THE PLAYBOOK

The Project

How to run a project with ten one-page artifacts: a charter and a kickoff to start it, a RAID log, risk register, RACI, and status report to run it, a decision log and stakeholder map to keep it honest, and a pre-mortem and closure to bracket it. Opinionated defaults, named failure modes, and the weekly rhythm that ties the ten together.

Paul Prado Pacardo

Executive Assistant · Project Manager · Operations   me.kabuhayan.app

Contents

THE TOOLS

Every artifact in this playbook exists as a free fill-in-the-browser tool at me.kabuhayan.app/tools, each with a worked example one click away and a printable blank version. This document is the connective tissue: it tells you when to reach for which sheet, and what goes wrong when people don't.

CHAPTER 0

How to use this playbook

Most project management writing is either a glossary or a religion. A glossary defines a RACI matrix and leaves you alone with it. A religion tells you that unless you adopt the whole framework, certification included, you are doing it wrong. This is neither. It is a field guide to ten one-page artifacts and the order they earn their keep, written by someone who runs projects with them, not someone who sells the training.

The stance throughout is that defaults beat blank pages. Every artifact here comes with a recommended cadence and a recommended size: a kickoff runs 60 to 90 minutes, a RACI row gets exactly one Accountable, the top risks get re-scored weekly, a pre-mortem takes half an hour. These are practitioner conventions, not laws of physics. Change them the moment you understand what the default was protecting you from. Until then, use them, because the alternative is not a better number, it is a blank page and a meeting about the meeting.

The playbook follows a project's life in five phases: before it starts, starting it, running it week to week, when things wobble, and ending it. Each chapter opens with the phase's one job, slots in the artifacts that serve it, and names the failure mode in plain language. The failure modes have names on purpose. "Watermelon reporting" and "two Accountables" stick in a way that "suboptimal status communication practices" never will, and the thing you can name is the thing you can catch.

Who this is for

Anyone who has been handed a project and a deadline without being handed a method. Coordinators becoming project managers. Executive assistants who already run projects in everything but title, which is how I started. Founders running their own delivery. If you manage a portfolio of forty projects under formal governance, you have heavier machinery and a PMO to feed it; this playbook is for the project that has to work with one person keeping it honest.

> DO FIRST
>
> Open the wiring page in chapter 6 before reading anything else. Ten artifacts sound like bureaucracy until you see that each one feeds the next; after that, skipping one feels like leaving a gear out of a watch.

CHAPTER 1

Before it starts

The phase's one job: decide what this project is, who can sink it, and how it is most likely to die, before anyone books the kickoff. Three artifacts do that work: the `charter`, the `stakeholder map`, and the `pre-mortem`.

The charter: one page of authority

A charter authorizes the project and hands the project manager the right to spend people's time. It says why the project exists, what done means in numbers, what is explicitly out of scope, four to six milestones, and who signed. One to two pages, written by the PM, signed by the sponsor, on partial information, because a charter that waits for complete information is a history book.

The out-of-scope section is the load-bearing wall. PMI's 2018 Pulse of the Profession found 52 percent of projects experience scope creep, up from 43 percent five years earlier. The cure is not vigilance, it is a written exclusion list with a signature under it, because in month three you do not argue about scope, you reopen a page both of you signed.

> THE FAILURE MODE
>
> **The novel-length charter.** A charter that tries to be the project plan gets skimmed once and never reopened. The charter authorizes; the plan details. If a section wants to grow, it is telling you it belongs somewhere else.

The stakeholder map: who can sink this

Rate everyone who touches the project on two axes, power and interest, and the quadrant tells you the strategy: manage closely, keep satisfied, keep informed, or monitor. The framework goes back to Aubrey Mendelow's 1991 paper on stakeholder scanning; the four-box grid everyone recognizes was popularized later by Johnson and Scholes. The map's real product is your communication rhythm: who gets the weekly report, who gets a monthly line, who gets invited to demos.

> THE FAILURE MODE
>
> **The quiet director.** High power, low interest is the most under-managed quadrant. The director who ignored your project in April can own its budget in June, which is why the map carries a last-reviewed date and gets re-plotted at every phase change, not filed after kickoff.

