


Field Guide

The complete reference: every tool, every screen, and what each control does. Written for the people using Ceruleo, not for the people building it.




Ceruleo - User Guide & Knowledgebase

Welcome to **Ceruleo**, the platform for planning and producing experience-driven creative work - themed attractions, exhibits, events, activations, shows, and everything in between. This is the living user guide and knowledgebase for the platform: start here for orientation, and come back as a reference. You don't need to read it end-to-end; skim the headers and dive into what's relevant.

This is a living document. It is kept current as Ceruleo evolves - new and changed features are reflected here as they ship. If something you see in the app isn't described here yet, let the team know and it'll be added.

1. The big picture




All projects: your company's whole portfolio - starred first, then A to Z, each card wearing its status chip.

Ceruleo is organized into three nested ideas:

- **Organization (your company).** Everyone on your team belongs to one organization. Org-wide things - templates, members, and shared idea boards (Sparks) - live here.
- **Projects.** A project is a single production (a show, an attraction, an exhibit). Most of your day-to-day work happens inside a project.
- **Elements.** Inside a project, work is organized into a tree of **elements** - the rooms, areas, features, and parts that make up the production. More on this in section 4; it's the core concept worth understanding early.

A quick mental model: **Organization → Projects → Elements → the work attached to those elements** (flags, RFIs, notes, tasks, assets, files, and discussions).

2. Getting in



The sign-in screen at login.ceruleo.app - accounts come from your admin's invitation.


- Sign in with your email and password. If you were invited, accept the invitation from your email link and set your password.
- **Passwords** must be **at least 10 characters** and include an **uppercase letter**, a **lowercase letter**, and a **number** - the set-password screens show a live checklist that ticks off each requirement as you type. The

policy applies whenever a password is **set or changed** (accepting an invitation, or resetting a password); if your existing password predates it, you can keep signing in with it as-is and you'll simply meet the new rules at your next password reset.

- **Two-factor authentication (2FA)** - an optional second step at sign-in, using an authenticator app (Google Authenticator, 1Password, Authy, or similar). Turn it on yourself in **Settings → Profile → Two-factor authentication**: scan the QR code with your app, type the 6-digit code it shows, and it's on. From then on, signing in asks for your password and then a **6-digit code**, on the web and in the phone app. Your company can also **require** it - see *Templates & settings*. **There are no backup codes**. If you lose the phone with your authenticator app, nobody can get you in except a **company admin**, who can reset 2FA for you from **Company settings → Members**; you then set it up again on your new phone. Because of that, set it up on a phone you will still have, and if you use a password manager with a built-in authenticator, storing the code there means losing your phone doesn't lock you out.
- **Forgot your password?** Use **Forgot password** on the sign-in page and follow the emailed link to set a new one. Reset links are **single-use** and expire after a while - if a link has already been used or has gone stale, the page tells you so and offers to send a fresh one.
- **How long you stay signed in.** On the **web**, Ceruleo keeps you signed in through the night and signs you out after about **a day without using it**, so a browser left open on a borrowed or shared computer does not stay open indefinitely. When that happens you are returned to the sign-in page with a short note saying why, and signing back in takes you to the page you were heading for. On the **phone**, reopening the app checks your sign-in - if it's still good, the screen you're on simply **refreshes itself**; if it has truly expired, you're taken **straight to the sign-in screen** (no more mysteriously empty pages).
- **On your phone: get the app from the App Store.** The Ceruleo iPhone app is publicly available - search "**Ceruleo**" on the App Store and install it like any other app, then sign in with the **same email and password** you use on the web; it is all one account. The phone is built for the field: capture photos and notes to your **Pocket**, add **daily log** entries and photos on site, and work your **tasks, flags, and RFIs** on the go with push notifications. (If you were on the TestFlight build, switch to the App Store version - updates now arrive automatically, and the TestFlight copy can be removed.)
- You'll land on your **Home dashboard**. From there, the left sidebar is your main navigation: the **Ceruleo logo** returns you home, **Projects** lists everything you can access, and workspace-wide views (**All tasks, All flags, All schedules**) gather those items across every project you're in.

Legal agreements. On your first visit you accept the **Terms of Service** and **Privacy Policy** (always readable at /legal/terms and /legal/privacy). If a materially updated version is ever published, you'll be asked to accept it again on your next visit; small wording updates never interrupt you. Companies also have a **Customer Agreement**, accepted once on behalf of the whole company by a **company admin**. When that agreement is updated, admins see a banner and have **30 days** to review and accept before it becomes required; only truly critical updates ever pause the team's access, and in that case any company admin can clear it in one click.

3. Your Home dashboard



Your Home dashboard: momentum stats up top, then your recent work, today's tasks, and the activity pulse.

The dashboard is your "what's happening" view across all your projects:

- **Hero & momentum stats** - a quick pulse of active projects, tasks due this week, open flags, and unread notifications.
- **My recent work** - the five things **you** most recently worked on, newest first: sparks, rallies, sandboxes, and published assets, and now also **flags, RFIs, notes, and schedules** (editing a schedule's rows counts as working on that schedule). It ranks by *your* last touch - a teammate working on the same item after you doesn't erase your recency - and working in other tools (tasks, discussions, the library) surfaces elsewhere on the dashboard instead.
- **My recent work** - the things you most recently created or edited, so you can jump back in. It respects your current access: if your **Studio** (or **Library**) access to a project is removed, that project's sandboxes and published assets stop appearing here (and in the Activity pulse), so you're never shown work you can no longer open.
- **Activity pulse** - a cross-project feed of who did what (created/updated a flag, RFI, note, sandbox, rally, etc.). It **refreshes on its own** - new activity appears when you come back to the tab and every minute or so while you're looking - so you don't have to reload to see the latest.
- **Today** - your open tasks, due-date sorted.

- **Watching & Notifications** - items you're following and your alerts. The Notifications list works like the full Notifications page: click an alert to open the item it's about, **Mark all read** at the top, and hover any row for its quiet controls - a **circle** to mark a read alert **unread** again (it gets its dot back and counts in the bell badge) and an **x** to **clear** it from your notifications entirely. Everything stays in step with the bell.
- **Schedule lookahead** - a full-width, collapsible band showing what's due across *all* your projects over the **next 14 days**, drawn as the same graphic timeline you see on a project's Overview (milestones above the line, tasks below, each linking into the schedule it lives in). **Overdue** items aren't plotted on it - instead a red **"N overdue items"** chip sits in the corner; click it to open your **All schedules** view with just the overdue items listed. Collapse the band from its header when you don't need it.
- **Active projects** - cards for the living projects you're on the team of, ordered by **most recent activity**: any new log, comment, RFI, task, file, or schedule change on a project (by anyone on the team) moves it forward, so the shelf reflects what's actually in motion. **Starred** projects stay pinned to the front regardless. A project that's been quiet for two weeks or more says so on its card ("Last activity 3 weeks ago"), so the ordering is never a mystery. The shelf shows up to six; **See all** opens the All projects page.

Project cards. On Home and on the **All projects** page, each project's card is its cover image edge to edge - the framed 16:9 shot fills the whole card, with the project's identity set on a dimmed band along the bottom: company, project name, **location**, and **project dates** (the start and estimated end set in Project settings, shown as e.g. "Aug 2026 – Mar 2027"; if only one is set you'll see "From ..." or "Through ..."). On All projects, a small badge in the corner marks a project that's **on hold**, **completed**, or **archived** (active projects wear no badge), and the star button still sits top-right for pinning favorites.

All projects opens on the living ones. The status dropdown defaults to **Open projects** - everything except projects whose status behaves as *Completed* or *Archived* (custom statuses follow their underlying kind, so a *Canceled* status marked as Archived is hidden too). Pick **Everything** or a specific status to widen the view.


Filtering by program. If your company groups its projects into **programs** (set up by an admin in Company settings), a **program dropdown** appears on **All projects**, **All schedules**, **All flags**, and **All tasks** - pick one to narrow the view to just that program's work, or leave it on **All programs** (a project in several programs shows up under each). On **All projects** the list also offers **No program**, which is the quickest way to spot a project nobody has classified yet. The dropdown only shows programs that actually have something in the view you're looking at, and it stays hidden entirely until your company defines its first program.

Cover images. Projects, sparks, and sandboxes can each have a cover image you upload. If you haven't set one, Ceruleo shows a clean **automatic cover** - a brand-colored gradient with the item's initials, generated from its id so the same item always looks the same. Upload a cover any time to replace it.

Scout

Scout answers questions about your work in plain language. (If you knew this tool as **Ask** - same tool, new name; Scout is growing into a character of its own around here.) Press **⌘K** (**Ctrl K** on Windows) from

anywhere, or click **Scout** at the top of the sidebar, type a question, and press Enter.



Scout, one keystroke away: ask about your work, or ask how to put something into Ceruleo.

It is not a keyword box. Ask questions the way you would ask a colleague:

- *did we get a reply from the township on the permit*
- *what's overdue on Riverside*
- *what did we decide about the loading dock*
- *who's our contact at the fabricator*
- *which RFIs are still waiting on an answer*

What it searches

Everything you can already open, across **all of your projects at once**: flags, RFIs and their replies, notes, tasks, meetings and their minutes, daily log entries, discussions and comments, submittals, schedules, studio assets, sandboxes, the project library, the Organizer, the Directory, sparks, rallies, and (if your company uses it) opportunities and logged interactions.

Crucially it reads the **body** of those records, not just their titles. A phrase buried in the third paragraph of a meeting note is findable now. That was never true of any other search in Ceruleo.

You don't have to say which project

This matters more than it sounds. Scout searches across every project you're a member of and works out where the answer lives, so a question like *"did the township reply about the permit"* finds the record even if you would have guessed the wrong project. **Every answer names the project it came from**, so you can see straight away whether it found what you meant.

If you're inside a project when you ask, Scout leans toward that project - but it will still look everywhere rather than tell you "nothing found" when the record is filed next door. When strong matches genuinely sit in two projects, Scout says so, or asks which one you meant.

Ask it how to use Ceruleo

Scout also gives directions on the app itself. Tell it what you're trying to do and ask where it belongs:

- *I need to decide on a name for our VIP program - how should I put this into Ceruleo?*
- *should this be a flag or an RFI?*
- *where do I capture an idea that isn't ready to be work yet?*

Scout reads the Field Guide (this document) to answer, so its advice always matches how Ceruleo actually works. If the right answer depends on something it can't know - one project or company-wide, decision or FYI - it asks a short follow-up first, then commits to one concrete recommendation.

It's a conversation

You can reply. Answer Scout's follow-up question, narrow a search ("which of those was from the township?"), or push back - the exchange stays on screen and Scout remembers it until you press **New chat** or leave the page.

It can pre-fill a form for you

When Scout's advice is "start a rally" or "create a task," it ends with a link that opens the creation form **with the title already typed**. Nothing is created until you review the form and press the button yourself - Scout still cannot create, change, or send anything on its own.

Sources

Under each answer is a row of **source chips** - the actual records the answer drew on. Click one to open it. If an answer has no sources, treat it with suspicion. (Advice about how to use Ceruleo comes from the Field Guide rather than your records, so it may carry no chips.)

Before your first question

The first time you open Scout it shows a short statement of how it works, and asks you to confirm you have read it. It appears once per person. If the substance of that statement ever changes, you will be asked to read the new one.

What it can and cannot tell you

Scout only sees what is **recorded in Ceruleo**, and only the parts **you personally** have permission to open. A private flag you are not a collaborator on is invisible to Scout, exactly as it is in the rest of the app. Two people can ask the same question and correctly get different answers.

This is not a filter applied to the results after the fact. Scout runs every single search **as you**, using the same permission checks the rest of Ceruleo uses to decide what you can open. A record you are not entitled to see is never retrieved, so it cannot be quoted, summarized, or mentioned - Scout has no way to learn that it exists.

The important consequence is in the wording of a negative. If Scout says *"there is no reply recorded on that RFI"*, that means the reply is not in Ceruleo - **not** that nobody replied. They may have answered by email, by phone, or to a colleague. Scout is careful to phrase it that way, and you should read it that way.

Scout **reads only**. It cannot create, edit, send, assign, or delete anything. When you find what you were looking for, you take the action yourself.

Notes

- Answers stream in as they're written. **Stop** halts one mid-flight; **Esc** closes the panel.
- Quoted "exact phrases" work, as does -word to exclude something.
- There is an hourly cap on questions per person, generous enough that normal use will never reach it.
- Scout is available to everyone in Ceruleo - there is nothing to switch on.

Privacy, in detail

Scout is private by design. What you search for, and what it tells you, belongs to you alone. Nothing you ask becomes public, at any point, by any route.

Stated precisely, because this is the question clients ask:

Your questions and your results are confidential. Each question is recorded against your own account so the feature can be measured and improved. Nobody else in your company can read them - administrators included. There is no company-wide log of who asked what. They are never published, never shared or sold to anyone, and never appear in another person's results.

They are never used to train anything. Anthropic's commercial terms state that, by default, they do not use inputs or outputs from their commercial products to train their models. The exceptions they list are explicitly reporting feedback or bugs to Anthropic, or otherwise choosing to allow it. Ceruleo does neither.


What leaves Ceruleo, and what does not. To write an answer, Ceruleo sends your question and short excerpts of the records **you are allowed to see** to Anthropic, who make the Claude model that produces it. That is the only place it goes, and it is used only to answer that one question. It is not sent anywhere else, to any other company, or to any other customer.

We state that plainly rather than claiming your data never leaves the platform, because generating an answer is exactly what it leaves for. Everything else stays put: your files, drawings, images and records live only in Ceruleo's own private storage, and only the text excerpts needed for a specific question are sent, only at the moment you ask.

If a client agreement requires this in writing, or if your company needs a zero-retention arrangement with the model provider, raise it with your Ceruleo contact.

4. Elements: the project structure (the key concept)

Every project is a **hierarchy of elements**. Ceruleo uses a consistent underlying model, but the **labels are tailored to your industry**, so the words you see may differ from project to project.



The Organizer: the project's element tree, where structure is sketched and refined as the work decides it.

The hierarchy

There are six ordered levels:

1. **Level 1** (*renamable* - e.g. *Venue, Land, Wing, Collection*)
2. **Level 2** (*renamable* - e.g. *Zone, Area, Gallery, Sequence*)
3. **Level 3** (*renamable* - e.g. *Scene, Activation, Exhibit, Look*)
4. **Feature** (*fixed* - *cannot be renamed*)
5. **Component** (*fixed* - *cannot be renamed*)
6. **Part** (*fixed* - *cannot be renamed*)

The first three levels are **configurable** and can each have **multiple named types at the same tier** - for example, Level 2 might offer both a **Zone** and a **Scene** as peers.

The bottom three - **Feature, Component and Part** - are the **fixed levels**. They mean the same thing in every Ceruleo project, so **their names cannot be changed**, though you can still switch them on and off like any other level. In Project settings the three you can rename sit in editable boxes; the fixed three are plain text under a **Fixed levels** heading that says so.

"Element" is the generic word for *any* node in this tree. A Zone, a Feature, a Part - they're all elements.

Rules that always hold


- The order is fixed: a Level 3 can never sit above a Level 2, a Part can't contain a Feature, and so on.
- You can **skip levels** - a project can go straight from a Level 1 to a Feature if it doesn't need the middle tiers.
- Projects can be as shallow (a single Feature) or as deep (Level 1 → ... → Part) as they need.
- **You don't have to set the structure up before you start.** A level you have enabled but not yet given any named types to is still addable, under whatever that level is currently called ("Level 1", or "Zone" if you have renamed it). Naming types is an improvement you make when the shape becomes clear, not a step you have to finish first. If you later add named types to that level, the elements you already made stay exactly where they are, and they follow the level's name if you rename it.
- A level you have **switched off** in project settings stays off: it is never offered, whether or not it has named types. And because a project with no levels at all could never have an element, settings keeps **at least one level on**.

Elements don't have to be physical

The tree reads most naturally as a map of the physical build, but an element is really just a **home for a body of work that belongs together** - it doesn't need to be something you can point at on a site walk. Work that matters to the project but has no physical home deserves a branch of its own: a **Marketing** element for the launch campaign and press materials, an **Operations** element for staffing plans and run-of-show, a **Sponsor Deliverables** element collecting everything owed to a partner. These behave exactly like physical elements -

their detail pages gather every flag, RFI, task, note, and asset attached to that effort, at any level of the tree that fits.

Finding your way around



An element's page: color-edged counters that open each tool pre-filtered, published assets as imagery, and child elements at a glance.

- The **Project Navigator** (on the project Overview) shows the element tree. Click any element to open its **detail page**.
- An element's **detail page** opens with its **lineage rail** - the mustard edge carrying the element's path in the tree (each ancestor is a link), its type chip, and its status - then a row of **color-edged counter tiles**: your tasks, flags, RFIs, submittals, notes, published assets, and (when any are linked) rallies for this element, each showing a small **red urgency badge** when it has overdue or critical items. **Clicking a tile opens that tool already filtered to this element** - see 3 submittals here, click, and you land on the Submittals tab showing exactly those 3, with a clearable **Element** chip (clear it to see everything). Below the tiles: a **priorities split** - **Needs your action** (linked items assigned to or awaiting you) and **At risk** (linked items overdue or due soon that aren't yours) - alongside **Published here** (the element's published assets as image cards), **Inside this element** (its child elements), **recent activity**, and a full-width **Discussion** (topics, threaded replies, @mentions) at the bottom. Everything is scoped to *this* element, not the project at large.

- An element's **description** is a single, locked **paragraph** (up to **1034 characters**) that reads in full without a scrollbox. It can only be edited by the element's **creator, a project admin, or an org admin** - everyone else sees it read-only (an **Edit** button appears only for those who may change it).
- If you can manage elements, the **edit pencil** in the element header lets you **rename, reclassify, mark inactive, or delete** the element right from its detail page (deleting also removes its child elements; you're then taken back to the parent element or the project). You can also do all of this from the **Organizer** and the **Project Navigator** tree.

Active vs. Inactive elements

You can mark an element **Inactive** (via the edit pencil on the element). Inactive means:

- It's **dimmed** and tagged "Inactive."
- **No new work can be attached to it** (it won't be selectable when you create a new asset, RFI, flag, task, or note).
- It **stays fully accessible** for reference, and its existing work can still be opened, edited, and **moved to an active element**.

Deactivation **cascades**: cutting a parent (say, a whole Gallery) makes its children inactive too - but **moving a child under an active parent reactivates it automatically**. This is the intended way to "cut" something while salvaging and relocating its work.

Reorganizing

The **Organizer** (the "Organize" link in the Project Navigator) lets admins drag elements to restructure the tree, rename/reclassify them, add children, and toggle active/inactive - all with the ordering rules enforced. (*Access: project admins, org admins, and studio admins.*)

New elements are added to the bottom of their parent's list, so a structure built in order stays in that order as you go. Drag any of them afterwards if you want a different arrangement.

Elements on the phone

The mobile app used to know about elements only when you were **tagging** one - picking which element a flag, RFI, note, task or rally belonged to. You could attach work to an element but never look at one. Now you can.

The project screen carries an **Elements** tile alongside Tasks, Flags and the rest. It opens the project's tree: tap a triangle to fold a branch, and once a project has more than a handful of elements a search box appears (searching shows every match as a flat list, so nothing stays hidden inside a folded branch).

Tap any element to open it. You get its **lineage** at the top - each ancestor tappable, so you can walk back up - the element's type and description, then a row of counts: **Tasks, Flags, RFIs, Notes, Submittals, Library and Rallies** for that element. **Tap a count and you land in that tool, filtered to this element**, with a chip naming

it; tap the chip to clear the filter and see everything on the project. Below that, **Inside this element** lists its children so you can keep walking down.

The phone is **read-only** on structure: creating elements, renaming, reclassifying and reorganizing the tree stay on the web, where dragging things around actually works. Everything you can already do on the phone - raising a flag against an element, filing a photo to one - is unchanged.

Copying an element to several places at once

Areas, zones and scenes very often carry the **same** features: lighting, graphics, audio. Adding "Lighting" by hand under thirty scenes is work nobody should do twice, so hover any element in the Organizer and choose **Copy**.

Pick every element that should also have one, and each gets its own copy. What that means precisely:

- **It copies the name**, along with the element's type and description. The original keeps its flags, RFIs, notes, tasks and files - those are not duplicated.
- **Nothing is moved, and nothing gets two parents**. Each target receives a **separate new element** that happens to share a name. They lead independent lives from that moment: renaming one does not rename the others.
- **Only places it can legally go are offered**. A Feature is offered the elements that may contain a Feature, and nothing below it in the order.
- **Somewhere that already has one is shown but greyed out**, saying so. That includes the element's own current home, so running the copy a second time can't create duplicates.
- If the element **has sub-elements**, a checkbox offers to bring them along; it is off by default, because usually the name is the whole point. Ticking it copies the entire branch beneath the element into each target.

Copies are added to the **bottom** of each target's list, as any new element is.

Building the structure from a spreadsheet

Plenty of projects arrive with the elements already decided, worked out in a creative concept process before anyone opened Ceruleo. When that list already lives in a spreadsheet, you don't have to retype it: **Import from spreadsheet**, at the top of the Organizer, loads the whole structure at once.

Start with **Download the template CSV**. The template comes out already using **your project's own level names** as its column headings, so there is nothing to translate.

One column per level, left to right, top of the hierarchy first, then a Description column. Each row is a **full path** from the top:

ZONE	AREA	SCENE	FEATURE
Main Hall			

ZONE	AREA	SCENE	FEATURE
Main Hall	Lobby	Welcome Moment	Lighting
Main Hall	Lobby	Welcome Moment	Graphics
Main Hall	Cafe	Coffee Bar	Lighting

A few things follow from that shape:

- **Repeat the parent names on every row.** Naming "Main Hall" four times doesn't create it four times - a name that already exists is reused, which is exactly how one feature ends up under many parents. The table above gives Lighting to two different scenes.
- **Leave a cell blank to skip a level.** A row can hang a Scene straight off a Zone.
- **A short row just creates the ancestors.** The first row above adds Main Hall and nothing else.
- **The Description belongs to the last element named in the row.**

The header row names your levels. Change a heading and that level is renamed when you import, so the same file that carries the elements can also set up your vocabulary. Levels 1 to 3 can be renamed this way; **Feature, Component and Part keep their names**, as they do everywhere else. Renaming a level needs project admin rights - without them the elements still import and the headings are ignored.

If your file uses a level the project doesn't have switched on yet, importing **turns it on**. An import never turns a level off and never deletes or edits anything that is already there.

Nothing is written until you say so. After you choose a file you get a preview: how each of your columns was read, how many elements will be added at each level, any level being renamed or switched on, and any rows being skipped with the reason. Check the column list if you have added or removed columns - it is the quickest way to spot a file whose columns have shifted.


Re-importing the same file is safe. Everything in it already exists, so nothing happens. That also means you can send out an updated spreadsheet with more elements in it and import the whole thing again: only the new rows are added.

5. The core tools inside a project

The project bar is how you get around inside a project. On every page in a project - the project's own dashboard, a sandbox, a rally, a schedule, any item's page - a slim bar stays pinned at the top: **Overview**, **Studio**, **Workshop**, then the tools you have access to (a tool set to **none** for you simply doesn't appear). It's always in the same place, so the next tool is always one click - never a trip back through the dashboard. At the bar's right end, past the tools, a **+** button raises something new **in the project you're already in**: an **RFI**, a **Flag**, a **Task**, or a **Note**. The form opens **on top of whatever you're working on**, and closing it puts you straight back there - you never lose your place, so you can log an RFI in the middle of reviewing a drawing and

carry on. It's the same form you'd get from that tool's own New button, so nothing is lost by taking the shortcut, and because it is the project's own bar there is never any question about which project the new item belongs to. You'll only be offered the kinds you're allowed to create. Clicking the **project name** on its left opens a jump menu: your **recent items in this project** (to get back to the thing you were just working in), **your other projects**, an **Also in this project** list, and **Back to dashboard**. The **Organizer** lives in that **Also in this project** list rather than on the bar itself - it is setup work rather than something you reach for daily, and keeping it off the bar leaves more room for the tools you use every day. It still shows only if you can reorganize. The projects list is A to Z, has a **search box** at the top (it focuses as soon as the menu opens - just start typing the project's name), and **scrolls** when the list is long, so no project is ever out of reach. Each core tool has its own section below, opening with an **at-a-glance** answer for when you're in the field and just need the quick version.

The **Workshop** (rallies) and the **Studio** (sandboxes and assets) are their own areas rather than list tabs - reach them from the bar or from the prominent cards near the top of the project Overview; they're covered in their own sections. The Overview's colored **count tiles** (Tasks, Flags, RFIs, and so on) still jump to each tool too - they're the at-a-glance summary, with counts and urgency the bar doesn't carry.



A project's Overview: Studio and Workshop up top, the colored count tiles below, and the schedule lookahead.


Most items support **attachments**, **discussions**, and **watching** (see section 7). You can **attach files right in the create form** for **flags**, **RFIs**, **notes**, **tasks**, and **log entries** - drag and drop or click to add them as you fill

the form, and they're saved with the new item (no need to create it first and re-open it to attach). (For log entries this is available on entry types that have **Allow file attachments** turned on.)

You can **download a flag, RFI, note, or log as a PDF** - open the item (or a day's log) and click **PDF** in its header. It produces a clean, branded document (not a screenshot): a header showing your **company name**, the **project name**, and the item's **number and title** (plus the **project icon**, if the project has one set in Project settings), and its **status / priority / type** badges; a **detail block** (due date, assignee, linked elements, created-by, and any **custom fields** for that type); the item's **rich-text body**; and, for an **RFI**, its **question** followed by **every response** in thread order with the **official answer** marked. A **log's** PDF instead lists that day's **entries** (grouped by type, with their fields and notes) and its **follow-up flags**. Any **images attached** to the item are embedded as a **gallery**; a **log's photos** appear **under the entry each one belongs to**, numbered ("Photo 1", "Photo 2", ...) so they're easy to reference, with any photos not tied to an entry in an **Additional photos** section at the end - the **distribution email** lays its photos out the same way (each entry says which photo numbers it carries, and the numbering matches the PDF). Any attached **PDF files** download alongside as their own documents (exactly as uploaded, without the header/footer). Every page carries a **CONFIDENTIAL** footer with the project · item name and **"Page X of Y"**, and long items paginate across pages.

Which tool when - the cheat sheet

At a glance - Ideas are **sparks**, decisions are **rallies**, making is **sandboxes**, problems are **flags**, questions-for-the-record are **RFIs**, information is **notes**, action is **tasks**, approvals are **submittals**, the plan is a **schedule**, the record of a conversation is a **meeting**, the record of a day is a **log**.



Every tool has a home on the project Overview - this cheat sheet is about picking the right one.

Ceruleo gives you a lot of tools, and picking the right one is what keeps each of them useful. This section is the decision guide: what each tool is *for*, and the mix-ups that happen most. It is also available as a printable one-pager - click **Quick Reference Tool Guide (PDF)** at the top of this Field Guide and pin it to the shop wall.

The tools in one pass

- **Spark** (*company board / Studio*) - An idea worth keeping, before anyone commits to it. Capture it, gather thoughts and reactions, let it wait until a project wants it. Not for working an active problem: when an idea needs a verdict, take it to a rally; when a project commits, move it into a sandbox.
- **Rally** (*Workshop*) - The team gathered around one question, to decide something. Use it when you need many heads on one thing fast - finish options, a layout call, a vendor choice, how to fix what a flag raised. Contributions come in, a direction emerges, action items capture what happens next. Not for producing the work itself, and not for questions an outside party must answer on the record (that is an RFI).
- **Sandbox** (*Studio*) - A workspace where creative work is developed until it is published. Assets get made here: concept work, drawings, documents, models - versions accumulate, reviews render verdicts, approved work publishes. Its **Questions** panel takes the creative, subjective asks that come up while making ("does this read as aged?") - informal by design; nothing there becomes a formal record. Not for a quick group decision (rally) or for storing finished reference files (Library).

