USER MANUAL

Vera AI

The complete guide to the platform · every screen, every workflow

VendorBenchmark · 2026-08-06

00 · Getting started & orientation

This chapter gets you signed in and oriented: how to log in securely, how the app is laid out, and the handful of screens you will return to every day. If you read only one chapter first, read this one.

> **What Vera AI is for.** Vera AI, by VendorBenchmark, helps you buy software at the right price: upload vendor proposals for expert benchmarking, run instant price checks, track your agreements and renewals, and work negotiations to a close. Everything else in this manual is a specific way of doing one of those four things.

THE SIX IDEAS BEHIND THE PLATFORM

The platform is large, but it is not complicated: nearly every screen is one of six ideas applied somewhere specific. Read these once and the rest of the manual becomes navigation rather than learning.

1. **Proposals become delivered reports.** You upload a vendor quote or contract as a proposal; our analysts benchmark it against the market; a delivered report comes back that you can forward. This is the expert loop, and it is the platform's oldest promise.
2. **Instant benchmarks answer the three questions every buyer has.** What discount you are being offered, what comparable customers get for the same deal as a range, and where your deal ranks. Instant benchmarks answer them self-serve in minutes, for over 1,280 vendors and product lines, no upload needed.
3. **Your estate is watched, not stored.** Agreements carry expiry dates and notice periods; Renewals sorts every deadline; Your Archive reads whole folders and files them by vendor. The point of filing paper here is that dates get watched and windows never shut silently, on the alert terms you choose.
4. **Vera is the analyst on call.** Press Ctrl/⌘ J on any screen and ask about a price, a clause, or your own portfolio. Vera's answers are grounded in your data and the benchmark evidence, and cited, so a number she gives you is a number you can stand behind.
5. **The market comes to you.** Terms watch diffs every tracked vendor's price list and prices the change at your seat counts; the Monday Weekly report files the week's licensing and pricing news; the Vera Wire carries the signals in between. You should not have to hunt for the fact that changes your renewal.
6. **Deals are worked to a close, not just analyzed.** The Renewal playbook builds the plan, a tracked Negotiation holds the mandate and the concession ledger, and Cases let Vera run a negotiation under decisions you approve. Analysis that never becomes a position is a report; the platform is built to close.

When you cannot place a feature, ask which idea it serves, or ask the direct question in Which tool do I use?

CREATING YOUR ACCOUNT

Where: `/get-started` · `/request-access`
Who: anyone with a corporate email

Two doors into the platform, both free to start, and both reviewed by our team. Access is always granted by a person, never automatically: whichever door you use, an analyst approves the request and sets up your workspace, usually within 1 business day.

Quick request (self-serve form). On `/get-started`, enter your name, corporate email, and organization, and tick the Terms of Service. That files a request for review; we email you a confirmation right away, and your sign-in link and temporary password follow once an analyst approves it (usually within 1 business day). A free 30 day trial with the sample portfolio loaded is waiting when you sign in.

Guided request. Prefer to add context, or want a specific plan? Request access on `/request-access` and a named analyst reviews the request and sets the workspace up for you within 1 business day. It is the same review, with room for a note.

While your request is reviewed. After you confirm your email, the waiting screen says plainly that your workspace is being prepared, and it is not a locked door. One click on **Explore the live demo** drops you into a working demo workspace immediately: delivered benchmark reports, renewals on the clock, and the AI analyst, all on a seeded sample portfolio. The demo is shared with other visitors and nothing you do in it is saved: the session runs for 30 minutes, a countdown at the bottom of the screen shows how long is left, and at the end you are signed out and the whole demo workspace is deleted and rebuilt from scratch for the next visitor. Click anything, and keep your own contracts out of it. Your approval continues unchanged in the background: when the invitation email arrives (typically within 1 business day), sign back in with your own credentials and your real workspace opens. The waiting screen also shows the dossier our analysts are preparing from your request, a get-ready checklist, and a preview of the benchmark catalog waiting inside.

Tips

- If your company already runs Vera AI and has domain capture on, either door routes you to your team's workspace admin instead of creating a second workspace.
- Personal email domains (Gmail, Outlook, Yahoo, and the like) are not accepted; use your work address. Corporate domains of the software vendors and resellers we benchmark are declined too: the platform is licensed to software buyers only (Terms of Service, Section 2).
- Every signup includes ticking agreement to the Terms of Service: the platform is licensed to software buyers only, you register with your real identity and corporate email, and benchmark data and playbooks are confidential to your organization. Teammates tick the same box when they accept an invite or finish account setup, and accounts created before the Terms shipped are asked once at their next sign-in.

SIGNING IN

Where: `/login`
Who: everyone

You sign in with your work email. The screen asks for your email first, then for the credential your account uses:

How to use it

1. Go to the sign-in page and enter your work email, then choose **Continue**.
2. If two-factor authentication is on for your account, you are asked for the six-digit code from your authenticator app; that code is your sign-in, no password step. Otherwise, type your password, or choose **Email me a link** to get a one-time magic link instead.
3. **Remember this device for 30 days** is ticked by default: this device stays signed in for up to 30 days. Untick it on a shared or borrowed machine and your session ends as soon as the browser closes.

Tips

- Invited by a teammate? Your invite email links to `/signup`, where you set your name and password once, then land in your organization already set up.
- Forgot your password? Use **Reset password** on the sign-in page; you will get a link to set a new one.
- Lost your authenticator? Choose **Email me a link**, sign in with the link, and use one of your recovery codes at the two-factor prompt.

TWO-FACTOR AUTHENTICATION

Where: `/mfa` (challenge) · set up under Settings › Security
Who: everyone (required for elevated access)

Two-factor authentication registers a rotating six-digit code from an authenticator app (Google Authenticator, 1Password, Authy, and the like) against your account. Once it is on, the sign-in page asks for that code in place of your password, and sign-ins that arrive another way (a magic link, single sign-on) are asked for the code at the `/mfa` challenge before the workspace opens. This protects your account even if your email inbox or password leaks. You turn it on and register your device in Settings › Security.

YOUR FIRST SESSION

Knowing what the first sign-in does, and deliberately does not do, makes it calmer. In order:

1. **You land on a result, not a tour.** A brand-new account is walked into your first verdict: drop a document of your own and get a defensible verdict in four visible steps, usually under five minutes. On plans without AI, the example report shows what a deliverable contains instead. Nothing about this step is throwaway; the verdict is saved with your other runs.
2. **The offers wait for you.** After your first win, Home offers the platform introduction (a six page picture book) and the desk survey (six questions that arrange the rail around your work). Both are links, never gates: skip them and they wait; take them whenever you like.
3. **Sample data is there if your workspace is empty.** Load the sample portfolio from the dashboard so every screen has something to show while you explore, and remove it before real data lands. A trial workspace starts with it already loaded.

4. **The workspace asks its questions once.** Where money-critical alerts should reach you, and whether you want the weekly briefings by email: asked one time each, strictly opt-in, always adjustable later under Settings › Notifications.
5. **Three keys are worth learning in the first minute. Ctrl/⌘ K** jumps to any page, tool, or vendor. **Ctrl/⌘ J** opens Vera over whatever you are looking at. The **Guide** door in the top bar opens this manual at the section for the screen you are on.

From here, Chapter 15 turns the first session into a first month.

THE APP SHELL

Every screen shares the same frame: a dark navigation rail on the left, a white utility bar across the top, and your content in the middle. One behavior rides every screen: any chart, and in the playbooks any card of text, enlarges on click into a near full screen reading view, with hover tooltips giving the exact figure under the pointer (see the note in Tooling for the details).

The left rail

The rail is your main map of the product. It groups destinations by the job you are doing:

- **Contracts & documents** — the paper itself: proposals, agreements, clauses, and document search.
- **Renewals & deals** — the deal calendar: renewals, negotiations, the rider, and the watchdogs.
- **Benchmarks & market** — price checks, delivered reports, research, and the buyer community.
- **Spend & savings** — the money: portfolio analytics, spend, invoices, opportunities, and savings proof.
- **Tooling** — upload-driven optimizers and the vendor playbooks.
- **Team & AI** — tasks, approvals, and Vera, the AI analyst.

On the screen

- **Recent** — a section at the top of the rail listing the last four screens you visited, so the pages you actually use stay one click away in either rail mode.
- **Used often** — under Recent, the pages you open most over the past month, marked "pinned by usage", once you have built a real habit. It never crowds out the fixed navigation.
- **Your views** — the saved list filters you return to, each with a live count. Owners can share a view with the whole organization; a shared view carries a **Shared** marker, is pinned in everyone's Your views rail, and stays owner-managed, so a teammate cannot rename, unshare, or delete it.
- **Group headers** — click one to fold or unfold it. Your folds are remembered per device. Until you fold something yourself, your workspace preset (or the role you picked in the getting-started course) decides which groups start open.
- **Core vs full rail** — by default the rail shows the 15 core destinations, every group keeping its flagships. The rest live in the **More** catalog and the command palette, so the rail stays short. You can switch to the full rail from the control at the foot of the rail; your workspace preset sets the starting tier until you choose.

Throughout this manual, a section whose destination sits outside the core rail carries a "via **More** or Ctrl/⌘ K" marker on its **Where** line, so you know not to hunt for it in the short rail.

- **"N more" rows (full rail)** — even the full rail no longer lists everything at once. Each group leads with the destinations a buying team uses weekly, about 30 across the rail, and folds its long tail behind a quiet **N more** row at the foot of the group. One press opens that group's tail, the choice is remembered per group, and the page you are on always stays visible even when it lives in the tail. Hover the row to see exactly which pages it holds.
- **Workspace summary (foot of rail)** — your organization and your role are always shown here, so you know which tenant you are in. Your workspace preset appears underneath with a **Change** link back to the picker.
- **Collapse button** — narrows the rail to icons only when you want more room.
- **On the phone** — the menu button opens the same rail as a drawer. Groups collapse and expand there exactly like on the desktop, share the same remembered folds, and honor the same **N more** rows, so the drawer opens short instead of scrolling through the whole catalog.

The top bar

The bar keeps to a handful of controls; everything less frequent lives in your account menu.

On the screen

- **Breadcrumb** — shows where you are (Vera AI › current view).
- **Apps and Dashboard switch** — the mode switch between your two homes: Apps opens the launcher, Dashboard opens the launcher with the feed beneath it, and the active mode is highlighted. It rides every screen, so either home is always one press away.
- **Add to apps** — on any feature page, this button adds the page to your launcher grid, so the app you use every week sits with the rest of them. To keep an app at the very front, press the star on its launcher row and it moves into Favourites. **Customization is one choice, not four.** Adding an app here, starring it on the classic launcher, and choosing it in the Estate Report editor all write to the same list, so a page you pin on one home is pinned on the others too. Pins you made before this became one list are carried over the first time you change anything.
- **Waiting on you** — the count each home carries beside your greeting: tasks assigned to you, approvals pending in the organization, and unread mentions. It is the same figure as the badge on **My work** in the left rail, and it links there. Renewals are not part of the count.
- **Jump to… (Ctrl/⌘ K)** — the command palette. Start typing any page, tool, or vendor and jump straight to it, and now run actions from the same box too: new proposal, start a benchmark, generate a negotiation brief, start a playbook, or ask Vera, each one keystroke away. This is the fastest way to move around; you rarely need to hunt in the rail. A query that matches nothing always ends in an Ask Vera row carrying your exact words, so a dead end still gets an answer.
- **Keyboard grammar** — press **g** then a letter to jump without the palette: **g d** dashboard, **g r** renewals, **g c** proposals, **g p** portfolio, **g s** spend, **g t** tooling, **g w** watch, **g b** benchmarking. Press **?** anywhere for the shortcuts overlay with every jump grouped.
- **The Vera cluster (Ctrl/⌘ J)** — one control with two modes: the text chip opens Vera, the AI analyst, over whatever you are looking at, and the voice segment beside it starts a spoken conversation. See Vera AI.

- **The inbox** — the bell opens one triage inbox folding notifications, watchtower changes, and money-critical signals into a single list, with a needs-decision filter and per-signal mute. Money items are flagged calmly, never alarmist. Team chat lives here too: the bell's badge counts unread chat along with everything else, and a pinned **Team chat** row at the top of the panel opens the full chat screen.
- **New menu** — the one place to start something: upload a proposal, add an agreement, and the other create actions.
- **Refer & earn** — the gold pill beside the bells. Like the platform? Invite a colleague from here: the dialog shows your personal referral code, sends up to five personal invites, and links to the Referrals page where you track who signed up and what you have earned (10% of what a referred company spends for their first 12 months).
- **Light and dark switch** — the moon (or sun) button beside your account chip flips the whole workspace between light and dark, on every screen, revealed as a circle of the new theme growing out from the button. The choice is your account's appearance setting: it follows you across devices and applies before first paint on your next sign-in. The three-way picker, including the System option that follows your operating system, lives in the account menu under Theme.
- **Guide** — the book beside the theme switch opens this user manual in a new tab, landed on the section that documents the screen you are standing on. A screen the manual does not yet cover in depth opens its chapter instead, so the door is never missing. The same handoff sits at the foot of the screen assistant's panel as **Read the full guide**.
- **Org chip and account chip** — switch organization (if you belong to more than one) and open your account menu. The account menu also carries What's new, Session activity, and the help and feedback actions (Platform introduction, Refer & earn, Report a bug, Suggest a feature), each keeping its unread dot, so the bar itself stays calm.

Help and feedback (in the account menu)

- **Platform introduction** — reopens the short picture book of the desk (see The platform introduction) any time, it is a good way to show a new colleague around.
- **Report a bug** — pick how badly it hurts (blocking, annoying, or minor), describe what happened and what you expected, and send. The page you were on, your browser, viewport, and timezone are attached automatically and shown to you before sending, so nothing is captured invisibly. You get a reference code to quote in follow-ups; bugs are typically fixed within 24 hours.
- **Suggest a feature** — tell the product team what to build: the area it belongs to, the idea, the problem it solves, and how much it would matter to your work. Every suggestion is read, and shipped ideas are announced in What's new.

FIVE MINUTES TO ORIENTED

The shell in prose, for reading once rather than referencing. The dark rail on the left is the map, grouped by the job you are doing, with your recent pages and saved views floated to the top; when a page seems missing, it lives in the **More** catalog or one keystroke away in the palette, never gone. The white bar across the top holds exactly one of each thing that matters: one search (**Jump to…**, which is the command palette wearing a field), one AI door (the Vera cluster), one inbox (the bell), one **New** menu for starting things, one **Guide** door into this manual, and your organization and account chips so you always know which tenant you are in. Everything rarer lives in the account menu.

Moving around rewards two habits. **Ctrl/⌘ K** then typing beats any amount of rail-hunting: pages, tools, vendors, and actions all answer to it. And the keyboard grammar, **g** then a letter, jumps without the palette (**g r** renewals, **g d** dashboard, **g b** benchmarking; press **?** anywhere for the full card). When a screen itself confuses you, its help is layered: the round assistant button explains the screen in place, the **Guide** door opens this manual at that screen's section, and Ctrl/⌘ J puts the question to Vera directly. The order to try them when stuck is written up as Get unstuck.

If you remember three things from this page: **Ctrl/⌘ K** to go anywhere, **Ctrl/⌘ J** to ask anything, **Guide** to read about wherever you are.

THE SCREENS YOU WILL USE MOST

The platform introduction

Where: offered on Home after your first win · reopen from your account menu, **Platform introduction**, or the Help page · `/welcome/intro`
Who: everyone, optional

A six page picture book of the desk, written for the moment it is offered: after your first verdict, not before it. It never blocks a sign-in. Your first result comes first, and the introduction waits as a link on Home until you take it or dismiss it.

The six pages

#	PAGE	WHAT IT COVERS
1	Welcome	The cover, over a slow moving photographic reel
2	One desk, not one tool	What the whole desk is, with the live count of benchmark datasets your price is measured against and the number of apps reading your data
3	What the desk does with a document	Two real screens side by side, the Verdict Bar and the renewal calendar
4	Vera remembers everything	Who Vera is, what she keeps, the Ctrl/⌘ J shortcut, and a plain note that no analyst, human or AI, is 100 percent accurate
5	Security is the first feature	The six controls that seal your data to your organization
6	Three doors worth opening first	The recap, then three next moves you can press

On the screen

- **Chapter rail (left)** — the pages grouped by chapter. Click any page to jump; pages you have read carry a check.
- **The page** — the explanation on the left, and beside it either a screenshot in a browser frame, live stat tiles, or the security grid. Click any screenshot (or its **Enlarge** control) to view it near full screen; Escape or a click returns to the page.
- **Progress bar and pager** — the footer tracks where you are; **Back** and **Next** turn the pages, and the arrow keys do the same.
- **Skip for now** — a permanent answer: skip once and the introduction is never offered again. It stays one click away in your account menu.
- **The three doors** — on the last page, **Hand Vera a contract** (drop a document on the Verdict Bar and get the report in about five minutes), **Ask Vera something hard**, and **Explore with sample data** (loads a worked sample portfolio, delivered reports included, one click to remove later). Pressing one records the introduction exactly as the checkbox asks and takes you to that screen.
- **Do not offer this again** — the checkbox on the last page, ticked by default. Leave it ticked and the offer never appears again, whether you leave by a door or by **Enter your workspace**; untick it and Home keeps offering the introduction on a later sign-in.