The pre-mortem: fail it on paper first

After the plan exists and before work starts, put the team in a room and say: it is six months from now, the project has failed, it was a fiasco. Why? Everyone writes alone for two minutes, then you go around the table, one reason each, leader first. Then two more minutes each on what they personally can do so their reason never happens, and every preventive action leaves with an owner and a date.

The method is Gary Klein's, from his 2007 Harvard Business Review piece on the project premortem, and it rests on research from 1989 by Mitchell, Russo, and Pennington on prospective hindsight: imagining an event as having already happened improved people's ability to correctly identify reasons for outcomes by around 30 percent. Be precise about what that number is. It measures explanation accuracy in a study, not a 30 percent jump in project success, and a playbook that inflates its own evidence has already told you how it handles status reports.

Why it beats a risk brainstorm: asking "what could go wrong" invites diplomacy. Stipulating failure grants permission. The concern most worth hearing is usually held by the person least likely to volunteer it, and "the project already failed, explain it" is the only prompt that reliably gets it out of them.

The whole session runs 20 to 30 minutes. Its output does not get admired, it gets typed straight into the risk register as the opening set of scored risks, which is the first piece of wiring between the artifacts.

> THE FAILURE MODE
>
> **The softened prompt.** "What are some risks we see?" produces the same three polite answers as every other meeting. The prompt is "it failed, why," said with a straight face. And if the most senior person in the room speaks first and at length, the exercise is over; go round-robin, one item at a time.

> DO FIRST
>
> If you only adopt one thing from this chapter, run the pre-mortem. It is thirty minutes, it needs no software, and it front-loads the honesty that every other artifact depends on.

CHAPTER 2

Starting it

The phase's one job: convert a signed charter into shared understanding. Four artifacts: the `kickoff checklist`, the `RACI matrix`, and the opening entries of the `RAID log` and `risk register`.

The kickoff: 60 to 90 minutes, prepared

A kickoff is not where scope gets discussed, it is where scope gets confirmed. The checklist has three parts because the meeting is the middle, not the whole: prepare (charter drafted, stakeholders mapped, agenda out two to three days ahead, every agenda item with a named leader), run (sponsor delivers the why in person, scope in and out, roles said out loud, timeline, working agreements, park anything that threatens the clock), and follow up (summary with actions, owners, and dates within 24 to 48 hours). The follow-up is the real kickoff; the summary is what people actually work from.

> THE FAILURE MODE
>
> **The discussion kickoff.** If the meeting opens questions the charter should have closed, you are holding planning with an audience. Park it, assign it, and protect the room; a kickoff that ends on time with five confirmed outcomes beats one that solves everything and drains everyone.

The RACI: one Accountable, always

Tasks down the side, roles across the top, and each cell gets one letter: Responsible does the work, Accountable owns the outcome and signs off, Consulted gives input before, Informed hears after. One rule carries the entire tool: **exactly one A per row.** Two Accountables is how a deliverable falls between two chairs while both owners assume the other has it. And most people should be I, not C, because a matrix where everyone is Consulted is a meeting generator with a grid on it.

Keep it to one page and only the danger-zone rows: client-facing deadlines, cross-team handoffs, approvals, budget decisions. Use role titles, not names, so the matrix survives the people.

Standing up the logs

The RAID log and risk register open now, seeded from the pre-mortem, and started rather than exhaustive. Ten to fifteen scored risks is a working register; sixty is a document nobody will ever update again. Every entry gets an ID, R-001, I-003, D-002, because "the vendor thing" cannot be tracked and R-003 can.

CHAPTER 3

The weekly heartbeat

The phase's one job: keep a predictable rhythm so problems surface on schedule instead of by ambush. Two artifacts carry it, the `status report` and the `decision log`, with the `RAID log` and `risk register` maintained underneath.

Day	What the PM touches
`Monday`	Set the week's three to five outcomes with owners. Scan the RAID log for anything due. Check the top risks. Start the week with a plan, not an inbox sort.
`Midweek`	Standup or async update: progress, blockers, dependencies. Log decisions the day they happen. New RAID items go in immediately, not at review time.
`Thursday`	Re-score the top risks. Update RAID owners and mitigations. Confirm next week's handoffs first, they anchor the plan. Stakeholder touchpoints per the map's cadence.
`Friday`	Close the week and send the status report while it is still true. Pre-plan Monday so future-you starts moving instead of remembering.