- **Flag** (*project tool*) - Something is wrong, and now it has an owner. Raise it when a problem threatens the work - damage, a conflict, a blocker, a safety issue - then shepherd it to resolution. If the fix needs group thinking, start a rally from the flag. Not for questions (RFI), to-dos (task), or FYIs (note): a flag that is really a question sits unresolved because nobody can "fix" it.
- **RFI** (*project tool*) - A formal question that needs an official answer on the record: a spec clarification, a jurisdiction question, a client decision. Reply as an **Answer** when you are answering the question; reply as a **Comment** to discuss without touching the official record. If nobody outside the room needs the answer in writing, it is probably not an RFI.
- **Note** (*project tool*) - Information the project should keep, or someone should receive: decisions, instructions, site conditions, records. Type it with your company's note types and send it **TO** people when they need to see it. A note asks nothing of anyone - if someone must do something, make it a task.
- **Task** (*project tool*) - A to-do with a name on it. Someone needs to do a specific thing, ideally by a date; assign it and it lands in their My Tasks until it is done. A task with no clear "done" is a note or a flag in disguise.
- **Submittal** (*project tool*) - Something needs formal approval before it is used or built: material samples, shop drawings, product data. Route it through review to an approval that stands as a record. Design feedback on work still forming belongs in a Studio review instead.
- **Schedule** (*project tool*) - The plan over time. Work that unfolds across weeks needs dates, owners, dependencies, and progress; give the schedule an owner and they run it. A handful of one-off to-dos does not need one - plain tasks carry those fine.
- **Meeting** (*project tool*) - An agenda, a room, minutes, and action items. For a scheduled conversation that needs a record: agenda ahead, roll call, per-topic notes, decisions, issued minutes. For decisions that do not need a meeting, a rally gathers the same input asynchronously.
- **Log** (*project tool*) - What actually happened today, from the site or the shop: entries by type, photos under each entry, follow-up flags for anything that needs an owner. A log states; it never asks.
- **Scout** (*everywhere, Cmd+K*) - A question about what is already **in** Ceruleo: "did we get a reply from the township?", "what's overdue on Riverside?". Scout reads what your permissions allow and links its sources. New questions for humans are RFIs, rallies, or discussions.

The common mix-ups

- **Rally or sandbox?** Rally to *decide*, sandbox to *make*. A rally is a conversation with an end; a sandbox is a workspace with a lifespan. Decide the direction in a rally, then develop the work in a sandbox.
- **Flag or RFI?** A flag says "something is wrong and it is ours to fix." An RFI says "we need an answer from someone, in writing." If resolving it means someone answers a question, it is an RFI.
- **Sandbox Question or RFI?** A sandbox Question is the room thinking out loud - subjective, creative, answered when the asker says it is, and **never memorialized as a formal record**. An RFI is a question the record must remember, with an official answer. "Does the brass read as aged?" is a Question; "what is the rated load of the truss?" is an RFI.


- **When is a flag appropriate?** When a problem threatens the work and needs an owner *now*. Not for questions, not for to-dos, not for FYIs. If you cannot say what "resolved" looks like, it is not a flag.
- **Note or task?** Information is a note; action is a task. "The paint arrived damaged" is a note (or a flag, if someone must deal with it); "reorder the paint" is a task.
- **RFI Answer or Comment?** Answer when you are answering the question - it goes on the official record and clears the review. Comment to discuss, ask for clarification, or think out loud - the record stays clean.
- **Meeting or rally?** If it needs a time slot, roll call, and issued minutes, it is a meeting. If the same input can arrive over a day or two while everyone stays where they are, it is a rally.
- **Spark or rally?** A spark keeps an idea alive for later; a rally works a question now. Sparks have no deadline and no verdict. When a spark is ready to be decided on or built, move it along.

The one-page PDF

Click **Quick Reference Tool Guide (PDF)** at the top of the Field Guide to download this section as a compact, printable PDF - the same content, laid out to pin up or hand out. It is your company's copy to share with the team however works best.

5.1 Flags

At a glance - Raise an issue that must get resolved. Attach it to elements, discuss it, own it; once resolved or closed the record locks (reopen to edit). A flag has one **owner** (you, when you raise it) and no assignees - work that needs assigning becomes a **rally**.



The Flags tab: open issues with urgency dots, type chips, and overdue markers.

Quick answers

Raise a flag

1. From the **Flags tab** (or the element's page, which pre-fills the placement), click **New flag**.
2. Pick a **type** for its custom fields, or "**Add without a type**" for a plain flag - a project with no flag types set up can still create flags.
3. Give it a title, urgency, due date, and element placement; **attach files right in the create form** (drag and drop works).
4. **You own it** - the person who raises a flag is its owner and shepherds it to resolution. Pull others in as **collaborators** for visibility and @mentions; if the fix needs people actually assigned to work, **start a rally from the flag** (its elements carry over).

Settle a flag

1. Change its status to **Resolved**, **Closed**, or **Void**.
2. The record **locks** - title, description, urgency, due date, custom fields, and privacy can't be changed (the Edit button reads "Locked - reopen to edit"), so the settled record stands as written.
3. Need to amend it? **Reopening (changing its status) is the unlock**. The discussion stays open throughout, locked or not.

Read a flag's page

A flag's page reads as a clean record, like a task or an RFI: a **header card** carries what it is (the red FLAG tag, its number, its type, and the title) with the **status control, privacy, and Start a rally right beside it** - and a **fact grid** below where every fact keeps the same slot every time: Owner, Raised by, Urgency, Due, Elements, Raised date, and the collaborators with the watch toggle. A rally started from the flag shows as a **blue rally chip in the header** that links through. Below the header, the flag's substance (custom fields, description, and attachments, with the **PDF export on the card itself**) sits beside its **discussion, now on the page** instead of behind the edge tab.

Who can do what

ACTION	WHO
See the Flags tab	Anyone with Flags access view or higher on the project
Create / edit a flag	Flags access edit or higher
Resolve, close, void, reopen	Flags access edit or higher
Delete a flag	Its creator, a Flags tool admin, a project admin, or an org admin
Transfer a flag's owner	Admins (the new owner is added as a collaborator automatically)
See a private flag	Its creator, collaborators on it, and admins


Worth knowing

- A flag's **status is its lifecycle**: open flags nag by existing; settled flags stand as written.
- Flags placed **on elements** appear on that element's page forever - place at creation, it's one click.
- **Download as PDF** from the flag's header: a clean branded document with the detail block, custom fields, body, and any attached images as a gallery.
- Flags can also be **raised from a daily log** (when the log allows flags) and **spun into a Rally** when the fix turns out to be bigger than one person - the rally inherits the flag's elements. That rally is where assignment lives: its **action plan** creates real assigned tasks.
- **On mobile**, the New flag form attaches photos and files **while you're creating the flag**: take photos with the camera (it stays open shot after shot - cancel when you're done), pick several from your photo library at once, choose a file, or pull in **captures from your Pocket** (tap the ones you want, then **Attach**). Everything lands on the flag the moment it's created - each attachment files into the flag's Library folder, same as a web attachment, and a Pocket capture leaves your Pocket once it's safely on the flag. Mobile offers the same four urgencies as the web: Critical, High, Medium, Low.
- **Also on mobile**, the New flag form picks the **owner and collaborators** up front: you own the flag by default, or hand it to a teammate as you raise it - they're added as a collaborator automatically (so a private flag stays visible to them) and get notified. Collaborators picked at creation are notified the same way.

- Flags assigned before August 2026 keep those assignments (and their tasks) - the model changed going forward, nothing was taken away.

5.2 RFIs

At a glance - Ask a question that needs an official answer on the record. The asker is notified of answers and marks one official, which closes it.



The RFIs tab: every question and where it stands - who owes the answer, and what's overdue.

Quick answers

Raise an RFI

1. From the **RFIs tab** (or an element's page, which pre-fills the placement), click **New RFI**.
2. Write the question, set priority and due date, place it on its elements, and **attach files right in the create form**.
3. Submit it - the people who owe the answer see it routed to them.

Save it as a draft now, send it later

The create form submits by default - its button says **Create and submit**, because an RFI exists to be asked. If you want to finish the question later, tick **Save as draft instead** and the button changes to **Save as draft**. A draft is yours alone: it has a number and a home, but nobody has been asked anything yet.

To send it, open the RFI and click **Submit RFI** at the top right of the page, next to the status chip. A draft says so in plain words in the body too, so you can tell at a glance whether a question has actually gone out. The same spot is how an RFI moves on later (**Mark under review, Close, Reopen**) - the buttons beside the chip are always the moves available from the current status. They appear when you are the RFI's creator or an admin, and are hidden while you are editing - finish your edit, press Save, then move it.

Answers and comments - what each one means

Every reply on an RFI is one of two things, and you choose which when you write it:

- An **Answer** is a proposed answer to the question. Posting one notifies the requester and gives them a **review task** in My Tasks, and an answer is the only kind of reply that can be **marked official**. If you are the person the RFI is assigned to, your reply box starts on Answer.
- A **Comment** is everything else - discussion, context, a follow-up question, "checking with the vendor". Comments create **no review task and no review notification**, and a comment can never be the official answer. For everyone except the assignee, the reply box starts on Comment.

Replies support **@mentions**: type @ and pick a teammate (or **@team** for everyone on the RFI) and they're notified that you mentioned them - people who can't see the RFI are silently skipped, so a private RFI never leaks. Picked wrong? A comment can be turned into an answer with **Make this an answer** (by its author, the requester, or an admin) - it then enters review exactly as if it had been posted as an answer. In the thread, every reply carries a **colored stripe down its left edge** so the kind reads at a glance: **answers blue, comments gray**, and the **official answer green** - the same left-edge signal used on spark contributions and unread notifications - with a small answer/comment word beside each reply's timestamp - hover a reply for its Edit, Delete, and promote actions. An **Answers only** filter at the top of the thread hides the discussion when you just want the answers. The downloadable RFI PDF labels comments too, so the record stays readable.

Answer and close one

1. Post an **Answer** on the RFI. **The requester is notified** - a bell notification (it also flows into the daily digest) plus a **review task** in My Tasks, so answered RFIs don't sit forgotten.
2. The requester reviews the answer and **marks it official**, which **closes the RFI**.
3. The review task clears itself once the requester acts - marks an answer official, replies (an answer or a comment - either counts as acting), or closes the RFI. **If the answer doesn't actually answer the question**, reply asking for what you need: the task clears, the RFI stays open, and the next answer raises a fresh one. You can also just **tick the review task off** and leave the RFI open, for when you have read the answer and have nothing to add.

Who can do what

ACTION	WHO
See the RFIs tab	Anyone with RFIs access view or higher
Create, reply, mark official, close	RFIs access edit or higher
Reply to an RFI (answer or comment)	The person it's assigned to, the person who asked it, anyone added as a collaborator, or a project admin
Mark an answer official	The person it's assigned to, the person who asked it, or a project admin - and only an Answer can be marked, never a comment
Turn a comment into an answer	The comment's author, the requester, or an admin
Change the requester (whose question it is)	An RFIs tool admin, a project admin, or an org admin
Delete an RFI	Its creator, an RFIs tool admin, a project admin, or an org admin - its replies, links, and auto-generated tasks go with it
See a private RFI	Its creator, collaborators on it, and admins

Worth knowing

- **Sending a question outside Ceruleo?** Use **Official PDF** on the RFI page: it builds a formal Request for Information - your company and project up top, To/From, date issued, response-requested-by date, the question, enclosures listed, reference photos included, and a blank response-and-signature block - ready to email or hand to a building official or jurisdiction who will never have an account. No links back into Ceruleo appear on it. When their answer comes back, paste it into the RFI as a reply and mark it the official answer.
- **Categories come from a list you control.** Set the categories your company files RFIs under in **Company settings > RFI categories**; a project can add its own in **Project settings > RFI categories**, and its RFI form offers the company list plus the project's own. Picking one is then a dropdown rather than free typing, so the same category is spelled the same way every time. You can change an RFI's category later from its page (**Edit**, then the Category dropdown) - it used to be fixed at creation.

Two things worth knowing. **Until you add a category, nothing changes:** the RFI form keeps its free-text Category box, so a project nobody has set up yet still works. And **removing a category never rewrites history** - RFIs already filed under it keep the name they were given, shown in the dropdown as "*(not in the list)*" so you can see it is a leftover, and it stays put until someone deliberately changes it.


Company categories are managed by **company admins**; a project's own by its **project admins** or anyone with **admin-level RFIs access** on that project.

Filter the RFI list by category with the dropdown beside the priority filter on the RFIs tab. It lists the categories actually in use on that project's RFIs - including any typed by hand before the category list existed - plus **No category** for the ones never filed under anything. Two spellings of the same word (*Structural* and *structural*) count as one category rather than splitting the list. The count line reads "12 of 40" while a filter is narrowing things.

- **Hand an RFI to a new requester.** Every RFI has a **Requester** - the person whose question it is. The requester gets the answer notifications and the review task, their reply counts as "this answer is not enough", and they decide what gets marked official. It starts as whoever submitted the question. If the question really belongs to someone else, or the original asker leaves the project, an admin (RFIs tool admin, project admin, or org admin) can reassign it: open the RFI, click **Edit**, and pick a new name in the **Requester** row. Any open review task moves to the new requester with it, and on a private RFI the new requester is added as a collaborator so they can see their own question. The original creator stays on the record as the author and can still reply and mark answers official.
- **Collaborators are the people who should see the RFI** without being the one who owes the answer - and they have a voice in it: a collaborator can **reply**, with the composer starting on **Comment** and **Answer** one click away when they know the answer. Marking an answer official stays with the assignee, the requester, and admins. Add them in the create form under **More options**, or on the RFI's page any time. The **assignee is automatically a collaborator** too, so you do not need to add them twice.
- You can **assign an RFI to yourself** - you appear in the assignee list as "(me)". Owning the chase for an answer is a legitimate assignment.
- When an answer sparks a decision bigger than one person, **start a rally from the RFI** - the button sits at the top of the RFI's page, the rally links back to the question (each shows the other), and the RFI's elements carry over.
- An RFI is a **question with a paper trail**: every response is kept in thread order, and the official answer is marked as such in the record and on the PDF.
- **Flag vs RFI vs Note**: fixed → flag; answered → RFI; read → note. If nobody owes you an answer, it isn't an RFI.
- **Download as PDF** from the **Question** section (the **PDF** button beside **Official PDF**): the question, then every response in thread order with the official answer marked. The exports live on the question itself because they print the document - moving the RFI's status is a separate thing, done at the top right.

5.3 Notes

At a glance - Write knowledge down on clean paper: typed or untyped, with attachments straight from your Pocket.



The Notes tab: typed and untyped knowledge, written down where the next person can find it.

Quick answers

Write a note

1. From the **Notes tab** (or an element's page, which pre-fills the placement), click **New note**.
2. Pick a **note type** if your company defines them (a typed note carries that type's custom fields, worth it when you'll query the answers later), or write an untyped note - just a clean page.
3. Write on the **paper editor** - headings, lists, and formatting read back exactly as written, on the web and in the mobile app.
4. **Attach** as you go: **Add file** to upload, or **From Pocket** to pull a capture straight out of your mobile Pocket (it copies the file in and offers to remove the original from your Pocket).

Send a note to people (a transmittal)

Some note types carry a **people field** - a "TO" line. It works like an address, not a task:

1. Create the note as usual and pick the type that has the field (a **Project Transmittal**, say).
2. Click the field and choose the people it goes to. **Select all** addresses everyone at once, and you can search the list. It only offers people who have Notes access on that project.

3. Save. Everyone on the TO line is **notified that you sent them the note** ("Kelly sent you a project note"), and it buzzes their phone.
- Recipients can **open the note even if it is private** - being addressed is what grants them access.
 - Nobody is **sent it twice**. Editing the note later, or adding a name to the TO line, only notifies the people who are **newly** added.
 - Removing a name takes the note back: they stop being a recipient, and lose the access that came with it.
 - It is **not a task**. A transmitted note never appears in anyone's task list or as assigned work - it is delivered, not owed.
 - **Collaborators still exist alongside it** and mean something different: collaborators are the **cc** and the visibility list, recipients are the people the note is actually **addressed to**.
 - Note types **without** a people field are just plain notes, exactly as before.

You can do all of this **from the phone** too, on the New note form and in a note's **Edit details**.

Who can do what


ACTION	WHO
See the Notes tab	Anyone with Notes access view or higher
Create / edit a note	Notes access edit or higher
Delete a note	Its creator, a Notes tool admin, a project admin, or an org admin
See a private note	Its creator, collaborators on it, anyone it was sent to , and admins
Send a note / change its TO line	Its creator, a Notes tool admin, a project admin, or an org admin

Worth knowing

- **Flag vs RFI vs Note**: a note demands nothing from anyone - it informs whoever finds it. If it needs fixing, it's a flag; if someone owes an answer, it's an RFI.
- Place notes **on their elements** so a site walk finds them - "what did we agree about this thing?" should be one click from the thing.
- **Download as PDF** from the header: the branded page with the note's detail block, custom fields, body, and attached images as a gallery.

5.4 Tasks

At a glance - One plate per person: everything assigned to you, whether created by hand or raised by flags, RFIs, reviews, rallies, and schedules.



The project Tasks tab: everyone's open work, filterable by status - with auto tasks and manual tasks side by side.

Quick answers

See your plate

1. Open **My Tasks** (or **All tasks** inside a project). Everything assigned to you is there - **auto tasks** raised by other tools (an RFI reply to review, a flag or schedule item assigned to you, a rally action item) alongside **manual tasks** people created by hand.
2. Use the **"I assigned"** view for tasks you delegated to others.

Act on a task

1. Click it. It opens **the item it came from** (the flag, RFI, rally, schedule item) so you act in context.
2. No access to the source? The task opens in its **own view** where you can read it, join its **discussion**, and **mark it complete or reopen it** - and when you *can* see the source, that view links straight through.


Park a task, or write it off honestly

A task has four states: **Open**, **On hold**, **Void**, and **Complete**.

1. **On hold** - the work is paused on purpose and expected to come back. Open the task and click the **Hold** button (the raised hand) next to Mark complete; on a rally or sandbox action plan, hover the item and tap

the **raised hand** there. A note on *why* is required, and it shows wherever the task appears. Anyone who could complete the task can **Resume** it later.

2. **Void** - you've decided **not** to do it: circumstances changed, the task was overtaken by events, it's no longer needed. Choose **Void** and say why (required). The task leaves every open list but **stays on the record as not done** - crossed out with an X, never a lying green checkmark. That's the honest alternative to marking a dead task "done" just to clear the list.
3. Only person-made tasks can be held or voided (manual tasks, and rally, meeting, and sandbox action items). Auto tasks from flags, RFIs, and reviews keep following their source.



My Tasks: your one plate - everything assigned to you across projects, auto and manual tasks together.

Who can do what

ACTION	WHO
Create a manual task	Anyone on the project
Complete / reopen a task	Its assignee, its creator, or an admin
Put on hold / void / resume	Same people as completing - and only on manual tasks and rally, meeting, or sandbox action items

ACTION	WHO
Delete a manual or rally task	Its creator, or an admin
Delete an auto task	Nobody directly - it clears itself when its source item is deleted or settled

Worth knowing


- **Auto tasks tidy up after themselves:** delete the flag, RFI, or schedule item they came from and the task is removed automatically - no orphaned to-dos.
- **Don't close-and-recreate a task that slipped:** edit it instead (move the due date, reword it, reassign it) or put it **on hold** - re-typing it daily fills the Done pile with tasks that were never actually done.
- A held or voided task **stops counting as open** everywhere: dashboards, project tiles, overdue counts, and the daily nudge emails all leave it alone.
- An **unplaced** task (no element) means "applies to the whole project" - that's a meaning, not a mistake.
- **Prefer raising work inside its context** (a meeting action item, a rally action item) over bare manual tasks - the container's record then shows what it produced.

Completed tasks. Ticking a task off makes it leave your list, so there is a **View completed tasks** link under the list (and a **Completed** link on the dashboard beside *See all*) that opens everything you have finished, newest first. If you ticked something by mistake, **Reopen** puts it straight back. The list loads fifty at a time with **Show more**, so nothing ever becomes unreachable.

Some tasks close themselves. When an RFI is assigned to you, the task that appears is a reflection of that RFI - it clears when you answer. Those tasks have no checkbox, because ticking one would hide the reminder without answering the question. Open the record and do the work, and the task clears on its own. They still show in your completed list so the record of what you did is complete, but they cannot be reopened by hand.

5.5 Logs

At a glance - The site record, one dated log at a time: configurable entry types, attachments, follow-up flags, distribution, and honest backfill.



The Logs page: the day's entries by type, follow-up flags, and the distribute-at-night flow.

Quick answers

Record the day

1. Open the project's **Logs** and today's log (or **Start a log for this day**).
2. **Add entries** by type - Work performed, Weather, whatever your log's entry types are. Cadence is up to the log type: daily, weekly, or any rhythm - not all logs are daily.
3. **Attach photos and files as you fill the Add-entry form** - **Add files** from your device or **From Pocket** to pull captures from your phone; they upload the moment you add the entry, and you can add more afterward. (Pocket offers to keep or remove the originals.)
4. Issue / distribute per your log's setup - the calendar view shows what's been logged and distributed.

Raise a follow-up flag from the log

1. On a log with **Allow flags** turned on (Log settings), click **Add flag**.
2. Assign it - up to two people, limited to teammates with Flags access. The flag carries the log as its provenance.

Backfill a missed day

1. Open the past date - the log **calendar**, or the < **Previous day** arrow.

2. Click **Start a log for this day** and add entries like any other day. The log carries the correct **date**, so the record stays accurate. (Future days can't be started.)

Finding Logs. Logs sits in the project bar with the other tools, so it is one click away from any page in the project - not just the project dashboard.

Who can do what

ACTION	WHO
See / read logs	Logs access view or higher (only View-or-higher people can join a distribution list)
Add entries	Logs access edit or higher
Edit or remove an entry	The person who wrote it (or a log manager), while the log is still a draft
Issue, distribute, reopen	Log managers (issuing is web-only)
Configure entry types & options	Project/org admins (Log settings); company-wide defaults live on the company log template


Who counts as a log manager? Project admins, company admins, and any member whose **Logs** access is set to **Admin** (Project settings → Team members). That last one is the point: the person who runs the daily log - often a superintendent or site lead - can be given full charge of distributing, reopening, and tidying the log **without** being made a project admin.

Worth knowing

- **Attachments per entry type** are switched on once - on the **company log template** (Company settings → Log templates) so every project inherits it, or per project in **Log settings**.
- **Step through a day's photos in place:** click any photo to open it full screen, then use the **arrows** (on screen or on your keyboard) to move to the next or previous photo without closing - a counter shows where you are, and Escape closes.
- An **issued log is sealed**; reopening for a correction is possible (log managers) and stays on the record.
- **Download a day as PDF:** entries grouped by type with their fields and notes, follow-up flags, and photos placed under the entry each belongs to, numbered to match the distribution email.

5.6 Submittals

At a glance - The formal submission-review-approval record: a Submittal is permanent, each Submission is a version inside it.



The **Submittals** tab: each row a permanent record, with its current version, due dates, and reviewers.

The scheduled submission, review, approval, revision, and permanent record of work products (drawings, design packages, specs, renderings, PDFs, images, video/audio, supporting documents).

A **Submittal** is the permanent record; each **Submission** is a version inside it, so a submittal accumulates revisions over its lifecycle. (Full details in its own section below.)

Quick answers

Submit a version

1. Open the submittal from the **Submittals** tab (or create one: number prefix like **SUB-1** , title, elements, submitters and reviewers).
2. Add your **Submission** - it becomes the next numbered version inside the submittal; earlier versions stay on file, read-only, forever.
3. Reviewers are notified and review the current version; when a revision is requested, submit the next version - the record accumulates.

Find one fast

1. The **search box** matches number, title, description, and the names of files on it (the row tells you when the match came from a file or the description).

2. Narrow with the **status**, **role** (I submit / I review / I created), and **element** filters.

Who can do what

ACTION	WHO
See the Submittals tab	Submittals access view or higher
Create / manage a submittal	Submittals access edit or higher
Submit a version	The submittal's designated submitters
Review a version	Its designated reviewers (the review page is for the manager, reviewers, and recipients)
Delete	Nobody - submittals are archived (recoverable), never permanently deleted

The full lifecycle

A **Submittal** is the permanent record of a work product as it moves through submission, review, approval, and revision. Each **Submission** is a numbered **version** inside the submittal - when a revision is requested, a new version is created and the previous one stays on file, read-only, forever. Open the **Submittals** tab inside a project to see the list; each row shows the auto-number (an editable prefix like SUB-1), the title, the status, the current version, and how many reviews are in. **Search** the list with the search box at the top - it matches the submittal **number**, **title**, **description**, and the **names of files** on it (when a match comes from a file or the description, the row tells you so) - and narrow it further with the **status**, **role** (I submit / I review / I created), and **element** filters.

Permissions. Submittals is a standard per-project tool with a permission level (None / View / Edit / Admin) set per member, just like RFIs or Schedules. **View** lets you see submittals and be named as a submitter or reviewer; **Edit** lets you create and manage them; **Admin** (and the record's creator) can override outcomes, close, and archive.

Creating one. Click **New Submittal** and set: the **title** and an editable number **prefix**; a **description**; the **submission due date**; the **submitter(s)** (who owes the work); the **reviewers**; whether reviewers work as one **shared** review group (everyone sees each other's notes) or in **named, separate groups** (you name each group and assign reviewers to it; each group only ever sees its own notes); the **number of business days** after submission that reviews are due; any **linked project elements**; and whether it's **private**. You can drag and drop files right into the form. The submittal starts in **Scheduled** (awaiting submission).

Editing after creation. The creator and anyone with **Edit/Admin** can change a submittal after it's made - click **Edit** in the submittal's header. You can always update the **title**, **prefix**, **description**, **linked elements**, and **private** setting. While the current version is still **Awaiting submission**, you can also adjust the **submission due date** and **review window**, change the **submitter(s)**, and **add, remove, or regroup reviewers** (including switching between shared and separate groups and renaming groups). Once a version has been **submitted**

for review, the window and submitters lock - but you can still edit the title, description, elements, and privacy, and you can still **add or remove reviewers** right on the submission's Reviewers list (an added reviewer gets a review task with the remaining window; removals ask for confirmation). Changes are recorded in the audit trail.

Submitting. Each submitter gets a **"Submit ..."** task in **My Tasks**. Open the submittal, add or confirm the files, and click **Submit for review**. That locks the version, completes the submit task, sets every reviewer's **review due date** automatically (the business-day count from the submission date), and creates a **"Review ..."** task for each reviewer. The record moves to **Pending Review**.

While a version is **Awaiting submission**, the **file submission area is private to the designated submitter(s) and the creator/admins** - reviewers and other members can't see the files yet, just a clear "Awaiting submission" state. Naming **more than one submitter** doesn't split the work: it means **any one of them can submit** (whoever's ready), not that each must submit separately. Once it's submitted, the files become visible to everyone who can see the submittal so reviewers can review them.

Reviewing. A reviewer opens the submittal and chooses one decision - **Approved**, **Approved as Noted**, **Revise**, or **Rejected** - with a **required narrative explanation**. Their review task clears automatically once submitted. Each reviewer's decision and **narrative** appear on the submittal under the reviewer roster - but **group-walled**: in a **shared** review, reviewers see each other's responses; in **separate** groups, each reviewer sees only their own group's responses; the creator and project admins always see every review; submitters see none until they're distributed.