How to use it

1. Turn the pages with **Next** or the right arrow key; jump around from the chapter rail if you prefer.
2. On page 3, press **Enlarge** on either screen to read it properly before you go there yourself.
3. On the last page, decide whether Home should keep offering it, then either press one of the three doors or **Enter your workspace**.
4. To revisit it later, open your account menu (top right) and choose **Platform introduction**.

Tips

- The introduction is a picture book, not a lecture: every claim on every page is something you can click into minutes later.
- The screenshots inside it are the same captures used throughout this manual, so what you see in the book is what the product actually looks like today.

The desk survey

Where: offered after your first verdict and on Home · retake from Getting started · `/welcome/survey`
Who: everyone, optional

Six quick questions that arrange your desk around your work. The survey never blocks a sign-in: it is offered once your first result is on the table (on the first verdict summary and as a link on Home) and waits until you take it. Answer once and it never repeats unless you retake it.

On the screen

- **Step 1, the six questions** — mark everything that applies: benchmarking, contract terms review, which vendors you are interested in (tap the chips, or type any vendor and press Add), IT cost optimization, cloud optimization, and optimization tools. The vendors you pick go on watch: they light up on your desk with live benchmark coverage and their fiscal calendars feed your alerts.
- **Step 2, the three views** — three previews of your navigation rail, each built live from your own answers (gold sections are the ones you lit up). Pick how the desk should treat everything else:

- **Focused desk** — hides what is not relevant; one press on **Show all sections** in the rail brings everything back.
- **Relevant first** — your sections lead the rail, open; the rest stays visible but folded.
- **The full desk** — every section open from the start; you fold what you do not need.

How to use it

1. Mark your interests and vendors, then press **Continue**.
2. Compare the three previews, press the one you want, then **Continue to your workspace**.

Tips

- Nothing is ever locked away: whichever view you choose, the command palette (Ctrl/⌘ K) and the More catalog always search the whole platform.
- Change your mind later from Getting started, **Retake the desk survey**, or just use the rail's own **Show all sections**.

Your home screen

Where: sign-in lands here · Workspace › Home · the **Apps** button in the top bar · /app
Who: everyone

Your front door since July 2026: a two panel home built around picking up where you stopped. The left panel is a navy briefing where Vera opens with what you were doing last time, one press continues it; the right panel is every app on the platform as a calm catalog, grouped by job. A brand-new account is first walked through its first verdict (see Your first verdict); after that, every sign-in opens here. The classic home this screen replaced as the landing stays one press away, rebuilt as a daily report on your estate (see The classic home screen). Choosing the dashboard as your sign-in home changes where you land, nothing else: the **Apps** button stays in the top bar of every screen, so this home and app mode are always one press away whichever home you picked.

Every home screen's footer now states its job, links the other homes, and carries the same **make this my sign-in home** choice; the setting follows your account across devices.

On the screen

- **Vera · welcome back** — the letter under the gold orb says what you were last doing, in her words: the question you asked, the paper you were reading, the benchmark or decode you ran, or the contract you were working, with when. The gold button beneath continues it in one press. Nothing fresh in the last 14 days? She suggests one corner of the platform you have not tried instead, and a brand-new desk simply gets its doors.

- **Your last actions** — your 20 most recent runs across the platform, newest first and scrollable: every contract, SOW, and MSA decode, every Verdict Bar run, and every benchmark tool you worked in. A decode or verdict row reopens the exact saved run; a benchmark row reopens its tool. The list is yours alone and scoped to the workspace you are in. Until you have run something it shows your last four apps instead, a device-local trail that never shows a made-up suggestion.
- **Saved & shared** — your newest bookmarked papers and reports with a **Browse all** link to the full Saved page. Press the bookmark on any paper or report anywhere on the platform and it lands here.
- **Workspace pulse** — three live numbers at the panel's foot: contracts on file, renewals closing inside 90 days (it turns warm when something is on the clock), and unread alerts. Each is a link straight to its page.
- **Waiting on you** — tasks assigned to you plus approvals pending in the organization plus unread mentions, the same figure as the **My work** badge in the rail, linking there. The classic home screen and the control room carry the same count.
- **On the wire** — one line from the market feed in the briefing panel, naming the vendor and the headline, linking the Estate Report. It appears only when the wire has something on file.
- **The catalog (right)** — every app as a row: icon, name, one line of what it does, grouped under Start with a result, Negotiate, Watch the market, and Operate. Each section's first door leads as a featured card with its full description and an Open button: Vera AI, Vera AI-run Negotiation, the IT spend optimizer, and AI agents & workflows. The gold icons are the AI flagship doors. Research, the analyst library, sits under Watch the market beside the Buyer Network. Typing in the search field filters the catalog live; renewals carry a live count.
- **Edit apps** — remove any door, add removed ones back from the Removed shelf, or press **Restore default** to bring the original set back. There is one customization behind every home: this editor, the star on the classic launcher, the Estate Report's **Edit this screen**, and every tool's own **Add to apps** button in the top bar all write the same list, so a change made on any of them shows on all of them. Pages you add that way appear under Your additions.
- **Light and dark switch (bottom left)** — the sun and moon capsule flips your whole account theme, the same setting as the account menu's Theme control.
- **Classic home screen · Dashboard** — the panel's footer links: the springboard home, and the full command view.

How to use it

1. Read Vera's line. If it names the thing you came back to do, the gold button is your fastest path.
2. Otherwise press a recent app, a saved item, or search the catalog and press the door.
3. Shape it once: **Edit apps**, remove what you never use, **Done**. Restore default undoes everything.

Tips

- Apps open in app mode: the tool owns the screen under the navy app band, the same dress as this screen's briefing panel; press **Apps** or **Back to apps** in that band to come back here.

- Apps opened this way lead with their work, not their title: the search bar, the dropper, the ops floor, or the booking desk sits at the top, and the page title and explainers follow below it. Navigation stays put, so any tab bar a page carries is still at the top. Opened from the dashboard instead, every page keeps its familiar title-first layout.
- Rooms have links. When a tool opens a full screen room (a case file, a theatre, a workroom, the sourcing control room) the address bar carries it, so you can bookmark it or send it and the person who opens it lands in the same room, on the same document. Closing a room clears it from the address, so coming back later lands you on the page rather than reopening what you dismissed. A link to a room with nothing in it yet lands on the tool instead.
- The pulse numbers are links: a warm renewals count is usually your first stop.

The classic home screen (the Estate Report)

Where: the Classic home screen link on your home screen · `/welcome`
Who: everyone (kept alongside the new home)

The screen that was the sign-in landing before July 2026, kept one press away and rebuilt in July 2026 as a daily report on your estate. Signing in now opens Your home screen, and the **New home screen** link in the masthead leads back there.

It reads like a report because that is what it is. The wire leads it, since the market is the only part of this screen that changed since yesterday: the lead signal gets a headline and Vera's read of it. Your workspace figures sit beneath. Every one of your apps runs down the right hand side as a directory, so nothing is hidden below the fold and any tool is one press away. Use an app, press **Home** in the rail (or the brand mark in the corner), and you are back here.

On the screen

- **Masthead** — Vera AI over **Estate Report**, then your organization, today's date, and who the report is prepared for. It stays put while you read: on a desktop screen the masthead never scrolls away, so the way through to your dashboard is always in the same place.
- **The lead signal** — the most significant thing on the wire today, with its headline, Vera's read of what it means for you, and a byline naming the category, the vendor, when it was filed, and how many of today's signals touch a vendor you actually hold. Press the headline to open it in the Wire.
- **Key figures** — four live numbers in a ruled band: contracts on file, renewals closing inside 90 days, unread alerts, and the number of tools on your desk. Each is a link straight to its page; the renewals figure turns gold when something is on the clock.
- **Signals against your vendors** — the rest of today's wire under the figures, each named by vendor with its category, the time it was filed, a Caution or Risk tag where one applies, and a **Touches you** marker on the ones that concern a vendor in your estate. **Open the Wire** leads to the full feed.
- **The directory (right)** — every app on your home screen as a ruled list, grouped under Your apps, each with what it does and a live count where one applies (Renewal calendar carries what is closing in). Every door opens in **app mode**: the navigation rail disappears and the tool owns the screen, with Back to apps

on the left and the app's name as a gold label. The tool keeps your appearance setting, light stays light and dark stays dark. Hover any entry (or tab to it) and a small card tells you when to use that app.

- **Run the introduction** — the cream strip at the top of the directory starts First Light, the guided walkthrough that builds the desk around you. It is a first-days affordance, not furniture: it retires on its own once you run the introduction, or 30 days after your workspace opens, whichever comes first. Hover it and it says exactly that.
- **Light and dark switch (masthead)** — the sun and moon capsule flips the whole workspace between light and dark, revealed as a circle of the new theme growing out from the switch. The choice is your account's appearance setting: every app follows it, it travels across devices, and it applies before first paint on your next sign-in. The same setting lives in the account menu under Theme, alongside the System option that follows your operating system.
- **Dashboard link (masthead)** — a standing shortcut to the full command view.
- **Edit this screen** — compose the directory itself: press an app to remove it, press Add apps to put any tool or page on the platform here (searchable), and Restore defaults brings the original set back. Every tool also carries an **Add to apps** button in its top bar.
- **Choose dashboard apps** — flip the directory into a picker, highlight the apps you want, and your dashboard carries exactly that clean set as a Your apps band. No customization means all of them; a Change link on the dashboard band leads back here.
- **Apps and Dashboard switch (top bar)** — every screen carries the mode switch in the top bar: Apps opens this screen, Dashboard opens the full command view, and the active mode is highlighted. Either one can be your sign-in home.
- **What about security?** — a control at the foot of the directory that opens the security answers in a popup right here: tenant isolation, file handling, the AI data flow, deletion, sign-in security, and SOC 2 status, taken word for word from the full security FAQ. A button inside opens all 60 answers in a new tab, so you never lose your place.
- **The two minute tour** — a quiet line under the report opens the desk as a short staged show in a centered dialog: thirteen stops over labeled sample figures showing how a contract becomes a verdict, the benchmark opening as a case file in the workroom, two order forms on one stage with AI redlines at the exact clause, the negotiation playbook, a cited answer typed live by Vera, the renewal radar counting down to each notice deadline, a worked case the desk runs end to end, the standing agent roster coming on duty, an interactive reel of the toolkit's headline screens with real product captures, the executive brief that leaves the building, and the ground rules that keep your data yours. A large Next button names the next stop, Skip tour sits top right, the progress rail up top jumps anywhere, arrow keys move between stops, Escape leaves at any point, and the final stop opens your dashboard. Every product capture in the show carries an **Enlarge** control: click the frame and the screenshot opens near full screen, Escape returns to the show.
- **Optional follow-ups** — until taken or dismissed, this screen offers the platform introduction and the desk survey as quiet links under the report. Links, never gates: they wait as long as you like.
- **Footer choice** — "Prefer the dashboard? Make it your sign-in home" turns this screen off as your landing: from then on, sign-ins go straight to the dashboard. If you have already done that, the same footer offers to make this screen your home again.

- **One-time asks** — the screen asks two questions, once each and never together: where money-critical alerts should reach you, and whether you want the weekly briefings by email (price list updates per vendor, vendor licensing news, the knowledge summary). Both are strictly opt-in, both stay adjustable in Settings › Notifications, and every briefing email carries a one-click unsubscribe.
- **On a narrow screen** — below a laptop width the report becomes a single column: masthead, the lead signal and figures, then the directory underneath as a grid. Nothing is dropped.

How to use it

1. Read the lead signal and the byline. If it says signals touch your estate, that is usually your first stop.
2. Check the figures: a gold renewals number means something is on the clock.
3. Press the app you came to use in the directory. When you are done, press **Home** in the rail to come back.
4. Want the full picture instead? The dashboard is in the masthead, or **g d** on the keyboard.

Tips

- Your very first visit looks different: before your first result, this screen is a springboard rather than a report, and the door you choose there also arranges your workspace. Once you have a first win it becomes the Estate Report described above.
- If you turn this screen off as your landing and miss it, the command palette (Ctrl/⌘ K) still finds "Home".
- Holding a document and want proof before you explore anything else? Your first verdict turns it into a defensible verdict and your first plan moves in four steps.

Your first verdict

Where: from Home (the launcher) · `/welcome/first-verdict`
Who: requires an AI plan
You get: a verdict on your own document plus the first drafted moves of a renewal plan · under five minutes

The summary card now ends with **Send it on**: an editable follow up email drafted from the verdict's own figures, for you to copy into your own mail. Nothing is sent automatically.

One continuous flow from a fresh sign-in to a defensible verdict on your own document, in four visible steps: Upload, Decode, Benchmark, Plan. Under the hood it is a normal Verdict Bar run, so nothing about it is throwaway: the verdict it produces is a real one, saved with your other runs.

On the screen

- **The step rail** — Upload, Decode, Benchmark, and Plan across the top with a progress bar beneath. Each finished step collapses to its key result lines, so the story of your document builds as the run works.
- **The dropper** — drop a quote, contract, or invoice, or click to pick one. PDF or DOCX up to 30 MB, photos up to 3.5 MB. A document read as a contract saves to your workspace and joins your benchmarking queue; photos and invoice captures are read in memory and never stored.
- **Or run it on a document you already have** — if your organization already holds documents, the most recent ones are listed here so you can run the flow without uploading anything new.

- **No document at hand?** — the quiet card under the dropper loads the sample workspace instead: a worked portfolio with delivered benchmark reports, seeded into your organization so every screen has something to show while you explore, and one click to remove from the dashboard later. Choosing it exits the flow like a skip, so you are never routed back here uninvited.
- **The kind chooser** — if the run cannot tell what the document is, it stops and asks: a quote or priced offer, a contract or terms, or an invoice. Choosing costs nothing; the same file re-runs straight through the right engine.
- **The summary card** — once the fourth step lands: the one-line verdict, what grounded it, each step's key results, and the doors onward. Open the full verdict, benchmark the vendor, start the drafted renewal plan, or go to your desk.
- **Skip to the dashboard** — the top-right link out at any time; nothing here is mandatory.

How to use it

1. Drop the document, or pick one from the list of documents you already have.
2. Watch the steps land: Upload verifies the file and classifies it, Decode reads every clause, line, and figure, Benchmark places the deal against comparable ones and the peer corridor, and Plan drafts the first moves of your renewal plan from the matching playbook.
3. If the run asks what the document is, pick the kind and it continues on the same file.
4. From the summary card, open the full verdict for every figure, or follow one of the doors onward.

Tips

- A refresh mid-run resumes exactly where the run left off. The one exception: if the run stopped to ask what the document is and you reload before answering, drop the file once more, because the original never left your browser.
- Every run after your first lives in the Verdict Bar under Recent runs; this page is simply the guided first pass over the same machinery.

Related: Verdict Bar · Renewal playbook

The example report

Where: linked from onboarding and research surfaces · `/example-report`
Who: everyone

A complete, illustrative benchmark report on a fictional Workday agreement, so you can see exactly what an analyst deliverable contains before you upload anything of your own. The **Present** button in its header plays the report screen by screen, from the masthead to the bottom line, which is the fastest way to show a colleague what a delivered report looks like; the show starts from the first screen on every open.

Related: Delivered reports · The benchmark report

Dashboard

Where: Workspace › Dashboard · `/dashboard`
Who: everyone

Your home page, and two things only. The launcher fills the first screen: your greeting, Ask Vera, and every application one press away. One scroll down there is the feed, a single browsable stream of everything that changed in the last 30 days.

On the screen

- **Your greeting** — the hello on the launcher masthead, beside the Ask Vera kicker, with today's date on the right. It changes with the day and the hour.
- **Ask Vera** — the AI front door, directly under the greeting: ask a question about your portfolio, search every application by name, or attach a contract. Typing here also filters the application grid below, so the same field is both the question and the search.
- **Front of desk** — the band at the top of the applications: the eight instruments most work starts from (Vera AI, Compare, Instant benchmarks, Document search, Research, Contract decode, SOW Decoder, MSA Decoder), each with its own mark, a one line description, and its family label. Hover one and a small explanation of what it does pops up; click to open it. The band steps aside while you search or filter, so results are never pushed down the page.
- **The applications** — below the band, every app you have as a ledger: each family under its gold kicker and navy rule, and every app on its own row with its icon, its name, and a one line description of its job. Hover any row and a tooltip explains the app in full before you spend a click. A gold badge on a row is a live count of things that need you (proposals in review, open tasks, renewals inside 90 days). Press the star on a row to keep the app under Favourites, and use the group pills across the masthead (Favourites, All, Contracts, Renewals, Benchmarks, Spend, Tools, Team) to narrow the ledger. Arrow keys move through the rows, Enter opens one.
- **Your context** — the rail on the right, in four tabs: **Recent** (the pages you last opened), **Saved** (research and reports you bookmarked), **Wire** (the newest market signals), and **New** (what shipped and what was published).
- **The feed** — one scroll down, the last 30 days in one stream: what we shipped, the research we published, and the licensing, pricing, and vendor news the market made. See below.

