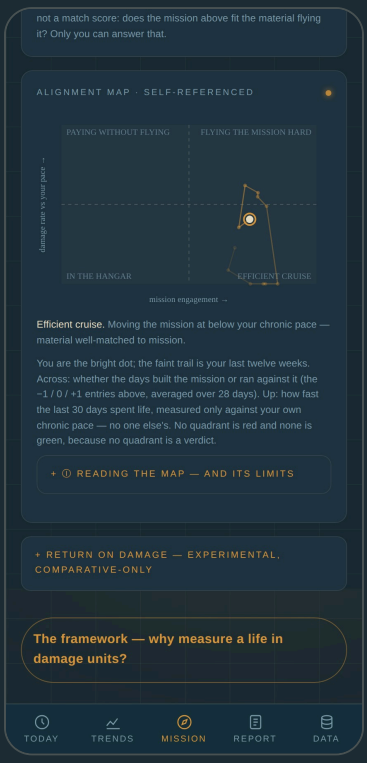

The service log lays your declared missions under the D(t) curve as amber bands — *Mission 3 — 0.31 life-units over 2.1 years* — a service log, not a diary. Every mission flown stays reviewable; sometimes the hangar was exactly the right call, and the log should say so without apology.

The alignment map is a 2×2, measured only against yourself. Across: 28-day mission engagement (the chips). Up: 30-day damage rate against your own chronic pace. You are the bright dot; **the faint trail is your mission's days.**

The trail is daily now — points from the very first chip, thinning to weekly averages as the record grows. Backfill chips from a journal and it fills in immediately.

No quadrant is red and none is green, because no quadrant is a verdict — and the tab stays silent when your last two weeks have been heavy.

The pacer, and return on damage.

The mission pacer. Edit a mission and mark a **reference stretch** — a spell when you were flying it well — and a Strava-style strip appears on the current mission: your last month's damage rate against your own reference rate, **behind / on pace / ahead**.

Mismatch is the signal in both directions: above-pace on a Depot or Sailplane refit is the alarm ("the refit is running hot"); below-pace on a build mission means it isn't being flown. Pace is never a score to maximize, and the target only ever comes from your own record — never a lookup table.

Return on damage (experimental) folds at the bottom of the tab and starts reading within days of your first chips: engagement bought per life-unit spent over an adaptive trailing window, shown only against **your own trailing median**. There is no good number — only your own history.

Not all damage is spending. Illness, loss, and weather accrue on every airframe. Return-on-damage is a lens for chosen missions — never a judgment on what happened to you.

MISSION INSTRUMENTS · NEW IN V2.6, EXPERIMENTAL

Resonance — the natural-frequency check.

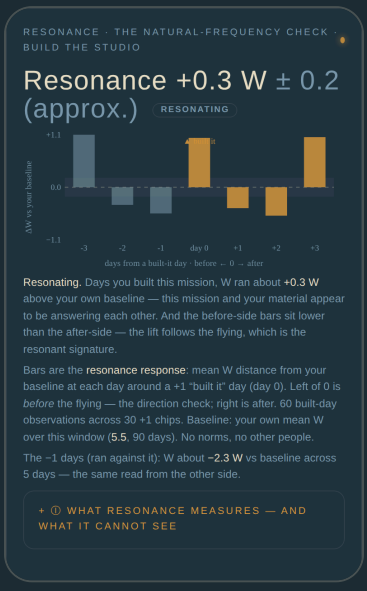

The instrument's third leg. D measures what the flying **costs**; the chips measure whether the mission **moved**; resonance asks what neither can: does flying this mission **feed the airframe**? On the days you tapped +1, did W run above or below your own baseline?

The headline is a mean ΔW — always with an approximate ± band, never a bare number. Beneath it, the **resonance response**: your average W distance from baseline at each day around a built-it day, three days either side. The bars *before* day 0 are the honest direction check — if W is already up before you engage, the record reads "you fly when well" at least as much as "flying makes you well," and the caption says which side your record favors.

Three reads, and the band must clear zero first: **Resonating** — the mission and your material are answering each other. **Reading level** — no clear response yet. **Holding on too tight?** — engagement running against your own wellness. A question by design: some missions are worth what they cost, and only the pilot prices that.

Not all damage is spending. Resonance is a lens for chosen missions — never a judgment on what happened to you, and never a verdict on whether a mission is worth its cost. No new inputs, no new storage: it only reads what the dial and the chips already hold.

TRENDS · STRUCTURE — NEW IN V2.3, EXPERIMENTAL

A second opinion from fracture mechanics.