Marking up & commenting (the review surface). From an open submittal, reviewers and the creator/admins click **Open review & markup** to enter a full-screen review workspace over the submission's files. Pick a file on the left - a multi-page PDF lists its **pages right underneath it** in that same left rail, each with its note count, so jumping to page 3 is one click (the < > arrows in the toolbar still work too). Dropped a **pin** in the wrong spot, or on top of the very note it points at? **Drag your own pins** to a better position any time while the review is open - the number and the note come along. The center shows the file with the right tools for its type:

- **PDFs & images** - drop **pins** (numbered, each anchoring a note), draw **boxes, circles, arrows, freehand lines, and text**, zoom in/out, and page through multi-page PDFs. With the **Select** tool you can click any shape to **select** it, then **drag to move** it or use the corner handles to **resize** it; a selected shape also shows a **Comment** button that pins a note to it (the numbered note rides with the shape and is removed if you delete the shape), and a **Delete selected** button (or just press **Delete**). Pins are numbered in order and **re-sequence automatically** when one is removed, so there are never gaps. Pick a **markup color** from the toolbar's color dot (a compact palette plus a custom color picker; your last color is **remembered** so you don't reselect it each time) - picking a color while a shape is selected also **recolors that shape**, and **new pins take the picked color too** (pins placed before this feature keep the classic yellow, a **resolved** pin always turns green, and the pin's number flips to white on dark colors so it stays readable). Pin colors carry into the **downloaded marked-up PDF** as well. To get a clean view of the design, use the **eye at the top of the Notes panel** to **hide/show all markups** at once, or the **eye on an individual note card** to hide/show just that note's pin (and the shape it's attached to) - the note's own text tells you exactly what you're

toggling. These show/hide controls are available to **anyone viewing** the document, including recipients after distribution.

- **Video & audio** - scrub the timeline and leave **time-stamped notes** that appear as markers on the scrubber; click one to jump to that moment.
- **Office & other files** - open the file and leave **narrative notes**.

Every file has a **Notes** panel on the right, with the **add-a-note** box pinned to the top so it's always in reach: threaded comments you can **reply** to (press **Esc** or **Cancel** to back out of a reply), **@mention** people in, **resolve**, edit, or delete, plus **supporting files** you can attach **while writing the note** or to an existing note or the review in general (handy for a marked-up PDF or a reference). Attached images show as **thumbnails**, and clicking one opens it in a built-in **image viewer** (other file types open in a new tab). All of this is **group-walled** - reviewers see only their own group's pins, markup, notes, and attachments, and never another group's; the creator and admins can switch between groups to see everything. (Submitters still don't see review notes here by default; they only see what the creator/admin chooses to **distribute** to them - see **Distributing review notes** below.) Everything you add is saved as you go; click **Done** (top-right) or **Back** to return to the submittal whenever you like - "Done" just means done for now, you can always reopen the review and add more.

Outcomes. Once every reviewer has responded (or been removed), Ceruleo resolves the submission: **all Approved / Approved as Noted → approved; any Revise → a new version is required; any Rejected → rejected**. From there the creator can **Request a new version** (which re-opens scheduling with reviewers you can adjust), **Close** an approved submittal, or **Override & close** with a recorded reason. When you **Close** a submittal whose files were **approved** (every reviewer chose **Approved** or **Approved as Noted**), Ceruleo asks whether you'd like to **save the approved files to the project Library**. You choose **which** of the files to keep (not every submittal's files belong in the Library), can **rename** each one, and pick where they go - an **existing folder** or a **new folder** you name on the spot - or **close without saving**. Right in the same dialog you can also set **sharing**: mark the saved files **private** and/or pick who they're **shared with** (needed if someone with limited/**explicit** Library access should be able to see them). Saved files become normal **Library items** (downloadable by anyone on the project, and tagged as having come from the submittal), and they are automatically **linked to the same project elements** as the submittal - so the approved files show up under each of those elements' **Library** sections. (A Library item can be linked to **multiple** elements, the same way a submittal can.) Closed submittals can be **archived**, or **reopened** - if a submittal is closed by mistake, the creator/admin can click **Reopen submittal** and it returns to its active state. While a submittal is **closed**, the whole review surface is **frozen**: no one - not even the creator/admins - can add, edit, or delete markups or notes. **Reopen it first** to make any changes. Reopening also re-opens the review itself, so **every reviewer can revise their decision and narrative** - click **Edit** next to your recorded review, change your response, and **Update review**. A reopened review **stays open** (it won't snap shut after the first edit), so all reviewers can revise in one sitting; when everyone's done, the creator/admin clicks **Finalize review** to re-resolve the outcome from the updated decisions.

Distributing the review. The review stays private to the reviewers and creator/admins until it's **distributed** - and you can only distribute once **all reviews are in** (the submission has been resolved) or you've **overridden & closed** the submittal; the Distribute button isn't available on a pending review. The creator/admin opens

Distribute review and: (1) picks **whose review to share** - you share a reviewer's review *as a whole* (all of their notes, markups, and decision narrative), so you might share one reviewer's review and not another's; (2) sets the **recipients** - defaults to reviewers + submitters, and you can add or remove **anyone on the project**; (3) adds an optional **cover message**. When you share, Ceruleo **emails every recipient a transmittal PDF** - a clean cover sheet listing the shared reviewers' **decisions and narratives**, the **names of the files that were under review**, your cover message, and a link back to the submittal. (The individual markups/notes are deliberately left off the PDF - without the file open around them they lack context; the in-app Shared review is where recipients with access see the actual markups.) The recipients split into two kinds: anyone **with Submittals access** also gets a notification and a **read-only "Shared review"** in the app showing exactly those reviewers' notes and markups on the files, plus their decision narratives on the submittal; anyone **without Submittals access** still receives the **emailed PDF** but has no in-app access to the submittal. You can **distribute again** later with a different selection or recipients (an in-app recipient sees the combination of everything shared with them), and a distribution can be **revoked**. The audit trail records each distribution.

Deleting a submittal. If a submittal turns out to be unnecessary - a test, a duplicate, work that was cancelled - the creator or an admin can **Delete** it from the header (next to Edit). Deleting removes the whole record: every version, file, review, and its history, with a confirmation first, and it cannot be undone. For a submittal that ran its course and should stay on the record, **Close** or **Archive** it instead.

Late or unresponsive reviewers. Reviews can't be finalized until everyone has responded or been **removed**. The creator or an admin can remove a reviewer (with a confirmation); the removal is recorded in the audit trail.

Notifications. People are notified when a submission is submitted, a review is assigned, all reviews are complete, a revision is requested, and when a submission is approved, rejected, or the submittal is closed.


Watch a submittal to opt into the loud feed: watchers are also notified as **each reviewer's decision lands**, with a progress count ("2 of 3 reviews in") - everyone else hears only the milestones. Reviewers who have not decided yet are never told how others voted, even if they watch.

Audit trail. Every submittal keeps a permanent, append-only **audit trail** - creation, submissions, reviews and decisions, reviewer removals, overrides, and closure - each entry stamped with who did it and when.

(Saving approved files to the Project Library on close, and full-text search are planned follow-ups.)

5.7 Schedules

At a glance - Trees of tasks and milestones with dependencies and working calendars; dates recalculate themselves. Several schedules per project is normal.



A schedule's Gantt view: dependencies drawn between bars, milestones as diamonds, the critical path visible.

Plan the work as a tree of **tasks** and **milestones** with **dependencies** and a **working calendar**.

Each project can have several schedules.

Dates recalculate automatically as you change durations, dependencies, or the calendar. (Full details in the next section.)

Quick answers

Complete your item from the field

1. Open the schedule (web or phone - phones read schedules and complete items).
2. Find your item (your assigned items also sit in **My Tasks**) and mark it **done**. Downstream dates re-flow on their own.

Move a date

1. If the item has a **dependency**, don't fight it - its dates are derived. Change the predecessor, the lag, or the dependency itself (dependency editing lives in the **Cards** view).
2. If it's a **manual** item, edit its dates directly. Everything that depends on it recalculates automatically.

Read the schedule's health

1. The analytics strip on top: health score, critical-path health, overdue count, completed count. An overdue item surfaced honestly is the cheapest problem you'll have all week.

Who can do what

ACTION	WHO
See a schedule	Schedule access view or higher (drafts hidden from collaborators unless the schedule shares them)
Build / edit items and dependencies	Schedule access edit or higher
Approve or publish a version	The schedule's approver (the creator when none is set), its owner, or an admin - and in the lighter approval modes, direct publishing is open to the creator, owner, and admins (never plain edit access). See <i>Status, publishing & approval</i> below
Complete an assigned item	Its assignee (and editors) - phone-friendly
Update an item's progress on the phone	The item's assignee, the schedule's creator or owner, or an admin - tap the item and pick Not started / In progress / Blocked / Done
Attach files to an item	Anyone who can edit the schedule - files also appear in the Library under Schedules
Save a baseline	The schedule's creator, the schedule's owner, or an admin
Assign or change the schedule's owner	A schedules tool admin, project admin, or company admin
Change the approval mode	The schedule's creator, a schedules tool admin, a project admin, or a company admin - deliberately NOT the owner
Delete a schedule	Its creator or an admin - and the owner may also delete a draft that has never been published . Once a schedule has published versions, deletion stays with the creator and admins

Schedule owners

Every schedule can have an **owner** - the person who runs that plan day to day. A project manager might create a *Creative* schedule and hand it to the VP of Creative: from then on the owner has **admin controls over that one schedule** - editing items and settings, assigning people, approving, saving baselines, correcting completion dates - **everything except deleting it**. Ownership is per schedule, so someone can own *Creative* without gaining anything on *Fabrication*.

Only a **schedules tool admin, project admin, or company admin** can assign or change the owner (in the schedule's **Settings**, under Owner). The new owner is notified when they're handed a schedule.

Owners are watched out for, too: while a schedule has items that are **past their end date with nobody assigned**, its owner gets **one notification a day** naming the schedule and the count - so unowned overdue work can't sit silently. Assign the items (or move their dates) and the reminders stop.

Files on schedule items

Open any item's details and use **Files** to attach documents, drawings, or photos to it (attaching is part of editing, so enter **Edit schedule** first). Attached files also appear in the **Library** under the **Schedules** folder, in a subfolder named for the item - the same paper trail as Flags, RFIs, and Notes.

See a schedule in your own calendar

Every schedule can be **subscribed to from Google Calendar, Apple Calendar, or Outlook**, so the plan sits alongside your meetings without opening Ceruleo.

On the schedule page, open **Export** and use **Copy calendar link** under the Calendar heading. Then, in your calendar app, subscribe to that URL - in Google Calendar it's **Other calendars → From URL**; Apple Calendar and Outlook have the same "subscribe by URL" option. The schedule appears as its own calendar: each task shows as an all-day span across its dates, milestones show with a diamond, and the events never block your free/busy availability.

Worth knowing:

- The feed shows the **published** schedule - the approved plan your dashboards read - never unpublished draft edits. Publish a new version and the calendar follows.
- **Calendar apps refresh subscribed feeds on their own clock** - Google typically once or twice a day. A date change appears within a day, not within minutes.
- The link works **without signing in** - that's what lets a calendar app read it - so treat it like a shareable key to this schedule's dates. Anyone who can edit the schedule (or its owner or an admin) can **Reset link** in the same menu, which mints a fresh link and disables every previously shared one at once.

Progress from the phone

On the mobile app, open a schedule and **tap any item** you're allowed to update to set its progress: **Not started, In progress, Blocked, or Done**. Done works exactly like the web - including the "completing after the planned finish" question when relevant - and the swipe shortcut to complete or reopen an item still works.

The full picture

A **schedule** is a plan of work - a tree of tasks and milestones, scheduled across a working calendar. A project can have **multiple schedules** (e.g. *Design & Build, Install*), each independent. (How many depends on your subscription tier: **Base** includes one schedule per project, **Studio** includes three, and **Professional** and **Enterprise** are unlimited. Your tier also sets how many **active projects** your company can have at once - Base

3, Studio 10, Professional and Enterprise unlimited. If you hit a limit, archive or complete something, or ask your company admin about upgrading. Members are never limited on any tier.)

Before you can schedule: the project needs a **start date** and an **estimated end date** (set them in **Project settings**). Every task and milestone must fall inside that window.

Archiving & deleting a schedule. On the **Schedules** list, each schedule has a **...** **menu** (visible to the schedule's **creator** or an **admin**) with two options:

- **Archive** - tucks an old or finished schedule out of the active list while **keeping everything** (tasks, milestones, baselines, history). Archived schedules collapse under a **"Show archived"** toggle at the bottom of the list, and you can **Restore** one any time from the same menu. Use this for plans you're done with but want to keep for the record.
- **Delete** - permanently removes the schedule and all of its tasks, milestones, and assignments (a confirmation asks first). This can't be undone from the app, so if you're unsure, **Archive** instead.

Creating a schedule. Give it a title, an optional description, choose its **working days** and work hours, pick an **approver** (defaults to you), pick a **color** (or leave it on **Auto**), add **collaborators**, and optionally make it **private**. Under **More options** you can also set everything the schedule's Settings offers, right at creation: **Show draft to collaborators**, **Lock against incoming dependencies**, and the **holidays & non-working days** the schedule should skip. You can change all of these later from **Settings** (the gear on the schedule page). You can also **build a schedule from a CSV file** instead of starting blank - see *Importing from CSV* below.

View mode vs. Edit mode. A schedule opens in a calm, read-only **View mode** so anyone with access can scan it quickly. If you have edit rights, click **Edit schedule** (top-right, just above the table) to make changes, and **Done editing** to return to the clean view. *Even in View mode*, the people responsible for a task - its **assignees**, the schedule's **creator**, and **admins** - can update that task's **status** without entering Edit mode.

Tasks, milestones, and the tree. In Edit mode, add **tasks** (work with a duration) and **milestones** (a single dated moment, shown as a diamond). Nest items to build a work breakdown - **drag the grip handle** on a row to move it (in the Cards view, Edit mode), or use the **indent / outdent / reorder** row controls. **Dragging moves the whole subtree** - a parent carries its children with it. Drop **between two rows** to place it there as a sibling, or drop **onto a row** to nest it as that row's child; to pull a child out of its parent, drop it in a gap at a higher level. (You can't drop a row into its own descendants.) A parent becomes a **summary** row whose dates and status **roll up from its children** - its span is always exactly the earliest start to the latest end of its sub-items, and its **status reflects its children**: it reads **Done only when every sub-item is Done**, **Blocked** if any is blocked, otherwise **In progress** - and because it spans its children, if a sub-item is **overdue** the parent reads overdue too. (You don't set a summary's status directly.) If a move would change a parent's dates to span a new child, Ceruleo **asks you to confirm** (showing the old → new dates) before it drops; **Cancel** leaves everything where it was.

Milestone status & health. A milestone is a gate - it's simply **Pending** or **Achieved** (no "in progress"). Mark it **Achieved** when it's hit. Its **health** is then shown automatically by the **color of its diamond** (in the Cards view, Calendar, and Gantt) and a Cards-view badge: **green = Achieved**, **red = Overdue** (its date passed without

being hit), **amber = Due soon** (within 14 days), **teal = On track** (further out). So you can see at a glance which key dates are safe and which are slipping.

Durations in days or weeks. Each task has a duration you can enter in **days** or **weeks** - handy for long stretches (type "10 wk" instead of doing the math). A "week" means the schedule's number of working days (e.g. 5 for a Mon–Fri schedule).

How dates are set. You never pick a scheduling mode - Ceruleo infers it from dependencies. A task **without dependencies is pinned to your dates**: give it a **start** and a **duration**, and the **end is calculated** (in working days, skipping holidays and non-working days). You can also type the **end date** directly - with a start set, that adjusts the duration; with no start yet, Ceruleo back-calculates the start from the duration. In the **Cards view's Edit mode**, every card carries labeled **Start / End** fields on its right edge, so you can date the plan as fast as you can type - no need to open each item. A task **with a dependency is placed automatically**: its **start follows its predecessor** (marked with a small **Auto** badge in the Cards view) and its **end still comes from the duration you already entered** - you never enter it twice. **Adding a dependency** hands the item's start over to the predecessor (your duration is kept); **removing its last dependency** pins the item right where it sits.

Dependencies. Link items so one drives another - **Finish-to-Start, Start-to-Start, Finish-to-Finish, Start-to-Finish**, each with optional **lag** or lead. Dependencies can even cross schedules in the same project. When a predecessor moves, everything downstream recalculates. (Circular links are blocked, and a predecessor on a private schedule you can't see shows as a neutral placeholder.) You can manage them from **any view**: an item's **details panel** (opened by clicking the item on the Cards view, Calendar, or Gantt) shows a **"Waits on"** section listing its predecessors - with **Manage** (or **+ Add**) opening the dependency editor right there - and the Cards view's rows also keep their **link chip** shortcut. In the editor, an existing link's **type and lag/lead are editable in place** (pick a new relationship, or type a lag - negative for lead - and dates recalculate immediately); no need to remove and re-add it.

Lock against incoming dependencies. Don't want other schedules building on yours? In **Settings**, turn on **Lock against incoming dependencies**. While it's on, items in *other* schedules can't create a dependency on this schedule's items (those items simply won't appear as options in their dependency picker, and the link is refused if attempted). Dependencies *within* this schedule are unaffected, and any existing cross-schedule links stay. Turn it off if you want to wire a cross-schedule link from this schedule yourself.

Heads-up before you affect another schedule. Because dependencies can cross schedules, a date change here can shift items elsewhere. If you move, resize, or re-time an item - or delete one - that feeds into another schedule, Ceruleo shows a quick **confirmation naming the affected schedule(s)** before it applies, so you don't shuffle someone else's plan by surprise. Changes with no cross-schedule effect just apply normally.

The project window bounds every date. A schedule item can never sit before the project's **start date** or after its **estimated end** (both set in Project settings) - that's the outer frame every schedule in the project plans within. If a change would land an item outside it - a typed date, a drag, or a dependency change that back-schedules something (a Start-to-Finish link can pull an item well before its predecessor) - Ceruleo declines it with a calm notice naming the window, and **nothing is changed**: the link, dates, and everything

downstream stay exactly as they were. If you genuinely need to plan there, widen the project's start or estimated end in Project settings first.

Working calendar & holidays. In **Settings**, set the schedule's **working days**, **work hours**, and a list of **holidays / non-working days** (a date plus an optional label, e.g. "Site closed"). Work is never scheduled on non-working days, and changing the calendar re-spreads the whole schedule.

Per-item working-day override. Need one task on a different rhythm - say an **installation that runs Mon-Sat** while the rest of the project is Mon-Fri? In that item's detail panel, switch **Working days** to **Custom** and pick its days. A small calendar icon marks overriding items, and their durations count in their own working week.

Item details. Click the document icon on any row to open the item's panel: name, status, mode, duration, working-days override, and **Notes** - free-form details that live only here (they never clutter the table or other views). The panel shows a **live end-date** as you adjust things.

People & links. Assign one or more people (or crews) to a task - assignees get it on their **My Tasks** (and on the dashboard). You can only assign people who have **Schedule access** on the project, and assigning someone **gives them access to that schedule** (they join its team), so a private schedule stays visible to the people actually doing its work. Clicking that task opens its schedule and **highlights the item**, so you go straight to it. You can also link a task to project **Elements** so the plan ties back to the structure. Open an **Element** and its **Schedule** section lists every task linked to it across the schedules you can see - with status and dates, and a click to jump straight to the item in its schedule. Each row shows the item's **start-end dates**. **Overdue** items are flagged in **red** (with a ⚠ and an "Overdue" badge); items coming up within 14 days show a smart hint - **"starts in N days"** if the work hasn't begun yet, or **"due in N days"** once it's underway (milestones, being a single date, just read "in N days"). The section header shows an **"N overdue"** count so problems stand out at a glance. **Completed** items are hidden by default to keep the focus on remaining work - a **"Show completed (N)"** toggle reveals them.

The project Overview. Open a project and the **Overview** greets you with a cover banner carrying the project's **at-a-glance counts** - open **Tasks**, **Flags**, **RFIs**, **Submittals**, **Notes**, and **Rallies** - each a button that jumps to that tool (Rallies opens the Workshop). (The banner is the single place those counts live, so the page below never repeats them.) A tile with anything **overdue or critical/urgent** - say a submittal past its due date or a critical flag - shows a small **red urgency flag** with the count, so you can spot trouble at a glance. Each tile only appears if you have access to that tool. Below the banner, the Overview is about **priorities**, split two ways: **Needs your action** - things assigned to or awaiting *you* (tasks and flags/RFIs assigned to you, plus submittals where you're a pending reviewer or an awaiting submitter) - and **At risk** - project-wide items you can see that are overdue or due within 14 days but *aren't* yours (so nothing is listed twice). A graphic **schedule timeline** shows the next 14 days as a band: a dot for each item with a small flag naming it and how soon it's due (milestones as diamonds above the line, tasks as dots below), and anything **overdue** in red at the left edge; same-day items stack so a busy day reads clearly. The timeline only appears if you have **Schedules** access. **Recent activity** and the **Project Navigator** (the element tree) round out the page. The **whole-project Discussion** lives behind the slim teal **Discussions tab on the right edge** of the project page (wearing an

unread badge whenever there's conversation you haven't seen) - click it and the conversation slides out **over** the page, scrolling within its own panel, available from **every tab**. The same teal Discussions tab carries the discussion on flags, tasks, notes, sandboxes, and published assets - one color, one spot, everywhere a conversation lives. It's the same project-level conversation the mobile app's hero Discussion pill opens, so topics started on either surface meet. Everything here respects your permissions - a **private** item you can't see never appears, in the counts or the lists. (Your own assigned deadlines also live in **My Tasks**.)

Status, publishing & approval. The schedule you're editing and the schedule your team relies on are no longer the same thing. Approving (or publishing - same act) captures a **published version**: a frozen copy of the plan as it stands, numbered Version 1, Version 2, and so on. The team-facing surfaces - the dashboard's two-week lookahead, the project Overview timeline, and All schedules - read the **published version**, while editors keep working on the live copy underneath. Only published schedules feed those surfaces: a draft you're still sketching never appears on anyone's dashboard. The **Schedules tab calls the state out on every card**: a **Draft** chip on schedules that haven't been published, **Approved** on published ones, and an amber **Unpublished changes** chip when a published schedule is carrying edits the team can't see yet - so one glance at the list tells you where revision work is in flight.

- **Editing never un-approves.** Change dates or structure on a published schedule and it stays Approved - the header gains an **Unpublished changes since {date}** chip with a summary of the difference ("4 moved · 2 added · 1 removed vs published version"), and the Gantt shows each item's **published dates as ghost bars** behind your working bars, so you can see exactly what you're proposing. The team keeps seeing the published plan, undisturbed, until the revision is published.
- **Approval modes.** Each schedule chooses how much sign-off publishing needs, in **Settings**, under **Approval** (changeable by the creator or an admin - deliberately not the owner):
 - **Approval for every publish** (the default) - each new version needs the approver's sign-off. The approver sees **Approve** (or **Approve & publish changes** on a revision); everyone else routes through **Request approval**.
 - **Approval for the first version only** - version 1 goes through the approval flow; after that, the creator, owner, or an admin presses **Publish changes** directly.
 - **No approval** - no sign-off step at all. The button is plain **Publish**, and the badge reads **Published** instead of Approved.
- **The approver is optional.** When none is set, the schedule's **creator** is the approver. **Request approval** works as before - the approver gets a **task** in My Tasks and a **notification** - and it now also works on a published schedule that has unpublished changes. The approver (or an admin) can **Decline** with an optional note (the request clears; the working copy keeps its changes), and the requester can **Cancel request**. Direct publishing, where the mode allows it, is reserved for the **creator, the owner, and admins** - edit access alone doesn't publish.
- **Every publish is automatically a baseline.** Each version lands in the **Baseline** menu (Version 1, Version 2, ...) the moment it's published - compare against it, rename it, delete old ones. The current published

version can't be deleted, and drift comparisons default to it. No more remembering to save a baseline at approval; it happens by itself.

- **Unpublish** takes a published schedule back to a hidden Draft (its version history stays in the Baseline menu). **Drafts are hidden until first published:** while a schedule has never been published, only its **creator, owner, approver, and admins** (schedule tool admins, project admins, company admins) can see it - everyone else, including collaborators and members with schedule edit access, sees it once version 1 is published. To share a draft early, turn on **Show draft to collaborators** in **Settings** (or in the create form's **More options**).

Privacy. Like other items, a schedule can be **Private** - visible only to its collaborators (the **Access** list) and admins.

Project Overview timeline. On the **Schedules** tab, switch from **List** to **Overview** to see *all* the schedules you can access on one read-only timeline. Expand a schedule to reveal its tasks (bars) and milestones (diamonds), or leave it collapsed to see just its overall span. **Zoom** from **Day** through **Week, Month, Quarter, and Year** (or hit **Fit** to frame the whole project window), filter to **Milestones only**, and read the **Today** line and shaded non-working days at a glance. Use the **schedule chips** above the timeline to **hide or show** individual schedules. Each schedule has its own **color** - shown on its **List-view card**, its **span bar**, its **row rail**, and its **chip dot** - so the same schedule reads the same color everywhere. (Task bars keep their **status** colors.) It's reference-only - nothing here changes your plan.

All schedules (across projects). The left sidebar has an **All schedules** view: every schedule you can see, from *all* your projects, overlaid on one timeline - grouped **by project, then by schedule**. Use the **Filter** button to choose exactly what to show: a searchable, scrollable checklist (check a whole **project** or individual **schedules**, with **Select all / Clear**) - built to stay manageable even with many projects. A compact "**Showing X of Y schedules**" summary and a few removable chips sit next to it. So you can line up (say) the install schedules from several productions and compare them directly. Each project expands to show its schedules' bars; expand a schedule to drill into its tasks and milestones. Same zoom controls, colors, and Today line as the per-project Overview. **Your last view is remembered** (which schedules are shown, the zoom, and the milestones-only toggle) so you come back to exactly what you were comparing - and it's now saved to **your account**, so the same view **follows you across devices** (set it up on your laptop, open it on another machine, and it's already the way you left it). Read-only. **Overdue focus:** clicking the "**N overdue items**" chip on your Home dashboard's Schedule lookahead opens this view with an **Overdue items** panel above the timeline - every overdue schedule item across your projects, each showing how many days late it is and linking straight to it; **Clear filter** returns to the full view.

Four views (Cards / Grid / Calendar / Gantt). Switch with the toggle at the top of a schedule; your choice is remembered per schedule.

- **Cards** (formerly "Table") - the work breakdown as an indented outline. Each item is a **two-line row**: the **name** (with its task dot or milestone diamond and status) on top, and a clean metadata line beneath - **duration, assigned team, linked element, dependencies**, and (when a baseline is active) its **Δ** variance. The **start → end dates** sit in their own **right-aligned column** down the right edge, so they're easy to scan

no matter how long the names are. The metadata only shows what an item actually has, so rows stay readable no matter how much you attach. **Search, filter & sort.** A **search** box filters the list to items whose name matches, keeping each match's **parent chain visible (greyed) for context** so you see where it lives. The **Filter** menu narrows the same way by **Status, Needs attention** (overdue / at-risk), **Milestones only**, and **Assignee** (including a one-tap **Me**) - combine them freely; a count shows how many are active. A **Sort** menu reorders by **Due date, Start date, Status, or Name** - sorting happens **within the tree** (items reorder inside each parent; summaries reorder among their siblings by their rolled-up dates), so the work breakdown stays intact. Sort, search, and filters are **display-only**; your manual order is preserved and restored by choosing **Manual order** (structural editing - drag, indent, reorder - is available only in Manual order with search/filters cleared). In **Edit mode** the metadata line becomes editable (rename, set duration/dates, assign team, link elements, add dependencies), with reorder / indent / add / delete actions on hover, plus a quick **Auto / Manual** toggle by each item's dates (switch to **Manual** to pin a start date you type in, **Auto** to let durations + dependencies place it). To **add items quickly**, use the **quick-add row at the bottom**: type a task name and press **Enter** - it's added and the cursor stays put so you can keep typing the next one. (Use **Add milestone** for milestones.)