The feed

- **Lenses** — the filter row across the head: Everything, New in the product, Research, Licensing, Pricing, Vendor news. Each carries its own count for the window, so you can see what is in there before spending a click. A lens with nothing in it is not shown.
- **Your vendors** — a gold filter beside the lenses that narrows the feed to items touching a vendor you hold a contract with. It appears only when at least one item does.
- **Filter the feed** — free text search across every title, summary, and vendor in the window.
- **The rows** — grouped by day, newest first, each with its lens colour, the vendor, a **Your estate** flag where it applies, and the wire's verdict (Favorable, Caution, Risk) where the item carries one. Click any row to open it: the summary in full, the one line worth acting on (Vera's read on a market item, the "try it" on a new feature), the source where the item is a verified market signal, and a link to the item itself.

- **Show more** — the feed loads 25 rows at a time and tells you how many are left. At the foot, links to the complete record: every release note, the research library, and the full wire.

How to use it

1. Start at the top: ask Vera a question, or press the application you came for.
2. Star the four or five applications you open every day so Favourites becomes your working set.
3. Scroll once to the feed and read it with the **Your vendors** filter on. That is the shortest version of "does anything this month affect my contracts".
4. Switch to the **New in the product** lens when you want to catch up on what we shipped, and open a row for the one thing to try.

Tips

- The Apps and Dashboard switch in the top bar moves between this page and your home screen, either of which can be your sign-in landing.
- On a brand-new account, load sample data first so every screen has something to show while you explore. Remove it before you start entering real contracts.

Control room

Where: Workspace › Control room · `/dashboard/control` · via **More** or Ctrl/⌘ K
Who: everyone

The second reading of the same workspace: where the Dashboard is built to start work, the control room is built to supervise it. Open it when you want to know what is about to stop being possible, what is waiting on a person, and what your agents are currently allowed to do. Your Dashboard is untouched and stays the default.

On the screen

- **The last 24 hours** — a single strip of counts across the top: documents filed, actions prepared, workflow runs finished, runs waiting on a person, and decisions taken. A zero is never coloured as movement, so a quiet night reads as quiet rather than as bad news.
- **Renewal pressure** — every notice window shutting inside the next 120 days on two axes: days until it shuts across, annual value up. Marker size follows the same money, so the renewal that is nearest and biggest is the largest circle furthest left. The red band on the left is the 14-day bracket, the fortnight inside which noticing a deadline later stops helping you. Only agreements that carry both an expiry date and a notice period appear, because only those have a window to miss.
- **Readout** — the panel beside the timeline, driven by whichever marker your cursor is on: annual value, the notice and renewal dates, the uplift cap or the absence of one, whether it renews itself, what the vendor has asked for as a percentage, and whether a benchmark exists. Click a marker to pin it so the readout survives the mouse leaving.

- **The board** — every open thread in four lanes, sorted by who it is waiting on rather than by what kind of thing it is. **Needs you** holds prepared actions the guardian escalated or blocked, actions awaiting your approval, and tasks assigned to you. **In flight** holds actions cleared to send. **On watch** holds ranked vendors with a real deadline and nothing owed today. **Cleared in 24h** holds what was sent or closed since yesterday.
- **Vera is working** — the workflow runs that have actually executed in the last seven days, newest and most urgent first. A run paused for a person is flagged in gold and will not continue on its own; a failed run shows its error.
- **Authority** — the mandates in force, the rope they give (advise, check with me, or just run it), the target, the largest single amount an agent may commit, the first expiry, and how many prepared actions are waiting on you. With no mandate signed it says so plainly: the desk can prepare work, but nothing reaches a vendor unless you send it.
- **The dock** — every app one press away at the foot, honouring the same picks as Choose dashboard apps on your home screen.

How to use it

1. Read the timeline left to right and stop at the first circle inside the red bracket: that is the deadline with the least room left.
2. Hover it to fill the readout, and check the uplift cap line. No cap on file means nothing bounds what they can ask for at renewal.
3. Click **Work \<vendor\> on the desk** to open it, or clear the Needs you lane first if something there is blocking a send.
4. Glance at Authority before you leave: if Waiting on you shows actions, nothing in that queue moves until you approve or decide.

Tips

- Nothing here is estimated. An empty control room means the rows are missing, not that nothing is happening: add an agreement with an expiry date and a notice period and the window appears on the timeline.
- The board reads your desk's own guardian verdicts, so an action lands in Needs you exactly when the mandate you signed says a person has to decide.

Related: Dashboard · Renewals · Workflow agents

Getting started (the course)

Where: Account › Getting started · `/start`
Who: everyone

The page opens with an onboarding spine: five ticks covering the first verdict, the platform introduction, the desk survey, the engagement letter, and first light, each linked, with one line naming your next undone step.

A guided tour matched to why you are here: benchmarking one contract, understanding a single contract, connecting your data sources, or running the full platform. Pick a path and tick steps off as you go; progress is saved to your account and syncs across your devices.

On the screen

- **Your workspace (preset picker)** — four arrangements of the same workspace: **Renewal desk** (one or two vendors, a renewal on the clock), **Deal check** (pressure-test a quote), **Portfolio command** (the whole estate), and **Explore everything**. Your pick decides which rail sections start open, whether the rail begins simple or full, which tools are pinned to your dashboard, and where the course begins. If you answered the goal question at signup, the matching preset carries a **Suggested** marker. A preset only changes emphasis: every feature stays available in search and the More catalog, and you can switch presets anytime here or from the foot of the rail.
- **Four doors** — benchmark one contract, understand a contract, connect your data, or run the full platform. Change paths anytime with **Change path**.
- **Role and use-case questions** — the full-platform path asks your role and what you want to get done, then builds a personal track from your answers.
- **Progress path** — the row of numbered dots shows your whole track: green dots are done, the blue one is your next step.
- **Next up card** — your single next step, promoted to the top with what you get out of it.
- **Your own vendor, carried through** — the vendors you named when you set up the workspace steer three steps. The instant benchmark step opens the grid already filtered to your vendor, the research step opens the library filtered the same way, and the Meet Vera step opens with the question about your renewal already written in the composer, ready to send or change. The step text is the same either way: only the destination arrives knowing who you buy from.

How to use it

1. Check the Your desk card: if the rail does not fit how you work anymore, retake the desk survey from there.
2. Pick the workspace preset that matches your situation; the sidebar and dashboard rearrange immediately.
3. Pick the door that matches why you are here.
4. Answer the role and use-case questions if prompted.
5. Do the "Next up" step, then mark it done. Repeat down the track.

My work

Where: Workspace › My work · `/my-work`
Who: everyone

Everything assigned to you personally: tasks, renewals you own, and approvals waiting on you. This is your answer to "what needs me?"

On the screen

- **Queues** — your open items grouped by type, so you can clear them in one sitting.

- **The all clear** — when nothing waits on you, the page shows one forward move: the next renewal to reach its notice window, with links to the renewal record and Ask Vera, or a price check prompt when none is upcoming.

How to use it

1. Open My work at the start of the day and clear each queue.
2. For team-wide to-dos rather than just yours, see Tasks.

The More catalog

Where: Workspace › More · `/more`
Who: everyone

The full, searchable catalog of every tool and page, including everything the short rail hides. When you cannot remember where something lives, this is the place to browse or search.

How to use it

1. Open More and search, or browse by group.
2. Faster still: press **Ctrl/⌘ K** anywhere and type what you want.

Search (the card catalog)

Where: Ctrl/⌘ K › search your records, or Home › the Instant search door · `/search`
Who: everyone

Global search across everything the platform holds for your organization: proposals, vendors, agreements and renewals, benchmark metrics, contract clauses, tasks, and AI deliverables. The command palette answers "take me to a page"; this screen answers "find me a record".

On the screen

- **The search toolbar** — the field, and under it a line telling you how many cards were pulled from how many drawers. Before you type, **Recently filed** lists the records you last touched as chips, so the proposals you were working on are one click away rather than a blank box.
- **Best match** — when one record clearly answers the query, it floats above the rest with a one line, fully grounded verdict.
- **Result drawers** — matches grouped by type (Vendors, Proposals, Agreements & renewals, Benchmark metrics, Clauses, Tasks, AI deliverables), each row linking straight to the record. A spine rail on the left shows how many of each type matched before you scroll.
- **Ask Vera about these results** — the line that closes the page. The records answer *where*; this asks what they mean, opening Vera with the question already written about that term across your agreements, and running it straight away. It shows even when nothing matched, because an empty result is often the most useful thing to ask about.

How to use it

1. Press **Ctrl/⌘ K**, type your term, and choose the search row; or open the **Instant search** door on Home.
2. Open the group that matches what you are hunting: a clause term lands in Clauses, a vendor name in Vendors and Proposals.
3. Nothing filed under your term? Press **Ask Vera about these results**, or try a vendor name, a contract title, or a clause term like "auto renewal" or "liability".

Tips

- For full-text search inside the documents themselves, use Document search; this screen searches the records the platform keeps about them.

Related: Document search · The More catalog

Migration Studio

Where: account menu or Ctrl/⌘ K · `/migration`
Who: org owners only · switched on per organization

Migration Studio brings an entire contract estate onto the platform in one guided run. Use it once, at the start, when you have hundreds or thousands of agreements to bring across. For a handful of documents use **Upload a proposal**, and for a folder you want read and filed straight away use **Your Archive**.

The part worth understanding before you begin: everything a migration loads is held in staging. While the run is in progress it is invisible to the rest of the platform. No renewal alert fires, no digest counts it, and no agent reads it. Your estate becomes real only when you sign it off.

Every step opens in the same frame, so nothing moves under you between them: a masthead across the top carrying the run's live figures, a numbered rail beneath it running Scope, Survey, Transfer, Register, Sign off, and a column down the side holding your plan, the quality bar you are measured against, and where the documents currently sit. The rail marks the step you are in and walks back to any step you have finished.

The walkthrough below follows one migration from the first question to go live. Every screenshot is from a real run; click any of them in the online manual to enlarge.

Before you begin

- You need the **owner** role in your organization, because a run writes thousands of records. Ask an owner to start the migration; they can hand you the work afterwards.
- Migration Studio is switched on per organization. If the account menu does not show it, ask your VendorBenchmark contact to enable it.
- Gather two things: the top folder of your contract archive (the whole tree, subfolders and all), and your contract register if one exists (the master spreadsheet, CSV or XLSX). The register is optional; documents alone still produce a portfolio.

Step 1 · Answer the six scoping questions

Open **Migration Studio** from the account menu or Ctrl/⌘ K. The run starts with six questions about your estate: roughly how many agreements you hold, roughly what you spend a year, which vendors matter most, where the documents live today, who owns your contract register, and the date you need this useful by. Rough numbers are fine. They are what the run is reconciled against at the end, so an honest guess beats a blank.

Press **Generate the plan**. Nothing you load is visible anywhere else on the platform until you sign it off at the very end.

Step 2 · Read your plan and the quality bar

The plan is generated from your answers rather than fixed, so a 300 document estate does not get the same schedule as a 5,000 document one. Each phase carries a date range, counted in working days from today. Under it sits the quality bar: the completeness, exception rate, and spend variance this run will be measured against. Both are yours to challenge before anything moves, and both stay in the column beside you for the rest of the run rather than scrolling away.

Step 3 · Survey the folder, on your own machine

Press **Choose your contract folder** and select the top folder of your archive. Include everything; we filter. The survey runs entirely in your browser: no file leaves your machine, which is why you can run it as many times as you like. It reports the files found and their total size, how many we can read (PDF and DOCX), every file type we cannot read named rather than silently dropped, how many look like duplicates, how many are over the 40 MB limit, and the biggest folders. Folder names become the first guess at each contract's vendor, so an archive filed as vendor then year needs far less correcting later.

Step 4 · Start the transfer

Press **Start the transfer**. Files go straight from your machine to encrypted storage, six at a time, and the console keeps a live count: transferred, remaining, duplicates skipped, and anything that could not transfer, each with its reason. Pause and resume whenever you like, or close the tab entirely; the work list is held on our side. If you return in a new tab, press **Re-pick the folder** so the browser can find your files again, and **Re-queue stalled files** if a previous run was interrupted mid-upload.

When the queue empties while you are watching, the studio moves you on to the register by itself. Open a run whose transfer finished earlier and it does the opposite: it holds on the result, so you read what was transferred, what was skipped as a duplicate and what could not be taken, and move on with **Continue to the register** when you have.

Step 5 · Map your contract register

When the transfer completes, upload your register (or press **Skip the register**). Each of its columns is shown with three real values from your own file, what we think the column means, and how sure we are. Change anything that is wrong. If nothing in a date column proves whether 01/04/2024 means 1 April or 4 January, the column is flagged and the dry run stays locked until you choose, because guessing moves every date in your register by up to eleven months.

Press **Dry run** to see what your mapping would do: how many contracts would be created, how many updated, and how many rows cannot be used with the reason for each, all computed without writing anything. The panel beside it shows what your first rows become, so you read your own data as we understood it. Rejected rows download as a CSV you can fix and upload again. When it reads right, press **Apply mapping**, then **Continue to reconciliation**.

Step 6 · Read the documents

Press **Start reading**. Each document is opened and read for the terms it states: vendor, dates, values, seats, notice. Reading runs in the background and continues after you close the tab; a large estate takes several passes over hours, and the counters keep score: read, waiting, could not be read. A document is read for what it says, not for what the contract ends up saying, which is why an amendment that changes only the notice period leaves every other term standing.

Step 7 · Decide the reconciliation proposals

Press **Run reconciliation**. Every judgement the machine made is a proposal waiting for you, grouped by kind: vendors that look like the same company, documents that are byte-identical duplicates, and agreements that supersede one another. Each carries the evidence it was based on and a confidence, and you accept or reject it with one click. This is also where binding happens: a master agreement and its amendments become one contract, whose terms in force are computed across the whole stack, newest document winning on each term separately, ordered by signature date.

Then clear **What needs a person**: everything the machine could not settle, sorted by annual value so the top of the list is what matters. Resolve each line or waive it; a scan with no text layer, for example, is reported as needing OCR rather than silently filed as a contract with no terms.

Step 8 · Compute the scorecard and go live

Press **Compute** on the scorecard. Every line is computed from your stored data and measured against the bar you set in step 1: completeness, exceptions, and your reported spend reconciled against what was found, with every part of any gap named. A line that misses the bar can be waived, but the written reason you give is stored with the acceptance record.

Then the one irreversible button. **Go live** promotes the whole estate out of staging in a single transaction: your contracts become visible everywhere on the platform, renewal tracking switches on, and the acceptance is written to your audit log with the scorecard attached. Until you press it, nothing here has fired an alert, sent an email, or been seen by anyone outside this run.

Tips

- Duplicates are detected by content, not by file name, so re-exporting the same archive twice is safe. Identical documents are counted and skipped, never imported twice. Two order forms with the same name a week apart are two real order forms, and we leave them alone.
- Only PDF and DOCX are imported, and each file must be under 40 MB. Everything else is reported in the survey so you know what was left behind.
- A browser cannot remember your files across a reload. The work list survives, but re-picking the folder is what lets a new tab carry on uploading.
- A column headed something like Contract Value is read as the total over the term, not the annual figure, and divided by the term to get the annual value. If it is already annual, map it to **Annual value** instead.
- Columns we do not recognise are kept on the contract as custom fields, exactly as you wrote them. Nothing in your register is thrown away.
- Binding is what makes a renewal date correct. A master agreement often says the term ends years before the amendment that extended it, and only the bound contract shows the real date. Binding is reversible: move a document to a different contract at any time and the terms in force recompute immediately.
- Overlapping agreements with one vendor are never proposed as superseded, because enterprises legitimately run several at once.

Related: Your Archive · Upload a proposal · Plans, allowances & limits

What's new

Where: account menu (your avatar, top right) · `/whats-new`
Who: everyone

Release notes in two levels: curated highlights in the journal, and the full changelog of every small change behind the Full changelog button in the header. The latest release leads as a full-width plate, older releases group by month, and on large screens a month index on the left jumps to any month. Each entry links to the feature it describes and closes with a Try it line: one concrete first move to make with the feature, the control to press or the document to drop. The full changelog reads as a full width ledger: each ship day holds its date and change count in a left rail, with every change on one line showing its category, title, and detail, and a summary band on top counts the changes, ship days, and areas covered.

Next: Chapter 15 · Your first 30 days for the guided ramp, or Chapter 01 · Contracts & renewals to start the screen reference.