The whole rhythm costs a couple of hours a week. The point is a heartbeat, not more meetings.

The status report: four lights, not one

Rate schedule, budget, scope, and risk separately, red, amber, or green, each with one line of why. A project can be green on time and red on budget, and a single overall light exists to hide exactly that. Define the colors before you need them: green is on track, amber is "will go red unless the named thing happens," red is "senior intervention required." Lead with the decisions you need from the reader, because executives skim and the ask should not live on line forty. One page, sent at the end of the period, on a cadence, even when nothing changed, because a skipped report reads as a hidden problem.

The decision log: the why is the payload

Everyone remembers what was decided. Nobody remembers why, and the why is what stops month-three re-litigating. Each entry: ID, date, the decision with its rationale in one line, who made it, what was considered and rejected. Rejected options are the working part: "we considered X and chose Y because Z" ends the "did anyone think of X" meeting before it is booked. Reversed decisions get marked superseded, never deleted; the trail is the project's memory.

CHAPTER 4

When things wobble

The phase's one job: make bad news travel fast and upward. The artifacts do not change; what changes is whether you let them tell the truth.

Watermelon reporting

Green on the outside, red on the inside. It is the most studied failure in project communication: research on software project status reporting by Snow, Keil, and Wallace found reports were biased about 60 percent of the time, with the bias twice as likely to be optimistic as pessimistic, and follow-up work in MIT Sloan Management Review confirmed how routinely bad news gets filtered on the way up. The caveat belongs in the open: that research is about IT projects, with a modest sample. But every PM who has watched a project jump from green straight to red, no amber in between, knows the watermelon did not grow overnight.

The fix is structural, not moral. People filter bad news when honesty is punished. So: define amber as an early-warning color and treat its use as competence, not confession. Substantiate every color with the line of why. Never tie a red status to a performance review, because the strongest predictor of dishonest reporting is a culture where honest reporting hurts. And walk the floor; a status color you only receive is a status color someone chose for you.

When a risk lands

A risk that happens stops being a risk. Close it in the register, open an issue in the RAID log with an owner and a next action, and say so in the status report with the decision you need. The paper trail from R-004 to I-002 to "decision needed by Friday" is what escalation looks like when it works: no drama, just routing.

Re-score on review, not on memory. A risk scored likelihood 2 in week one can be a 4 by week four, and a register whose scores never move is either mis-scored or unread. High scores get a mitigation and a contingency; low scores get watched, not worked.

> THE FAILURE MODE
>
> **The always-amber report.** If everything is amber every week, amber means nothing and the reader learns to skim past it. Amber is a commitment: name the thing that must happen, and the date by which its absence turns the light red. An amber without a named condition is a hedge wearing a color.

CHAPTER 5

Ending on purpose

The phase's one job: make the project stop existing, formally, and make its lessons load into the next one. One artifact: the `closure sheet` with its lessons-learned table.

Book the closure meeting at kickoff. It is the meeting most likely to be skipped, because by the end the team has mentally moved on, and it is the cheapest hour the next project will ever buy. Closure means: deliverables verified and accepted in writing, every RAID item at a final status, contracts and invoices actually closed, documentation and credentials handed over, people released with a thank-you and an end date instead of quietly evaporating. Unclosed projects haunt; they generate "quick asks" for months because nobody said done out loud.

Plan versus actual: the honest table

Three columns, no adjectives: area, planned, actual. Launch Nov 30, launched Dec 2. Budget envelope, 4 percent under. The gap between the columns is not a verdict, it is next project's estimate, corrected, and it is the fastest-compounding page in this entire playbook.

Lessons that get reused

The central failure of lessons learned is that they are written and never read: captured into a repository at closure, searched by no one, ever. The fix is to change where they live. A lesson is not finished when it is written down; it is finished when it appears as a line item in the next project's kickoff checklist. "Communicate better" teaches nothing. "The content deadline goes into the contract, with a placeholder fallback agreed in writing" changes the next kickoff, which is why the closure sheet asks what you do differently next time and what it applies to, not how you feel.

Ten specific lessons beat a hundred observations, and a facilitator who is not the PM gets more honest ones.