- **Grid** - the schedule as a **formal spreadsheet**: dense uniform rows with real gridlines and columns for the **outline number, name, status, start, end, duration (working days), dependencies, and assignees**, in outline order (parents bolded, milestones marked **◆**) - a **chevron on each parent row collapses or expands its sub-items** (numbering never changes; it's just hidden), in view and edit mode alike. In view mode it reads like a printed schedule. In **Edit mode** the cells go live - **click any cell and type: Tab** moves across cells, **Enter** commits and moves **down the same column** (spreadsheet style), **Esc** reverts the cell. The **Duration** cell carries the same **d / wk unit toggle** as the Cards view (the number reads and edits in whichever unit the item uses). The **Waits on** cell opens the **dependency editor** and the **Who** cell (small profile pictures + names) opens the **assignee picker** - the same editors as the Cards view. Dates follow the same rules as everywhere else: a typed **Start** pins the item, a typed **End** back-solves its duration, an item that **follows a dependency** shows its start read-only with an **auto** tag (move its predecessor instead), a **summary row's dates roll up** from its children, and **Done** items keep their dates locked. A **quick-add row at the bottom** appends items as fast as you can type - a **◆ Milestone** toggle beside it switches between adding tasks and milestones, **Enter** adds the row and stays put for the next name, **Tab** adds the row and moves **into its cells** so you can fill the whole row out in one pass. **Moving rows works by editing the outline number itself**: type where the row should live in its **#** cell - **5** makes it the fifth top-level item, **2.3** makes it the **third sub-item of item 2**, **2.3.1** nests a level deeper - and the row (with its whole subtree) moves there when you commit. Numbers are **insert-at**: taking an occupied number slots the row in at that position and everything below shifts down and **renumbers automatically**, so there's never a duplicate. The parent part of the number you type (the **2** in **2.3**) refers to the numbering **as you see it** on screen; nesting under a milestone or inside the row's own sub-items is refused with a plain message, and a number past the end of a family just lands the row last. A parent's dates and status keep rolling up from its children, exactly as in Cards (where drag-and-drop reordering still lives). **Deleting** an item lives in the **Cards** view; editing rights are exactly the same as every other view (edit rights + Edit mode).

- **Calendar** - a **Month** or **Week** grid. Tasks show as bars across their days, milestones as diamonds, and non-working days/holidays are shaded. Click any item to open its details. In **Edit mode** (if you have edit rights), **drag an item to a new day** to reschedule it - dates recalculate automatically. An item that **follows a dependency isn't draggable** (its start comes from its predecessor - move the predecessor, or remove the dependency from the Cards view, to place it yourself); its tooltip says so. You can also **add** here: use **+ Task / + Milestone**, or hover a day and click its **+** to drop a task right on that date - the new item opens so you can name it and set details.
- **Gantt** - the classic timeline: the task tree on the left, bars/diamonds/summary brackets on the right, with **dependency arrows** (all four link types) that always match the current links and **follow the bars live while you drag**, a **today** line, non-working shading, and **zoom** from Day through Year (or **Fit**). A **Details** toggle beside the zoom controls adds **Start, End, Who (assignees), and Status columns** to the task list on the left (remembered in your browser) - the timeline plus the hard numbers, without switching views. In **Edit mode**, **add** a **+ Task / + Milestone**, **drag a bar to move it**, or **drag its right edge to resize** (change its duration) - all write back and recompute. An item that **follows a dependency can't be move-dragged** (its start comes from its predecessor - drag the predecessor instead, or remove the dependency; the tooltip explains), but its **right-edge resize still works**. **Click any bar, diamond, or task name** to open the item's details - a clean click opens it, a drag moves it.

Assign from any view. An item's **details panel** (click the item on Cards, Calendar, or Gantt) now shows **Assigned to** - the current assignees, with a **Manage / + Assign** button in Edit mode that opens the assignee picker right there. No more returning to the Cards or Grid view just to hand someone a task. (Assigning still works exactly as before: the person joins the schedule's team, gets the task in My Tasks, and private-schedule pickers only offer people with access.)

Adding (from any view) opens a **draft** - the item is only created when you press **Create**, so closing the form leaves nothing behind. Editing a schedule from the Calendar or Gantt is gated the same way as the Cards view: you need edit rights (creator, schedule/tool admin, project admin, or org admin) and you must turn on **Edit mode**. Otherwise both are read-only. To **delete** an item from any view, open its detail panel in Edit mode and choose **Delete** (a summary deletes its sub-items too).

Overdue items are flagged. Any task or milestone whose planned finish has passed without being marked **Done** gets an **overdue flag** everywhere it appears - a red **⚠ marker** by its name, a red **outline** on its Calendar/Gantt bar (or milestone diamond), and a red row tint with red dates in the Cards view. Crucially, the item **keeps its status color** underneath, so an overdue item is distinct from one with the **Blocked** status (which is solid red) - and an item that's both blocked *and* overdue shows both. Mark it **Done** (or reschedule it) and the flag clears.

Marking work done (after the planned finish). When you set a task or milestone to **Done** *after* its planned finish - on the web or the phone - Ceruleo asks how to handle the difference:

- **Keep the planned dates** - just record it as done. The plan and baseline drift are unchanged, and the item simply shows **Marked Completed on [date]** - a record of when it was marked done, not a judgment on the work (marking Friday's work complete in Monday's meeting doesn't make the work late).

- **Update the end to the actual date** - the item's end moves to the completion date, dependent items shift accordingly, and **baseline drift grows** to reflect the slip. If that pushes past the project's estimated end date, the project end **extends automatically** to fit.
- **It was finished on another day** - record the date the work **actually** finished, rather than today. Useful when work is done in the field and written up days later. The item's end moves to that date, so the plan and the record agree, and drift is measured from it. **This option is limited to schedule tool admins, project admins and company admins**, and it is enforced by the platform, not just hidden - anyone else marking an item Done records it as completed today.

Correcting a completion date afterwards. The prompt above only appears when an item is completed **after** its planned finish, so an on-time completion is recorded as today without asking. To fix either kind afterwards, open the item's **Details** in **Edit mode**: a **Marked completed on** date sits under the dates and can be changed. Unlike the option above, this **only corrects the record** - the planned start and end are left exactly as they are, and are edited separately in the same panel if you want them to move too. This is limited to the same three admin roles; everyone else sees the completion date as read-only text.

Completing an item - even when it updates dates - does **not** count as an unpublished change or disturb the published version; publishing is reserved for deliberate re-planning of future work, and completions are progress, not re-planning. On-time completions are recorded without a prompt. (This is why a completion after the planned finish can now show up as drift - previously it changed nothing, because only the planned dates were stored.) Marking an item **Done** - and reopening it - is reserved for the item's **assignee(s)**, the schedule's **creator**, and **admins** (schedule-tool, project, or company), on the web and the phone alike.

Completed items are locked. Once a task or milestone is marked **Done**, its dates are protected everywhere - you can't drag or resize it on the Gantt, drag it on the Calendar, or edit its start/duration in the Cards view or detail panel. To change a completed item's dates, **reopen it first** by setting its status to something other than Done. (Its status, name, and notes stay editable, and dependency changes elsewhere can still shift it automatically.)

Baselines & drift. Save a **baseline** - a named snapshot of the schedule's planned dates at a moment in time - and then track how the live plan drifts from it. Use the **Baseline** menu (top of the schedule, next to the view toggle): **Save current as baseline** captures the plan (name it, or accept **Baseline 1**, **Baseline 2**, ...), and you can keep **several** baselines and **rename** or **delete** them. **Published versions appear in the same menu automatically** (Version 1, Version 2, ..., marked with a version chip; the current published one wears a green **Published** chip and can't be deleted), so a published schedule always has something honest to measure drift from. Only the schedule's **creator, a tool admin, a project admin, or an org admin** can save or manage baselines; anyone who can see the schedule can compare against them. Pick which baseline to compare with from the same menu (or turn it **Off**); your choice is remembered per schedule. When a baseline is active:

- A **progress-vs-baseline** panel appears at the top with three readouts: **Schedule drift** - how far the *current plan's finish* has moved from the *baseline's finish* (*days behind* in red, *days ahead* in green, or *on track*); note this compares plan to plan, so it reads "on track" until you actually change dates, not because work is or isn't getting done. **Completed** - % of tasks done, with a marker for what the baseline expected done **by**

today (flagged "behind pace" if your actual completion trails it). **Overdue** - tasks whose planned finish has already passed and aren't Done yet (with the worst lateness), which is the readout that tells you you're behind on *execution* even when the plan hasn't drifted.

- The **Gantt** shows a faded **ghost bar** (or ghost diamond for milestones) at each item's baseline dates, beneath the current bar, with a small day-variance label.
- The **Cards view** adds **Baseline** (baseline finish) and Δ (signed day variance) columns. New tasks added since the baseline read as *new*.

Tip: you rarely need to save a baseline by hand anymore - every publish captures a version automatically, and the comparison defaults to the current published version, so drift is always measured from the plan everyone signed off on. Manual baselines remain useful as extra reference points (say, a what-if you want to keep).

Exporting (PDF & CSV). Use the **Export** menu (top of the schedule, next to the view toggle):

- **Download CSV** - every task and milestone with all its data (WBS outline number, type, status, mode, start/end, duration, working-day overrides, assignees, linked elements, dependencies, notes, actual completion date, and - when a baseline is active - baseline start/end + variance). Opens straight in Excel or Sheets.
- **Download PDF - {current view}** - generates a clean, branded PDF of the view you're on (a real document, not a screenshot). Every page carries a header with the **project name, schedule title, approval status** (Draft / Pending approval / Approved), the **published version number** when there is one, and the **generated date**. A **Paper size** picker in the menu chooses the sheet - **Letter, Tabloid, A4, or A3** - and the Download button confirms it ("Download PDF - Gantt · Tabloid"). Until you pick, exports use your project's **paper size** (US Letter or A4, set in **Project settings**) with the **Gantt on Tabloid**; a pick is remembered in your browser. Layouts per view: **Cards** portrait, **Calendar** landscape (one page per month), and the **Gantt** landscape with **Start / End / Status columns** beside each task name and a **Gantt layout** choice: **One page** (the default - everything scaled to fit a single sheet; a big schedule stays whole, just smaller) or **Paginate** (full-scale pages that tile **across the timeline, then down by rows** - the task columns repeat on every page so each sheet reads standalone, every page states its **date range, row range, and tile number**, and bars crossing a page boundary continue on the next sheet). When you export the **Calendar**, a quick dialog lets you **choose the first and last month** to include (it defaults to the schedule's full span) so you can print just the months you need instead of the whole timeline. The file downloads straight away - no print dialog.

Every file Ceruleo generates for you downloads with a **consistent, sortable name**:

yyyyymmdd_title_type.extension - for example 20260614_kitchen-remodel_table.pdf , 20260614_kitchen-remodel_calendar.pdf , or 20260614_kitchen-remodel_schedule.csv . The date is when you downloaded it, so exports line up chronologically in your downloads folder. (Files you uploaded yourself keep their original names when you download them back.)

Importing from CSV. On the **Schedules** list, click **Import CSV** to build a **brand-new schedule** from a spreadsheet. The importer expects the same columns Ceruleo's **Download CSV** produces (**WBS, Name, Type, Status, Start, End, Duration, Predecessors, Notes**, and the rest), so the simplest workflow is: **export a schedule, edit it in Excel or Sheets, then import it back** as a new schedule. If you're starting from scratch


instead, use **Download a template CSV** (the link under the file picker) to grab a ready-to-fill file with the exact column headings and a handful of example rows - a summary with sub-tasks, a milestone, a dependency, and a working-days override - showing how each column should be filled in. It rebuilds the **task and milestone tree** from the WBS outline numbers and re-links **dependencies** by task name. Before it commits anything it shows a **preview** with the counts (tasks, milestones, summaries, dependencies) and any warnings. A few things to know:

- Every imported date must fall **inside the project window** (its start and estimated end). If any rows fall outside it, the import is **blocked** and lists the offending items - **widen the project dates in Settings, then re-import**.
- Imported tasks come in **manually scheduled** (pinned to their imported dates); switch any to **auto** afterward if you want dependencies to drive them.
- **Assignees and elements aren't imported** - add them after - and any predecessor that can't be matched to a task name is skipped (the preview tells you how many).
- Building a schedule this way counts against your plan's **schedule limit**, exactly like **New schedule**.

Watching & notifications. Click the **popcorn-bucket** icon beside the schedule's **Access** faces to **watch** it. While you're watching, you'll be notified whenever someone **changes a task's dates** in that schedule - and the notification takes you straight to the task. You're also notified automatically when **someone adds you** to a schedule (as a collaborator) or **assigns you** to one of its tasks, and - if you watch an individual task - when its **status** changes. Notifications are deliberately quiet: one change is one alert (no duplicates if you happen to be watching both the task and its schedule), and the routine date-shuffling that happens downstream when one task moves doesn't spam you - only the task someone actually edited does.

5.8 Project contacts

At a glance - The project's people directory: team members flagged as contacts plus external contacts (with privacy) for everyone else you deal with.



Project contacts: the people you deal with on this project - team members and externals alike.

Quick answers

Find someone's number on site

1. Open the project's **Contacts tab**. Two lists: **Project team** (members flagged as contacts, with contact role and email) and **Other contacts** (people not in Ceruleo at all - a vendor, a client, a city utility, an executive).

Add a contact

1. **A teammate**: flag them as a contact from **Project settings → Team members** - they appear under Project team.
2. **An outside person**: a project or org admin clicks **+ Add contact** - name, role/title, company, phone, email, notes.
3. Mark an external contact **Private** to share it with only the specific project members who need it (admins always see it).

Who can do what

ACTION	WHO
See non-private contacts	Everyone on the project (read-only)

ACTION	WHO
See a private contact	The members it's shared with, plus admins
Add / edit external contacts	Project or org admins
Flag a member as a contact	Project settings → Team members (admin)


Worth knowing

- Project contacts draw on the company **Directory** where one exists - company users with contacts access manage that from **Directory** in the sidebar.
- Keep the vendor you'll call at 7 AM in **Other contacts** with their cell number in notes - that's what this list is for.

5a. Meetings

Meetings turns a project's recurring coordination meetings (and one-off sit-downs) into a single, disciplined record: build the agenda ahead of time, send it out with a calendar invite, take notes topic by topic during the meeting, capture action items as real tasks, then issue the finished notes to everyone who attended - including people outside your company.

Find it on the **Meetings** tab inside a project.



The Meetings tab: issued records and upcoming agendas, numbered as a series.

At a glance - Agenda in advance, calendar invite out, notes topic by topic, action items as real tasks, minutes issued and sealed. Series carry open topics forward with their numbering.

Quick answers

Run the meeting cycle

1. **Before:** build the agenda on the meeting page (any member can add topics; only authors edit their own), then **Send agenda** - attendees get the PDF and a calendar file, external guests included, by email.
2. **During: Start meeting,** take roll call, write notes per topic, record **decisions** with who made them, add ad-hoc topics, and raise **action items** - they land as real tasks on real plates.
3. **After: Complete and issue** - the minutes PDF goes to every attendee and the record locks. Amendments are explicit **revisions**, never quiet edits.
4. **Next time (series): Create next meeting** carries open and standing topics forward with their numbering intact (3.01 = meeting 3, topic 1) and each topic's history attached.

One-off meetings and recurring series

New meeting creates either:

- A **one-off meeting** - a single agenda, run once, issued once.
- A **recurring project meeting** - tick *Recurring project meeting* and the meeting starts a **series** (weekly, every two weeks, monthly, or custom). Each occurrence is numbered (**Meeting 1, Meeting 2, ...**) and, when one wraps up, **Create next meeting** spins up the next occurrence with the roster and the unfinished business carried forward automatically.

Every meeting has a title, a date and time, a duration, a location and/or a **join link**, and an **organizer** (who runs it - by default, you). A meeting can also be marked **private** - then only its attendees (plus the organizer and admins) can see it at all; everyone else loses it from the Meetings tab entirely. Toggle privacy when creating the meeting or any time before the notes are issued, and a private series stays private as new occurrences are created.

Attendees - including people outside Ceruleo

The organizer builds the roster:

- **Team members** - anyone on the project with Meetings access. They get in-app notifications and can open the meeting in Ceruleo.
- **External attendees** - **search your Directory by name or email** (start typing and the list filters; only people appear - a company can't attend a meeting), or just type a name and email. Externals participate **entirely by email**: they receive the agenda and the issued notes as PDF attachments, and never need an account.
- **Timezones** - a meeting is scheduled **in its own timezone**: pick the zone alongside the date and time (it defaults to yours), and the meeting page, agenda, notes, and calendar invite all show that zone. A 10 AM Las Vegas walkthrough reads 10 AM PDT for everyone, wherever they are.

The organizer can also mark one or more attendees as **note takers** - see *Who can do what* below.

The agenda

Before the meeting, anyone on the project can **add a topic** to the agenda - but each person can only edit **their own** topics (the organizer and admins can edit anything). A topic can optionally be **assigned to an attendee** - internal or external - marking whose topic it is; the name shows on the agenda, in the meeting, and on both PDFs. Topics can also carry a **category** (Design, Operations, whatever fits - the picker suggests categories already in use): once any topic has one, the agenda groups topics under **category headings**, on the page and on both PDFs, with uncategorized topics gathered under *General*. Topics get a classic minutes number the moment they're created - **3.01** means *meeting 3, topic 1* - and **keep that number forever**, even when they carry to later meetings. The organizer can reorder the agenda.

Attachments ride along with the agenda: attach files to the **meeting itself** (the Attachments section under the roster) or to a **specific topic** (the + Attach control on the topic). When the agenda is emailed, every one of those files goes out **as a real email attachment**, and the agenda PDF **lists them** - meeting files in their own Attachments section, topic files on an "Attached:" line under their topic - so it's always clear what was sent. Anything too large to email stays listed and lives on the meeting in the app.

Send agenda emails every attendee the agenda as a **PDF**, with a **calendar invite (.ics)** attached so the meeting lands in Outlook / Google / Apple Calendar with one click. The meeting's **join link rides inside the calendar event itself** - it leads the event description, and when the meeting has no physical location it doubles as the event's location - so external guests join straight from their calendar without ever opening Ceruleo. Edited the agenda after sending? Send again - the email is marked *updated*, and a rescheduled date updates the calendar event rather than duplicating it.

A recurring series can also opt into **auto-send**: pick a lead time (say, 24 hours before the meeting) and the agenda goes out on its own - no one has to remember.

Running the meeting

Start meeting converts the agenda into live meeting notes:

- **Roll call** - tap attendees present or absent; attendance appears in the issued notes.
- **Notes per topic** - the note takers (plus organizer and admins) write rich-text discussion notes under each topic. For a topic carried from earlier meetings, the **previous discussion** from each prior meeting sits collapsed right above - the topic's whole history in one place.
- **New business** - topics that come up mid-meeting are added on the spot and numbered under the current meeting. An **Add topic** button sits both at the top of the agenda and **below the topic list**, so a note-taker deep in a long meeting never scrolls back up to record something new.
- **Action items** - under any topic, add an action item with an assignee and due date. These are **real tasks**: the assignee is notified and sees it in **My Tasks**, and its open/done state shows wherever the topic appears. Completing one from either side stays in sync. **Tick it off from anywhere it appears** - the meeting page, the project's **Tasks** tab, or **My Tasks** - and **issuing the notes does not close it off**: action items are the work that happens after the meeting, so they stay tickable on an issued meeting. Who can tick one: its **assignee**, whoever **added** it, or a **meetings admin**. If it isn't yours the checkbox is there but greyed, and hovering it tells you why.
- **Decisions** - each topic has a highlighted **Decision** field. When something is settled, record it there and pick **who decided** from the attendee list: decisions stand out in gold in the notes with the decision maker's name, follow the topic through its history, and the issued notes PDF opens with a **Decisions** summary - the meeting's decision log at a glance.
- **A side conversation on the meeting** - every meeting carries the same slim teal **Discussions tab** as the rest of Ceruleo, on the right edge of its page. Use it for the back-and-forth that isn't a topic note: chasing a missing attachment, agreeing a follow-up date, asking a question after the fact. It works on **past meetings** too, so a conversation can carry on after the notes are issued, and it stays separate from the minutes.
- **Finish later** - notes and decisions **autosave as you type**. Hit **Save & finish later** (or just close the page) and the meeting holds as a draft; come back whenever and **Complete & issue** when the notes are done.

Finishing a topic - or keeping it around

- **Complete** a topic when it's resolved - it stops appearing on future agendas. Its record stays put in every meeting where it was discussed (nothing is erased); it simply isn't carried into new occurrences. Completing also **unpins** a standing topic - Complete always means *retire from future agendas*.
- **Keep on agenda** pins a topic as a **standing item** - it stays on every future agenda until you unpin or complete it (think *Safety review* or *Schedule check* on a recurring meeting).
- Anything still **open** carries to the next meeting automatically as old business, number intact.

Issuing the notes

Complete & issue locks the meeting and emails every attendee the full record as a **PDF** - a **Decisions** summary up top, attendance, every topic with its notes, decisions, and history, and each action item with its owner, due date, and status. Internal attendees also get an in-app notification. You can also grab the same document yourself any time with **Download PDF** on the meeting page - the agenda before the meeting starts, the notes after.

Issued notes are **locked** - that's what makes them the record. Spot an error later? **Amend & reissue** reopens the notes for correction; issuing again sends everyone a clearly-marked **revision**. Reopened by accident?

Close without re-sending puts the notes straight back to issued - no email, no revision - so an accidental Amend never spams the room.

Who can do what

ACTION	WHO
Create meetings, add agenda topics	Any project member with Meetings access
Edit a topic's title/description	Its author - plus the organizer, project admins, org admins
Reorder, complete, or pin topics; manage attendees; send agenda; issue and amend notes	Organizer, project admins, org admins
Write discussion notes, take roll call	Designated note taker(s), the organizer, project admins, org admins
Add action items	Any project member in the meeting
Complete an action item	Its assignee, whoever added it, or a meetings admin
Attach files to the meeting or its topics	Organizer/admins (meeting files); a topic's author, the organizer, admins (topic files)
Delete a meeting (any status)	Company admins only - deleting removes the record, so it is deliberately a company-admin action


Meetings has its own column in the project **permission grid** (none / view / edit / admin), just like Flags or RFIs.

On mobile

The mobile app shows meetings read-only: the **Meetings** tile on a project opens the list, and a meeting shows its details, attendance, numbered topics with their notes, and action items. Running the meeting - agendas, sends, issuing - is desktop work.

Schedule Analytics & Health Metrics

At a glance - The analytics strip on every schedule: a health score, critical-path health, overdue and completed counts. Diagnostics, not grades - overdue surfaced early is the cheap problem.



The analytics strip rides the top of every schedule: health score, critical path, overdue and completed counts.

At the top of every schedule is a **Schedule analytics** dashboard - a set of metric cards that summarize how the plan is tracking. Every metric is fully documented: click the small "i" on any card to see its plain-language **definition**, the exact **calculation**, **why it matters**, a **worked example**, and the **data source** it reads. The goal is complete transparency: you can always see exactly how a number was produced and why a schedule is getting a particular health score.

How metrics are scored. Scored metrics use a 0–100 scale with three status bands, shown in green/yellow/red with a plain-language label: **On track (green, 90–100)**, **At risk (yellow, 70–89)**, and

Behind (red, below 70). A metric reads **N/A** when it can't be computed yet (for example, before a baseline is saved, or before any milestone is completed).

Baselines

A **baseline** is a named snapshot of the schedule's planned dates, captured at a point in time. **Publishing a schedule captures one automatically** (Version 1, Version 2, ... in the Baseline menu), and the comparison defaults to the current published version - so a published schedule always has a baseline to measure from. You can still save extra ones by hand from the **Baseline** control. Baselines power the drift, forecast, and baseline-variance metrics below - until one exists, those cards show a calm "Activate a baseline" hint and are left out of the Ceruleo Schedule Health Score.

Schedule Drift

How far the **projected finish** has moved from the **baseline finish** since the baseline was captured - in **either** direction. Calculation: projected finish minus baseline finish, in calendar days; $\text{score} = 100 - \min(100, | \text{days} | / \text{project span} \times 150)$. Moving earlier is also penalized, because a schedule that keeps moving is harder to coordinate around than a stable one. *(Needs a baseline.)*

Forecast Completion Variance

Whether the project is forecast to **finish on time** versus the baseline. Calculation: projected finish minus baseline finish; $\text{score} = 100 - \min(100, \max(0, \text{days late}) / \text{project span} \times 200)$. Only **lateness** is penalized - finishing on time or early scores 100. *(Needs a baseline.)*

Critical Path

The **critical path** is the chain of dependent activities - found with a full Critical Path Method pass (forward/backward through the dependency graph, all four dependency types plus lag, counted in working days) - that drives the finish date and has **zero scheduling slack**. **Critical Path Health** scores how much of that chain has slipped: $100 - (\text{fraction of critical activities overdue} \times 60) - \min(40, \text{average days late} \times 2)$. Slippage on the critical path moves the whole finish date, so this isolates the activities that actually matter for delivery.

Milestone Reliability

How reliably the project hits its **key milestones** (major gates like Concept Approval, Design Freeze, Fabrication Release, Soft Opening, Grand Opening). You can mark **any** schedule item as a milestone, at any level of the plan, and it's shown with a **diamond**. For each **completed** milestone, the score is based on how many calendar days late it finished: **on time = 100; 1-3 days = 90; 4-7 days = 75; 8-14 days = 50; 15+ days = 0**. The metric is the **average across all completed milestones**. For experience, exhibit, event, and production projects, these gates are often more meaningful than individual task timing.

The dashboard also shows the **number completed on time, the number completed late, the number upcoming, and the next critical milestone** (via the **Upcoming Milestones** card).

Overdue & Late Completion

- **Overdue Activities** - how many activities are past their planned finish but not yet Done (a live, today-vs-plan signal).
- **Late Completion Variance** - among activities already finished, how many completed **after** their planned end, and by how much on average (an execution-quality signal that predicts future slippage).

Ceruleo Schedule Health Score

A single 0–100 score for overall delivery readiness, intentionally weighted toward **executive outcomes** over individual task performance:

FACTOR	WEIGHT
Critical Path Health	30%
Forecast Completion Variance	20%
Milestone Reliability	20%
Overdue Activities	10%
Late Completion Variance	10%
Schedule Drift	10%

Any factor that isn't available yet (for instance, before a baseline is saved, or before any milestone is completed) is **dropped and the remaining weights are renormalized to 100%**, so a missing input never unfairly drags the score down. The card shows how many of the six signals contributed ("Based on N of 6 signals").

Customizing your dashboard

Different stakeholders care about different indicators, so the dashboard is yours to arrange. Click **Customize** (top-right of the dashboard) to:

- **Show / hide** any metric with its checkbox.
- **Drag to reorder** the cards.
- **Expand / collapse** an individual card (the chevron on each card) to show just its headline number or its full detail.


Your layout is **saved to your account for that schedule**, so it's there next time and on your other devices - and your choices don't affect what anyone else sees. Click **Reset to default** to return to the standard layout. (Available metrics today: Ceruleo Schedule Health Score, Schedule Drift, Forecast Completion Variance,

Critical Path Health, Overdue Activities, Late Completion Variance, Milestone Reliability, Open Activities, Completed Activities, Upcoming Milestones, and Baseline Variance Summary.)


Coming soon: a **Schedule Performance Trend** metric (health over time) and **project-level dashboard templates** (a default layout an administrator sets for everyone joining the project, which you can then personalize).

6. Studio, Sparks, and the Library

At a glance - Ideas wait as Sparks, take shape in Sandboxes, harden into Assets, and publish onto elements. The Library is the project's shelf; published assets file themselves there.



The Studio: the project's creative workspace - sandboxes as explorations, assets maturing inside them.



Inside a sandbox: the studio floor (sparks, assets, references) with the Questions and Action plan pinboard rail.

Quick answers

Take an idea to published work

1. **Spark it** - on the company board, or **New spark** inside a sandbox for one that belongs to this project from birth.
2. **Move it into a sandbox** when the project commits (**Move into a project** on a company spark; it *moves*, nothing is copied, and the audience narrows to the sandbox's team - the move screen names anyone losing access first).
3. **Develop assets** in the sandbox - documents, files, models - versioning as thinking hardens; request **reviews** when a version needs a verdict.
4. **Publish** the approved asset onto its element(s). Published assets also appear in the project Library under **Studio Assets**, so Library access is enough to find them.