01 · Contracts & renewals

This is the paper-trail heart of the platform: where you upload vendor quotes for expert benchmarking, keep your signed agreements, and watch every expiry, notice deadline, and price-list change so none of them catches you out. It is also where you work a live deal to a close, from setting a mandate through to signature. If a screen touches a contract, a vendor, or a renewal date, it lives in this chapter.

PROPOSALS

Where: Contracts & documents › Proposals · `/contracts`
Who: everyone

The related records strip states whether this vendor's signed agreement is on the ledger; when none exists it says so and offers to add one so the notice date is tracked.

Proposals holds the vendor quotes and offers you put in front of an analyst, from upload to delivered report. Signed paper lives on the agreement ledger, and a whole folder at once goes to Your Archive, which the empty state now links directly.

Vendor quotes and contracts you have submitted for benchmarking. Each one moves along a path: Uploaded, In review, Benchmarked, then Delivered when your report is ready.

In plain words: the tracker for everything you sent to the analysts, and the door to each finished report.

When to reach for this

- "Where is my benchmark?" answered by the status pill and the expected delivery date.
- The report landed and you want to read it: the Report tab is on the proposal's own page.
- A quote deserves the expert treatment rather than an instant read.

On the screen

- **Status pills** — where each proposal sits in the review pipeline. Every step talks to you: an alert and an email when an analyst picks the proposal up (with the expected delivery date), when the benchmark is built and moves into final quality review, and when the report is delivered. If a benchmark ever slips past the promised turnaround, you get one proactive email owning the miss with a revised delivery date. Tune these under Settings › Notifications (Benchmark status updates and Report delivered).
- **Proposal page** — open a row to see its status path, the files you attached, and, once delivered, the benchmark report. While a proposal is anywhere before Delivered, the page refreshes itself within half a minute of a status change, so you never need to reload; delivery announces itself with a toast that opens the Report tab.

- **Benchmark request card** — an upload normally enters the analyst queue by itself. When your plan's benchmark allowance is used up it cannot, and the proposal page says so plainly: the card states your usage (for example 4 of 4 requests this month) and resolves it in place. An organization owner adds a benchmark credit ($5,000 for one full analyst benchmark, billed by invoice, no card) in a dialog; the request lands on our desk as a support ticket, our team confirms and invoices it, and the credit is added to your account. A teammate who is not an owner asks their owner in one click; the owner gets the request with a link back to the proposal. While a credit request is in flight the card shows it, with a link to follow the ticket under Support.
- **AI review** — on a proposal page you can run an AI redline scan of the vendor PDF that flags risky clauses with quotes from the document.

How to use it

1. Upload a proposal: vendor, contract type, your question, and the files (PDF or DOCX).
2. Open a proposal to track its status, or wait for the emails; every stage change reaches you either way.
3. Read the report once it is delivered.
4. About 30 days after delivery, answer the What happened? question, from the one-touch email or the card on the Report tab. A won reduction, with its annual figure, lands in your savings proof; if the contract has no stored annual value the card asks for the pre-reduction baseline so the saving can be reconciled against invoices.

If something looks wrong

- The proposal reads **Saved, not queued**: your plan's benchmark allowance is used up, and the page resolves it in place with a benchmark credit; see the card described above.
- The status emails never arrive: check the Benchmark status updates toggle under Notifications and Troubleshooting.

Related: Upload a proposal · Delivered reports

UPLOAD A PROPOSAL

Where: Contracts & documents › Proposals › New · `/contracts/new`
Who: everyone

One drop zone sends vendor quotes and contracts to our analysts for benchmarking. Drop one file or twenty; each document is read as it lands, its details fill in as editable chips, and you confirm the whole batch once. You will be emailed when each report is ready.

In plain words: drop the paper, check the chips, confirm once.

A worked example. Drop a renewal quote like the sample Contoso Analytics one and watch the chips fill: vendor Contoso Analytics, 120 seats, $450 per seat, $54,000 annual value, the term dates, and, tellingly, no uplift cap, because the quote names none. A missing or low-confidence value shades gold or red; that is your cue to check the paper, and the benchmark focus box is where you tell the analysts what you actually want answered ("the quote carries no uplift cap; what does that cost us?").

On the screen

- **The drop zone** — drop PDF or DOCX files, browse, choose a whole folder, or paste files from the clipboard. Every file is verified by its contents (not just its name) the moment it lands, and stored privately: your organization and our analysts only.
- **File cards and chips** — vendor, title, category, annual value, currency, seats, term dates, and uplift cap fill in per file, shaded by how confident the read is: plain for high confidence, gold for worth a glance, red for please check, blue for values you set yourself. Click any chip to correct it; the AI pass never overwrites your edits.
- **Discount band** — once the vendor is known, the card shows the discount band buyers like you land with that vendor at your deal size, with sample sizes, so you know the stakes before you submit.
- **Notes and benchmark focus** — optional per file: notes for our analysts and what you want benchmarked, with a one-click starter template matched to the proposal category.

How to use it

1. Drop your files (up to 20 at a time). Each card starts reading immediately: rules first, then an AI pass that upgrades anything you have not edited.
2. Check the chips, correct anything gold or red, and add your benchmark focus.
3. Confirm once (Ctrl+Enter also works). A batch shows a progress bar while it submits, and each card ends as Queued with a link to its proposal page.
4. Track progress afterwards on the Proposals list.

Tips

- Files are limited to PDF and DOCX and capped at 40 MB each. Files upload straight from your browser to storage.
- Your edits survive a stray navigation: chip corrections, notes, and the benchmark focus are restored when you re-add the same files in the same browser session.
- If your plan's benchmark allowance is used up, the upload still lands but ends as **Saved, not queued** instead of Queued. Open the proposal page to fix it in place; see Proposals below.

Related: Proposals

YOUR ARCHIVE

Where: Contracts & documents › Your Archive · `/archive`
Who: everyone

In plain words: drop your whole contract folder and get back a filed, judged estate.

When to reach for this

- Week one: the folder your organization already keeps is the fastest way to a working workspace (Chapter 18 compares the doors).

- Any time paper accumulates: the archive is a standing door, not a one-time import.
- When someone asks what the estate actually holds: the intake brief is the forwardable answer.

Your Archive is every contract document you hold, filed under its vendor. Renewal dates and notice windows live on the agreement ledger, and a single quote you want benchmarked goes to Proposals. A renewal candidate the scan finds that is already on the ledger now links straight to it, and says that confirming it advances that row rather than duplicating the vendor. The overnight portfolio import at `/contracts/import` fills the same store from a ZIP of up to 40 documents at half the AI cost; Your Archive takes up to 300 files a drop and files them while you watch.

Your whole contract estate as one console, led by the ledger. A navy mast states the estate in a single reading: vendors and agreements in force, the annual run rate, the top vendor's share, what still awaits a first review, and the nearest renewal. Four numbered views sit under it, and the vendor ledger is the working surface. Drop files or entire folders and Vera reads every document, fills in its terms, files it under its vendor, and judges whether each agreement is current, due for renewal, expired, superseded, or a perpetual license. The judgments respect how vendors actually work: a perpetual Oracle license from years ago is a current asset, and only your newest Microsoft Enterprise Agreement counts as live.

On the screen

- **The mast** — the estate's verdict in one row of figures: annual run rate, the top vendor's share, agreements awaiting first review, and the nearest renewal, with Add documents and the review action beside it. A meter on the tab band tracks how much of the estate has been reviewed.
- **The five views** — numbered tabs on the mast, the console's one navigation: Vendors, Agreements, Renewals, the Review queue, and the Explorer.
- **Vendors** — the ledger: every vendor ranked by annual spend with its share of estate drawn in the row, the exact annual value, the next renewal with months out (inside 6 months reads in gold), and what still needs review. An accounting rule closes the table with the estate total. Select a row to focus its vendor file; the chevron or a double press opens the full file.
- **The vendor file** — the focused vendor as a cover sheet beside the ledger: annual value and share, agreements, what awaits review, the next renewal, the largest agreements as press-to-open rows, and the analyst read in plain prose, with Open the vendor file and the review action beneath.
- **Agreements** — every document in the estate, largest first: title and filename, vendor, the term drawn against a today line, annual value, and the lifecycle state. Press a row for the document peek.
- **Renewals** — the next 24 months, soonest first, opened by the runway read: how many renewals sit inside 18 months and what they are worth, or the all clear.
- **Review queue** — the documents Vera could not fully read, each with the reason, and Review all one press away. Your correction always outranks the automatic reading.
- **The Explorer** — the estate as a file library: one folder per vendor, each with its document count, annual value, and a meter showing how much of it Vera has indexed. Open a folder and the documents sit as a dense details list, one row per paper: name and file name, the indexed chip, file size, last updated, and annual value. Drag a row onto another vendor in the folder rail to refile it by hand, exactly like a review-

queue correction; the hover actions read the paper on the Document Desk, start a renewal project, or delete it. Drop files from your desktop on a folder or anywhere on the view and a confirm dialog asks before anything is uploaded or indexed (indexing uses your plan's scan allowance). Deleting is permanent, stored file and contract record both, with a deletion certificate written for your records.

- **Index with Vera** — documents filed before scanning existed (or imported without a read) show as Not indexed. The Explorer's toolbar offers **Index with Vera** for the whole archive or the open folder, and the Document Desk offers it for the single paper in front of you. You confirm first; the run uses your plan's scan allowance, the same reading pipeline as the intake fills in vendors, terms, values, and lifecycle judgments, and your own hand corrections always outrank the reading.
- **Search, filters, and sort** — find anything by vendor, title, or filename; filter agreements to Current, Renewal due, History, or Needs a look; sort vendors by biggest spend, due first, or A to Z.
- **What next** — the console closes with the moves the estate itself suggests: give the estate its first read, get ahead of the nearest renewal, open the top vendor's room, or ask Vera where the risk is.
- **Adding documents** — drag PDF or DOCX files or whole folders anywhere on the page, or press Add documents. Up to 40 MB per file and 300 files per run; every file is verified by its contents, not its name, and stored privately.
- **The filing ticker** — during a scan a live feed narrates each document as Vera reads it: filename, vendor, value, renewal date. Leave any time; a floating chip follows you on every screen with the percent and a notification brings you back when the archive is ready, even if you closed the tab mid-scan.
- **Next moves** — after a scan, a numbered rail with the money named on each decision, and nothing happens without your confirmation. Renewals: Vera lists only the contracts genuinely coming due (never ten-year-old paper) and you choose which go on the calendar; vendors you already track are advanced, never duplicated. Spend: your top 20 vendors by annual spend, ready for the Spend Optimizer. Benchmarks: a card wall with vendor logos and a quota dial to pick which contracts our analysts should benchmark. Each move ticks off as you answer it.
- **The finishing ceremony** — when a scan completes while you are on the page, a short full screen reveal stamps in the vendors and contracts, runs the spend odometer, and names the renewals due with their money. Skip it any time; it plays once per scan and lands instantly under reduced motion.
- **The Estate Intake Brief** — every scan produces an executive document in the house brief style: the vendor table, the renewals inside the window, the lifecycle findings, and a bottom line. Open it from the ceremony, the Next moves rail, or the button above the floor; it prints clean for a CFO forward.
- **Timeline and peek** — inside a vendor file, the Timeline tab lays every agreement on a time axis against a today line, and clicking any contract opens a side peek with its terms, the lifecycle reasoning, the verbatim passages each date and value was read from, View the document, and the full record one click away.
- **The Document Desk** — View the document opens the paper in a near full screen reading window over the archive instead of a bare browser tab. The document renders in the middle (PDFs natively with zoom controls in the top bar, 50 to 300 percent or Fit; Word documents through the Office viewer), the left rail lists your whole archive grouped by vendor with a search box, and the arrow keys step from paper to paper. The right rail shows the scan's read with its verbatim passages, a what if lever that prices a modeled discount against the annual value on record (per year and through term end), and the ways out: Send it to Vera AI (the Vera door opens with the document already attached, you only choose the read), Start a

renewal project (onto the agreement ledger and into the renewal room, never duplicating a vendor), the Deal Room, the contract decoder, self play, and second opinion. Forward sends the document to a teammate as an in-app heads up plus an email linking the contract record; the file itself never leaves the workspace, and every open uses a short lived signed link.
- **Needs a look** — documents the scan could not read wait behind the review coverage panel and the ledger's gold chip. Fill in a vendor, end date, or value in the review queue and the contract refiles instantly; your correction always outranks the automatic reading.

How to use it

1. Drop your contract folder, then leave the screen if you like; the chip and the completion notification keep you posted.
2. When the scan finishes, work through the three follow-up asks: add the due renewals, review the top spend, choose benchmarks.
3. Read the console any time: select the biggest row, read the vendor file beside it, and open the full file for every agreement and its evidence.

Tips

- Files that fail verification are skipped with the reason on the file chip; the rest of the batch continues.
- Uploads land as contracts, so everything you drop also feeds Document search, the spend portfolio, and the vendor pages automatically.
- A reloaded tab cannot resume files that never finished uploading, but everything already uploaded still scans and gets filed.

A worked example. A folder holding the four sample agreements comes back as four vendor folders worth $331,200 a year together, each contract judged: the Tailspin endpoint agreement (400 endpoints, $38,400 a year) carries a 90 day notice period, so it is the kind of row the Next moves rail asks you to put on the calendar, while an expired predecessor in the same folder files as history without polluting anything. The three questions at the end (renewals, top spend, benchmarks) turn the filing into a working desk in one sitting.

If something looks wrong

- Files were skipped: the reason is on each file chip, and the survey never drops anything silently. Only PDF and DOCX are read, 40 MB each; Troubleshooting covers the rest.
- A contract reads with no terms: it is probably a scan with no text layer, queued under **Needs a look** rather than filed wrong. Fill in the vendor and dates and it refiles instantly.

Related: Proposals · Renewals · Spend optimization

COMPARE

Where: Contracts & documents › Compare · `/contracts/compare` · via **More** or Ctrl/⌘ K
Who: everyone

A saved comparison says when the paper moved: on the contract record and on your Comparisons shelf it shows the date it ran and names the compared document that has a newer version since, with a **Compare again** link that reopens this desk with the same two contracts selected. Nothing re-runs on its own.

Side-by-side document comparison with an AI difference scan: two proposals or versions next to each other with the changes called out. The result opens with a verdict band (terms compared, which side is stronger where, watch-outs), then one card per term showing both sides' language with the stronger position highlighted, a **Your options** card laying out your next steps, and an amber watch-out panel closing the read.

On the screen

- **Document panes** — the two documents side by side.
- **Difference scan** — the AI's read of what changed and what it means commercially.
- **Your options** — two to four concrete next steps drawn from the scan, numbered, with exactly one marked **Recommended**, so the read ends in a decision rather than a table. The same block appears on the compare card each contract's record page keeps, told from that contract's own perspective.
- **Saved automatically** — every completed scan files itself on both contract records and on the Comparisons shelf of your Saved page; the green note under the result confirms it.
- **What next** — a finished comparison closes with a What next panel: seed a case or a renewal playbook on the weaker contract, put an Ask Vera question on the sharpest difference, or copy the verdict for a chat thread.

How to use it

1. Pick two documents and run the scan. While it reads, a progress bar counts toward done and names each step as it happens; a full read usually lands inside a minute.
2. Read the verdict and pick from **Your options**. There is no save step: the finished scan is already on both contract records and on the Comparisons shelf of your Saved page, so you can reopen the pair any time.

Tips

- Use this to check a renewal quote against last year's paper, or a redlined draft against the version you sent. For a year-over-year read on the same vendor, see the Time machine tool.

Related: Saved

VENDORS

Where: Contracts & documents › Vendors · `/vendors`
Who: everyone

Every vendor across your proposals and agreements, one page per vendor, so you can see your whole relationship in one place.

In plain words: the whole relationship with one vendor, on one page.

When to reach for this

- Before any vendor conversation: contracts, benchmarks, market intel, and notes in one read.
- Taking over a colleague's book: this page per vendor is most of the handover (the recipe).
- Opening the case file when the meeting deserves the full executive-brief treatment.

On the screen

- **Your position** — the masthead states the estate in one sentence under a gold rule: how many vendors you hold, their stated annual value, how many renew inside 90 days, and how many are priced above the market median.
- **The briefing band** — three reads before the rows. **Where the money is** draws your spend concentration as one bar across the top five vendors, with the share each holds. **The renewal runway** plots every tracked renewal across the next twelve months, soonest first, and counts the vendors with no renewal date on file. **Needs attention** names up to four profiles worth opening today, each with the reason: renewing inside the window, priced above market, or carrying proposals with no stated value.
- **The register** — vendors grouped by the role they play in your estate: anchor vendors (the top five by stated value), the working portfolio, and the profiles with no stated value. Each row carries the vendor, a one line position statement, a verdict in plain colors, the annual value, and the renewal countdown. Choosing a sort (spend, renewal urgency, recently active) flattens the groups into one ranked list.
- **Vendor page** — open a vendor for its contracts, benchmarks, market intel, and notes.
- **Open the case** — a navy band under the related records opens the vendor as a case file: a full screen file of numbered executive brief documents in a grouped left rail, paged with the arrow keys or the rail, with a read meter tracking your progress and Esc to close. Up to six documents build from what the vendor's record holds: your agreements ledger, the pricing band at your deal size, the vendor's terms and conduct intel, a negotiation angles brief fitted to your footprint, and a closing What now sheet whose every move runs on your live records. Drafts open in your own inbox, copies go to your clipboard, and nothing ever sends from a case. Several records across the platform open the same way; each section below says what its case holds.