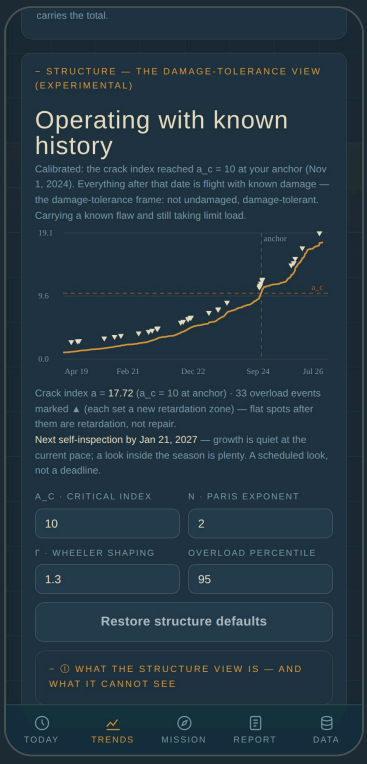

A fold at the bottom of Trends runs a second life-model on the same days: a crack index a(t), grown **cycle-by-cycle in time order** (Paris' law with Wheeler retardation). Where the damage sum is sequence-blind, this track is sequence-dependent by construction — the same days in a different order give a different answer.

Hard stretches — your own worst 5% of swings, marked ▲ — temporarily slow the growth that follows. **Retardation, not repair: nothing heals.** A hard event changes how the ordinary cycles after it bite, for a finite while.

Set your anchor and the curve calibrates; everything after that date reads **"Operating with known history"** — the damage-tolerance frame: not undamaged, damage-tolerant. Log a repair and the growth restarts fresh while every superseded timeframe stays one tap away — and a COMPARE view states plainly what the repairs bought.

"Next self-inspection by [date]" is the app's only forward-looking line, and it is maintenance language — when to look inside, the way airframes get scheduled inspections. Never a countdown: spent is knowable; remaining is not, and **neither track ever shows remaining life.**

TAB 5 · REPORT

The month, read like a service log.

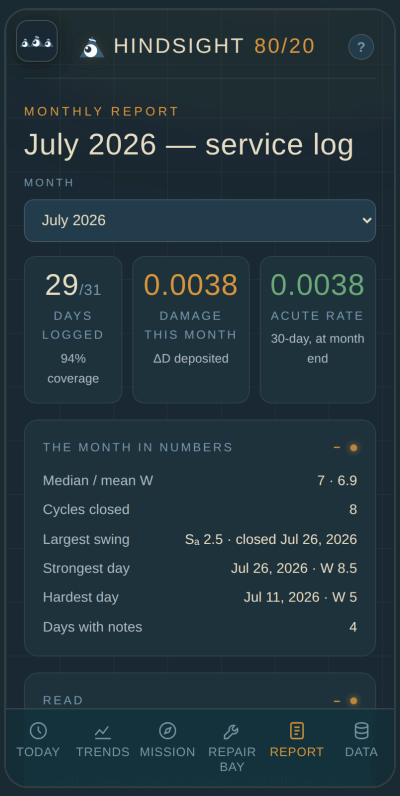

Every calendar month builds its own report: coverage, damage deposited, the acute rate at month end, cycles closed, largest swing, hardest and strongest days, and a plain-language read.

1. Pick the month from the selector.
2. **Save as PDF** opens your phone's print sheet — choose "Save as PDF" and keep or share the file.
3. **Month CSV** exports just that month's rows.

The read to carry forward: does next month's acute rate rise or ease against this one?

TAB 6 · DATA

It lives here. Nowhere else.

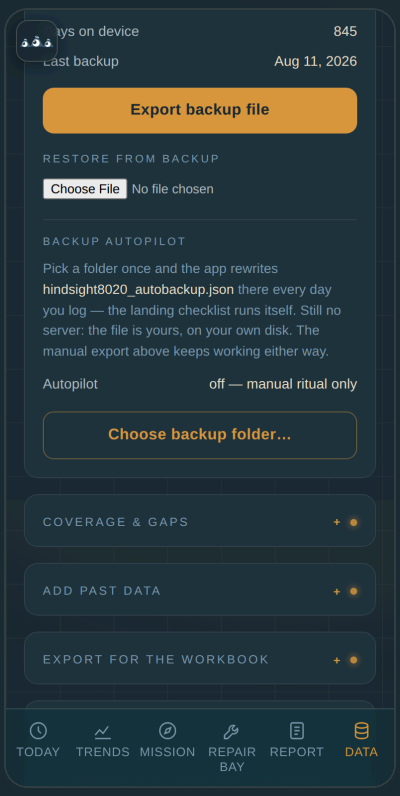

The monthly ritual — one tap. Your record exists only on this device. Once a month, **Export backup file** (or Share backup to email or cloud). A banner nags when a backup is missing or stale.