Find published work without Studio access

1. Project **Library** → **Studio Assets** (read-only) - anyone with Library view can open published work.
- **Studio** is where visual/creative work lives:

- **Sandboxes** - working spaces for developing assets and ideas, with image **galleries** (Pinterest-style, with reactions and comments), file/link attachments, and collaborators. A sandbox's **References, Resources, and Tools** show as an **image-forward tile grid**: image files display their picture, galleries show a cover image with a count badge, and embeds, links, non-image files, and referenced **library items** (which render like their own type) show as icon tiles. Whatever **note** was written when a reference was added reads **on the tile**, under its name - a line or two on what it is and why it's there - so the reason travels with the reference. **Click a reference and it opens full screen**, with a column on the right carrying its **title**, the full **description**, what kind of thing it is, and who added it and when - and that's where you **edit** them, in place. A **link** opens as a live page inside that view where the site allows it; many sites refuse to be displayed inside another page, and for those you get a tidy card naming the page with **Open in a new tab** instead of a browser error. **CAD drawings** (.dwg, .dxf, .dwf) attached to a sandbox open in the same full-screen **drawing viewer with measuring and calibration tools** as on a rally. You can **add a CAD drawing directly as a reference** - the Add reference form's **File** option takes .dwg, .dxf, and .dwf alongside images and PDFs - so the base building drawings the team measures against can sit in References without becoming assets. **Working files from design tools** - .procreate, .psd, .psb, .sketch, .fig, .indd, and Affinity documents - can be added the same way, so the editable source sits beside the flattened export it came from. Ceruleo cannot display those formats (no web browser can), so opening one shows a short note - *This file type is not compatible within Ceruleo. Download and view externally* - with a **Download** button beside it; on a phone the same note offers **Open in another app**, which hands the file to the app that made it. The same note appears for any file Ceruleo cannot display, so you always get an answer instead of a blank screen. Working files can be very large, so if one is better kept in your own cloud storage, add it as a **Link** instead and the sandbox points at it without storing a copy. Hovering a tile reveals the same three actions it always had, now as icons: the **pencil** (its details), **archive**, and **delete** - and moving a reference to another sandbox lives inside the details, under the description. Galleries still open into the gallery itself; their pencil opens the details, with **Open gallery** a click away. **Galleries are collaborative**: anyone on the sandbox's collaborator list can open one and **+ Add images**, not only the person who created it, so a shared mood board can actually be built by the group. You can remove **images you added yourself**; the sandbox's owner and admins can remove any of them. (Images added before September 14 2026 have no record of who added them, so those stay with the owner and admins.) The Add reference form also takes a plain **Note** - a title and any text worth keeping with the sandbox (a vendor call summary, a decision, a field reminder); it shows as a small paper tile in the grid and opens full screen to read or edit like everything else. And a sandbox has its **own email address**: the **Email records** row at the bottom of the References card reveals it, and any exchange CC'd there lands on the sandbox as a readable chip (see *Working together* → *Documenting email in Ceruleo*). A sandbox's **owner** or an **admin** can **Archive** it to tuck it away (it's recoverable - **Restore** it anytime); archived sandboxes drop out of the Studio list and dashboards but are still there under the **Archived** filter. (A company admin can also set sandboxes to **auto-archive** after a stretch of inactivity - see *Templates & settings* → *Lifecycle & auto-archive*.)
- **The sandbox is a workspace** - the page is laid out as a studio floor with a pinboard. The **work column** runs Project Sparks, then Assets, then References; a **pinboard rail** on the right holds the sandbox's

Questions and **Action plan**, always in view while you work. An **empty section collapses to a one-line strip** (its name, a hint, and its add-button still in place) instead of a tall blank card, and grows back to the full card the moment it has content - so a fresh sandbox is compact, and the sections never change order. Four things to know:


- **The Brief.** Every sandbox can carry a proper brief - rich text, as long as it needs to be - behind the slim **gold BRIEF ribbon** under the sandbox's name, which shows the brief's opening line at a glance; click anywhere on it to read (or write) the whole thing, so the full text never crowds the workspace. The one-line **Summary** under the name (set in Edit) stays as the quick orientation; the brief is the deep context: what this sandbox is making, what done looks like, what constraints matter. The **owner or an admin** writes it. When the brief changes, the sandbox's **collaborators and watchers are notified**, and the Brief button wears a small **gold dot** for anyone who hasn't read the new version yet - open it once and the dot clears. The ribbon is on your phone too: the mobile sandbox screen shows the same BRIEF line with its dot, and tapping it opens the full brief to read (writing the brief stays on the web).
- **Questions.** Creative questions - the subjective kind that would feel absurd as a formal RFI ("does this material read as aged under warm light?") - get their own panel. Ask **the room** or direct a question **to a person** (they're notified). A question is a discussion under the hood, so clicking one opens its thread with full replies and @mentions. When something answers it, the **asker** (or an admin) marks it **Answered** - you judge what answered you - and it folds into the answered tail. **A Question is deliberately not an RFI:** no numbers, no transmittals, no official answer - nothing here is memorialized as a formal record. When a question needs an answer on the record (a spec, a jurisdiction, a client decision), raise an **RFI** instead - the panel says so right on the ask form.
- **The Action plan.** Real tasks, attached to the sandbox - same timeline, assignees, due dates, and the On hold / Void lifecycle as a rally's action plan. They show up in My Tasks and dashboards like any task and deep-link back to the sandbox.
- Ambient conversation stays in the sandbox's Discussions - the slim teal **Discussions** tab on the right edge of the page.
- **Project Sparks** - a sandbox's ideas, in their own format. A spark that moves into a project **doesn't get reformatted:** it stays a spark - same cover, same idea cards, same resources, same discussion, same page - it just lives in the sandbox now, at the top of the page above the assets. Open one and it works exactly as it did on the company board (the breadcrumb walks back to its sandbox). There are two ways one gets there. **Start it here:** **New spark** in the Project Sparks section creates a spark that belongs to this sandbox from birth - it never appears on the company Sparks board. **Move an existing one in:** on a company spark, **Move into a project** asks for the project and a destination (an existing sandbox or a new one), then moves it. The spark's own **creator** can move it, as can a **spark admin** or a **company/project admin**. Either way, nothing is copied. The important part is **who can see it:** once a spark lives in a sandbox, its audience is that **sandbox's team** - not the company. A spark everyone at the company could read becomes visible only to the people on that sandbox, and company-level **spark admins** no longer have any special claim on it (the project's admins do). Because that's a real loss of access, the

move screen **names the people who will lose it** before you commit, so you can add them as sandbox collaborators if they should keep it. Moving a spark is **not** a way to give someone project access - add people to the sandbox or project the normal way.

- **Assets** - when work is ready, it's **published** as an asset (documents, files, embeds, even 3D/GLB models with a viewer). File assets preview in-app - images and PDFs, and now **Adobe Illustrator (.ai)** files, which render (and can be **marked up** with pins/notes) just like a PDF, paging through multiple artboards. To remove a markup, select it and press **Delete** (or use **Delete selected** in the bottom bar). Asset **cards** show a page-1 thumbnail for .ai/PDF files too, so a sandbox is scannable at a glance. **AutoCAD (.dwg)** files can be uploaded and **downloaded** (in-app CAD preview is planned for later); anything that can't be previewed shows a download link. Published assets can be linked to elements and reviewed, and once published they also appear in the project **Library** under a read-only **Studio Assets** folder, so anyone with **Library** access (or Studio access) can find and open them - you no longer need Studio access to see published work. Every asset - including **video and audio** - has a **Details** tab (link it to project elements, add notes) alongside comments and reviews, so you don't have to open Edit just to assign an element. On video, audio, image, and PDF assets, **versions now live at the top of that Details tab**: the **Version History** card lists every version (switch, preview, restore) and carries the **Upload new version** button, replacing the old Versions button that hid in the bar under the asset. Document assets keep their own **Versions** tab in the same right panel. On a **document** asset, the **Edit** control sits right on the document, **next to Download PDF** (so it's clear you're editing the document itself, not the asset's details); the **Download PDF** produces a proper multi-page document - **portrait at your project's paper size** (Settings → US Letter / A4), with a full header on page 1, a running header on the following pages, and a footer on every page (project/document name, date, a **CONFIDENTIAL** mark, and "**Page X of Y**"), and you can **collapse the right-hand panel** (details / versions / comments / reviews) with the **Hide** button to work on just the page - click the slim **Details** tab on the right edge to bring it back. In fact, **every asset type's right panel collapses the same way** - video, audio, files/images, and 3D models all have a › button at the end of their tab bar that tucks the panel away so the content gets the full width, with the same **Show panel** rail to reopen it. The **Close** button is in the **top-right, after Publish**.
- **Placeholders** - a **place for an asset before the asset exists**. When the need is clear but the work hasn't started, add an asset and choose **Placeholder** as its kind: it takes a name and an **owner**, and deliberately asks for **no file type**, because the format is often still an open question. The card shows the owner's name and avatar with an **Awaiting work** chip, so a sandbox reads as what's planned as well as what's done. Open it and there's one thing to do: **Upload the first file** (or drag one onto the panel). **That first upload decides the format** - drop in a PDF and it becomes a file asset, an MP4 and it becomes a video - and it becomes **version 1**. From then on it behaves exactly like any other asset of that kind, and later uploads are **new versions**, not a change of format. It stays the **same asset** throughout, so the discussion, the element links, and the history all survive the transition - a placeholder grows into the asset rather than being replaced by one. The **owner**, alongside the creator, collaborators and admins, can upload that first file and new versions. Because there's no work in it yet, a placeholder **can't be published, moved out of Draft, or sent for review** - the controls say so and name the reason. If the

asset needs a **date**, put it on the sandbox's **Action plan**; placeholders deliberately carry no dates of their own.

- **Reviews** - request a review on an asset, then assigned reviewers open it to **comment and record a decision** (Approved / Approved with comments / Reviewed / Rejected). This works on **every kind of asset** - uploaded files, video and audio, embeds and Google Docs, rich-text documents, and 3D models
 - open the asset, go to its **Reviews** tab, and click a review to give your feedback - the comment box **grows as you type**, so a lengthy review is comfortable to write. Each reviewer also gets a "**Review: ...**" **task** in **My tasks** (home and project dashboards) until they record their decision - it clears too if the review is closed or they're removed. For an **explicit**-scoped member, a review invitation opens **only the asset under review and that review's thread**, not the rest of the sandbox. On **file and image** assets a reviewer can **draw markups** as part of their review (they go on the reviewer's own review layer), and on **video/audio** their **timeline comments** attach to the review the same way - reviewers see their own feedback plus co-reviewers' on the same request (each person's toggleable on/off - just for your own view), the requester sees it all, and the asset's regular working markups stay separate. **On your phone**, an asset's Reviews section shows each review with whose turn it is; open one to read the brief, see every reviewer's decision, add notes, and record or change your own decision - the same rules as the web (anything other than plain Approved needs a note). A closed review says so and can be reopened by its manager. Markup stays a desktop tool. When you request a review you can add **Instructions** - a short brief saying what to look at, what to ignore, and what finished looks like. It shows at the top of the review, above the reviewer list, so everyone reads the same ask before they start. Keep the **title** short and name-like; the instructions box is where the detail belongs.
- **Timeline comments** (video & audio) - open a video or audio asset and click the scrubber to leave a **time-stamped comment** pinned to that exact moment; the pins show on the timeline (red = open, green = resolved) and in the **Timeline** panel on the right. Open a comment there to **reply** (comments carry a full reply thread, and each row shows a **reply count**), **@mention** teammates in the comment or a reply (they're notified), and **mark it resolved** or reopen it. Anyone who can view the asset can comment and reply. The player's control bar has a **mute button and volume slider** next to the time readout - your volume level is remembered between assets and sessions.
- **Markups** - annotate an asset (rectangles, circles, arrows, clouds, text, and labels) right on the document or image. Your markups are color-coded as your own **markup set**, and you can create **multiple named markup sets** of your own (e.g. one per round of feedback) via **New markup set**; toggle each set's visibility to compare them. The eye on **someone else's** set hides it **just on your screen** (their view - and everyone else's - is untouched); the eye on your **own** set is the shared toggle. Each of your sets can be marked **Reviewable** or **Internal**: an **Internal** set is for the working team only and is **never shown to reviewers** - mark it Reviewable when you want reviewers to see it.




The Sparks board: the company idea pool - active, shelved, and converted ideas side by side.

- Sparks** - lightweight, **org-wide idea boards**. Capture an idea on "**The Idea**" card, and when one card isn't enough, build a **stack**: click **+ Add card** to add more cards behind it for the details, give each one an optional **heading** (leave it blank and it just shows its page number), and **page through the stack with the arrow buttons**. You can **reorder** the extra cards (the ← → move buttons) or **delete** them; The Idea always stays first. Only the people who can edit The Idea (the spark's creator, a spark admin, or an org admin) can add, edit, reorder, or remove cards. **Export PDF** produces a designed, presentation-ready document: page 1 opens with a **dark header band** carrying your **company logo and name** (light-colored logos read properly on the dark ground; SVG logos work too), a prominent **SPARK** mark in the upper-right with the **spark's owner** beneath it (their profile picture when they have one, otherwise their name), and the spark's **title** large in white - with the **cover image as a compact photo card docked on the band's edge** (breaking onto the page with an offset navy slab behind it; shown exactly as framed at upload, never re-cropped), then **The Idea** (its first paragraph set slightly larger, as a lede) and every card in order, each section opened by a small **gold accent bar** and its heading. Later pages carry a slim version of the dark band, and every page has a **CONFIDENTIAL** mark and page numbers in the footer. There's no spark number and no Ceruleo branding on it - it's your company's document. Gather **contributions** (thoughts, questions, concerns, opportunities, feedback) and reactions, build a mood-board, export "The Idea" as a PDF, and **move a spark into a project** when it's ready to develop - it becomes a **Project Spark** inside a sandbox (see Studio → Project Sparks), keeping everything it already has. Sparks that started at company level stay on this board afterwards, marked **Converted**, and open into the sandbox they moved to. As a team rallies around a spark, **filter the contributions by type** - tap the type chips above the stream (All,

Thoughts, Questions, Concerns, Opportunities, each showing a count) to focus on, say, just the concerns or just the opportunities instead of one long list. Tap more than one to combine them. When someone **replies to your contribution** or **@mentions** you on a spark, the notification takes you **straight to that feedback and briefly highlights it**, and any contribution you've been alerted to but haven't seen yet is flagged **New** until it's been on your screen for a moment. **Reply to any contribution - or to a reply**: the thread stays flat under the contribution that started it, in time order, and a reply that answers another reply carries a quiet **L**, **name** marker - click it to jump to the post it answers, briefly highlighted. On a contribution **or reply** you can **attach files and links directly** (the **+ File / + Link** buttons) and/or **reference an existing resource** - directly-attached items stay **on the post** as clickable, previewable chips and are **never** added to the spark's Resources (so casual posts don't clutter it), while a reference points at a real resource. Each of a spark's **Resources** can carry a **note** - what it is, why it matters, what to look at - asked for when you add a **link** or an **embed**, and addable to anything later. **Galleries are collaborative**: anyone who can contribute to the spark can open a gallery and **+ Add images**, even one somebody else started, so a mood board gets built by the group rather than by its creator alone. You can remove **images you added yourself**; the person who created the gallery, and admins, can remove any of them. (Images added before September 14 2026 have no record of who added them, so those stay with the gallery's creator and admins.) **Click any resource and it opens full screen**: the image, PDF, video or audio fills the screen, and the column on the right carries its **title**, the full **note**, its type and size, and when it was added - with **Download** and (on a rally) **Mark up** as actions. **Edit them right there**, in that column. A **link** or an **embed** opens as a live page inside the same view where the site allows it; many sites refuse to be displayed inside another page, and for those you get a tidy card naming the page with **Open in a new tab** instead. To go straight to the details, hover any reference - an image tile, a gallery, a file row, an embed - and click the **pencil**; it's the same pencil, in the same place, on every surface. An image also carries its note **on the tile**. Adding a **single** file opens its details immediately so you can write the reason there and then. Click a chip to **preview in-app** (images, video, audio, PDFs) or open links/other files in a new tab; image attachments also show a thumbnail. You can **edit your own contribution or reply** any time (the quiet **Edit** under it - reword it and, via the type chips atop the composer, change its **type**) and **delete** it - admins can delete anyone's; deleting a contribution removes its replies and its directly-attached files (resources it merely references stay). Star sparks to keep your favorites on top. Sparks can be sorted into **categories** (defined by an admin in **Company settings → Spark categories**): pick one when you **create** a spark, or set/change it later from the spark's **Settings**. The category shows on the spark's card (an uncategorized spark reads **"Idea"**). On the Sparks board, filter by **Category** and by **Owner** right next to the search box - alongside the existing status filter (All / Active / Shelved / Archived / Converted) and the sort options.

- **Shelving a spark** - between active and archived there's the **shelf**: **Shelve** (in the spark's header) parks an idea deliberately - not now, but worth keeping - without burying it in the archive. A shelved spark freezes like an archived one and wears a gold **Shelved** badge, but coming back is one click: **Take off the shelf** (the spark's creator or an admin - unlike unarchiving, which takes an admin). You can also archive straight from the shelf when a parked idea is truly done. The shelf is for company sparks only - a Project Spark already lives in its sandbox.



The project Library: your folders plus the automatic system folders - every attachment files itself.

- Library** - the project's curated files and documents, organized into folders you create. Anyone with **Edit** access to the Library can create folders and subfolders; you can **rename or delete the folders you created** yourself, while someone else's folders (and the automatic **system** folders) take a Library, project, or org admin. Deleting a folder never deletes the items inside - they move back to **Uncategorized**. (The company-wide library under Company settings is managed by **company admins** and by members given the **Manage library** permission - granted per person in Company settings > Members. That permission covers the whole job: creating and renaming folders, adding files, and deleting - in the company library only, never in any project's library.) Files you attach to a flag, RFI, note, or log live on that item **and** file themselves into the Library automatically: each tool has a read-only **system folder** (Flags, RFIs, Notes, Logs) holding a subfolder per item - "Flag #12 - Cracked panel" and so on - so the project's full file record is browsable in one place. These automatic folders can't be renamed, moved, or added to by hand (attach to the item instead), and privacy carries through: a **private** item's folder, its files, and even its number and title appear only to the people who can see that item - everyone else simply doesn't see the folder. Items show a **thumbnail** (a preview for images - and now a rendered page-1 preview for **PDFs** and **Illustrator .ai** files, and a **play mark** on anything that opens in a player - a type icon otherwise) so you can scan a folder at a glance. Open a file and you get an **in-app preview**: images and PDFs as before, and now **Adobe Illustrator files (.ai)** render right in the app (they're read as PDFs, so you can page through **multiple artboards**) - no need to download and open Illustrator just to look. **Video and audio play right there in the item, with nothing to download first**: open the file, press play, and you can scrub straight to any point instead of waiting for the whole file to arrive. Nothing starts on its own; the player waits for you. This covers the

formats browsers can play (.mp4, .mov, .m4v, .webm for video; .mp3, .m4a, .wav, .aac, .ogg for audio). A format no browser can play (.mkv, .avi, .wmv and the like) still shows a download button, and so does a file whose contents your particular browser can't decode - which tells you so rather than sitting there blank.

A video uploaded from the web app is prepared for streaming on the way up, which is what lets a long file start playing immediately; videos added before this, or added from the phone, may take a moment longer to start. The **Download** button at the top right is unchanged if you'd rather keep your own copy. (On the **phone**, a Library video still downloads to the device before it plays.) **AutoCAD drawings (.dwg)** can be stored and **downloaded** here today; an in-app CAD preview (and measuring) is planned for later. If an unusual .ai can't be rendered, you'll get a clean download link instead. Every item has a **Shared with** list (open the item, or set it when you add one). For a **private** item, that list is exactly who can see it - only those people, plus you and admins. For a normal item, everyone with Library access already sees it, so the list is how you make sure someone with **limited (explicit) access** to the Library can see it too - add them and they will. Only people with Library access can be added either way. The item's **Usage & details** note - a place to record how it's used, versions, source, who to ask - is the same field whether you fill it in when adding the item or later from the item's details. Approved files you **save from a submittal** land in the Library folder you choose and behave like any other Library document. Published **studio assets** also surface here automatically, in a read-only **Studio Assets** folder - open one to see **just the asset itself** - its **preview** (images, PDFs, embeds, **video/audio that play inline**, and **in-app documents rendered read-only with a Download PDF at your project's paper size**), a **download** for other file types, and its **linked elements**. Its **previous versions, discussions, comments, and reviews stay in Studio** and are visible only to **studio members** - so a Library-only teammate sees the finished asset, not the studio workspace around it. You can't rename, move, or delete these from the Library (manage them in Studio). The folder shows to anyone with **Library** access; whether you can actually open an asset follows **Library-or-Studio** access. A **privately published** asset (one whose audience was limited to specific people when it was published) appears here **only for its audience** - marked with a small lock - and stays completely hidden from everyone else. You can browse the Library two ways with the **Folders / All items** toggle: the **Folders** tree (as above), or an **All items** view that lists everything across every folder - real files and Studio Assets together - and lets you **group** by category, linked **element**, or **kind**, **filter** by element and kind (tap the chips), and **sort** by name or newest. The **Add item** form takes **several files at once** - drag and drop them onto it (or click to pick several, and keep adding more before you save); each file becomes its own Library item **named after its file**, and the description, Private setting, and Shared-with choices apply to the whole batch. This works in the project Library and the company library alike. **Working on several files at once:** in a project Library folder, hover a file and a **checkbox** appears - tick a few, or use **Select all**, and an action strip above the list lets you **Move** them to another folder, **Share** them with people in one action, **Make private** or **Make visible** to the project, or **Delete** them together (recoverable from the **Delete log** for 7 days, like any deletion). Checkboxes appear only on the files you can change - the ones you created, or all of them if you administer the Library - so a batch never partly applies. Changing who files are **shared with** needs to be the file's creator or a project/company admin; if some of your selection doesn't qualify, the panel says which and leaves those files alone. The automatic system folders and published Studio assets can't be batch-edited, and files can't be moved into the company-shared library

this way. **On a single file**, the same two changes live in the file's own details panel: **Location** carries a **Move** link listing the project's folders, and **Visibility** reads *Private* or *Visible to the project* with a **Make private / Make visible** switch. Both are for the file's creator or a Library / project / company admin.

Who can edit what (across every tool). Having a tool's **Edit** permission lets you **create** items and **edit the ones you created** - it does **not** let you change someone else's. Editing another person's item (its details/settings) requires a **tool admin, project admin, or org admin**. This is consistent everywhere - submittals, schedules, rallies, library items, studio assets, flags, RFIs, notes, sparks, and so on. Being an **assignee** or **collaborator** doesn't grant edit of an item's details - but it never takes away your ability to **participate**: you can always comment, react, take part in discussions, and (if you're an assignee) change status / mark your assigned item done.


Studio specifics. In a sandbox, the **sandbox owner**, an **asset's creator**, and **admins** (a **Studio-tool admin**, project admin, or org admin) can edit asset/sandbox **details, change an asset's status, and publish it**. Before an asset can be **published** it must be **Approved** and **linked to at least one element** - the **Publish** button stays disabled until both are true. A **placeholder** sits below all of this until its first file arrives: it can't be approved, published, or sent for review, because there is nothing yet to judge. **Publishing now asks who the audience is: Everyone on the project** (the default - anyone with Library or Studio access, as before) or **Specific people only** (a private publish: pick exactly who can see it; a lock marks it in the Library). Either way, the **working team** (the asset's creator, the sandbox owner and collaborators), **reviewers on it**, and **admins** always keep access - a private publish never hides work from the people making or reviewing it. The same picker is also how a member with **explicit (need-to-know) access** gets a *public* published asset into their view: add them and they'll see it. The audience is **editable afterwards** from the asset's **Details tab → Published audience** (creator or admin), and unpublishing **remembers** the audience for the next publish. Because an **Approved** asset is frozen, the copy people see in the **Library** never changes underneath them. A published asset stays **pinned at Approved** (its status dropdown is disabled while published) and must stay linked; to make **any** change - edit the content, change the status, or swap its element - **unpublish** it first, then move the status off Approved, edit, re-approve, and republish. (There's no separate "approver" role - that was removed for being confusing; setting the status, including marking something **Approved**, is just part of managing the asset.) **Approving an asset locks its content**: once its status is **Approved**, the asset is a signed-off record - its **document body, new-version uploads, markups**, and **new review requests** are all frozen (its name, notes, element links, and published audience stay editable, and it can still be published). A reviewer on a review that was **already open** can still finish their review - including their review markups and timeline comments. To change an approved asset, a manager moves its status **off Approved**, edits, then re-approves. (Publishing controls only Library visibility - it doesn't lock anything.) **Collaborators can add their own assets** to the sandbox (and then manage only those), can **take part in the sandbox discussion**, and - for a **document** asset - can **edit the document body (the text)**, but not its details (name, type, element links, publishing). A collaborator can also **upload a new version** of an asset **only if "Let sandbox collaborators upload new versions" is turned on** for it - otherwise new versions stay with the creator, sandbox owner, and admins. Set it with the checkbox on the **Add-asset** form, and a manager can **toggle it on or off any time afterward** from the asset's **Details** tab. Discussions stay open to collaborators and admins, and when an item is public, to anyone who can see it.

(**Workshop** lives inside each project - see section 5, "The core tools inside a project".)

7. The Workshop and Rallies

The **Workshop** is where a project's team works a problem out together. A rally can be **linked to one or more project elements** - pick them (optionally) when starting the rally or later from its **Settings** - and linked rallies appear on those elements' pages alongside their flags and RFIs. A rally with no links wears a quiet **No element** tag as a reminder; linking is never required. Where **Studio** is for producing the work itself (drawings, documents, media, models), the Workshop is for the thinking around it: operational planning, production coordination, and solving whatever's in the way. The two sit side by side as the project's creative spaces.

Finding it. The Workshop is its **own area**, not a tab. Open a project and it's one of the two prominent cards near the top - **Studio** and **Workshop** - each showing what's inside (sandboxes and assets, or active rallies) and taking you straight in. The Workshop wears a distinct warm-toned background, so you always know where you are.



The Workshop: the project's active rallies, each one a cause the team is working out together.

At a glance - A rally is a focused push on one question: a brief, typed contributions, an action plan of real

tasks, resources, and a close. Link it to its elements so the answer stays findable.

Quick answers

Start a rally

1. From the Workshop board, **New rally** - title, **brief (required)**, optional cover, optional **elements** (linked rallies show on those elements' pages; never required, but a rally about a specific thing should link it). The brief is a short statement of what this is about and why it matters - a rally without one is hard to join, so it can't be skipped. Refine it any time on the rally page.
2. Or spin one **from a flag** when the fix outgrows one person - it inherits the flag's elements.


Work it

1. Post **contributions** by type (update, question, blocker, idea, feedback); reply, react, attach files or reference resources.
2. Build the **Action plan** - who's doing what, when; assigned items land in the assignee's tasks, unassigned ones read as events.
3. Gather **Resources**: files, links, embeds, video, audio, 3D models, or documents written right on the rally.

Close it

1. When the question is answered, **Close** the rally with its outcome. The record - and the element pages it's linked to - keep the whole story.

Rallies



A rally: brief ribbon up top, the conversation front and center, and the action plan and resources in the rail.

A **Rally** is the Workshop's record: one topic the team is rallying around. A rally has a **cover image**, a **brief** (the rich-text statement of what this is about), an **owner**, and a **collaboration team**. The page leads with the conversation: a slim header (cover thumbnail, number and title, status, people - and, when the rally came from a **flag** or **RFI**, a **red or blue chip right in the header** linking to it), then the **gold BRIEF ribbon** showing the brief's opening line (click it to read or edit the whole thing), then the **contribution stream** front and center with the composer on top. The **Action plan** (its full date timeline) and **Resources** sit in a rail on the right, always in view while you read - the same floor-and-pinboard layout as a Studio sandbox. It can be **private** - visible only to its owner, its invited collaborators, and admins - or open to everyone on the project. **Star** a rally to keep it near the top; **archive** it when it's done (admins can restore it).