How to use it

1. Open a vendor to see everything you have with them.
2. Run an instant benchmark for a vendor from here.

WATCH

Where: Renewals & deals › Watch, and the Market Desk's **Your watchtower** tab · `/wire/watch`
Who: everyone

One watchtower, four rooms. The **Terms** tab watches the vendor side: every change to a tracked vendor's published price list, diffed and priced at your seat counts. The **Billing** tab watches your side: what you are billed against what you are entitled to against what the market pays, checked daily. The **Price list check** tab is the machinery behind the first tab on one screen: every price list the platform tracks, the live state of the automated verification sweep, and the controls that decide which vendors' changes reach your inbox. The **Archive by vendor** tab holds every price-list change ever detected, re-sorted by vendor. The four sit as a pill row in the page header, under the Market Desk strip, so you can move between them from any of them, including the archive. The old addresses (`/vendors/watch` for Terms watch, `/spend/watch` for Billing watch) redirect to the right room, so old links, bookmarks, and emails keep working.

All three working rooms open with a navy masthead that states the room's live figures, keep their context in a column beside the work, and carry the same navy band: **Open the workroom** raises the watchtower file, a four document case that reads the whole tower for the last 90 days. The watch brief states the period in tiles and one narrative, the change ledger lists every change on your vendors with both published prices and your exposure, On your paper folds the period per vendor with the renewal windows to act, and What now carries the stakeholder memo, the copy moves, and the rooms. ⚡ **What now** and ✦ **Ask Vera**, grounded on the file, ride beside every paper, and the same file opens from any of the three rooms.

Terms watch

> **Where:** Your watchtower › Terms · `/wire/watch`
> **Who:** everyone

Every change to a tracked vendor's published price list, diffed on each refresh and priced at your own seat counts. This is the vendor side of the watch: what a vendor moved, and what that move would cost you at renewal. Every change is kept, so you can read a vendor's full price history any time.

In plain words: when a vendor moves its published prices, you find out, priced at your own seat counts.

When to reach for this

- Setting up the watch on your top vendors, once, in the week the estate lands.
- A renewal is inside a year: the vendor's recent moves are your negotiation file.
- A price rise just landed mid-term: the change's case file carries the response play.

Your position sits in the masthead: price changes on your vendors in the last 90 days, vendors tracked and how many are in your estate, your annualized exposure at the next renewals, and the biggest single move in the window. The vendor follow switch and the other standing watches sit in the column beside the strips.

On the screen

- **The watchtower workroom** — the navy band above the floor opens the watchtower file described at the top of this section: the whole period across Terms, Billing, and the check, as four executive papers with the moves and Ask Vera docked beside them.
- **The floor toolbar** — one instrument row over the strips: scope the board to your vendors, the tracked market, or both; cut it to increases, decreases, new SKUs, or retirements; and read the last 90 days or the full year. The count at the right restates what the filter left, so a narrowed board never reads as a quiet market.
- **Email me when price lists change** — a switch in the side column turns on the price-change digest; nothing is sent until you turn it on, and every email carries a one-click unsubscribe that turns off only this email.

Below it, one chip per vendor company (Microsoft covers Microsoft 365, Azure, and Dynamics 365 in a single chip; Oracle and Google merge the same way): pick the vendors you care about, or choose **All tracked vendors** to hear about every tracked change. The list here and on the Price list check tab are the same roster, so the two tabs always agree.

- **Your vendors** — a seismograph strip per vendor you hold: increases strike upward in brick, decreases downward in green, new SKUs in gold.
- **Open the change workroom** — selecting a change shows a second navy band under the diff panel; it opens that one change as a five document workroom: the cover sheet, the change itself with both published list entries side by side, an impact brief priced at your own seat counts, a response sheet weighing the grandfathering, renewal, and acceptance options, and a What now sheet. ⚡ **What now** drafts five letters from the change itself, the grandfathering request, the claim on a new list price, a clarification request on a new SKU, a continuity request on a retired SKU, and the internal note to stakeholders, then copies the change summary or the one line version and opens your renewals or the full price history. Both published prices are quoted rather than derived. ✦ **Ask Vera** answers from this change and your own exposure to it; if you do not hold the vendor, the room says so and treats the move as market intelligence rather than exposure.
- **Market watch** — tracked vendors outside your estate. Changes here are context, not exposure, because nothing below touches a contract you hold.
- **Also watching for you** — cross-links to the Billing tab and to standing answers that changed.

How to use it

1. Read your vendors first: any strip with a recent brick strike is a price rise heading for a renewal you hold.
2. Use the switch to have price-list changes emailed to you, and follow the vendors you want to hear about.
3. Open a change to see the old and new price, the delta, and what it costs at your seat counts.
4. Switch to the Billing tab to see whether your invoices already reflect a move.

If something looks wrong

- The digest emails are not arriving: the switch here is the source of truth, and nothing sends until it is on; Troubleshooting covers the delivery cases.
- Your vendor's strip is missing: following is per vendor company, so check the chip roster on the Price list check tab, and request the vendor from its card if it is not tracked yet.

Related: Archive by vendor · Billing watch · Weekly report

Terms watch archive by vendor

Where: Your watchtower › Archive by vendor · `/wire/watch/archive`
Who: everyone

Every published price-list change Terms Watch has ever detected, kept and re-sorted by vendor. It opens with the read: how many vendors have moved, increases against cuts, and which vendor moves its list most. The index lists every vendor whose price list has moved, with how many changes it has, a breakdown of increases, cuts, new SKUs, and retirements, and when it last moved. Open a vendor to read its full history, newest first:

each change shows the old and new price, the delta, and the grounded analyst note (what changed, why it likely changed, and the one action to take before renewal).

How to use it

1. Press **Archive by vendor** in the watchtower's pill row, from any of its rooms.
2. Scan the index (most-changed vendors first) or find the vendor you are about to renew.
3. Open the vendor to read its full price history and the note attached to each change.

Billing watch

Where: Your watchtower › Billing · `/wire/watch/billing`
Who: everyone

Every day, per vendor per SKU: what you are billed against what you are entitled to against what the market pays. Drift arrives with its corrective draft attached, and a correction only counts once later invoices confirm it. This is the customer side of the same seismograph as the Terms tab.

On the screen

- **The masthead** — the headline number lives in the navy masthead: how much a month is leaking, across how many vendors, with live bleed, open findings, vendors leaking, and corrections landed. **Recompute** sits beside it.
- **The watchtower workroom** — the navy band above the board opens the shared watchtower file (see the top of this section), which reads the billing drift beside the period's price-list changes.
- **Leak board** — a row per leaking vendor with the SKUs where billed, entitled, and market prices disagree.
- **The side column** — how the watch reads (billed against entitled against market), entitlements waiting for your confirmation, and the rest of the tower: the Terms tab, invoice disputes, proven corrections, and standing answers that changed.

How to use it

1. Read the verdict band for the live monthly bleed.
2. Open a leaking vendor to see the SKU-level findings and the corrective draft attached to each.
3. If the board is empty, confirm entitlements from a contract or connect a license source so the daily check has something to price against.

Tips

- An approved correction is proven or disproven against later invoices on Savings proof; mischarge disputes are tracked on Invoices.

Related: Terms watch

Price list check

Where: Your watchtower › Price list check · `/wire/watch/price-lists`
Who: everyone

The verification machine on one screen. The book of published list prices refreshes weekly, an automated check re-reads each vendor's own public pricing page and compares it against the anchors on file, walking the whole roster inside every week, and an analyst confirms any detected change before it is published to the Terms tab. This tab shows all of it live, and doubles as the place you choose what you are notified about.

On the screen

- **The watchtower workroom** — the navy band above the desk opens the shared watchtower file (see the top of this section).
- **The sweep band** — the dark panel at the top of the desk states the live numbers: price lists under watch, published prices on file, checks against vendor pages in the last 7 days, and detections currently with an analyst. Beneath them, a heartbeat marks every automated check this week: green verified clean, gold a detected change, red a page that could not be read and retries.
- **Notification bar** — the email switch, the **All tracked vendors** chip, and your follow count. The roster below is the picker: the bell on each row follows or unfollows that vendor and saves instantly, no separate save step. Following is one chip per vendor company, so all the Microsoft lists ride one follow.
- **Latest verified changes** — the most recent analyst-confirmed changes with their deltas, linking to the Terms tab where each is priced at your seat counts.
- **Search and lenses** — find a vendor by name, or cut the roster to your vendors, lists with changes, stale prices, or lists needing attention.
- **The roster** — grouped by category, a row per price list: the vendors it covers, how many SKUs are on file, a freshness bar for the published prices, the list's recent check history, and the latest result. A gold dot marks vendors in your estate.
- **The drawer** — open any row for the full detail: every anchor's list price, unit, publication date, and confidence, a price-history trend per SKU, the latest check's verdict for each anchor (matches the page, change detected, or not printed publicly) with a link to the vendor's own pricing page.
- **The side column** — how the check works in three steps (the book refreshes, the check reads the vendor's page, an analyst verifies), and **Request a vendor**: buying from someone we do not track? Ask and the list joins the roster with its own weekly check.

How to use it

1. Read the sweep band: if detections sit with an analyst, a price move may be coming to Terms soon.
2. Turn on the email switch, then either keep **All tracked vendors** or click bells to follow just the vendors you hold.
3. Use the **Your vendors** lens to confirm your estate is covered; anything green was recently re-verified against the vendor's own page.

4. Open a row before a renewal to read the vendor's current anchors and their history, straight from the book our analysts maintain.

Related: Terms watch · Archive by vendor

WEEKLY REPORT

Where: Benchmarks & market › Weekly report, and the Market Desk's **Weekly report** tab · `/wire/weekly`
Who: everyone

Two rooms under the tab: **This week** and **Archive by vendor**, as a pill row in the page header.

Every Monday the agent sweeps the week's enterprise software news for pricing and licensing changes, anything that moves costs or negotiating position, and files this wire. It is short on quiet weeks and longer only when something major breaks; items on vendors you hold are marked.

The page opens with a plain-language rundown of what each Monday update covers: price and discount moves, licensing and packaging changes, support, audit and end-of-life shifts, and the M&A or earnings news that changes your leverage, each with a read on what to do about it.

On the screen

- **What each week covers** — a short panel at the top listing the kinds of change the agent files each week.
- **The wire** — the week's report as an executive brief, built to scan: the executive summary opens with a lede line and bullets the rest of the week, and every item separates the facts (a short bulleted list) from a **For buyers** list with one move per bullet. Items touching a vendor you hold carry an "in your estate" eyebrow.
- **Open the workroom** — on a week with items, a navy band above the paper opens the wire as a four document workroom: a cover sheet with the week at a glance and the item to read first, the week's ledger with every item as a row, a priority brief that reorders the wire for your desk (your vendors and major items first), and a What now sheet. ⚡ **What now** forwards the week's read as a drafted note, copies the week summary for chat or the one line version, and opens your renewals, the Market Desk, or the archive by vendor. ✦ **Ask Vera** reads this week's own items, each with its magnitude, its source and its date, so what in the week affects you is answered from the wire itself rather than from a general summary of the market. The room recomposes from the live report on every open; quiet weeks carry no band.
- **Subscribe card** — a switch at the top left turns on the Monday email; nothing is sent until you turn it on, and every email carries a one-click unsubscribe that turns off only this email. Below it, pick the vendors you care about so their items lead your copy.
- **Past weeks** — the archive strip below the brief; a gold dot marks a week where something major broke.
- **Browse the full archive by vendor** — a link into the per-vendor archive (see below).

How to use it

1. Read the latest wire, starting with the items marked as in your estate.
2. Use the switch in the subscribe card to have it emailed to you each Monday.

3. Open a past week from the archive strip to catch up, or open the archive by vendor to read one vendor's full history.

Tips

- Your weekly emails now arrive as one Monday brief: the vendor market wire (this report), renewals coming up, watchtower changes, and savings and wins in a single send. Turn any section off while keeping the rest under Settings › Notifications.

Related: Archive by vendor · Terms watch

Weekly report archive by vendor

Where: Weekly report › Archive by vendor · `/wire/weekly/archive`
Who: everyone

Every weekly update the agent has ever filed, kept and re-sorted by vendor so nothing drops off. It opens with the read: how many vendors carry a major market-moving update, how many touch pricing, and which vendor the desk keeps writing about. The index lists every vendor the agent has covered, with how many updates it has and when it was last mentioned; a gold dot marks a vendor with a major update on record. Open a vendor to read its full history: every price move, licensing change, and market shift on it, newest first, each linking back to the week it ran in.

Commission the desk sits beside the archive: gold chips that turn what is on file into a written document rather than leaving you to act on a feed. A **renewal strategy brief** on the vendor the desk covers most, offered only where your workspace actually holds a contract for that vendor; a **Monday briefing** that states how many notice deadlines fall inside 45 days; and a closing question that opens Vera on the same subject. Each commission writes a full document to your Deliverables and keeps writing even if you close the dialog. The same rail rides the price watch archive, where the chips are the renewal strategy brief on the vendor with the most published price moves and talking points for the call on that vendor.

How to use it

1. Press **Archive by vendor** in the weekly report's pill row.
2. Scan the index (most-covered vendors first) or find the vendor you are about to renew.
3. Open the vendor to read its timeline, and click any **Week of** eyebrow to jump to the full weekly brief it came from.
4. Press a **Commission the desk** chip to turn the coverage into a document you can forward.

AGREEMENTS

Where: Contracts & documents › Agreements · `/agreements` · via **More** or Ctrl/⌘ K
Who: everyone

The ledger is signed paper only, one row per vendor. The full document estate is in Your Archive, and a quote you want priced against the market goes to Proposals.

Signed contracts you are tracking for spend and renewals, worked as a ledger and an inspector. The table groups every vendor folder by what binds you first; selecting a row opens that folder in full beside the table. Set the expiry and notice period once, and the platform watches the deadlines for you.

In plain words: the ledger of what you signed, with the dates watched so windows never shut silently.

When to reach for this

- Filing a newly signed contract, the day it is signed, while the dates are at hand.
- Checking what reaches renewal inside 90 days and what vendors are asking above your caps.
- Auditing inherited agreements for the two fields that matter: expiry and notice.

A worked example. The sample Northwind CRM agreement carries $142,800 a year, an expiry with a 60 day notice period, and a 7 percent uplift cap. The vendor's renewal proposal on file asks $162,792, a 14 percent increase, which is exactly the kind of gap the masthead's "asking above your caps" figure counts and Anomalies flags. One ledger row, set up once, is what makes that breach visible without anyone rereading the contract.

On the screen

- **The ledger** — one row per vendor, grouped by urgency: notice window open, past renewal, renews inside 90 days, later this cycle, and no date yet. Each group states its count and annual value, and the runway column shows how much of the final year has run with the notice deadline ticked on it.
- **The inspector** — select a row and the panel beside the table works that folder in full: the term drawn as a timeline, the editable terms, the uplift verdict against your cap, and up to five documents. On a narrow screen it stacks under the table. Days remaining are counted on your organization timezone's calendar day, the same one deadlines and alerts use, which owners set under Renewals.
- **Chips, search and views** — filter to the notice windows, the next 90 days, past renewals or undated folders in one press, search the ledger by vendor, or switch to the **Runway** tab to see every renewal drawn on a timeline. You can track up to forty vendor agreements.

How to use it

1. Add an agreement with **Add vendor** or **Add with AI**: vendor, expiry date, notice period, and the signed files.
2. Press the **In notice** chip and work the top row in the inspector: confirm the dates, enter the vendor's proposed price against your cap, and file the paper.
3. See them sorted by deadline on Renewals.

If something looks wrong

- An agreement never alerts: check it carries both an expiry date and a notice period; without both there is no window to watch. And remember the alerts themselves are opt-in under Notifications.
- You are out of room: the ledger tracks up to 40 vendor agreements with 5 files each. Past that scale, the estate belongs in Your Archive or Migration Studio.

Related: Renewals

THE DESK (BETA)

Where: Renewals & deals › The Desk · `/desk`
Who: everyone. Signing or revoking a mandate needs edit access.

The buy side as one desk file. The navy stage states the book in figures and opens the file, the morning brief is the landing paper, everything waiting on a person in priority order, and the console below holds the live controls for the selected supplier as tabs: Position, Mandate, Execute, Close.

What beta means here. Live on your data: the morning brief, the ranked queue, the leverage behind every row, the connected feeds it reads, signed mandates, real outbound actions with the guardian ruling on each one, and closed deals handed to Savings Proof. Still to come: the desk never picks who a message reaches, and it does not yet open events on its own schedule.