> DO FIRST
>
> Before your next kickoff, read the closure sheet of the last comparable project, yours or anyone's, and carry two lessons in as named checklist items. That single habit is the loop that separates teams that learn from teams that archive.

CHAPTER 6

The wiring and the defaults

Ten artifacts, one system. Read the arrows and the bureaucracy disappears: each sheet exists to feed the next.

> **Charter feeds kickoff.** The kickoff confirms what the charter authorized; it never re-opens it.

> **Charter and kickoff feed the stakeholder map,** and the map sets the communication rhythm: who gets what, how often.

> **The pre-mortem seeds the risk register.** Failure stories become scored risks with owners.

> **The register and the RAID log mirror each other.** The register is the scored deep dive; its top three sit on the RAID page so the one-page view stays honest.

> **RAID issues and Thursday's re-scoring set Friday's colors.** The status report is the week's verdict, not a separate chore.

> **The status report cites the decision log,** and the log holds the why so the report can stay one page.

> **Closure lessons feed the next charter and kickoff.** The loop closes, and the next project starts smarter.

Artifact	Default that matters	Cadence
Charter	One to two pages; out-of-scope list signed	Once; reopened at phase gates
Stakeholder map	Power x interest sets the strategy	Kickoff, then every phase change
Pre-mortem	"It failed, why"; leader speaks first, once	20 to 30 min, after plan, before work
Kickoff	Prepare, run, follow up in 48 hours	60 to 90 min, once
RACI	Exactly one Accountable per row	Set at start; review at phase changes
RAID log	IDs on everything; one owner per item	Working review weekly
Risk register	Score = likelihood x impact, cause-and-effect wording	Top risks re-scored weekly
Status report	Four lights, decisions first, one page	Weekly, end of period
Decision log	Rationale and rejected options recorded	The day the decision is made
Closure	Plan vs actual; lessons land in the next kickoff	Booked at kickoff, run at the end

Small team or a four-week project? Keep a half-page charter, a light RAID, and the weekly report; drop the rest until the project earns it.

CHAPTER 7

Sources, caveats included

The claims in this playbook rest on the following. Where a number is contested or narrower than it sounds, that is stated here and in the text, because a playbook about honest reporting should model it.

Status-report bias. Snow, Keil, and Wallace, "The Effects of Optimistic and Pessimistic Biasing on Software Project Status Reporting," Information and Management 44(2), 2007: reports biased about 60 percent of the time, optimism twice as likely as pessimism. Caveat: software projects, modest sample. See also Keil, Smith, Iacovou, and Thompson, "The Pitfalls of Project Status Reporting," MIT Sloan Management Review 55(3), 2014.

Pre-mortem. Gary Klein, "Performing a Project Premortem," Harvard Business Review, September 2007, drawing on Mitchell, Russo, and Pennington, "Back to the Future: Temporal Perspective in the Explanation of Events," Journal of Behavioral Decision Making, 1989. The circa 30 percent figure measures improved explanation accuracy under prospective hindsight, not project-success uplift.

Scope creep and waste. PMI, Pulse of the Profession 2018: 52 percent of projects experienced scope creep, up from 43 percent five years earlier; organizations wasted an average 9.9 percent of every project dollar. Caveat: industry self-survey, and headline extrapolations from it are modeled, not measured.

Stakeholder mapping. Aubrey Mendelow, "Environmental Scanning: The Impact of the Stakeholder Concept," Proceedings of the Second International Conference on Information Systems, 1991. The four-quadrant presentation was popularized by Johnson and Scholes in Exploring Corporate Strategy.

Project size and failure. The Standish Group's CHAOS reports are widely quoted and formally criticized; Eveleens and Verhoef, IEEE Software, 2010, concluded the success figures are unreliable for benchmarking. This playbook uses them only for the direction everyone's data agrees on: small projects succeed far more often than grand ones.

Cadence defaults. Kickoff length, weekly re-scoring, and review intervals are practitioner conventions drawn from the working guidance of Atlassian, Smartsheet, and The Digital Project Manager, among others. They are sensible starting points, not findings.

THE TOOLS, AGAIN

All ten artifacts are free, fill-in-the-browser, and printable at me.kabuhayan.app/tools. Nothing you type is sent anywhere. If this playbook helped, that page is where it becomes practice.