A rally can also be started **from a flag** - when something raised on the project needs a team to solve it rather than one assignee. The rally then shows that flag as a card you can open, and the flag lists the rallies working on it. A rally can be linked to project **elements**, shown as chips in its header; a rally started from a flag inherits that flag's elements.

Contributions: the conversation

The heart of a rally is its **contribution stream**. Post an **update**, a **question**, a **blocker**, or an **idea**, and reply to anyone else's - including **to a reply**. The thread never burrows: every reply sits at one indent under the contribution that started it, in time order, and a reply that answers another reply carries a quiet L, **name**

marker naming who it answers - click it to jump to that post, briefly highlighted. Contributions are typed on purpose: tap the type chips above the stream to focus on just the blockers, or just the questions, instead of reading one long thread. Tap more than one to combine.

React to any contribution - 👍 Agree, 💡 Builds On This, 🎯 Important, 🔥 Love It, 😞 Needs Work, 🚀 Explore Further - as a quick way to register agreement without another post.

You can **edit your own** contribution or reply at any time (the quiet **Edit** beneath it - reword it, and use the type chips to re-type it: an Update can become a Question), and **delete** your own. Admins can delete anyone's. Deleting a contribution takes its replies and its directly attached files with it.

@mention teammates to pull them in, and type **@team** to mention everyone on the rally at once. When someone replies to your contribution or mentions you, the notification takes you **straight to that post and briefly highlights it**. Anything you've been alerted to but haven't seen yet is flagged **New** until it's been on your screen for a moment.

You can **attach files and links straight to a post (+ File / + Link)** or **reference an existing resource**. Attached items stay **on the post** as clickable chips and are deliberately kept **out of** the rally's curated Resources list, so a casual post doesn't clutter it. Click a chip to preview images, video, audio, and PDFs in-app; other files and links open in a new tab. Deleting a post never deletes a resource it merely referenced.

The Action plan

Every rally has an **Action plan** - a lightweight **timeline** of *who's doing what, when*. Each action item has a title, an optional **assignee** and **date**, and free-form **notes**. Items group by date, and an assigned item shows up in that person's **All tasks** like any other task, so the plan doesn't live in a silo.

An action item left **unassigned** reads as an **event** - something happening on that date with nothing to complete (it carries a calendar marker). Give it an assignee and it becomes a task to check off.

It's intentionally simple: no dependencies, no milestones, no Gantt chart. When you need those, that's what the **Schedule** tool is for.

Two rules worth knowing.

Action items can only be assigned to people with **Workshop access** on the project, so an assignee can always open the rally they're working in. And an item can be **completed or edited** (title, who, when, notes) by its **creator, its assignee, or an admin**, while the **creator or an admin** can delete it.

Resources

A rally's **Resources** section holds the reference material the team gathers. Attach **files** (images preview inline as a grid; other files list as rows - click one to open it, or hover for **Download**, which saves it to your computer), **links** (click one to read it in place), **embeds** (paste a **Google Doc, Sheet, Slides deck, Form or Drive file**, or a YouTube, Vimeo, Figma, Miro or Loom URL, and it renders right on the page - a Google file is only visible to people it has been shared with), **3D models** (.glb/.gltf, in an interactive rotate/zoom viewer), and

video or **audio** that play in place. You can also write rich-text **documents** directly on the rally. Uploading auto-detects the type. **CAD drawings** (.dwg, .dxf, .dwf) preview too: click one and use **Preview drawing** to open it in a full drawing viewer with **measuring tools** - distances, areas, and angles, plus **Calibrate** (click a known dimension and enter its true length to set the scale; both live in the toolbar's ruler menu). The **first** preview of a drawing prepares it for up to a minute; after that it opens instantly for everyone. The drawing viewer also does **markups**: click **Mark up** (top left) and redline the drawing - freehand, arrows, boxes, circles, revision clouds, and text, in your pick of colors - then **Save markups**. Everyone who can open the drawing can mark it up, and **each person draws on their own layer**: you see everyone's redlines (each person's name is a chip you can show or hide), but you can only change your own. Markups belong to **that uploaded file** - upload a revised drawing and it starts clean while the old file keeps its redlines - and they never alter the drawing itself: the .dwg you download is always the untouched original. The same drawing viewer opens CAD files **everywhere they attach** across Ceruleo: studio sandbox references and assets (a drawing asset keeps its discussions, reviews, versions, and details alongside the viewer), the project library, flags, RFIs, notes, tasks, meetings, submittals, daily logs, discussion posts, company contacts, opportunities, and your Pocket. Drawing preview is powered by an Autodesk-based rendering service built into Ceruleo - there is nothing to set up, and the file always downloads normally as well. **Clicking any resource opens it full screen**, with a column on the right carrying its **title**, its **note**, what it is and how big, and when it was added - plus **Download** and **Mark up** where they apply. The details column **collapses**: click the slim rail on its edge (the chevron) to tuck it away and give the preview the full width, and click the vertical **Details** tab to bring it back. **Edit the title and note right there**, in that column; to go straight to them, hover the reference and click the **pencil** (the same pencil on image tiles, file rows, embeds, video, audio, and 3D models - the last three had no way to be named at all before). A **link** or an **embed** opens as a live page inside that same view where the site allows it; many sites refuse to be displayed inside another page, and for those you get a tidy card naming the page with **Open in a new tab** instead. Every resource takes an optional **note** - a line on why it's here and what the team should look at - which the **Add link** and **Add embed** forms ask for up front, and which you can add to or reword on anything later. The note also reads under the resource and on an image's tile, so the reason is there for whoever looks next. Adding a **single** file opens its details straight away.

Mark up an image or a PDF. Reference material is easier to talk about when you can point at it. Hover a file (or open it in the preview) and choose **Mark up** to open it full screen with drawing tools - **Box**, **Circle**, **Arrow**, **Draw**, **Text**, and **Pin** - plus a colour picker, zoom, **Fit**, and page controls on a multi-page PDF. A **Pin** is a marker with a **note** attached: drop it on the spot you mean and write why it matters. The **Notes** list on the right is numbered to match the pins on the page, and you can reword or remove your own at any time. Everyone works on **their own layer in their own colour** - the **Markup by** panel lists everyone who has marked the file up, and hiding or showing any of them changes only **your** view of it. **Share to the rally** posts your markup into the contribution stream as an Idea, with your notes written out and the file attached. **Download marked-up PDF** (top right) saves the whole file as a PDF with the markups you can currently see baked onto every page - numbered pins, notes colors and all, plus a header naming the project, the rally, the file, and the date. What you see is what downloads: hide someone's layer in **Markup by** first to leave their marks out. The original file is never altered, and the button is there even in read-only view, so a closed rally's redlines can still be taken to a meeting. Anyone who can take part in the rally can mark up; you manage your own marks and a Workshop

admin can clear anyone's. A closed or archived rally is read-only. This one is web only - the phone still opens the file itself, but doesn't draw on it.

Anyone who can see the rally can add resources - the same people who can take part in it; you don't need edit access. You can edit or delete anything you added, and admins can manage any resource.

Who can do what

Workshop is a **per-project tool with its own permission level** (None / View / Edit / Admin), set per member exactly like Studio or Schedule - so someone can have Workshop access on one project but not another, independent of their other tools.

- **View** is enough to **take part**: read a rally, post contributions and replies, react, add resources, and **mark up** the images and PDFs on it.
- **Edit** lets you **create** rallies and manage the ones you created.
- **Admin** manages any rally in the project.

The **collaboration team** on a rally is its notify list and, for a **private** rally, exactly who can see it.

Working in a rally: the full detail

a space inside the project for operational planning, production coordination, and solution development, where the team rallies around a problem together. Its records are **Rallies**: each with a **cover image**, a **brief**, an **owner**, a **collaboration team**, a typed **contribution stream**, an **Action plan**, and its own **Resources**. It's its **own area** (not a tab) - open it from the prominent **Studio** and **Workshop** cards near the top of the project page.

See the Workshop and Rallies section for how to use it. The rest of this bullet covers the details:

post **contributions** (updates, questions, blockers, ideas) and reply with feedback, **react** to anyone's contribution (👍 Agree, 💡 Builds On This, 🎯 Important, 🔥 Love It, 😞 Needs Work, 🚀 Explore Further), filter the stream by type.

You can **edit your own contribution or reply** any time (the quiet **Edit** under it - reword it and, via the type chips atop the composer, change its **type**) and **delete** it - admins can delete anyone's; deleting a contribution removes its replies and its directly-attached files (resources it merely references stay).

When someone **replies to your contribution** or **@mentions** you here, the notification takes you **straight to that feedback and briefly highlights it** (no more hunting for it), and any contribution you've been alerted to but haven't seen yet is flagged **New** until it's been on your screen for a moment.

On a contribution **or reply** you can **attach files and links directly (+ File / + Link)** and/or **reference an existing resource** - directly-attached items stay **on the post** as clickable, previewable chips and are **kept out of** the rally's curated **Resources** list, while a reference points at a real resource (deleting the post never

deletes a referenced resource). Click a chip to **preview in-app** (images, video, audio, PDFs) or open links/other files in a new tab.

Build an **Action plan** – a simple **timeline** of action items laying out *who's doing what, when*: each item has a title, an optional **assignee** and **date**, and free-form **notes** so everyone knows the plan. Items are grouped by date in the timeline and appear in the assignee's **All tasks** like any other task.

An action item left **unassigned** is treated as an **event** (shown with a calendar marker – there's nothing to complete, it's just something happening on that date); give it an assignee and it becomes a task to check off. It's intentionally lightweight (no dependencies, milestones, or Gantt charts – just the plan).


Action items can only be assigned to people who have **Workshop access** on the project (so the assignee can actually open the rally), and an item can be **completed or edited (title / who / when / notes) by its creator, its assignee, or an admin**; the **creator or an admin can delete** an item.

A Rally also has a **Resources** section for the reference material the team gathers: attach **files** (images preview inline as a grid; other files list as rows with **View** to preview in-app and **Download** to save), **links, embeds** (paste a **Google Doc, Sheet, Slides deck, Form or Drive file**, or a YouTube, Vimeo, Figma, Miro or Loom URL, and it renders inline), **3D models** (.glb/.gltf, shown in an interactive rotate/zoom viewer), **video** and **audio** (play right on the page), and write rich-text **documents** directly on the Rally. Uploading auto-detects the type (video, audio, and 3D models are recognized from the file).

Anyone who can see the rally can add resources (the same people who can take part in it – you don't need edit access); you can edit or delete anything you added, and admins can manage any resource.

8. Working together

These run throughout the platform:



The collaboration layer on one page: a task's discussion beside its facts, collaborators, and the watch toggle.

- Discussions** - most items (flags, RFIs, notes, tasks, sandboxes, assets) have a discussion panel for threaded conversation. **Anyone who can see the item can start a topic and comment** on it; a topic can be marked **private** to restrict it to specific collaborators. **Adding someone to a topic notifies them**, and the notification opens that exact topic rather than just the item it hangs off. You can **reply to any comment - including a reply**; the thread stays **flat** (every reply at one indent under the comment that started it, in time order), and a reply that answers another reply carries a quiet **L name** marker naming who it answers - click it to jump to that comment, briefly highlighted. Deleting a comment takes its whole reply chain with it. When commenting you can attach files (drag in, **Attach from Pocket** to pull a capture from your mobile Pocket, or **Add a GIF** to search and drop in a GIF). The panel opens to your **Active** topics by default; use the filter tabs (Active / Stale / Archived / All) to see the rest. The drawer looks after a comment in progress: if the window is too short for the whole composer it **scrolls** so the Post button is always reachable, the click that merely brings your browser window back to the front never closes it, and closing it on purpose (backdrop, **x**, or Escape) with unsaved writing asks **Discard your draft?** before anything is lost. **Managing a topic** - editing its title, **archiving** (and **unarchiving** - archiving is never permanent), or **deleting** it outright (removing its comments too) - is reserved for the topic's **creator**, its **manager**, and project/company **admins**. Commenting on a topic does not grant these controls.
- Privacy** - marking an item or a discussion topic **private** (so only its collaborators and admins can see it) is reserved for the item's **creator** and project/company **admins**. Having edit access to an item does not by itself let you change its privacy; non-owners simply see a read-only **Private** badge.

- **@mentions** - type @ in a comment to mention a teammate; they get notified. Mentions show as a **highlighted teal chip** so they stand out from the surrounding text everywhere they appear. (Creators and collaborators of an item are always mentionable.) You can also @-mention inside **document inline comments and their replies** (see Studio assets below); there the picker only lists people who can actually see that document, so you can't pull someone into a comment on a private sandbox they can't access.
- **#links - point at a thing, not a person** - anywhere you can @mention in a **project discussion** or a **rally contribution**, type # to link an item from the same project: **flags, RFIs, notes, rallies, sandboxes, submittals, and schedule items** (type a couple of letters and pick from the list). The link lands in your comment as a quiet **teal chip** - so you can write "there's an RFI for this at **#Structural steel**" and readers click straight through. A few things to know: the chip keeps the item's **title as it was when you linked it** (a later rename doesn't rewrite old comments, but the link still opens the right item); #links **don't notify anyone** (that's what @mentions are for); and the picker only ever offers items **you** can see - a reader who can't see the item is stopped at its page rather than let in. Company-level spark conversations don't offer the # picker (they have no project to search).
- **@team - mention everyone on the item** - in a **discussion** comment (or reply) and in a **rally or spark contribution**, the @ picker offers a special **team** option first. Pick it to mention **everyone on that item at once** - its **creator, its assignees, and its collaborators** (the same people shown in the item's avatar stack) - instead of tagging each person by hand. Each of them gets a normal mention notification; you're never notified about your own post, and on a **private** topic @team still only reaches the people who can already see it. (By design it doesn't pull in project or org admins who aren't otherwise on the item, and it isn't offered on Studio asset comments.)
- **Document inline comments** - on a **document asset**, select any text and add an **inline comment** anchored to that passage; others can reply in a thread. @-mention a teammate in a comment or reply and they're notified with a link that opens the document and jumps to that exact comment. Replying also notifies the original commenter.
- **Watching** - click the **popcorn-bucket** icon on any item to **watch** it and get notified when it changes. The popcorn always sits **next to the item's collaborator faces** (watching is a people thing - it deliberately keeps its distance from edit, archive, and delete buttons); grey means not watching, gold and full means you are. On your phone the popcorn sits beside the collaborator faces too (tasks, flags, RFIs, rallies, sparks; notes, sandboxes, and schedules keep it top-right) - tap to watch or unwatch, no menu digging. Your dashboard **Watching** list keeps following an item even after it's finished - completed tasks, resolved/closed flags, and closed RFIs stay in the list with a status badge (and sink to the bottom) rather than vanishing, so nothing disappears on you. Unwatch whenever you're done with it.
- **Notifications** - the bell keeps you current; mark items read, or clear all. **Every notification says which project it came from**, on its own line above the headline, so you can tell at a glance whether something needs you now without opening it. A few things are not tied to a project at all - a company spark, an organisation-wide announcement - and those read **Company** instead. Clicking a notification now takes you **straight to the item** it's about - and for a discussion comment, it opens that exact topic and scrolls to the comment (briefly highlighting it). Read something but not ready to deal with it? Hover a **read** notification - in the bell or on the full Notifications page - and a small **circle button** appears next to the x: click it to **mark**

the notification unread again. It gets its unread dot back, counts in the bell badge, and shows under the page's **Unread** filter until you return to it.

- **Daily digest email** - Ceruleo can send you a calm summary: your top unread notifications, tasks to keep an eye on, and a short activity log, with a **Log in now** button. If there's nothing new, you'll just get a quick all-clear (with a little something to make you smile). It's **on by default** and, out of the box, arrives **every day around 10am your local time** - but you can **set your own schedule**. Under **Settings → Profile → Email → Daily digest**, with the digest on, choose **which days** it sends (**Send on** - tap the day buttons **Su-Sa** to include or skip each one) and **what time (At** - any whole hour, interpreted in your profile **timezone**). Turn the digest off entirely there too. The digest only goes to people who are a member of at least one **active** project; if you're not on any active project, it pauses automatically until you are.

On the go (the Ceruleo mobile app)

The mobile app is your pocket companion to the web platform. Every item page shares the same top line: < **Back**, the page title centered, and - when you have edit rights - a **round pencil button** that opens the item's edit form directly, or a small menu when more actions are yours to use (**Archive / Delete**, and on a rally, **Close rally**):

- **Your home dashboard** - the home screen is a single **scrolling dashboard**. At the top, a greeting and an **At a Glance** row of tap-through tiles - **Projects, Tasks, Flags**, and **RFIs** - each opening that full list. Below that: **My recent work** (imagery cards of what you last touched across your projects - sparks, rallies, sandboxes, and published assets), an **Activity pulse** feed (who did what across your projects; tap to open the item), your **Watching** list (items you're following), and a **Schedule lookahead** card (below). Pull down to refresh. The full lists behind the tiles (and the tool lists in **More**) let you **filter by status** with the chips at the top (e.g. Tasks: *Open / Completed / All*; Flags: *Open / In Progress / Resolved / Closed*; RFIs by their stage; Rallies and Sparks by theirs) to find closed or completed items; tools you don't have access to never appear. Your **Alerts** live on the bottom bar's bell (with its unread badge); the top-right corner holds the **search icon** (opens the search sheet with the keyboard ready) and your **profile picture** for your profile and settings.
- **Getting around - the bottom bar, Alerts, and More** - the bottom bar is **Home · Pocket · + (capture) · Alerts · More**. **Alerts** is your notifications list, one tap from anywhere, and its bell wears a **live unread badge** so you always know when something's waiting. **More** opens with a **search box** at the top (it never grabs the keyboard - tap into it when you want to search) above a grid of every tool you can reach, grouped *Your work* (Tasks, Flags, RFIs, Notes, Logs, Schedules), *Spaces* (Projects, Rallies, Sparks, Sandboxes), and *You* (Pocket, Profile, Alerts) - tools you don't have access to are hidden - with **Sign out** and the app version quietly at the bottom. Your **Profile** is also reachable from the Home avatar.
- **Search - jump to any item** - type into the search box at the top of **More**, or tap the **search icon in the Home header** to open the Compass sheet with the keyboard ready. Search by **name or #number** to find any **task, flag, RFI, rally, spark, sandbox, project, or schedule item** you can see; each result shows its tool icon, name, and project, and tapping it opens that item. (Search only ever shows things you have access to.)

- **Capture to your Pocket** – snap a photo, pick an image, or **Jot** a quick snippet and it lands in your **Pocket** to organize later (place it from your phone or from the web Pocket page). Capturing happens on the phone; the web Pocket page is where you review and file. For a quick text note on the web, add a **Note** to a sandbox's references, or jot from the phone and it's waiting in your Pocket on both. Your home screen's **Quick Actions** fan (the floating **+**) is customizable (choose up to **4** so the labels stay readable) – **and you don't lose the rest**: open the fan and **swipe up from the +** (or tap the little **up-arrow** in the middle) to open a **New** sheet listing **every** action you can use, grouped **Create** (Task, Flag, RFI, Note, Log) and **Capture to Pocket** (Camera, Jot, Link, Photos, File). So a tool you didn't put on your 4 favorites is always one swipe away.
- **Create on the spot** – start a **Task, Flag, RFI, Note, or Log entry** right from Quick Actions (you only see the ones whose tool you have access to).
- **Dates and times are always picked, never typed** – every date field on the phone opens a **month calendar**: tap the field, tap the day. That goes for a **rally action item's** due date, a **flag's** or **RFI's** due date when you edit it, and any **date field your template defines** on a note, flag, or log entry. Time fields open a matching **time picker** (hour, minute, AM/PM). A date you've set shows in full – *Sep 1, 2026* rather than 2026-09-01 – with **Clear** beside it to empty the field again.
- **Notes** – browse the **Notes** from your projects (a **Notes** tile in the Compass, or under **More**): the list has a **search box** (matches a note's title, #number, type, or project), **type chips** to filter to one note type (or **No type**), and a **Newest / Type / Title** sort – Type groups the list by note type. Open a note to read its **type**, custom **fields**, rich-text **body**, and **attachments** – plus jump into its **Discussion**. **Write a new note** from Quick Actions or the Compass: pick a project and note type, fill in the type's fields and body, optionally link **elements**, mark it **private**, and choose whether to **notify the team**. You can **edit notes you created** right from your phone, in two clean pieces that each save once: tap **Edit** > on the note's **Body** paper sheet (or **Write the note** > if it's empty) to open the full-screen formatting editor – its **Save writes immediately** – and use **Edit details** in the header for the title, custom fields, **collaborators**, and Private toggle. Backing out of a details edit with unsaved changes asks whether to **save, discard, or keep editing** (the body editor's Cancel likewise confirms before discarding changed content). A note's **collaborators** are picked while creating the note or in **Edit details** (creator/admin; the picker offers people with Notes access on that project, and added people are notified), and they show as **chips near the top of the note's page**. Collaborators are how a **private** note is shared, same as the web. If the note type has a **TO** field, the phone offers the same people picker (with **Select all**) on the New note form and in **Edit details**, and sending the note notifies its recipients exactly as the web does – a **Dropdown** field likewise reads as a row of tappable choices.
- **Rich text editing on the phone** – a **note's body**, a **spark's idea**, a **rally's brief**, and a **document asset's body** are all edited straight from their **paper sheet**: tap **Edit** > on the sheet to open a **full-screen formatting editor** with the **same controls as the web editor** – headings (H1–H3), **bold / italic / underline / strikethrough**, bullet and numbered **lists**, **quotes**, **alignment**, **highlight**, and inline **code** – and its **Save writes immediately** (one save; the item's Edit-details form handles title and other details separately). Formatting written on either surface reads identically on the other. **Rich-text custom fields** (on a note or flag whose template defines them) use the same editor: tap the field in the create form or the note's **Edit**

details to open it - the formatting saves for real instead of flattening to plain text, and the fields read fully formatted on the flag page, as they already did on notes. A document asset that's **Approved or published stays locked** on the phone exactly as on the web (reopen its status to edit). A spark's extra **idea cards** read on the phone as a **swipeable stack** - swipe right/left between **The Idea** and each card (the label and "2 of 5" counter track where you are; tap a dot to jump), in the same order as the PDF export - and the spark's creator/admins can **manage the whole stack from the phone**: **Edit** > writes whichever card you're on (The Idea included), **+ Add card** sits right under the deck (and under The Idea before any cards exist), and an extra card's **⋮ renames** it (its title becomes the card's heading) or **deletes** it - the web's card stack, pocket-sized. A spark's **Resources** also show on the phone - the same **image-forward tile grid** as sandbox references (pictures fill their tiles, a **gallery** unfolds its thumbnail strip on tap, links/files as icon tiles; rallies too). The spark's creator/admins can **add links and files** from the section header, and **press and hold** a tile to rename or delete it, like rallies.

- **Logs** - browse your projects' **Logs** (a **Logs** tile on the project's dashboard - its list grouped by **log type**, newest first within each - or the cross-project list in the Compass / under **More**) and open any day's log - it leads with its **log type, number, and date** plus a **Draft / Issued** chip and entry/flag counts - to see its **entries** grouped by type - each with its field values, who added it and when - along with any **photos** and **linked flags** (tap a flag to open it). On entry types that allow it, **attach a photo or file** to an entry straight from your phone. (Adding a log entry from **+** already shipped; distributing or reopening a log stays on the web.)
- **Issue and answer RFIs** - raise an **RFI** (a request for information) to a teammate, then track it through its lifecycle. Whoever it's assigned to (or a project admin) can **post a response - with files attached (+ Attach**: photo library, camera, or a document; stage several before posting), and an admin or the person who raised it can mark a reply the **official response**, which closes the RFI. Responses show their attachments as tappable thumbnails and file chips, and the RFI's own attached files appear in an **Attachments** section. Reach RFIs from the dashboard's **Open RFIs** list, from a project's **RFIs** section, or by tapping an RFI notification.
- **Track submittals** - open a project's **Submittals** tile to scan its records (each with its status color), and open one for a clean record page: the submittal's number, title, and status with the current submission's **version, state, due date, and submitted date** right in the identity card, its **files** (tap to preview in-app), and the **review** - each reviewer with their decision and narrative, grouped by review group, with a "2 of 3 decided" progress line. (Creating, editing, and reviewing submittals stays on the web.)
- **See your schedules** - open a project's **Schedules** to view its timeline: the schedule opens with its **approval state** (Draft / Pending review / Approved) and the numbers that matter - **items, overdue count, and the next deadline** - in its identity card, then tasks and **milestones** in outline order with their dates and status, and at-a-glance **health** (overdue work and at-risk milestones are flagged) so you can keep an eye on deadlines from your phone. To mark an item **complete** (or reopen a finished one), **swipe its row left** and tap the revealed action - a deliberate gesture, so scrolling the list can't trigger it (available to the schedule's creator, admins, and the item's assignee). Completing an item **after its planned finish** asks the same question as the web - keep the planned dates, or move the end to today - and a done item shows **when it was completed**. The home **Schedule lookahead** card pulls this together across *every* schedule you can

see: a single **soonest-first list of upcoming items** - anything **overdue** plus what's due in the **next 14 days** - each tagged with its project and schedule and marked as a **milestone** (diamond) or **task** (dot), with overdue items flagged red. It shows up to **14** items at a time; if there are more, it tells you (e.g. *14 of 23 items*) and **View all** opens the full pooled list. Tap any item to jump straight into its schedule. (This is the timeline radar; the schedule work that's actually **assigned to you** also appears on your **Tasks** tile/list like any other task.)

- **Join rallies and sparks** - open a project's **Rallies**, or your company's **Sparks** from your Profile, to read the conversation and **participate**: post a contribution (an update/question/blocker/idea on a rally; a thought/question/concern/opportunity on a spark), **reply** to others - including **to a reply**: the thread stays flat under the contribution that started it, and a reply that answers another reply shows a small L, **name** marker naming who it answers - and **react** to any contribution with the **same six emoji buttons as the web** (👍 💡 🎯 🔥 😊 🚀 - counts ride on the chip, yours is tinted gold; tap again to take it back). **@mention teammates** in a contribution or reply - type @ or tap the @ **button** to pick from a photo list (a spark offers your company, a rally its project team), and mentioned people are notified. Tap an attached resource to open it. Your own posts carry quiet **Edit** and **Delete** links beside Reply: **Edit** reopens your words in the composer with the **type chips** above it - so you can re-type the post (an Update can become a Question) as well as reword it (your @mentions survive; a post written with rich formatting on the web saves back as plain text from the phone) - and **Delete** asks first - a deleted contribution takes its replies and directly-attached files with it. Admins can delete anyone's contribution; editing is always author-only.
- **Browse sandbox assets** - reach a project's **Sandboxes** from the **Studio** button on its dashboard. A sandbox opens with a **cover hero** (its cover image dimmed behind the sandbox's name and asset count), the floating **discussion bubble** (topics and comments on the sandbox, shared with the web), and its design **assets** grouped by category on lifted cards - **image assets show real thumbnails** - each with its status. The sandbox's **References, Resources & Tools** show below the assets as the web's **image-forward tile grid** - pictures fill their tiles, a **gallery** tile shows its cover with a count badge (tap it to unfold the thumbnail strip, tap a thumbnail to view), and links/embeds/files are icon tiles (tap a file to preview it, a link to open it; archived ones sit behind a toggle). A **document** asset's page opens with an **identity card** (name, status, created by, latest version) and **Edit details** in the header (name + description, creator/admin), with the document rendered on its formatted **paper sheet** below and an **Edit** > on the sheet (Approved/published stay locked); an **image** shows at its true shape (tap for full size), a **PDF** previews inline, and video/audio play in the in-app player; the description and a collapsed version history sit quietly below, with the floating **discussion bubble** in the corner. With **Studio edit** access, tap **+ New** on the Sandboxes list to spin up a sandbox (name + description) and jump straight in. (Adding the assets inside a sandbox stays on the web.)
- **Browse the Library** - open a project's **Library** to scan its curated items, each with an **image thumbnail** (image files show a preview; other items show a type icon). **Search by name** to filter the list, and switch between **Recent** (newest first) and **Folder** (a folder browser: you see just the folders, each with its item count, and **tap one to open it** - subfolders nest as they do on the web, **Back** climbs out a level at a time, and the header names the folder you're in; **empty folders still show**, marked *Empty*, so a folder just created on the web is there to file into) with the toggle - the app **remembers your choice across every project's Library**. Open an item to see its details; an **image file shows a full preview** you can tap to open

in the built-in viewer, and links/embeds open out. With **Library edit** access, tap **+ Add** to file a new item (pick a folder, then add a **link**, upload **files/photos**, or write a **note**, and mark it **private** if needed) - **Photo library picks many photos at once**, and camera shots and documents keep adding to the same batch, so a day's worth of photos files in one save (each becomes its own item named after its file). Big batches are quick: photos are **resized before uploading** and go up **several at a time**, with **Preparing photos 12 of 30** and **Uploading 12 of 30** counting the progress - and the header's **Stop** button halts a batch mid-way, keeping what already uploaded and leaving the rest listed to finish with one more tap on **Add**; edit or delete an item you own from its detail screen. Organizing folders and pasting embeds stay on the web.