On the screen

- **The stage**: the desk's one navy surface. It states the book in figures, spend under the desk, money in the workable window (renewals 60 to 150 days out), what waits on a person, mandates in force and savings verified, and carries two doors: **Open the desk file** raises the full file, and **How the desk works** opens the whole loop written down, station by station.
- **The morning brief**: the landing paper, prepared from your paper and feeds when you open the desk. Section one is what needs you, in priority order: notice deadlines inside a week, guardian refusals to read, counters to approve, approved drafts that only need a recipient, mandates about to lapse, windows open with no authority, and closes leaking against their invoices. Every item states the supplier, the reason and the money, and its link lands the console on the right supplier and station. Then the week ahead, dated; what the desk is watching; and the feeds behind the figures. A feed you have not connected reads as unmeasured, never as zero, because something the desk cannot see must never look clean.
- **The queue**: 0 to 100 leverage per agreement, from money at stake, how workable the renewal window still is, and the levers your own paper gives you. The window matters both ways: a deadline 14 days out scores lower than one 94 days out, because there is no longer room to run a real negotiation. An auto-renewing agreement, a proposal above your uplift cap, a vendor you have never benchmarked, and an unowned renewal each add to the score. Rows carry the days left, a mandate pill when one is in force, and the money any cap breach is worth. The queue also states who owns each deal: a supplier already in a Vera-run case or a human-run negotiation is flagged on its row, with a link to the owning record.
- **Position tab**: the brief the desk wrote before you opened it. The renewal window drawn in time with the workable band shaded. What your paper says: current spend, the vendor's proposal, the uplift cap, and the money any breach is worth, recoverable without a market argument. What the feeds measured: unused seats and billing drift, each labelled unmeasured when its feed is not connected. The market corridor drawn as a ruler: the observed band from the median close to the upper quartile, with the three

mandate targets marked and the number of closes behind them. And the signal ledger showing exactly why the row ranks where it ranks.

- **Where the targets come from**: push gently, push hard and go for the best are not fixed percentages. They come from the corridor of reductions actually achieved on that vendor: the median close, the upper quartile, and one step past the observed range for the anchor. Each option shows its percentage and the closes behind it, and the note underneath says whether those closes are this vendor's own or borrowed from the all-vendor pool. With fewer than five closes, the figures are labelled house postures rather than measurements. Switching rows loads that vendor's own corridor and cohort.
- **Mandate tab**: the console on the left sets the terms; the instrument card on the right reads them back as the document you are signing: objective, disclosure, autonomy, authority, term, and the count of human decisions it implies, two under a fully autonomous mandate against 15 to 40 by hand. A signed mandate is never edited; changing your mind revokes it and signs a new one, so the record is a sequence of decisions each carrying a signer and a timestamp.
- **The desk file**: the whole desk as formal papers, opened from the stage, from the navy band below the brief, or by sharing a `?open=case` link. It lands on the morning brief and reads on: **the book**, every agreement ranked by leverage with its standing and one line of why, then the selected supplier's case as a workroom: the cover sheet, the position sheet with the corridor targets and their provenance, the signed mandate read back as a formal instrument, the guardian ruled execution ledger and the close package with its verification state once that data exists, and a What now sheet. ⚡ **What now** briefs a stakeholder on the mandate, reports the close, prepares a counter under the mandate, copies the case summary, jumps back to the stations, and opens the savings ledger once a close verifies. The drafts state the mandate's real basis, including how many observed closes it rests on, whether those closes are this vendor's own or borrowed from the all-vendor pool, and whether the number reaches one step past the observed range on purpose, so a target never reads as more measured than it is. ✦ **Ask Vera** answers from the papers in the file.
- **Guardian preflight**: judges three moves against the settings on screen before you sign, naming the rule behind each verdict. Move the disclosure setting and the verdicts change. The case worth understanding is a move your mandate permits but org policy forbids: it is rewritten and released, not blocked, with the diff kept in the audit record.
- **Execute tab**: the action ledger. Press **Prepare a counter** and the desk drafts one from your stored figures, reproducible and incapable of inventing a number. Every row carries the guardian's verdict, the rule that decided it, and a timestamp; drafts open in their own window. The scope toggle switches between this supplier and every supplier. *Ask me first* leaves the draft for you to send, *Check with me* waits for one click, *Just run it* clears it to send. You always type the recipient yourself, and nothing the guardian refused can be promoted.
- **Close tab**: type what you will pay under the new terms and the signature package computes before you commit. The saving in money and percent, whether you met the mandate you signed or came up short and by how much, the market check, the uplift cap, and what is on the record. The mandate line is the one to read: a saving alone does not tell you whether the desk did what you told it to. Figures that are not a close are refused with the reason. The verification bar alongside shows claimed against verified against

leaking, re-checked by Savings Proof against the invoices that follow, and the network panel shows how many separate buyers stand behind your position, behind the floor of five and the give-to-get rule.

How to use it

1. Read the morning brief. Its first item is the day's priority, and its link lands the console on the supplier and the station that can act on it.
2. Press a queue row and read its **Position** tab. Every figure carries its source.
3. Open **Mandate**, set how hard to push, what may be revealed, how much rope, the spend cap and the duration, then read the instrument and the guardian preflight before you sign.
4. Open **Execute**, press **Prepare a counter**, read the guardian's ruling, and send it to a recipient you type yourself.
5. When terms land, open **Close**, enter the new price, read the package against your mandate, and press **Sign off and record**.
6. Press **Open the desk file** on the stage to read the whole desk as papers, or to walk someone else through it.

Related: Renewals · Negotiations · Agreements

VERA OS (BETA)

Where: Renewals & deals › Vera OS · `/vera-os`
Who: everyone; in beta and included with your plan

Vera OS is the autonomous procurement companion: a separate workspace where agents work renewals and deals end to end under mandates you sign, grounded on your VendorBenchmark data. The page on this rail is its front door, with the Control Tower reading your own estate live. Because it runs as its own product alongside the platform, it is documented in full in the technical reference: Vera OS, the autonomous procurement companion.

Related: The Desk (beta) · Renewals

RENEWALS

Where: Renewals & deals › Renewals · `/renewals`
Who: everyone

One renewal at a time, argued in full. The queue on the left is every agreement ranked by how soon notice falls due; the case in the middle states the decision, the money on each option, and what to do next. Renewal alerts are opt-in: anyone who turned them on in their notification settings hears when a notice window opens.

In plain words: your deadlines, ranked, each argued as a decision with the money on every option.

When to reach for this

- The Monday pass: what shuts next, and what it is worth (the desk routine).

- A renewal is close enough to decide: the case states the options rather than just the date.
- Choosing where the quarter's negotiation effort goes, by money and time left rather than instinct.

A worked example. On the sample Northwind CRM renewal, the three outcomes read: renewing as offered, $162,792 under the vendor's proposal; holding today's price, $142,800; and pricing at the market median, which is only derivable because a benchmark was delivered, and therefore carries the **recommended** flag. The gap between the vendor's ask and the recommended outcome is the case in one line, and the clock beside it says how much sourcing runway is already gone. That is the reading pattern for every case: options, money, time.

On the screen

- **The queue** — every agreement grouped into Decide now, Inside the sourcing runway, and Ahead of the runway, with the countdown to its notice date in tabular figures down the left edge, so the column reads as a list of numbers without reading a single vendor name. The total value of the book, how many sit inside the notice window and how many carry no benchmark are stated at the top. Press `j` and `k` to step through it without the mouse.
- **The case** — opens with one sentence that states the decision and the money on it, not a summary of the row you just pressed.
- **The clock** — the life of the renewal drawn to scale and graduated every ten days, each graduation a real date. The hatched stretch is your sourcing runway, the time you are meant to be working; the solid stretch is the notice window, the time the contract is deciding for you. Today is marked on it, so how much runway is already behind you is one glance rather than a subtraction.
- **The three outcomes** — renewing as offered, holding today's price, and pricing at the market median, with the annual figure on each and the difference against the opening position. Renewing as offered uses the uplift cap recorded on your agreement where there is one, and says so when it has had to assume one instead. The outcome marked **recommended** is only ever the one your delivered benchmark supports; where no benchmark has been delivered the median reads *Not derivable* rather than a guess.
- **The argument** — why the renewal sits where it does, composed from figures stored against the agreement rather than written for you, plus a cream panel on what the vendor's own fiscal quarter close does to your timing. A vendor with a quarter to close has discount authority it does not have the week after.
- **What to do next** — the moves in order, each one a link to where the move happens, and a button that opens the full renewal record.
- **Engage the vendor** — the two doors at the foot of the case, and the point of the whole desk. **Have Vera run it** opens a case with this renewal attached, so Vera benchmarks the position, drafts the strategy, and negotiates by email with your approval on every send. **Track it yourself** opens a negotiation thread prefilled with the vendor, so every ask, counter, and commitment stays on the record and the writing stays yours. Neither door creates anything on its own: each one opens the form, and you decide there.
- **When a deal already exists** — if a case or a negotiation is already running on that vendor, the two doors collapse to a single **Open the case** or **Open the negotiation** link, so the work stays on one record instead

of splitting across two. A Vera-run case takes precedence over a thread on the same vendor.

- **The guided playbook line** — when the vendor is one of the nine publishers with a guided playbook (Microsoft, SAP, Oracle, IBM, Adobe, Salesforce, Anthropic, Workday, ServiceNow), one line names it and links it. The match is on the publisher family, so Oracle America, Inc. and Microsoft Ireland Operations Ltd both resolve to their playbook.
- **The dossier** — on the right: the agreement's terms with leader dots, where your unit price falls in its market cohort with the median marked, and the risk score with every reason behind it listed.
- **The sample book** — **See the sample book** in the header fills the desk with fourteen worked example agreements, ten of them benchmarked, so you can see the desk working before your own agreements are on file. It includes a notice date already missed on a contract that renews itself, one seven days out, two priced below the market median where the work is defending them, and four with no benchmark at all. Countdowns are live. Nothing in it is your data and nothing in it is saved, alerted on or exported; a banner says so the whole time it is on.
- **Calendar** — the year at a glance is one tab across: the month grid, the renewal horizon strip, the stat band, the quarter capacity planner, Present mode and the calendar governance settings.
- **War room** — once a renewal crosses into its notice window, its vendor record leads with a live countdown, this week's playbook step, the negotiation brief, and the owner and approvers pinned, and the deal jumps to the top of My work.
- **The Renewal Event Engine** — on the Calendar tab. Owners switch it on and renewals stop being calendar entries: at 90 days before each notice window the platform opens the event itself, the deal thread is created, a mandate skeleton is drafted from the market corridor and the agreement's own figures, and the owner is notified. Humans still approve every mandate; the engine only opens the door, on up to five renewals per day.

How to use it

1. Read the case at the top of the queue, which is always the one closest to its notice date.
2. Check which outcome carries the **recommended** flag, and what it is worth against the opening position.
3. Work down **What to do next**, or press `j` to move to the next case.
4. When you are ready to engage the vendor, pick a door under **Engage the vendor**: **Have Vera run it** to hand the cycle to a case, or **Track it yourself** to open a tracked negotiation.
5. Build the full plan with the Renewal playbook.

If something looks wrong

- The queue is empty: rows come from agreements carrying an expiry and a notice period, so an empty desk means missing dates, not a quiet year. **See the sample book** shows the desk working while you fix that.
- The median outcome reads *Not derivable*: no benchmark has been delivered for the deal, so the desk refuses to guess. An instant benchmark or an uploaded proposal is what turns it on.

Tips

- The sample book hides the engage doors, because a worked example has no real vendor to engage. Switch back to your own book to see them.

Related: Renewal calendar · Renewal playbook · Negotiations

RENEWAL CALENDAR

Where: Renewals & deals › Renewals › **Calendar** · `/renewals/calendar`
Who: everyone

The year at a glance: every notice date, expiry and vendor quarter close. Planning a year of capacity and taking one decision are different jobs, so the calendar keeps its own tab beside the case desk.

In plain words: the year's deadlines and the vendors' quarter ends on one grid, so timing stops being luck.

When to reach for this

- The annual planning sit-down (the recipe): where the capacity crunch quarters are.
- Timing an ask: a vendor with a quarter to close has discount authority it lacks the week after.
- The QBR: Present mode walks the calendar as the session itself.

On the screen

- **Renewal horizon** — the same 18-month strip your dashboard shows: every dated deal sized by spend, colored by risk, and led by its notice deadline. One timeline everywhere, so the dashboard and this page always agree.
- **Month grid** — twelve months with deadlines, expiries and vendor fiscal quarter closes layered on the same days.
- **Stat band and quarter planner** — decisions due, renewals inside the sourcing runway, spend at risk, the next deadline, and the next four quarters with the value and ownership coverage of each.
- **Present** — runs the calendar as a full screen QBR session: the horizon and quarter planner open, each upcoming quarter follows deal by deal with notice-window renewals leading, and the decisions due close the show. Arrow keys move, N shows presenter notes, the print control leaves the walked screens behind as a packet, and a shared link carries the show.
- **Get it managed** — every renewal card carries a control to hand the renewal to a human team as an add-on engagement. If your workspace is served by a partner firm, the request goes straight to them; otherwise it goes to the Vera AI team. One live request per renewal.
- **Calendar governance** — folded shut at the foot: calendar sync, alert timing, vendor aliases, and the approval and SLA policies your renewals run by. The alert card states your own delivery plainly: whether renewal alerts reach you and on which channels, with a link to Settings › Notifications to manage delivery. Alerts stay off until you turn them on; the line only states where you stand.

How to use it

1. Read the horizon for the shape of the year, then the quarter planner for where the capacity crunch is.

2. Press a quarter to narrow the list below it.
3. Subscribe your own calendar app from **Calendar governance** so the notice dates ride along with your diary.

Related: Renewals · Renewal playbook

RENEWAL PLAYBOOK

Where: Renewals & deals › Renewal playbook, or the **Playbook** tab on the renewals desk · `/renewals/playbook`
Who: everyone

Name what is renewing and the platform builds the whole plan: a work-back timeline, the questions to ask, the plays to win, and the negotiation talking points. Turn on the agent and it watches the deadline and keeps you on track.

In plain words: name the renewal and get the plan, worked backwards from the date.

When to reach for this

- A renewal crossed into your working horizon and needs to become a plan, not an entry.
- You know the vendor but not the moves: the plays and talking points are the education.
- You want a deadline watched by something that never gets busy: the agent.

On the screen

- **The work-back wall** — the whole plan on one engraved scale, redrawing live as you type the date and the notice period below it. The three delivery phases are drawn to scale worked back from the notice deadline, the stretch after the notice date is hatched because that is time the contract is deciding rather than time you are working, and the graduations are real dates rather than a progress bar. It is the same instrument the Renewals case desk uses for a single agreement.
- **Their calendar** — the gold dashes on the wall are the vendor own fiscal quarter closes. Discount authority loosens as a seller chases a number, so a decision timed into one of those is worth more than the same decision a week later.
- **The three steps** — name what is renewing, fill in the renewal details, then choose who works it. A completed step stays pressable, so you can change an earlier answer without losing the ones after it.
- **Wizard** — pick an existing contract or just name the vendor and fill in the details. Contracts that already have a playbook offer "open" instead of starting a new one.
- **The agent** — turn it on and it tracks the deadline and prompts you at each step.
- **Engage the vendor** — the same two doors that close the Renewals case, at the foot of the playbook: **Have Vera run it** opens a case with this renewal attached, **Track it yourself** opens a negotiation thread prefilled with the vendor. If a case or a thread is already open on that vendor, one link to it replaces both.

How to use it

1. Pick the contract that is renewing, or name the vendor and its details.

2. Let the platform build the timeline, the discovery questions, the plays, and the talking points.
3. Turn on the agent if you want it to watch the deadline for you.
4. When the plan is built, pick a door under **Engage the vendor** to start the actual work.

Tips

- The step-by-step plan and the agent draw on your AI plan. Without one you still get the plan structure and the work-back timeline.

Related: Renewals · Managed renewals

MANAGED RENEWALS

Where: Renewals & deals › Managed renewals, or the **Managed** tab on the renewals desk · `/renewals/managed`
Who: everyone · activating the program is org owners only

Hand over the calendar: Vera runs every covered renewal under mandates you approve, and the fee is a percentage of savings verified against vendor invoices. Claimed figures never bill.

Your position sits in the masthead: savings banked and invoice-verified, fees paid and accrued, your net position, and how many decisions wait in the queue. The renewals tab strip rides the masthead's foot, so the program stays one press from the calendar and the pipeline.