Two hulls (v4.1). Every save is now mirrored into a second on-device store; if the browser ever wipes the first copy, the app restores itself from the second on next open. **Backup autopilot** (shown left, where the browser supports it): pick a folder once and *hindsight8020_autobackup.json* is rewritten there every day you log — the landing checklist runs itself. Still no server; the file is yours.

The evening tap (v4.1) — an optional once-a-day reminder, only if the day isn't logged by your chosen hour. No streaks, no escalation. And on a brand-new record, **first flight** shows Today, Trends and Data first — the other tabs come online at seven logged days, or one tap.

Restore — feed a new phone the backup file. Everything returns: entries, notes, settings, missions, the structure parameters, and the whole Repair Bay ride in the same file. Old backups restore into new versions and vice-versa.

Coverage & Gaps · Add past data — span, days logged vs. missing, each gap listed; auto-fill short gaps with tagged, removable estimates. Import a CSV (date,W,note) or paste rows; logged days are skipped unless you tick Overwrite.

Advanced — the model parameters, plus one honest button: set your anchor date and **Calibrate** pins D = 1.0 to it. Until then, watch the rates.

TRIAL & UNLOCKING

Fourteen days, on the house.

The first fourteen days are free — the pill in the top corner counts them down. When the trial ends, the instrument locks; **your record does not**. The lock screen itself carries an export button, because your data is never the hostage.

1. **Unlock Hindsight Daily** takes you to secure checkout (Stripe).
2. After payment you land back in the app and it unlocks itself.
3. On a second device, or if anything hiccups: enter the email you paid with (or your code) under **Purchased before?** — it verifies online once, then never needs the network again.

One purchase, yours. No subscription, no account, no tracking.

WATCH INSTEAD

The walkthrough & the spot.

A silent, captioned walkthrough of the whole instrument — including everything in this guide's mission pages — lives at isightlabs.netlify.app/daily/howto (3:42). The 30-second brand film is linked from the welcome card and the Data tab.

THE FINE PRINT THAT ISN'T FINE

What this instrument is not.

It is not a diagnosis. D is a gauge of accumulated load, read the way engineers read a wing — it is not a screening score, a clinical judgment, or a prediction. It cannot see a crisis. That includes the pace alert, the alignment map, resonance, and the structure view: they compare rates, nothing more.

> If you are in one, the right instrument is a person. In the US, call or text **988** — the Suicide & Crisis Lifeline — any time, any reason.

PRIVACY	Every entry stays in this browser's storage on this device. Nothing is uploaded, synced, or seen by anyone — including the lab. The unlock check sends only a code or payment ID, never your record.
THE CATCH	That privacy is also the risk: the device is the only copy. Hence the monthly backup ritual — the app will nag, and it means well.
DECIMALS	Half-steps are honored everywhere. A 6.5 kind of day is a 6.5.
MISSED DAYS	Expected. Log what you can; the method was built for irregular records.
SECTIONS	Everything folds to its title — tap any header to expand it. Your layout is remembered on-device and rides in backups.
MISSIONS	Not goals, not scores. A mission is a name for what the spending is for — and "in the hangar" is a legitimate place for an airframe to be.
RESONANCE	Correlation on the days of one life, never causation — and never a verdict on whether a mission is worth its cost. Some are. Only the pilot prices that.
GO DEEPER	The Application Workbook runs the same pipeline with every formula exposed, and the full story is in *Insider Threat*, Volume II — "The Spar That Failed."

Hindsight 80/20 · iSight Laboratories LLC · Hyrum, Utah · isightlabs.netlify.app

How fast do you get back up?

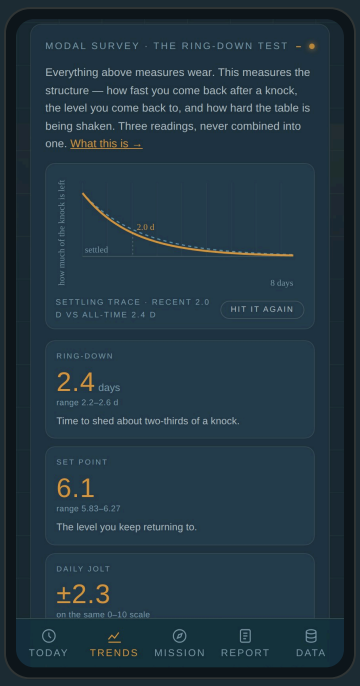

Every other panel on Trends measures **wear** — damage piling up on a structure whose properties are taken as given. This one measures the **structure**. The damage and crack tracks are *i*SIP Task V, force management; the Modal Survey is **i5.3.2.1**, Task III — the ground vibration test, where engineers thump a finished airframe and time how long it rings.