- **Browse your work** - the home dashboard (above) gathers your **Projects, Tasks, Flags, RFIs**, recent **work**, and **Activity**. Tap a project to open its **dashboard**: a cover **hero**, a prominent **Studio** and **Workshop** pair - the creative workspaces where the real work happens (Studio opens the project's sandboxes, Workshop its rallies) - then a scrolling row of **tool tiles** (Tasks, Flags, RFIs, Submittals, Schedules, Library, Notes, **Logs**, **Contacts** - each showing its count and opening that tool's list; tools you can't access are hidden), an **Activity pulse**, a **My tasks** list (the open tasks **assigned to you** in this project), and a **Schedule** summary. Open any task or flag for the full detail - tasks can be **completed, reopened, and edited** (title, notes, priority, due date, assignee, private) from your phone; on the **Tasks list** (and a rally's Action plan), **swipe a row left and tap** to complete or reopen it - the deliberate gesture used across the app - and tap the row to open it - backing out of an edit with unsaved changes prompts **Save / Discard / Keep editing**. The task page leads with its **project name** and the **assignee / due / completed** facts in the identity card, shows the task's **collaborators** (the creator or assignee can **Manage** them right there - added people are notified and become @mentionable in the task's discussion), and a task that was **generated automatically** (a rally action item, schedule assignment, submittal or review request) carries a card that opens **the item it came from** - shown only when you can see that item.
- **Project contacts** - a project's **Contacts** (from its dashboard) lists the **project team** (members flagged as contacts, with their role, pronouns, and email) and any **other contacts** - external people such as a vendor, utility company, or executive, with their role, company, phone, email, and notes. Tap a **phone number** to call or an **email** to write; private contacts stay hidden unless they've been shared with you. (Adding and editing contacts is done on the web.)
- **The Directory (company contacts)** - the company-wide address book, opened from **Directory** under Company tools in the sidebar (separate from a project's own Contacts tab above). It holds the vendors, partners, freelancers, and other people and businesses your company works with, whether or not they're on any project. A contact is an **individual** or a **business**: attach individuals to a business as its people, or keep them standalone (a freelancer). Each contact can carry one **category** - your company defines the list (Vendor, Partner, Freelancer, and so on) - plus any number of **tags** for the cross-cutting labels: tag your illustrators "illustrator" and clicking that tag chip filters the list to them instantly, alongside the search box, the **People / Businesses** toggle, and the category filter. **Everyone in the company can browse and search** the list; who can change it is granted per member (see Roles & permissions). Open a contact to see their full details (email, phone, website, address, notes - email and phone are click-to-write/call), the people at a business, and the **projects they're working on with you**. That projects list fills in **automatically**: when someone is added to a project using the same email as a contact, the project appears on that contact's

page - there's nothing to link by hand. A business shows the projects of the people attached to it, marked "via" that person. **Company admins can also load contacts in bulk:** click **Import CSV** on the Company contacts page, download the template, fill in one row per person or business (a person's *Business* cell attaches them to that business by name; tags separated by semicolons), and upload. A **preview** shows exactly what will happen before anything is written: categories you name that don't exist yet are created, a business a person names that isn't listed gets created with just its name (fill in its details later), and any row matching an existing contact - same email, or same name and kind - is **skipped**, so importing the same file twice never creates duplicates. **Businesses also carry compliance documents** - three dedicated slots on the business's page, above its general documents: an **NDA** (one file, recorded with the **year it was signed**), a **tax document** (one file, e.g. a W-9), and **insurance certificates** (as many as needed, each with its own **expiration date**; lapsed ones stay listed as history). The business's card wears the result at a glance: a green **NDA** chip with the year, a green **Insured** shield while any certificate is current, a red **Ins. expired** when they've all lapsed - and if there's no NDA on file, an amber **No NDA** warning that's always visible. Replacing an NDA or tax document keeps the old file in the business's general documents rather than deleting it. Uploading follows the same grants as editing contacts; individuals keep plain document attachments with none of this ceremony. The Directory's newer powers - multi-select **company types**, City/State/Country fields, Mobile on people, the **Active/Inactive** state, and the optional premium **Relationships** tool built on the same records - have their own section: see **Directory & Relationships**.

- **Collaborators on your items** - a **flag, RFI, task, or note** shows its **collaborators** on its page - and so do **sparks and rallies** - and the item's **creator** (or an admin; on tasks, the assignee too) can **Manage** them right there, on the web and on your phone - a multi-select picker offering only people with access to that tool, applying immediately and notifying whoever you add. **Collaborators you pick while creating the item are notified too** - it makes no difference whether you name someone on the create form or add them a minute later. (Uploading several files to the library at once sends **one** message naming the batch, not one per file.) Collaborators can be **@mentioned** in the item's discussion and are how a **private** item is shared.
- **Discussions + @mentions, anywhere** - on a **flag, task, note, RFI, sandbox, or asset**, tap the **floating teal speech bubble** in the lower-right of the page (on the web, the same drawer opens from the slim **Discussions tab** on the right edge) - it wears an **unread badge** whenever there's conversation on that item you haven't seen (a reply, a mention, or a new comment in a topic you've spoken in). The **project page** uses the same bubble for its whole-project discussions. Inside a Discussion: start topics (mark them **private** if needed), comment, and **reply** in threads - including replying **to a reply**; the thread stays flat, and a reply that answers another reply wears a small L **name** marker (tap it to light up the comment it answers). To mention a teammate, just **type @** in your comment (at the start of a word) - the picker pops up, like on the web - or tap the **@ button** beside the composer. The picker shows **profile pictures** and only lists people who can actually **see** that item, so you can't pull someone into something they can't access; it works on every discussion surface, including company-level **sparks** (where it used to do nothing). Tap **+** on a comment to attach a **photo** (library or **camera**), a **file/document**, a **GIF** (searchable), or a capture **from your Pocket** (shown with thumbnails so you can see what you're attaching), and **edit or delete** your own comments. If you started a topic, you can **Archive** it (which closes it to new comments) and **Unarchive** it later - archiving is never permanent. Everyone involved (people you mentioned, the person you replied to,

the topic's starter) gets notified, and tapping a discussion notification opens the thread **at the exact comment**.

- **Notifications that take you there** - tap a notification and it opens the item it's about (a mention or reply lands **on the exact comment**, briefly highlighted); **swipe an alert left** to flip its read state without opening it - an unread alert offers **Mark read**, a read one offers **Mark unread** (it gets its dot back and counts in the badge again). **High-signal events buzz your phone**: being **@mentioned**, getting a **reply**, being **assigned** something, and being asked to **review** arrive as real device notifications, shown with **who it's from** as the heading and what happened beneath (allow notifications when the app asks); everything else waits quietly on the **bell**. The **app icon carries your unread count** as a badge - each arriving push sets it, and it drains as you read your alerts.
 - **You choose what interrupts you, per channel** - under **Settings → Profile → Notifications** on the web, or on your phone's **Profile** screen, every category of notification - **Mentions, Replies, Assigned to you, Reviews, Watching, Added as collaborator, Status changes, Meetings** - has two switches: one for whether it **buzzes your phone**, one for whether it's **emailed to you** the moment it happens (grouped when several arrive together, delivered within about a minute). "all / none" shortcuts sit atop each column on the web. **Buzz** starts on for the high-signal categories (mentions, replies, assignments, reviews, meetings) and off for the ambient ones (watching, collaborator, status changes); **email starts entirely off** - it's opt-in, and separate from the **daily digest**, which keeps its own toggle and schedule. There's also a master switch for the **notification sound** the web app plays while a Ceruleo tab is open (with a "Hear it" preview). Turning anything off never hides it - every notification still lands on the **bell** and your notifications list; these switches only control what interrupts you.
-

Documenting email in Ceruleo

Real projects run on email too - a vendor quote, a township ruling, a client's sign-off. Ceruleo can keep those exchanges with the work, without ever pasting an email into a comment.

Every rally, spark, sandbox, RFI, and discussion topic has its own email address. Look for the small **"Email to this..."** control on the record (on an RFI it's in the details column under "Email records"; on rallies and sparks it sits with Resources; on a sandbox it's at the bottom of References, Resources, and Tools; a discussion topic shows ✉ **email this topic**). Click to reveal the address, click again to copy - then **CC or forward** any email to it, and the exchange appears on the record as a compact **email chip**: the subject, who's on it, how many messages, and when it last moved. Click the chip to read the whole conversation in a clean reader - quoted history folded away, attachments listed, and nothing in the email can run or phone home (remote images are blocked on purpose).

Replies keep landing by themselves. As long as the record's address stays CC'd, every reply joins the same chip - the count ticks up and the chip carries an "updated" time. In a discussion, a chip whose conversation just got a reply moves to the top of the thread, like any conversation with fresh activity. Ten replies are still one chip - email never floods a discussion.

Filed to the wrong place? Remove it. Open the email chip and use "**Remove from this record...**" at the bottom of the reader (a confirm step explains what happens). Removing deletes the email and its attachments from Ceruleo - the original email still exists in everyone's mailboxes, so nothing is truly lost. Who can remove: the **creator of the record** the email is filed on, a **tool admin** for that tool, a **project admin**, or a **company admin**.

Closed RFIs stop taking new email. Once an RFI is closed, the "Email to this..." and "From Pocket" controls go away, just like the reply box - a closed RFI is a sealed record. Everything already attached stays readable, and if late correspondence needs to join the record, reopen the RFI first. (If someone already had the RFI's address and replies after closure, that email still files itself - senders can't know your status - but nobody can newly attach material to a closed RFI from inside Ceruleo.)

Your Pocket has an address too - for the field. In **Settings → Profile → Email capture** you'll find your personal address (a memorable generated name like `creativepumpkin@in.ceruleo.app` - re-roll it until you like it, or choose your own). CC it when you don't know or can't find the right record's address in the moment; the email lands in your **Pocket**, and later you file it onto the rally, RFI, spark, or topic with **From Pocket**. Filing moves the *whole living conversation* - future replies follow it to the record automatically.

Email records read fine on your phone too. In the mobile app, an email captured to your Pocket shows as an envelope row (subject + who sent it) - tap it to read the whole conversation, attachments included. Rallies, sparks, sandboxes, and RFIs each carry an **Email records** card at the bottom of the screen listing the record's email threads, and **Email this record** reveals the record's address right there - long-press the address to copy it, or use **Share address** to send it straight into your mail app. Replying, filing a Pocket email onto a record, and removing a misfiled thread stay on the web for now.


Who can send: your own login email always works for your Pocket, and you can whitelist extra senders (a personal Gmail, an assistant) in the same settings block. A new conversation on a record's address must come from someone on your team - but once a teammate starts or files a thread, **replies from anyone already on that conversation** (the vendor, the township, the client) are accepted. Mail from strangers is quietly dropped.

And Scout knows. Email subjects and bodies are searchable - ask Scout "did we get a reply from the township?" and if the answer arrived by email, it's found, permissions respected as always.

Directory & Relationships

The company rolodex has grown into the **Directory**, and an optional premium company tool called **Relationships** now builds on top of it - relationship ownership, follow-up cadence, opportunities, pipeline, and forecasting for companies that pursue client work. The Directory is included for everyone; Relationships is an add-on your company can enable.

The Directory (everyone)



The Directory: the company rolodex - businesses and people, categories, and compliance at a glance.

Directory in the sidebar (under Company tools) is the same company-wide address book as before with a sharper name and richer companies. Everything you know still holds: individuals and businesses, people attached to a business, one category plus tags, compliance documents on businesses, CSV import, and the automatic projects list on every contact.

What's new:

- **Company types** - a business can now wear any number of **types** (Client, Prospect, Vendor, Partner, Venue, Architect, Fabricator, and so on - your company can edit the list). Types show on the company's card and page. They're descriptive, not permissions: a company's type never controls what anyone can do on a project.
- **Location fields** - businesses carry **City, State / region**, and **Country** alongside the street address; people gain a **Mobile** number next to phone and email.
- **Active / Inactive** - instead of deleting a company or person you're no longer working with, mark them **Inactive** in the edit form. Inactive contacts keep their whole history but drop out of the default list (a **Show inactive** toggle brings them back) and out of pickers.

Companies still connect to projects **only through people**: when someone whose email matches a contact joins a project as a member, that project appears on the contact's page - and on their company's page "via" them. There is nothing to link by hand, and there is no way to put a company itself on a project.

Relationships (premium add-on)

Relationships answers a different set of questions than the Directory: not just *who do we know*, but *who owns the relationship*, *when did we last really talk*, *when should we talk next*, *what work are we pursuing*, *what have we won*, and *when will it hit the business*. It's built for client-service companies - agencies, design firms, fabricators, consultants - and stays completely out of the way for everyone else: if the add-on is off, Ceruleo works exactly as before.

When it's on, a **Relationships** entry appears under Company tools, and every company page in the Directory gains a **Relationship** panel.

The relationship panel (on a company's page)

- **Relationship owner** - the person on your team who owns this relationship, plus an optional **relationship team** of others who work it. The owner is notified when handed a relationship.
- **Relationship status** - Active Client, Active Prospect, Strategic Partner, Past Client, Dormant (your admins can edit the list). Statuses can carry their own follow-up cadence defaults.
- **Strategic importance** - Standard, Important, or Strategic.
- **Relationship notes** - the context your team should carry into every conversation.

Follow-up cadence and the next touch

The heart of Relationships is staying in touch **on purpose**:

- **Cadence** is how often you want meaningful contact with a company - every 14 days, 30, 90, whatever fits. It resolves through a clear ladder, most specific wins: a **custom cadence set on the company** beats the **relationship status default**, which beats the **company type default** (if a company has several types, the shortest wins), which beats the **company-wide default** from Relationships settings. The panel always says where the current cadence came from ("Every 14 days - from Active Prospect default").
- **Next touch** is the specific date you expect to talk next. Logging a **meaningful interaction** sets it automatically from the cadence; you can also pin it by hand ("she said call back after Labor Day") without changing the cadence - the next meaningful interaction resumes the normal rhythm. **Skip** pushes it forward one cycle; **Pause** stops the clock entirely until you resume.
- Every managed company shows an attention state: **Overdue** (the next touch has passed), **Due soon** (inside the window - 7 days by default, adjustable in settings), **Current**, or **Paused**.

Logging interactions

Log interaction (on the company page and the Relationships overview) records what happened in one pass: the kind of interaction (Meeting, Call, Visit, Meal, BD conversation - admins can edit the list and which kinds count), what happened, who was involved, notes, an optional **follow-up**, and the recalculated next touch,

editable right there. Interaction kinds marked **meaningful** reset the relationship clock; the rest are just a record - adding a contact or uploading a file never counts as staying in touch.

Follow-ups are the specific to-dos of relationship work ("Send the revised proposal to Sarah by Friday") - each has an owner and a due date, lives on the company (and optionally an opportunity), and also appears in the owner's **All tasks** under Relationship follow-ups - and on the **home dashboard's Today list** once it's due, completable right there. They're deliberately separate from project tasks: no project required.

Opportunities and the pipeline

An **opportunity** is a piece of future work you're pursuing (or have won): a name and a company are enough to start; value, probability, dates, services, contacts, team, documents, and a delivery profile can come later.

- **Stages** are your company's own (Lead, Qualified, Proposal, Shortlisted, Preferred / Verbal, Awarded, Lost by default - rename, reorder, reprobability them in settings). Each stage carries a default **probability**; you can override it per opportunity.
- Separate from the stage name, every opportunity has a system **outcome**: **Open**, **Won**, or **Lost**. Whatever you call your final stage, Ceruleo knows what it means.
- **Weighted value** = value x probability for open work; won work counts at full value; lost work counts nothing. Lost opportunities keep their history (reason, competitor) but leave every forecast.
- The **Pipeline** page shows everything as a stage **board** or a **table**, filterable by company, owner, and outcome, with honest totals (full and weighted). Move an opportunity between stages right on its card.
- Timing is deliberately more than "when do we win": an opportunity tracks **proposal**, **decision**, **award**, **anticipated kickoff**, and **anticipated completion** dates separately, because winning in December and starting in April are different facts. The working rule: **award is when you expect ink; kick-off is the first day company resources are committed; completion is the last day they are free again**. For a touring kit that means shipped out to back in the warehouse - opening night is neither date. A deal through an external producing partner is still an opportunity (the partner is the counterparty); an in-house venture with no client is not - it belongs in Portfolio Planning as an item without an opportunity, since there is no award and never will be.

Won is not the same as converted

Winning an opportunity does **not** create a project. A won opportunity is **secured future work** - it shows that way (with its contract status: negotiating, pending signature, executed, notice to proceed, on hold) until you deliberately click **Convert to project**. Conversion opens a short guided step - project name prefilled, opportunity contacts offered as project invitations - and creates a real Ceruleo project exactly like the New project button would. Afterward the opportunity remains as the permanent commercial record, marked **Converted**, with an **Open project** link; the forecast moves its value from "won" to "confirmed projects" with no double counting.

Forecast

The **Forecast** page answers two executive questions by month, quarter, or year:

- **Commercial forecast** - *what are we likely to win?* Open pipeline and its weighted value, bucketed by expected award date, plus what was actually won in each period.
- **Delivery forecast** - *when does work hit the business?* Three bands per period, each opportunity in exactly one: **Confirmed projects** (converted), **Won, not yet converted**, and **Weighted pipeline** - bucketed by anticipated kickoff, or spread across a **delivery profile** if you've broken the opportunity into phases (Concept Design April to May, Production December to January, each with dates and a value or percentage).
- Admins can set a **target** per period right on the page; the bar shows the target line and the **gap**.

On the phone

The Directory and Relationships both live in the mobile app's Compass menu. Search companies and people, call or email in one tap, and see a company's relationship card - attention state, owner, status, cadence, last and next touch - with **Log interaction** built for the walk back to the car: pick the kind, dictate what happened, tag who was there, spin off a follow-up, and accept (or change) the calculated next touch, all in one sheet. Opportunities can be created, browsed by stage with totals, and updated - stage, value, probability, award and kickoff dates, next action - and the Forecast tab shows the quarterly bands and target gaps. The full pipeline board, delivery profiles, and configuration stay desktop things.


Who can do what

Two per-member grants in Company settings mirror the contacts pattern: **View relationships** (see everything, including values and forecasts - on by default) and **Manage relationships** (edit relationship fields, log interactions, manage follow-ups and opportunities, convert won work). Company admins can always do everything, and are the only ones who edit the configuration - statuses, interaction types, stages, services, cadence defaults, targets - under **Relationships → Settings**.

9. Templates & settings (mostly for admins)

Ceruleo adapts to how your company works through **templates** and **configurations**.

Company settings is laid out in two columns, like Project settings: **General** and **Templates** run down the main column, while **Subscription**, **Programs**, **Project statuses**, **Permission templates**, **Lifecycle & auto-archive** and **Delete log** sit in a narrower column beside them. **Members** spans the full width underneath, where the per-person permission rows have room. If your permissions only let you see some of these, the page falls back to a single column rather than leaving an empty gap, and on a narrow window everything stacks.



Company settings, Project statuses: rename, recolor, and add the statuses your projects wear.

Company (org) settings include:

- **Structure templates** - reusable project hierarchies. Define the level labels and the named types each level offers (e.g. a "Festival" template with Venue → Area → Activation → Feature). Templates are *starting points*: applying one copies it into a project, after which the project edits independently.
- **Flag / Log templates** - reusable typed-item definitions with custom fields. A **flag template** is a named **set** of flag types; it has two separate, clearly-labeled description fields, so keep them distinct: the **Template description** describes the whole set (e.g. "Standard construction site flags"), while each individual flag type inside it has its own **Description** (e.g. what that one flag type is for). Both are optional and are notes for whoever manages the template.
- **RFI categories** - the list of categories an RFI can be filed under, offered as a dropdown on the RFI form instead of a free-text box. Company categories apply to every project; a project can add its own in its Project settings, and sees both. Until at least one exists the RFI form keeps its free-text box, and removing a category leaves existing RFIs holding the name they were given.
- **Note categories & templates** - note **templates** (e.g. *Meeting*, *Transmittal*) are organized into **categories** (e.g. *Project Management Notes*). A category is just a grouping; the templates inside it are what you actually use to create notes. Every note automatically gets a **Title**, and turning on "rich-text Body" gives it a

formatted body for the main write-up - so you don't add Title or description fields yourself. **Custom fields** are only for extra structured data (dates, numbers, checkboxes, etc.).

Custom fields come in these types: **Text**, **Paragraph**, **Rich text**, **Number**, **Date**, **Dropdown**, **Checkbox**, and - on note types only - **People**.

Reorder the fields by dragging the grip handle at the left of each one. The order you leave them in is the order they appear on the form, on the item's page, and in its PDF. (You can also reorder with the keyboard: tab to a grip, press **space** to pick the field up, **arrow up/down** to move it, then **space** again to drop it, or **escape** to cancel.) The same grip is on log entry-type fields, and on flag types, at both company and project level. Reordering is saved with the rest of the type when you press Save.

- **Dropdown** gives the field a fixed list of choices that you write into the template: type them separated by commas ("Approved, Revise and resubmit, Rejected") and they appear in that order. Available on both note types and flag types, at company level and inside a project's configuration.
- **People** turns the field into a **TO line**. Whoever is picked in it is **sent the note** - they get a notification saying you sent it to them, and they can open the note even if it is private. Use it for a transmittal. It is not a task and creates no assigned work, and it is separate from **collaborators**, which stay the cc/visibility list. Notes only - a flag cannot be transmitted.
- **Project statuses** (org admins only) - the statuses a project can carry, in your company's own words. **Company settings → Project statuses** lets admins rename, recolor, and **add** statuses (e.g. *Design Development*, *Punch List*). Every status **behaves as** one of the four built-in kinds - Active, On hold, Completed, Archived - which is what counts, filters, and subscription limits go by, so a custom *Design Development* status marked as Active counts as an active project everywhere. Statuses show as colored chips on the **project cards** and in the project page header; deleting one leaves its projects showing the plain built-in kind.
- **Programs** (org admins only) - a **program** groups related projects under one name, like *NYC* or *2026 Retail*, so a company running many projects at once can still find things. Define them in **Company settings → Programs**: each has a name and an optional one-line description, and the list shows how many projects are currently in each. A project is put into programs from **Project settings → Programs** (by a project admin or a company admin) or from the **Programs** field on the New project form - a dropdown where you tick as many programs as apply (searchable once the list grows); a project can belong to **any number of programs** (an *NYC* build that's also part of *Holiday 2027* sits in both), and a project in no program is normal. A program is a *lens*, *not a container*: projects inside one keep their own team, permissions, structure, schedules, and every other setting, and nothing is shared or inherited between them. Once programs exist, a **program filter** appears on **All projects**, **All schedules**, **All flags**, and **All tasks**. Renaming a program updates it everywhere; **deleting** one leaves its projects completely intact, just no longer in a program (a deleted program is recoverable from the **Delete log** for 7 days, though its projects don't automatically rejoin it - re-pick the program on each one).
- **Spark categories** - reusable categories to organize and filter **Sparks**, managed just like Studio asset categories. Add, rename, or remove them in **Company settings → Spark categories**; deleting a category

simply leaves its sparks uncategorized (it never deletes a spark). Manageable by an org admin or a member with **Manage templates**.

- **Permission templates** - reusable permission grids to apply when adding or inviting project members. Each template sets a **level** (None / View / Edit / Admin) per tool **and** the per-tool **Explicit** modifier - an **Explicit** toggle sits on each tool row (shown only where it applies: not on Logs, and not when the level is None or Admin). Applying a template to a new member therefore carries the **Explicit** need-to-know scoping too, not just the levels, so there's no follow-up edit.
- **Lifecycle & auto-archive** (org admins only) - keep work from piling up quietly. Two kinds of protection, set per tool in **Company settings → Lifecycle & auto-archive**:
 - **Auto-archive from inactivity** - for **Sparks** and **Studio sandboxes**, set how many days of no activity mark the item **stale**, and how many mark it for **auto-archive**, plus a **warning lead**: that many days before *each* step, the item's **creator and collaborators** are notified so they always get advance notice and a chance to revive it - nothing ever archives silently. Any activity (a new contribution or asset, an edit, or **restoring** it from the archive) resets the clock and re-arms the warnings, so a revived item gets a fresh, full window. The warning lead can't be longer than the window it precedes. Leave a field blank to turn that step off.
 - **Attention safeguards** - for **RFIs, Tasks, Flags, and Submittals**, set a **nudge window**: an item left **open** past that many days with **no assignee and/or no due date** reminds its **creator and the tool's admins** to assign it, give it a due date, or close it - repeating until they do. (A **submittal** always has an owner - its creator - so it's nudged only when its current version has **no due date**.) Leave it blank to turn it off.
- **Two-factor authentication** (org admins only) - decide who must use a second step at sign-in: **Optional** (each person decides, the default), **Required for company admins**, or **Required for everyone**. Anyone it applies to is asked to set up an authenticator app before they can carry on using Ceruleo. Turning the requirement **off again never weakens anyone**: a person who has set 2FA up is always asked for their code, whatever the company setting says, until they remove it themselves. In **Members**, anyone with 2FA on carries a small **2FA** mark, so you can see how a rollout is going at a glance. Beside it is **Reset** - the recovery path when someone loses the phone with their authenticator app. It turns their 2FA off so they can sign in with their password and set it up again. Use it only when you are sure who you are talking to: for the moment in between, their account is protected by a password alone. **Every reset is recorded** - who was reset, who did it, and when - and that record cannot be edited or deleted by anyone.
- **Delete log** (org admins only) - the safety net for accidental deletions. **Company settings → Delete log** lists **everything deleted across your projects in the last 7 days** - flags, RFIs and their responses, tasks, notes, discussion topics and comments, spark/rally contributions and resources, library items, log entries, and **studio assets, their versions, sandboxes and sandbox attachments** - each row showing what it was, its project, **who deleted it and when**, and how long it has left. **Restore** brings the item back **to its most recent state**, and anything that was deleted along with it (a topic's comments, a contribution's replies) comes back in the same click. After **7 days** an entry is permanently unrecoverable, so act within the week. **Files come back too** - attachments on deleted content are kept for the same 7 days (they're only

permanently removed once the log entry expires), so a restored item returns with its file previews intact. The attachment records themselves are captured alongside their items - an RFI's or task's files, a comment's images, a log's photos - so a restore reassembles the whole item, files and all. **Deleting a sandbox takes every asset inside it**, so that restore brings back the sandbox, its assets, and each asset's versions together. In the list, an item's deleted files ride inside its row's **" + N contained"** count and come back with its Restore; a file deleted **on its own** (removed from an item that stayed) gets its own row and restores individually.

- **Members & roles** - manage who's in the org and what they can manage.

Adding company members. In **Company settings → Members**, enter an email, choose a **role (Member or Admin)**, and optionally pick projects to add the person to at the same time - the project list is a compact **searchable dropdown** (type to filter, tick as many as you need), with your picks shown as removable chips below it, so it stays manageable even with a hundred projects. What happens next depends on the person:

- **Already has a Ceruleo account** (for example, someone already on one of your projects) - they're added to the company immediately, no acceptance step.
- **New to Ceruleo** - they receive an email invitation; they join the company once they create their account and accept.