On the screen

- **The offer** — before you enroll: three numbered cards (Vera negotiates and discloses itself as an AI, you decide the moments that matter, the fee is a share of invoice-verified savings only) above the activation form.
- **The command queue** — leads the enrolled desk: every interruption typed (a mandate, a formal offer, a counter, a bounced contact), aged in place and flagged when it has waited five days, each row opening the deal it waits on.
- **The covered horizon** — every covered renewal on one time axis across the next fourteen months, with the open window shaded and dots sized by contract value. Press a state in the legend to isolate its lanes, hover a lane for the readout, and a lane with an open run clicks through to it.
- **The book of runs** — every run grouped by what it needs (waiting on you, Vera working, preparing, settled), with days to expiry on each line.
- **The fee ledger** — every fee line paired with the verified saving and the invoice that confirms it, with the program totals in the footer.
- **The pairing** — in the column beside the work: savings banked against fees paid, always shown net, with how the fee works and the program as signed beneath it.

How to use it

1. Read the offer and the fee terms.
2. An owner activates the program, which signs a fee agreement and sets the standing mandate defaults.

3. Approve each deal's mandate before anything reaches a vendor, and clear the command queue as Vera works; the horizon and the book of runs show everything else moving without you.

Tips

- Nothing reaches a vendor until you approve that deal's mandate, even with the standing defaults in place.

Related: Renewal playbook · Savings proof

NEGOTIATIONS

Where: Renewals & deals › Negotiations · `/negotiation` · via **More** or Ctrl/⌘ K
Who: everyone

A workspace for live deals. The room is **The Negotiation Clinic**: a lobby with Vera's move queue, one prioritized move per vendor, each carrying a six document case file. Document one, The move, is assembled from your own records: the evidence with each line's source, the commercial picture as a verdict table, a ready draft where the move sends something, and every figure cited to the record it came from. Documents two to six (vendor dossier, benchmark card, before you start, strategy, talking points) are written by Vera on request, at maximum effort, grounded only on what your workspace holds; anything not on file is named as missing, never invented. Opening a vendor opens its case file on the full case stage, the same reading room the rest of the platform uses. Approving a move files the decision and any draft to the deal log; nothing sends to a vendor from the clinic without the owner-set send rules. Each negotiation remains a tracked thread from mandate to signature, so nothing agreed on a call gets lost.

In plain words: the deal as a record: the mandate, every give and get, and what the vendor promised.

When to reach for this

- The moment a renewal turns into a live conversation with the vendor.
- A vendor email needs decoding: paste it and the tactics analysis names the play and the counter.
- Post-signature: the commitments list is what you check the final paper against.

A worked example. Standing up the sample Northwind CRM deal, the mandate writes itself from the benchmark: the target is the market median seat price of $504 (worth $16,750 a year against the $571 on paper), the discount rebuild to 22 percent is the second ask, and the walk-away is where your alternatives stop being credible. With the mandate set before the first call, every later concession in the ledger is measured against something, which is the difference between negotiating and drifting.

Starting one from a drafted email. Every vendor-facing draft in a workroom ends with one question: is this the start of a negotiation? Pressing **Save it to the log** saves that draft as your opening move on a tracked thread and puts the track on your dashboard, marked as waiting on the vendor. When they answer, open the track and paste their reply into the composer: Vera reads it against the whole thread and your signed mandate before naming the tactic, logging any concession, and drafting your response, so she is reading the conversation rather than one email. The prompt appears only on vendor-facing drafts, because an internal note to your CFO is not the opening of a negotiation. If a track is already open with that vendor the draft files onto it rather than starting a second one. Nothing sends from the platform either way: the draft still leaves from your own inbox, and the track records what you sent.

On the screen

- **The lobby** — the clinic opens on your vendors as cards, ranked by what actually presses: notice deadlines first, then missing mandates, stalled threads, timing plays, and deliberate holds. An urgent deadline in your agreements estate gets a card even when no deal thread exists yet, and entities of one vendor family show as a single relationship with the money aggregated. One press on a card opens that vendor's case file, with a short entrance while the documents are counted in. The lobby is also the whole workspace entry: the **All deals** directory beside the queue lists every thread with search, momentum, and the closed pile one fold away, with the Boardroom Brief and **New deal** beside the search box. Every thread states who owns the deal: a Vera-run deal carries a Vera case chip that links to its case record. A star on any move card shelves that case; **Your shelf** reopens starred and recent case files, and every case has a link of its own, so the dashboard's Back to the room lands on the exact case. A card below the queue offers to read the contract or proposal you are negotiating from, through the decoder, before you enter the room.
- **The case stage** — the vendor's case file takes the whole screen: seven documents numbered in the left rail with a read meter, one paper on stage at a time, the arrow keys to page, Escape to step back to the lobby. **Your brief to Vera**, under the document list, takes one line, what are we negotiating and what to focus on, and every document she writes then stays on that topic, written as someone who has negotiated that vendor's licensing for 25 years.
- **The Executive overview (document 1)** — the case opens on one question: are you negotiating from a proposal or contract? Upload it, PDF or DOCX up to 30 MB, and Vera reads it end to end and writes a two to three page assessment: the deal at a glance with every commercial fact exactly as the paper prints it, the high level read of what the vendor is really selling, the market position, a verdict table of the terms that will cost you, and the prioritized changes to ask for. Anything the document does not state is named as absent, never assumed. The document files with your contracts, the overview is kept for the whole team, and **Write it again** rereads the stored copy. **Skip to the case file** goes straight to the move when you are not working from a paper.
- **The case file** — the six documents of the focused move, numbered in the rail. All six are written by Vera at full depth, the move brief first in line overnight; until her version of the move exists, the clinic's instant read stands in with a **Have Vera write it** offer. While she writes, the page shows real progress: the stages checking off and a small page building, usually a few minutes. A written document is kept for the whole team and rereads instantly; **Write it again** fetches a fresh read when the picture changes. For workspaces that use the clinic, it pre-writes missing or stale documents overnight for the queue's top moves, so the case file is usually already on the desk at 06:00. Each document shows when it was written, with a note when it is more than a week old and the records may have moved.
- **The case rooms** — pills in the case's top band open the working tools on the stage without leaving the case, each scoped to this vendor: **Asks** (the ask ledger, what is open, in play, and landed), **Commitments** (capture a promise with the vendor prefilled, each one reconciled against the signed contract, chase drafts flagged), **Deal log** (add an entry in place; every clinic decision files there on its own), **Meeting notes** (the call note, dated), on deals with a live thread the **Live copilot** (paste the offer mid call, get the trade math

and the counter), and **Team**, where the case's collaboration lives (below). A capture in a room and a capture on the full page are the same write. Escape steps back to the papers.

- **The Team room** — bring colleagues into the case and talk it through on the record. Pick the teammates who should work the case with you: each gets an in-app note that opens this exact case, delivered according to their own notification settings, and nothing goes to the vendor. **Case discussion** is a threaded conversation pinned to the case: @name mentions notify the person mentioned, settled threads resolve out of the way, and author edits are kept honest in the audit log. A case with open threads shows the count on its lobby card and on the Team pill. When a teammate is reading the same case, their presence shows under the document rail, so you know who else is on the paper. On cases with a live deal thread, the room also links to the deal record, where the read-only vendor portal is minted and revoked.
- **Rewrite, export, email, save** — every written document ends with a rewrite line: type what should change, shorter and firmer, add our March date, drop section two, and Vera rewrites it with the current version as the base; empty, the button fetches a fresh read. The icon strip beside the stage downloads the whole case (or the open paper via its own export) as **PDF** or an editable **Word** document in the executive brief look, drafts a **Send it on** email from whichever document is open (vendor, CFO, CIO, or colleague; vendor drafts never name the platform), and a **star** saves the case to your Saved shelf, where one press reopens it.
- **Ask Vera, in the case** — the icon strip's Ask Vera panel answers on this case's own documents: the move's facts and every written paper, with a chip per document that jumps the stage to what she cites. The case also announces itself to the top bar's Ask Vera, so asking from anywhere carries the case as context.
- **Earlier versions** — every rewrite files the version it replaces, with the reason attached: your instruction, a refocus, or a fresh write. **Earlier versions** under any written paper reads them or restores one; a restore files the replaced version too, so nothing in the chain is ever destructive.
- **The Outcome room** — on families with a closed deal or landed asks, a scorecard in the case's top band: what landed per year in the ledger's own figures, what was dropped, the vendor's promises against the signed paper, **Run the post-mortem** on every closed deal, and a note that files what you would do differently straight to the deal log.
- **The case's own research and agents** — Put it to work names the specific research: up to three of the library's documents on this vendor open straight into the reading room. **Put an agent on it** points at the vendor agent that fits the move (work a renewal, prepare for a vendor call, do the homework) and at the fourteen-role negotiation staff.
- **Small comforts** — the read meter remembers which papers you have read on this browser; document mastheads name exactly what moved since they were written; printing a case prints the open paper clean on white; and Ctrl or Cmd with the up and down arrows walks between vendor cases without leaving the room. Chase reminders arrive as in-app notes when follow-up drafts are queued.
- **The runway** — the lobby's twelve-month spine: every vendor family's notice deadline as a red dot, their fiscal close as a gold diamond, month ticks behind them; one press on a lane opens that vendor's case. Inside a case, **The clock** in the left rail counts the same two dates down in days.
- **The decide panel** — opens beside the paper with the case: approve, edit the draft, hold, or skip. Every choice files to the deal log with its reasoning and persists across reloads and teammates. Holding takes a

date, for a week, for a month, or a day you pick, and the clinic raises the move again when it arrives. A notice or escalation draft is copied for sending from your own mail.

- **Vendor sends (optional, off by default)** — an organization owner can turn on vendor sending from the clinic's top bar. With it on, approving an **escalation** move queues the email for the next 09:00 window in your organization's own timezone, addressed to the vendor contact already on the deal, with you in CC; it sits visibly cancelable until it leaves, and the sent email files onto the deal thread. Vendor replies land back on the deal thread, where Vera reads the tactic and drafts the response like any other vendor email. Only escalations ever send: a legal notice always leaves from your own mail, whatever the toggle says. If the deal's contact changes before the window, the send cancels itself and asks for a fresh approval.
- **The wire out** — a panel beside the paper with every send the clinic has queued: its status, its window or sent time, and a cancel press on anything still waiting.
- **Put it to work** — a panel of next moves in five families, one press each. The deal: the deal file, the mandate, the agreement record, sign-off routing. Price the position: a benchmark run, a second opinion, proposal comparison, a competitive scenario. Research the vendor: the vendor room, the research library, a contract decode, Ask Vera. Tell people: copy a chat-ready summary or the one-line position, draft the vendor email in the ghost writer, invite a teammate. Plan the runway: the renewal calendar, the vendor calendar, the playbook library, and filing follow-ups as tasks.
- **Mandate** — your target price, walk-away point, and approvals, set before you engage the vendor.
- **Tactics analysis** — paste a vendor email and the AI reads the tactic being used and suggests the counter.
- **Concession ledger** — every give and get on both sides, with a running value.
- **War room** — the panel beside the thread on wide screens: the mandate against reality, the give and take math, momentum, open promises, and counters queued. It closes with its **What now** dock. **Draft the next counter** composes against the thread, the mandate, and the ledger, and lands the draft in the counter queue for you to approve; nothing reaches the vendor from the dock. **Brief the CFO** writes the finance facing read, where the money stands, what each side has given, and the decision in front of them, and renders it in the dock. **Copy the position summary** puts the target, the ledger, the momentum, and the queued counters on your clipboard as plain lines for a chat thread, with the walk-away deliberately left out. **Watch the mandate triggers** puts the deal on your notification wire, and like every alert here it is off until you turn it on.

- **Commitments** — what the vendor promised, captured so it makes it into the contract.
- **Competitive Quotes** — the desk's Competitive Quotes button opens autonomous multi-vendor quote runs (below).

How to use it

1. Work the clinic's queue first: approve what should move, hold what should wait, and let every choice file to the deal log.
2. Start a negotiation from a renewal or an opportunity, and set the mandate before you engage the vendor.
3. Analyze a vendor email for tactics, and log every give and get in the ledger.

4. Run a post-mortem when the deal closes.

Related: Renewals · Competitive Quotes · Ghost writer

THE DEAL FILE (PREVIEW)

Where: Renewals & deals › Negotiations › **The Deal File** · `/deal-file` · via Ctrl/⌘ K
Who: everyone

The whole negotiation rendered as one case file, in the same voice as a Second Opinion: a dark stage, a white executive brief per page, and a Put it to work rail. Six pages, walked in order or jumped by tab, with the arrow keys paging when nothing else has focus. This preview runs on the anonymized sample renewal; live runs on your own contract start from Second Opinion.

On the screen

- **The file rail (left)** — the six pages numbered in reading order with open and read markers, search across the file, and the seats at the table. **Offer a seat** brings a teammate into the same file on the same page.
- **The paper (center)** — one executive brief per page: the benchmark **Brief** with your standing against the peer band, the **Research** dossier where you tick your leverage directly in the table, the **Drafts** pack of five letters, the line by line **Compare** table with a strike button on every row, the **Redline** with adopt or hold on each clause, and the **Mandate** ready for sign-off.
- **The what-if desk (right)** — scope, term, and ambition dials plus three what-if pills. Every page reprices as you move it: the ask marker on the brief, the numbers in the letters, the mandate table.
- **Put it to work (right, below the desk)** — page-specific moves: the five drafts, copy the one line verdict, send the markup to legal, route the mandate for sign-off.

How to use it

1. Open the file from the negotiation desk and read the brief; the gold marker is the ask the file argues for.
2. Move the what-if desk until the target is one you would defend on a call, then tick your real leverage on page 2.
3. Pick the opening email on page 3, copy it, and send it from your own inbox; nothing sends from the platform.
4. Strike the indefensible lines on page 4, decide all six redlines on page 5, and route page 6 for sign-off.

The file autosaves in your browser: leave whenever, and it reopens on the same page with every number held.

Related: Negotiations · Second opinion · Ghost writer

VENDOR TWIN

Where: Renewals & deals › Vendor Twin · `/vendor-twin` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan

You get: a verdict on every ask in your set, in a minute or two

A rehearsal room. Write out the asks you intend to put to a vendor and the twin tells you what will happen to each one before you send anything, using a model of how that vendor is negotiating right now rather than how vendors negotiate in general.

In plain words: your ask list, graded before the vendor ever sees it.

When to reach for this

- The day before the call, with the ask list drafted (the rehearsal recipe).
- Deciding which asks to spend the meeting on: a Stalls verdict is a turn saved.
- Checking what an ask reveals: the Reveals verdict exists because some questions cost you information.

On the screen

- **Vendor** — the vendors you hold agreements with, nearest expiry first, because that is usually the one you came to rehearse.
- **Where they stand today** — the evidence the rehearsal runs on, assembled from six sources: how this publisher negotiates, where you sit in their fiscal calendar, what their last results said about the pressure they are under, what they changed in their standard paper this year, any viability or consolidation signal against them, and what other buyers verifiably landed with them.
- **The coverage strip** — the six sources, with the ones that have no data for this vendor struck through. A twin running on two sources and one running on six give answers that look alike, so the strip stays on screen next to the verdicts.
- **Your asks** — one per line, in the order you mean to raise them. Up to 12 are rehearsed.
- **Verdicts** — one per ask. **Lands** means expect to get it. **Traded away** means you will get it only by giving something back, so decide now what. **Stalls** means this vendor does not move here and the turn is better spent elsewhere. **Reveals** means asking it tells the vendor something you would rather they did not know.
- **Confidence** — **Evidenced** only when outcomes other buyers actually landed sit behind the reading. **Informed** when the vendor's live position and known behavior do. **Directional** when coverage is thin, which means treat the whole rehearsal as a prior rather than a prediction.
- **Play them in this order** — the sequence to raise the asks in, and the one change that would most improve the outcome.
- **What next** — after a rehearsal, a What next panel offers the tracked negotiation on that vendor, the vendor calendar timing check, and an Ask Vera question seeded on the sharpest verdict.

How to use it

1. Pick the vendor. The evidence panel loads without running anything, so you can read where they stand before you commit to a rehearsal.
2. Paste or type your asks, one per line, in the order you plan to raise them.
3. Press **Rehearse this ask set** and give it a minute or two.
4. Deal with any **Reveals** verdict first. Those are the only ones that cost you something you cannot get back.

5. Re-order your asks to match the suggested sequence, apply the single biggest fix, and rehearse again.

Tips

- Order matters to the result. The twin reads your list as the sequence you intend, so rehearsing a reordered set is a genuinely different question.
- Thin coverage is worth fixing before rehearsing. Running the earnings radar for that vendor is usually the fastest way to move a Directional reading up to Informed.
- The twin predicts the vendor; it never writes to them. Nothing here is sent, and your internal figures never appear in anything drafted for the vendor.

Related: Negotiations · Buying blocs · Renewal playbook

GHOST WRITER

Where: Renewals & deals › Ghost writer · `/ghost-writer` · via **More** or Ctrl/⌘ K
Who: everyone; drafting needs AI on your plan

The correspondence desk for every vendor conversation you are running, in the shape of the mail client the work actually happens in. A drafting box is right for one email and wrong for the eight weeks after it, where the negotiation is a thread: what they sent, what you sent, what you nearly sent, and what to send now.