Think of a marble in a bowl. Something knocks it up the side; the bowl pulls it back. Three separate things govern that, kept here as three separate readings and never averaged into one:

Ring-down, the days it takes to shed about two-thirds of a knock — how steep the bowl is. **Set point**, where the bottom of the bowl sits: the level you return to, whatever today happened to be. **Daily jolt**, how hard the table is being shaken. Each prints with a range, because a bare number would claim a precision the record hasn't got.

The second analogy is the design choice. A trainer is built stable — let go and it finds level on its own, predictably, slowly. A fighter is built with *relaxed* stability on purpose: it barely holds its own attitude, which is exactly what makes it quick, and it needs a control system running to stay flyable. Neither is the better aircraft.

Below the gauges, a **rolling survey** shows how the first two have moved — each point reading the 180 days *centred* on it, so the line lags real change by about three months. A **per-mission table** surveys each declared mission; short missions give wide ranges, and two numbers only differ meaningfully when their ranges don't overlap.

A short ring-down means you recover fast. It does not mean you are doing well. Read it beside the set point. Quiet and braced are opposite readings, not the same one: quiet is a small jolt, braced is a big jolt with a short ring-down.

It is not an early-warning alarm. Run against the worst event in the founder's own record, the ring-down was *improving* right up to that morning. A condition check, not a prediction, not a diagnosis — and one person's evidence so far. The full explanation: iSightLabs.netlify.app/modal_survey

What actually works on you?

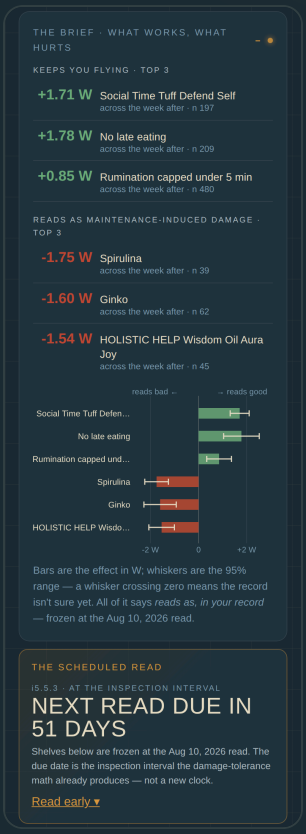

Every other tab measures the **damage**. The Repair Bay keeps the **maintenance record**: which recovery procedures earn their place in your record — and which well-meant ones read as **maintenance-induced damage**. The added burden is one tag on repairs you already log — what it was, its **dose**, and whether you did it *to recover*.

Three capped lists: **Certified** — works in your record (three slots), **Grounded** — well-meant maintenance that reads negative (three), **In Test** — candidates (five). A full list always asks who gets displaced; everything ever tried keeps its verdict in the **Service History** archive — sortable, filterable, renameable. Same-named studies pool into one all-time read, and Detail keeps the by-period breakdown, flagging when seasons disagree.

The gate: a procedure carries a **dose and an interval** or it is not certified. "Water" is not a procedure. "16 oz before noon, daily" is.

The Brief sits first on the tab: the top three that read good, the top three reading as maintenance-induced damage, and an effect ladder — signed bars in W with 95% whiskers, so size and certainty are one glance. Every verdict leads with the plain word — **reads good** or **reads bad** — and names its window in ordinary language: on the days you did it, the day after, across the week after.

Results recompute **only when the inspection interval comes due** — the same date the Trends panel shows. Between reads the shelves sit frozen under **NEXT READ DUE IN *n* DAYS**. That is deliberate: continuous self-inspection couples with the thing it inspects, and that hazard has its own task number in the program. You can read early; the bay says what it costs, and logs it.

Under the hood, a procedure is scored against the **residual** — the part of a day the Modal Survey's ring-down model did *not* already predict from yesterday — judged against matched days from the *same era* of your record, so neither momentum nor a hard year can masquerade as the procedure. Verdicts wait for **five occurrences**. A same-day-only effect gets flagged — *maybe something you do on good days* — and the remedy is a **coupon test**: declared before the run, falsifier in your own words, closed on the declared comparison only, with missed ON days tap-marked on the card and excluded at close — adherence prints beside the verdict.

It never says "causes." Every verdict reads *associated with, in your record*. Effects print in W units and damage-units per year — never a correlation coefficient — and the In Test list runs a false-discovery correction so a list of maybes can't masquerade as a list of findings.

If the reading is heavy. This bay tunes walks and bedtimes. Anything involving persistent depression, crisis, or harm routes outside the program — to a person. In the US: **988**, by call or text. Anything prescribed is a clinician's call; the bay records it as a finding and goes no further.