One company per person. A person can belong to **only one company at a time**. If you try to add someone who is already a member of another company, Ceruleo stops you with a message - they'd need to be removed from their current company first. (Being a collaborator on another company's *project* is fine and does not count as company membership.)

Member permissions. Admins always have full company access. For each non-admin **Member**, expand **Permissions** on their row to grant specific company powers: **Create projects, Manage templates, Permission templates**, and **Manage library**. Members only ever see what they've been granted:

- **Create projects** → a **New project** button on the **All projects** page.
- **Manage templates** or **Permission templates** → they can open **Company settings**, but it shows **only** the sections they're allowed (the template editors and/or permission templates). The company name, logo, subscription, and member list stay hidden.
- **Manage library** → the **Company library** item in the sidebar's *Company tools*.

A member with none of these sees only their projects and Sparks. Full **Company settings** (name, logo, subscription, and member management) is visible only to admins.

Setting up a new project

The **New project** form asks how you want to set it up:

- **Walk me through it** - a short guided run through the things a new project needs, one step at a time.
- **I'll set it up myself** - the project is created and you land on its settings page, where every section is laid out at once.

Either way the project is created the same way; the choice only decides where you go next, and you can change your mind by simply doing the rest in settings.

The wizard opens over your new project and covers five steps, in the order that makes each one easier than the last:

1. **Details** - dates, location, status, cover image and icon.
2. **Structure** - what this project's levels are called. This comes before the next step on purpose: these names become the vocabulary of the whole project, and the spreadsheet template in the next step is built from them.
3. **Elements** - import them from a spreadsheet, or leave it and build them in the Organizer later.
4. **Team** - add or invite people and set what each can reach.
5. **Tools** - adopt the company's flag and note templates for this project.

Next is the only button that saves. There is no separate Save or Done inside a step: pressing **Next** commits that step and moves on, and on the last step the button reads **Done** and closes the wizard on your new project.

Nothing is compulsory. **Skip** moves to the next step without saving that one, **Back** revisits anything, and closing the wizard keeps everything already saved. Whatever you skip is waiting in **Project settings**, which edits the very same things - so there is nothing you can only do inside the wizard, and nothing chases you afterwards to finish it.

Project settings is laid out in two columns: what you edit runs down the main column - **Details, Project structure, Flag and Note configurations** - while a narrower column beside it holds the things you set once and rarely revisit, **Visuals** and **Categories**. **Team members** spans the full width underneath, where its permission grid has room. On a narrow window everything stacks into a single column.

Project settings include:

- **Details** - the project's name, description, location, status, programs and dates, plus the paper size used for schedule PDFs.
- **Visuals** - the **cover image** shown behind the project card and the **project icon** shown in the header and sidebar.
- **Categories** - project-specific **RFI categories** and **Studio asset categories**, each added on top of the company's own list.
- **Project structure** - apply a company template, rename levels, enable/disable levels, and manage the named types for this project. Everything **auto-saves**. A level can't be switched off while elements of that level exist (move or delete them first), and **the last remaining level can't be switched off at all** - a project with no levels would have no kind of element left to add. The same rule applies to company structure templates, since those are applied to real projects.
- **Reordering fields in a project configuration** - a project can arrange a type's fields in whatever order suits it, by dragging the grips, without changing the company original. On a company type you have adopted, a

field's **label and type stay locked** but its **position is yours to set**.

- **Flag / Note configurations** - adopt and customize the company's templates for this project. A **template** (built in Company settings) is a named set of **types** - the categories a flag or note can be, each with its color and custom fields. A project **configuration** subscribes to one template and localizes it for that project: hide types you don't need, rename or reorder them, adjust a type's fields, or add **project-only types** - all without changing the company original, so other projects using the same template are untouched. The create forms offer the types from every **enabled** configuration, with your customizations applied. (**Log configurations** live in the **Logs** area itself - open Logs and use **Log settings** - so log types, entry formats, and distribution lists are all managed in one place.)
- **Team members** - add people, set roles, designate contacts, and grant tool permissions. Two ways to add someone:
 - **Add people from your company** - start typing a name or email and pick from the matches. The search only shows people who already have a Ceruleo account **and** are in your company - either a company member or someone on another of your company's projects. People from other companies never appear. Select several into a list, choose **one** role + permission set (or a permission template) for the whole batch, and **Add** them all at once. This adds them to *this project* only; it does not change anyone's company membership.
 - **Invite by email** - for someone who isn't in your company yet, enter their email to send an invitation (with the role and permissions you choose). They join when they accept. Better yet, **search your Directory** right in the invite form - start typing a name, company, or email and pick the person: their email fills in, their rolodex title becomes the suggested contact role, and they're flagged as a project contact by default (uncheckable). Typing an address by hand still works for anyone not in the rolodex.


Where you create a project, you can start **blank** or **from a company template**.

10. Roles & permissions (in brief)


- **Project admin / superuser** - full access within a project.
- **Org admin** - full access across every project in the organization, plus org-level management. A company admin is **implicitly on every project**: the **All projects** page lists the org's entire project portfolio for them (badged Admin), including projects created after they joined - no per-project add needed. Their **Home dashboard**, though, deliberately shows only the projects they've been added to as part of the **project team**, so their day-to-day view stays focused; adding an admin to a project team (via the Add-member form's picker or the project's Team page) is about the roster - appearing in assignee lists and contacts - not about access.
- **Members** - access is controlled per-tool. Each tool (Studio, Library, Flags, RFIs, Submittals, Notes, Logs, Schedule, Workshop) can be set to a level: **none**, **view**, **edit**, or **admin** (Studio also has a **reviewer** level for review-only access). For **Logs**, only people with **View or higher** can be added to a log **distribution list**,

and **Logs = Admin** makes a member a **log manager**: they can distribute, reopen, and manage the day's log (including editing anyone's entries) without holding the project-admin role. A tool set to **none** simply doesn't appear for that person on that project - its section is hidden on both the web app and the mobile project page (so the same person can have a tool on one project and not see it on another).

- **Collaborator** designation on a specific item grants access to that item even if it's private.
- **Directory (company contacts)** - everyone in the company can **view and search** the company contact list. Changing it is granted per member in **Company settings → Members: Edit contacts** lets a member add and edit contacts; **Manage contacts** additionally lets them delete contacts and manage the contact **category** list (Manage includes Edit). Company admins can always do everything.
- **Relationships** (when the premium add-on is enabled) - two per-member grants in **Company settings → Members: View relationships** (see relationship panels, opportunities, pipeline, and forecasts - on by default) and **Manage relationships** (edit relationship fields, log interactions, manage follow-ups and opportunities, convert won opportunities to projects; includes view). Configuration - statuses, interaction types, stages, services, cadence defaults, forecast targets - and the add-on itself are company-admin territory.
- **Deleting an item** - the same rule applies across the tools (flags, RFIs, notes, tasks, a day's log, library items, and so on): an item can be **deleted by its creator, a tool admin, a project admin, or an org admin**, and no one else. **Sparks, Rallies, and Submittals** are the exception - they're **archived** (recoverable), not permanently deleted. Archiving one automatically **clears its open action items and reviews** from everyone's task lists, My Tasks, and the dashboard/Overview counts, so nothing lingers as an outstanding to-do.
- **Explicit** designation - an optional, per-tool modifier (set by an admin) that *inverts* a person's visibility on that tool: rather than seeing everything the project shares, they see **only the items they created, are assigned to, or collaborate on**. It's separate from the level (you can be **View** or **Edit** *and* explicit), and it never applies to admins. You can also bake Explicit into a **permission template** (Company settings → Permission templates) so applying that template to a new member sets both the levels and the per-tool Explicit scoping at once. There's also an **explicit** setting for **sparks** at the company level - you can turn it on for a person when you **add or invite** them to the company (Company settings → Members), so it takes effect the moment they join, or change it later on an existing member. Explicit people also won't be @mentionable on, or notified about, anything they can't see. Being asked to **review** a Studio asset opens just **that asset and its review thread** to an explicit member - the rest of the sandbox (its other assets, Resources, and Discussions) stays hidden unless they're connected to the sandbox itself (creator, owner, approver, collaborator, or assignee).



What an external collaborator sees: their one project, and only the tools they've been granted.



A restricted member's view: schedule access at view level - the page works, the edit affordances simply aren't there.

If a tab or action isn't visible to you, it's almost always a permissions setting - ask a project or org admin.

11. A few orientation tips

12:28



ceruleo

POWERED BY SKAMBLE SYSTEMS



MR

Good afternoon, Morgan

Here's what's happening across your worlds.

1

PROJECTS

1

TASKS

0

FLAGS

0

RFIS

Activity pulse



Confirm LED wall pixel pitch

Morgan Reyes created

RFI · Aurora Pavilion

just now




Load-in logistics

Avery Ellison created

Rally · Aurora Pavilion

just now



Ceruleo rides along in your pocket - same account, same projects, capture in the field and it's on the wall at the office.

- **Drag files straight in.** On the web you can drop files onto the attachment area of a **task, flag, RFI, note, submittal, meeting, schedule item, contact record, rally, spark, library** or a **daily log entry** instead of clicking through a file picker. The area you are over lights up to show it will accept the drop.
- **Two-factor (2FA)** - a 6-digit code from an authenticator app, asked for after your password. Optional unless your company requires it. No backup codes, so a lost phone needs a company admin to reset it.
- **Placeholder** - an asset that has been planned but not yet produced: a name, an owner, and no file type until the first upload decides it.
- **Start at the project Overview**, then use the **Project Navigator** to understand the structure before diving into individual tools.
- **Click an element** to see everything tied to it in one place.

- **Watch** the items you care about so you're not chasing updates.
- **@mention** people rather than messaging out-of-band - it keeps the context with the work.
- The terminology you see (Zone, Scene, Venue, etc.) is **project-specific** - don't be thrown if two projects use different words for the same tier.
- **On a tablet, use the web app.** Ceruleo in your tablet's browser (login.ceruleo.app) is the full experience - markups, the CAD drawing viewer, Illustrator files, every tool at full size. The phone app is built for phones; your iPad or Android tablet gets more from the browser.
- The **Field Guide** is always a click away. Open it from the **Field Guide** link in the left sidebar for a searchable, always-current reference to everything in Ceruleo - browse the contents or search across every section. It's also **readable without signing in** - share the link with new teammates before their account exists. Not sure which tool a situation calls for? The **Which tool when** section is the cheat sheet, and the **Quick Reference Tool Guide (PDF)** button at the top of the guide turns it into a printable one-pager.

Mini-glossary

TERM	WHAT IT MEANS
Project bar	The slim bar pinned to the top of every page inside a project - the project name (click it for recent items and a project switcher) plus one-click links to every tool you can open
Scout	The ⌘K question box (formerly Ask) - ask in plain language and get an answer drawn from the records you can see, across all your projects, with links to its sources. Also answers "how should I put this into Ceruleo" questions from the Field Guide, holds a running conversation, and can open a creation form pre-filled
Element	Any node in a project's structure tree - a physical piece of the build, or a workstream like Marketing or Operations that needs a home for its work
Level 1–3	The renamable upper tiers (e.g. Venue / Zone / Scene)
Feature / Component / Part	The three fixed levels at the bottom of the hierarchy. They mean the same thing in every project, so their names cannot be changed
Flag	An issue or item needing attention
Transmittal	A note addressed to specific people - a note type with a People field on it. Recipients are notified that it was sent to them and can open it even when it is private; unlike a task, it asks nothing of them
Collaborator vs recipient	Collaborators are the cc on an item (visibility, and how a private item is shared - on an RFI they can also reply); recipients are the people a note is addressed to in its TO field
RFI	Request for information

TERM	WHAT IT MEANS
Requester (RFI)	The person whose question an RFI is - they get answer notifications and the review task, and mark the official answer. Admins can hand an RFI to a new requester
Answer vs Comment (RFI)	Two kinds of RFI reply: an Answer is a proposed answer - the requester is asked to review it, and only answers can be marked official; a Comment is discussion - no review task, never official
Schedule owner	The person who runs one specific schedule: admin controls over it (everything except deleting it), plus a daily heads-up while it has unassigned overdue items. Assigned by admins in the schedule's Settings
Log manager	Whoever runs a daily log's lifecycle - distribute, reopen, edit or tidy anyone's entries. Project and company admins, plus any member whose Logs access is set to Admin
On hold / Void (tasks)	The two honest ways off the open list besides completing: On hold pauses a task to come back to; Void records "we're deliberately not doing this." Both require a note saying why; a voided task stays on the record crossed out, never as a fake "done"
Submittal	The permanent record of a work product through submission, review, approval, and revision
Submission	A numbered version inside a submittal; revisions add new versions, old ones stay read-only
Review group	A named set of reviewers on a submission; members see each other's review notes/markup, but never another group's (the creator/admins see all)
Review surface	The full-screen markup workspace for a submission - pins/shapes on PDFs & images, time-stamped notes on video/audio, narrative notes elsewhere
Sandbox	A Studio working space for developing assets
Asset	A published deliverable (doc, file, model)
Spark	An idea board. Company sparks are org-wide; a Project Spark lives inside a project's sandbox and is seen by that sandbox's team
Workshop	A per-project tool for operational planning and production coordination; its records are Rallies
Rally	A single Workshop record - a brief, owner, and collaboration team, public to the project or private
Reference detail	The full-screen view of a single reference - the file (or a live page, for a link) on the left, its title, description, type, and origin in a column on the right, editable in place. A link that its own site refuses to have framed shows as a card with Open in a new tab . Reached by clicking the reference, or by the pencil on its tile
Markup	Drawing on a rally's image or PDF - boxes, arrows, freehand, and pins carrying notes. Each person marks up on their own coloured layer, which anyone can show or hide for themselves
Program	A company-level grouping of related projects (e.g. <i>NYC</i>) used purely to find and filter them; a project can belong to several programs, and projects in a program share nothing and inherit

TERM	WHAT IT MEANS
	nothing
Directory	The company-wide address book (formerly <i>Company contacts</i>) of the people and businesses your company works with - categorized, typed, taggable, and searchable by everyone in the company; a contact's page auto-lists the projects they're on
Company type	A label a Directory business can wear several of (Client, Vendor, Venue, Architect...) - descriptive only, editable by your company
Relationships	The optional premium company tool layered on the Directory - relationship owners, follow-up cadence, interactions, follow-ups, opportunities, pipeline, and forecasting
Cadence	How often a company should get meaningful contact (every N days) - inherited from status, company type, or the company default unless set per company
Next touch	The specific date you expect to talk to a company next - recalculated when a meaningful interaction is logged, or pinned by hand
Meaningful interaction	A logged interaction of a kind that counts as real contact (meeting, call, meal...) - it resets the relationship clock; other events never do
Opportunity	A piece of future work being pursued (or won) with a company - stage, probability, value, award and kickoff dates, and its own history
Kick-off / completion	The first and last days company resources are committed to a piece of work - what the delivery forecast and portfolio target dates run on. Opening night is neither
Weighted pipeline	Open opportunity value multiplied by probability - what the open pipeline is "really" worth
Won, not converted	An opportunity that's been won but not yet turned into a project - secured future work, counted at full value until conversion
Meeting	A project meeting record - agenda, roster, roll call, per-topic notes, action items - that walks Planned → In session → Issued (locked + emailed)
Meeting series	A recurring project meeting; each occurrence is numbered and open/standing topics carry forward automatically
Agenda topic	One numbered item on a meeting's agenda (3.01 = meeting 3, topic 1 - the number never changes); complete it to retire it, or keep it as a standing item
Note taker	An attendee the organizer designates to write the meeting's discussion notes and take roll call
Decision	A topic's settled outcome and who made it, recorded in the gold Decision field - highlighted in the notes and summarized at the top of the issued minutes PDF
Watching	Following an item for change notifications

TERM	WHAT IT MEANS
@team	A mention (in discussions and rally/spark contributions) that notifies everyone on the item at once - its creator, assignees, and collaborators
#link	A teal chip in a comment or rally contribution linking to a project item (flag, RFI, note, rally, sandbox, submittal, schedule item) - type # while writing to add one; it links without notifying
L, name	The marker on a reply that answers another reply (threads render flat, in time order) - click it to jump to the comment it answers
Explicit	A per-tool (or per-sparks) designation that limits a person to only the items they created, are assigned to, or collaborate on - even ones the team can otherwise see; never applies to admins
Field Guide	The searchable knowledgebase (this guide), opened from the left sidebar - readable even without signing in
Inactive element	A "cut" element - dimmed, no new work attaches, but still accessible and salvageable
Schedule	A plan of tasks and milestones across a working calendar
Milestone	A single dated moment in a schedule (no duration) - shown as a diamond
Dependency	A link making one schedule item drive another (FS / SS / FF / SF, with lag)
Approver	The person who signs off when a schedule's approval mode calls for it (the creator stands in when none is set; admins can always act)
Published version	The frozen copy of a schedule the team's dashboards and timelines read - capturing one is what Approve / Publish does. Versions stack up in the Baseline menu (Version 1, 2, ...) and double as baselines
Calendar link	A schedule's private subscription URL (Export → Copy calendar link) - add it to Google Calendar, Apple Calendar, or Outlook to see the published plan alongside your meetings; Reset link revokes shared copies
Approval mode	A schedule setting choosing how much sign-off publishing needs: every publish, the first version only, or none (a plain Publish button)
Unpublished changes	Edits to a published schedule's working copy that the team can't see yet - shown as a chip on the schedule (with ghost bars of the published plan) until they're published

Need a hand?

For anything this guide does not answer - billing, plan upgrades, account questions, or a problem you cannot solve in the app - email support@ceruleo.app. The Upgrade and Subscription buttons inside Ceruleo open a message to that same address.

- **Planning item** - a piece of work the company may plan around, in Portfolio Tools. It may link to a project or an opportunity, or stand alone.
- **Strategic resource** - anything scarce that work competes for (a kit, a crew, hours, capital), allocated across planning items in Portfolio Tools.
- **Allocation** - a claim on a strategic resource for a date range, graded committed / reserved / tentative / assumed.

Portfolio Tools

Available to organizations with Portfolio Tools enabled. Access is granted per person by a portfolio admin - it does not come from being on any project, and external project collaborators never see portfolio information at all.

Portfolio Tools sit above individual projects. Projects manage how work gets delivered; Portfolio Tools manage what the company intends to do - which future work might happen, how likely it is, and where current and future work compete for the same scarce resources. Find it in the sidebar under **Portfolio tools**: Relationships, Portfolio Planning, Reporting, and Portfolio Settings.

The idea in one sentence: **Relationships records what the company is learning, Portfolio Planning decides what the company intends, and projects deliver the work** - three jobs, kept deliberately separate.

The Overview

The **Overview** tab is the Monday-morning read. Eight counts - active, committed, probable and exploratory work, conflicts, decisions due within 30 days, stale signals, reviews awaiting project teams - and each one opens the thing it counts. Below them: a ranked **Needs attention** list, the current **Published Plan** with its assumption health (stale evidence, contradictions, overdue decisions, tentative overlaps - always itemized, never one unexplained score), and the decision calendar. Portfolio admins can hide cards for the whole organization with **Customize cards**.

Planning items - one population for all work

A **planning item** is any piece of work the company may plan around: an active project, a likely future job, an internal initiative, or an opportunity that is still just a conversation. They all live in one table on the **Planning** tab, so current and future work are always seen together and filtered - never split into separate systems.

Three rules worth knowing:

- **A planning item never requires a project.** Work that never becomes one - an opportunity that dies early, a capacity hold - lives and closes entirely here. When an item has a project, it shows the project's execution state; when it doesn't, that column is blank. A blank cell is honest.
- **One item per piece of work, forever.** An opportunity becomes a planning item once; if it later gains an evaluation project and then a delivery project, they all attach to the same item. Nothing is ever re-entered.

- **Three states, never collapsed into one.** The Relationships *pipeline stage* says where the deal is. The *strategic state* says how the company treats the work for planning (your own names - Exploratory, Probable, Committed and so on, configured in Settings). The project's *execution state* says where delivery stands. One piece of work can validly be Negotiation, Probable, and Evaluation at once.

To bring an opportunity into planning, open it in Relationships and choose **Start Portfolio Evaluation**. The linked planning item is created with the basics carried over, and an opportunity can only ever have one. The opportunity's award, kick-off and completion dates carry over as the item's **target award**, **target contract start** and **target contract completion**, so keep them honest by the rule: **kick-off is the first day company resources are committed, completion the last day they are free again** - for a touring kit, shipped out to back in the warehouse. Refurbishment time after a return lives on the resource's reset days, not in the dates. Within (or around) the window, add **critical milestones** to the item - load-in, opening night, strike - each a named date drawn as an amber tick on the Timeline. Opening night is a milestone, never the end date.

When an item needs real work to evaluate it - a feasibility study, drawings, a site review - open the item and choose **Create evaluation project**. That makes a full collaborative project workspace linked to the same item. It implies no commitment, advances nothing automatically, and creates no duplicate record. Items gathering lots of notes with no workspace show a quiet "evaluation?" hint - a suggestion, never an automatic project.

Strategic resources and allocations

Allocations normally sit inside the item's contract window. Save one that starts earlier or ends later and Ceruleo warns you once - save again to confirm. The override exists on purpose: internal claims like capital legitimately live outside the window; a kit that does is usually a date that needs fixing.

The **Resources** tab holds everything scarce that work competes for: a show kit, an install crew, fabrication hours, capital. A resource is either one discrete thing or a divisible capacity (like 4,000 hours per month), and can carry a **reset time** - the days it needs between uses for refurbishment or shipping.

An **allocation** ties a resource to a planning item for a date range, with one of four statuses - a ladder of firmness:

- **Committed** - the resource is assigned to this work
- **Reserved** - it is being held for expected work
- **Tentative** - the current plan assumes it, but that could change
- **Assumed** - the work needs it, but nothing has been set aside

The resource timeline draws each status differently - solid, outlined, dashed, dotted - so a glance tells you how firm every claim is. The **Timeline** tab does the same for the work itself: solid bars are happening, outlined are planned, dashed are hoped for, and items without dates are listed rather than drawn - no bar means no false precision. Beneath each bar, a thin rail spans the item's full **resource claims** - from the earliest allocation to the latest, including reset time - so a kit that stays committed long after opening night shows its whole footprint right on the work timeline. Hover the rail to see which resources and dates it covers.

Conflicts - computed, not curated

When two allocations claim the same resource at the same time (including its reset time), Ceruleo flags a conflict and grades it by the statuses involved: two committed claims are a **hard** conflict, committed against tentative is a **forecast** conflict, softer overlaps are **planning** conflicts or **pressure**. For divisible resources, demand is summed month by month against capacity, layered by firmness - "committed alone exceeds capacity" reads differently from "everything including assumed exceeds capacity".

Every conflict shows the rule that produced it. Resolving one means changing an allocation, not dismissing a warning.

Signals - the evidence behind the plan

A **signal** is a meaningful piece of information: "the client still wants September", "site possession slips to October 15", "the contract is signed". Each carries who said it, how authoritative they are, how confident the information is, and whether it supports the work, works against it, or is unresolved.

Click any item's **confidence** in the planning table to open its **Evidence** panel. The computed planning confidence sits at the top with every reason listed - each supporting or opposing signal, unresolved decisions, resource conflicts, and your recorded judgment. The label is never shown without its reasons, and when your judgment and the evidence disagree, both are shown - neither silently wins.

Three things signals do that notes cannot:

- **They age.** Each signal type has freshness thresholds (configurable in Settings) - a verbal target date goes stale in weeks, a signed contract never does. A stale signal counts for less, and says so. One tap **reconfirms** a signal with a fresh timestamp; every past confirmation is kept.
- **Contradictions coexist.** "September opening" and "possession not before October 15" both stay visible, marked as contradicting - the conflict lowers confidence instead of one statement silently overwriting the other. Superseding a signal keeps it, struck through; signals are never deleted.
- **They can be proposed from a note.** Paste what was said into the Evidence panel and Ceruleo proposes signals to accept or reject. With the AI option off, honest rule-matching does the proposing and nothing leaves the server.

Decisions live in the same panel: the external or internal calls the work is waiting on, with needed-by dates. An overdue decision drags confidence down until someone marks it decided.

Pipeline changes can move strategic state too: in Settings, each pipeline stage maps to **Automatic** (the state changes when the stage does, fully audited), **Suggest** (a confirmation appears at the top of the Planning tab - Update or Keep current), or **No relationship**. A newer stage change replaces an unanswered suggestion rather than stacking questions.

Scenarios and the Published Plan

A **scenario** asks "what if" without touching the plan. It is a layer of disagreements over the working plan - change an item's dates or state, move an allocation, add or remove one - and only those disagreements are stored. The Scenarios tab compares them side by side: feasibility (the same conflict math as everywhere else, run on each scenario), conflict counts, and value change. Duplicating a scenario copies just its changes.

Archive removes a scenario from the list while keeping it as history - the only way out for a published one - and portfolio admins can **Delete** an unpublished scenario outright. For a real decision between two resolutions - Mumbai gets the set, or Seoul does - the **Compare** button opens the two scenarios in columns: every item either one touches, working plan beside scenario A beside scenario B, changed cells tinted, override notes visible, resource moves listed, and each column's feasibility verdict on top. One screen for the leadership conversation, then publish the winner from its workspace. If you have never made one, the Scenarios tab shows a five-step walkthrough until your first scenario exists, and a new scenario explains itself until you make your first change - the guidance disappears the moment you no longer need it.

Publishing a scenario is the one deliberate act that changes the plan. It requires the Portfolio Admin role and a written reason, shows exactly what will change before you confirm, applies the scenario to the working plan, and records a numbered **Published Plan version** with a full snapshot - history is never overwritten. The scenario itself freezes.

The plan and your projects

Publishing never edits a project's schedule or tasks. If a published change touches work that has a project, that project's Overview shows a **review card**: what changed, why, and two honest responses - **Accept as target**, or **Propose alternative**, which sends a counter-proposal back to planning. The card exists only while unresolved.

Project teams can also push upward at any time: a **portfolio impact proposal** says "reality disagrees with the plan" - a release date slipping, a resource no longer workable - with a reason. Proposals land at the top of the Planning tab. Accepting one is a decision, not an automatic change: the reviewer makes the edit in planning, then accepts.

While the company is still deciding about a piece of work, its project shows a small **Portfolio** chip on the cover next to the status pill - click it for a read-only context panel with the strategic state, confidence, target window, and resource claims. Two things to know: it is **visible to company members only** (external collaborators see nothing at all), and it **retires once the work is confirmed** - a project in delivery just is. Company members can still reach the panel from the project menu's **Portfolio context** entry.

Portfolio alerts arrive as ordinary Ceruleo notifications - the bell on the web, the Alerts tab on your phone: plan changes touching your project, proposals arriving for review, answers to proposals you raised. Each says what changed and why it matters.

Reporting

Reporting aggregates what planning records: work by strategic state, pipeline stage against strategic state, demand versus capacity by quarter, open decisions with overdue marked, aging and stale signals, current

conflicts, and the published-plan history with every version's reason.

Portfolio Settings

Portfolio admins manage the organization's vocabulary and access here:

- **Strategic states** - names, colors, order, and what each state permits (whether it counts toward forecast demand, may reserve resources, requires a reason to enter). Every state maps to one fixed system category so reporting stays comparable however the names change.
- **People** - who holds portfolio access and at what level (Admin, User, or Viewer). Only organization admins can grant or revoke the Portfolio Admin role.
- **Pipeline mappings** and **signal freshness** thresholds, described above.
- **Assisted extraction (AI)** - off by default, and only an *organization* admin can enable it. When on: proposals are still reviewed by a person before anything is created, your information is processed only for your organization, never used to train any model, and not retained by the provider. Every AI exchange is written to the **AI activity log** - who triggered it, the exact text sent, the outcome - before the result is even returned, and the log is viewable here at any time.

Every administrative change, publication, allocation and state change is recorded in the portfolio audit log.