The three panes. Down the left are folders, worked out from the record rather than stored, so they can never disagree with the thread: **Needs a reply** (they spoke last), **All conversations**, **Waiting on them** (you spoke last), and **Closed**. Above them, type a vendor name to open a new conversation, or press one that already exists. The middle column lists the conversations, newest first, each showing the last thing said and a **Needs a reply** mark when the ball is with you. The right pane is the conversation itself: every message in order, yours and theirs, with **Vera's read** printed underneath each of their messages, the tactic named and the play she would run against it.

The composer. Underneath the thread. **Paste what they sent** adds their message to the conversation and reads it, against everything already said rather than on its own, so a number they are quietly repeating or a concession they are walking back is visible. Then pick an angle and Vera drafts: **Hold the position** (concede nothing, make them move next), **Ask, do not bid** (spend questions instead of money), **Bring the alternative forward**, **The budget ceiling**, **Offer two packages** (two shapes you value equally, so their choice reveals their constraint), or **Close it out** (narrow to what is open and collect the end-game asks).

The tone, on its own row under the angles, is a separate decision: the angle is what the email argues, the tone is how it sounds. **Direct** is the standard register, plain and businesslike. **Warm** is collegial, for a relationship you intend to keep; the warmth is in the wrapping only, so it concedes nothing a cold email would have held. **Formal** is measured and written to be read by people who were not in the conversation, for an email that will be forwarded, escalated, or attached to the file. **Firm** is cool and short, for a thread that has stalled: what is not working, what you need, and what happens if it does not arrive. Pressing one rewrites the draft in that register with the position unchanged, and the tone rides every rewrite and saves with the draft.

She works out where the negotiation stands from the record itself, probe, anchor, trade, or converge, and writes to that phase, and every draft is written against what is known about how that particular publisher negotiates and where its fiscal year sits. **What this draft is doing** appears above the email, for your eyes only: what it concedes, what it refuses, and the one thing to hold. **Ask Vera for a rewrite** takes plain English ("shorter and firmer", "hold the number", "ask for their final paper") and rewrites what is on screen, so edits you made by hand survive.

Drafts are kept. Save the draft keeps it as edited, so leaving the desk does not lose it. **Every draft Vera wrote** lists the whole history for that conversation, each with the angle and tone it took, when it was written, and a **sent** mark on the one that actually went, and **Put it back in the composer** restores any of them, in the voice it was written in. **I sent this** records that you sent it from your own inbox: it joins the thread as an outbound message, and every later draft reads against it.

Nothing sends from here. The platform never mails a vendor from this desk. You copy the draft or open it in your own email app, and the desk records what left. If you would rather work the thread from your inbox entirely, **Ghost writer by email** in the left rail sets up a private agent address you forward vendor mail to.

Related: Negotiations · The benchmark library · Workflow agents

CASES

Where: Renewals & deals › Cases · `/cases`
Who: everyone (mandate approval and offer decisions are owner only)

Case Mode: every agent-run deal as one complete record, with the worklist sorted needs-you first and one rule throughout: a case holds at most one pending human decision at any moment.

In plain words: Vera runs the negotiation; you approve the goal and decide every offer.

When to reach for this

- The tail vendors nobody has time to negotiate: Automated mode runs them without a strategy review each.
- A deal that matters but not enough for your own calendar: Guided mode, with the strategy approved by you.
- Checking on a running case: the worklist is sorted so what needs you is at the top.

Your position sits in the masthead: money in motion, decisions waiting on you, saved this year from certificates, and certificates booked. The operating rule and the worked example sit in the column beside the worklist rather than under it.

On the screen

- **The numbers**: money in motion across open cases, decisions waiting, savings booked from case certificates, and certificates issued. Every number clicks through to the cases behind it.

- **The worklist**: every case with its annual value, stage (Your decision, You have the wheel, Agent negotiating, Setup, Closed), who the case is waiting on and for how long ("Vendor · 3d", "You · 2h"), the predicted close date with its basis, and the next action. Filter by stage with the chips above the table. When a vendor has been silent for 3 days or more, the row grows a **Have Vera follow up** button: one click and she sends a polite chase, you on CC. The worklist also states who owns each deal: a row where you took the wheel links straight to its Negotiations thread.
- **The case record**: a status path across the top (Setup, Agent negotiating, Your decision, Closed; gold "You have the wheel" when you take over), the metadata grid, then **Your mandate**: the one panel that answers "what exactly did I authorize?" with the opening ask, target, walk-away, and leverage, a plain-language "she may / she may never" split, how she signs, who rides CC, and the exact AI disclosure sentence the vendor reads. Below it the market evidence and the **case file**: every email and every system action merged into one chronology (filter to Emails or Case log), sent messages edged in blue and vendor replies edged in gold, each with its addressing, timestamp, complete body, and Vera's read.
- **The decision panel**: when an offer stops the agent, the money leads: three tiles showing what you pay now, what the vendor offers, and what you would save or pay more, plus the vendor's offer trend across rounds so you can see the floor forming. Then the one decision: accept and close, approve and send Vera's prepared counter, or reject it and tell her what to push for so she counters the vendor again. A negotiation runs as many rounds as it takes. Accepting can carry **conditions**: state them and Vera sends the vendor a written acceptance confirmation, subject to your conditions, before the case closes, so the terms live in the thread. When an offer lands you also get an email that leads with the same numbers. Nothing goes to the vendor until you choose.
- **Take the wheel**: step in on any live case and the agent stands down entirely; work the thread yourself, then hand control back to Vera whenever you are done. Both moves are on the case log.
- **The certificate**: closing books a verified outcome certificate with the baseline, the accepted figure, the savings, any acceptance conditions, and the evidence counts, and it rolls into the Saved this year number. A closed case ends with next moves (start the next vendor, open the report, see the insights), never a dead end. It also shows on the savings ledger as a closed deal you can file with one press.
- **This deal elsewhere**: one line naming the other records this deal already lives on, so a single deal stops living in four places you have to remember. There are four possible surfaces, and it names only the ones that exist: **Negotiation** (you drive), **Case** (Vera drives), **Desk** (the standing mandate), and **Renewal record** (the paper). The record you are standing on is never listed against itself, and the same line rides the negotiation case file and the Desk supplier file.

Vera AI-run Negotiation (`/cases/new`)

The one intake for Case Mode, with two modes at the top. **Automated** is the easy path for the tail (tier-3 and tier-4 vendors): pick the vendor, give the contact email, and launch, and Vera benchmarks the discount, writes the opening email, and runs the negotiation on her own with no strategy to review. **Guided** is for the deals that matter: you review and approve the strategy, the asks, and the opening email per vendor before anything sends. Both copy you on every email and stop for your decision on any formal offer.

The first visit plays the Case Mode introduction (`/cases/intro`): an eight-scene, full-screen walkthrough of how a case runs, from intake to the certificate, with the control model spelled out (you approve the goal, she negotiates, you decide every offer). The product screenshots carry each page; click any of them (or the Enlarge control) to view it near full screen, and press Escape to return. Replay it any time from **How Case Mode works** in the Cases header. Finish with "Do not show this again" ticked and it never interrupts again.

In either mode you tick a vendor already on file (each row carries a colored vendor monogram so long lists scan fast), add one by hand, or drop new contracts (single files or whole folders; the drop zone accounts for every file: added, failed, or skipped for type), and supply the vendor's sales or renewals email by typing it or by pasting a signature or email thread. The paste box shows what it found ("Found dana@vendor.com · Dana Reyes") before you commit it. As each row lands, Vera reads the market for that vendor and states it in dollars: the published discount band, and what your target percent is worth per year against your current spend ("At your 10% target that is $32,000 a year off the table"), with a nudge when your target sits below what buyers typically get. A summary line under the guardrails totals the whole launch. The screen also shows the exact AI disclosure sentence the vendor will read, word for word, and warns you if a CC address sits on the vendor's own domain. Up to 10 cases launch at once; the counter is on the screen. Your mode, guardrails, sender name, and CC list are remembered for the next launch, and once your org has launched its first case the big explainer folds away into a one-line header.

The launch is live. After you press launch, each case plays its opening on screen: benchmarked, strategy set, first email in the vendor's inbox, then the actual email that went out, readable in full, with the sender name the vendor sees. The headline states the target on the table in dollars across everything you launched. Failed rows list their reason with a **Fix these and relaunch** shortcut. In Guided mode you then press **Prepare the strategy with Vera**: she benchmarks where you sit versus the market (the discount read), drafts a high-level negotiation strategy that uses real competitive tension (naming credible alternatives so the vendor has to stay price competitive), sets the opening ask, the internal target, and the walk-away, lists the prioritized asks as tick boxes (defaults are pre-ticked; untick what you do not want), and writes the opening email. The plan also states how she will play the rounds: the concession plan (steps that shrink round over round so the vendor reads the floor), what she finds out from the vendor before conceding anything, and, when talks would fail, the credible alternative as a costed path (If talks fail). Review and edit any of it, including the email itself, then approve; the approved plan becomes the mandate every one of her emails is checked against. The owner's authorization launches every prepared selection at once. How the email works is spelled out on the screen: Vera sends through VendorBenchmark's secure relay so vendor replies route back to her, the vendor sees your sender name (not VendorBenchmark), and you are copied at your real inbox on every message.

How Vera negotiates. Before every strategy and every counter she reads the VendorBenchmark research for the moment at hand (the vendor's Negotiation Toolkit, the Negotiation Craft doctrine, Renewal Outlooks, audit defense manuals, exit runbooks), so a quarter-end squeeze or an audit-as-leverage play gets the researched counter. She computes the vendor's offer trend across rounds (shrinking moves mean the floor is near), recognizes the rep from your prior threads and the plays they favor, knows the vendor's fiscal year and quarter closes and deliberately slows her pace into them, and never gives without getting: every concession is paired with an ask, on the shrinking curve you approved. Before any email goes out unattended, a seller-side check reads the draft the way the vendor's account team would; anything exploitable is parked in your approval queue with the findings attached. And a negotiation that reaches ten vendor rounds without landing stops and asks you to converge or walk, with a prepared next move waiting.

Editing what Vera sends. You can edit any email before it goes out: the opening email in the prepare step, and every counter Vera drafts (open the case, use Edit on the prepared counter, then approve to send exactly what you see). The AI-agent disclosure is always kept on the opening email even if you remove it.

If the first email has not gone out. A case only reads as "Agent negotiating" once the opening email has actually been sent. If a case shows "Setup · not started" with the next action "First email not sent yet", open it and use **Send the first email now**; if delivery fails you will see exactly why.

Insights (`/cases/insights`)

The management view leads with the number a leader repeats: verified savings a year across closed cases, on a navy banner. Below it: money in motion, decisions waiting, median cycle time, and the agent's email activity, with pipeline-by-stage and savings-by-month charts you can click into, and a table of every case with its predicted close date and basis. **Board deck** exports the whole portfolio as a three-slide PowerPoint in the executive-brief theme (the numbers, then every case), ready to drop into a board pack unedited; the portfolio report is one click away too.

Email that keeps you current. A weekly Your cases digest lands Monday morning when you have open cases: money in motion, decisions waiting, savings booked, and the worklist, decisions first. And the first time a vendor replies on a case you get a short note with the vendor's message, Vera's read, and her next step, so you never wonder whether the machine is working.

Reports
Every case has a report (`/cases/[id]/report`): the mandate, the full correspondence, Vera's reads, and the outcome, as an executive brief you can export and forward. The portfolio report (`/cases/report`) is everything Vera ran across every case in one brief.

How to use it

1. Open Cases and work top down; the list already puts what needs you first.
2. On a setup case, review the mandate table and approve it in place; that approval is the agent's written authorization.
3. When a decision is waiting, read the correspondence log first, then send the counter or accept and close.
4. To steer a live thread yourself, take the wheel on the case, work it, then hand it back to Vera.
5. Use Full deal thread on any case to drop into the complete negotiation workspace with every tool.

Tips

- Cases are the same deals as Negotiations, read as records; anything you do in either place shows in both.
- The sender name and CC list you choose at launch apply to every email Vera sends on that case, not just the first.

Related: Negotiations · Competitive Quotes · Renewals

COMPETITIVE QUOTES

Where: Renewals & deals › Negotiations › **Competitive Quotes** · `/negotiation/fleet`
Who: everyone (launching is owner only) · requires an AI plan

Autonomous multi-vendor quote runs: hand Vera a set of vendors, the tier-4 tail is the sweet spot, and she runs every quote conversation herself, so the vendors effectively compete for your business in parallel.

On the screen

- **The launcher** — up to 10 vendor rows, each with the vendor, the rep's contact email, a one-line scope ("25 seats, renews October"), and the current annual spend if you know it. Guardrails are set once for the run: the target reduction, the walk-away ceiling, and optional notes for the agent.
- **The authorization** — the launch checkbox is the written mandate approval the agent cites to vendors. Owners only, and the approved mandate is recorded per vendor in the audit log.
- **What Vera does** — introduces herself to each rep as your declared AI agent, requests a formal written quote, negotiates the back-and-forth inside the mandate, and politely chases silence. You are copied on every email she sends.
- **The offer gate** — the moment a vendor puts a formal offer on the table, the agent stops, notifies you, and waits. No offer is ever accepted without your explicit approval on the deal thread. A What next panel rides the moment: review the offer, escalate the run to a case, or ask Vera about the spread across the board.
- **The board** — every run with its state: quote requested, Vera negotiating, proposal awaiting your decision, paused, or closed, with sent and received counts and a link into each thread.

How to use it

1. Open Competitive Quotes from the negotiation desk and add your tail vendors with their rep emails.
2. Set the target reduction and the walk-away ceiling, tick the authorization, and launch.
3. Watch the CC'd emails arrive, and act only when the board says a proposal needs your decision.

Tips

- Give each row a real scope line and the current spend; the agent's opening ask is built word for word from what you enter, never generated.

Related: Negotiations · Approvals

BUYING BLOCS

Where: Renewals & deals › Buying blocs · `/blocs` · via **More** or Ctrl/⌘ K
Who: everyone · the app joins on the Enterprise plan (workspaces from before July 29, 2026 and trials keep access; see Plans, allowances & limits)

Buyers with the same vendor and the same renewal window, negotiating on one term sheet. You stay anonymous until mutual opt-in.

Your position sits in the masthead: invitations waiting on your answer, blocs you have joined, how many buyers are in them, and the combined annual value those buyers carry into the same window. That last figure is the whole argument for a bloc, so it leads.

On the screen

- **Matchmaking strip** — where your renewal calendar lines up with other buyers of the same vendor.
- **Invitations** — a bloc you have been invited to: you see the shape (buyers, combined ACV, window, term sheet version), never the members.
- **Your blocs** — blocs you have joined, with a room to coordinate the shared negotiation.

How to use it

1. Review the matchmaking strip and your invitations.
2. Open an invitation to read the bloc's shape and thesis.
3. Join a bloc to enter its room; identities unlock only on mutual opt-in after you join.

Tips

- Blocs form around vendor families where several buyers renew in the same quarter, built on a shared Fair Renewal Rider term sheet.

Related: Fair Renewal Rider · Renewals

CLAUSE LIBRARY

Where: Contracts & documents › Clause library · `/clause-library` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan

Your organization's negotiation positions per clause, worked like a desk: the playbook rail on the left, the selected position as a one-page dossier in the center, and the live contract's verdict on the right. These positions feed the AI engines, so scans and drafts argue for your standard instead of a generic one.

This is the same playbook you see under Settings › Clause positions, viewed from the negotiating side, and each page now states its role and links the other: this workspace compares contracts and drafts redlines, while Settings › Clause positions holds the must-have template Coverage sweeps. Every workspace starts with the VendorBenchmark starter library already in place, so the page has positions to grade against from day one; edit them to match your own standard.

In plain words: what your organization signs, written down once, argued for everywhere.

When to reach for this

- Setting the standard, once, so every later review argues for your language instead of generic language.
- Before signature: compare the draft against your positions and see the deviations named.
- When the same clause fight keeps repeating: a recorded fallback ends the relitigating.

A worked example. Set an uplift cap position (preferred: 4 percent indexed; fallback: 5 percent; walk-away: no cap, no signature) and compare the sample Contoso Analytics quote against it: the quote names no cap at all, so the comparison flags the silence and drafts your pre-approved language as the redline. One position, recorded once, catches the same gap in every future contract.

On the screen

- **The playbook** — the left rail lists every position grouped by area; its header states how many of the eight standard areas carry one, and names any area with no position, because a comparison cannot grade a clause you have taken no position on. Once a contract is compared, each clause carries its grade as a colored dot: green meets your standard, red weaker, gold missing, gray unclear.
- **The dossier** — the selected position as a one-page brief: your position, your preferred language, the acceptable fallback, the walk-away, and why the position holds. Every section carries a copy action, so the wording is one press from a redline.
- **Against the live deal** — pick a contract on the right and run the comparison. The rail grades the selected clause, quotes the language the AI found in the document, and drafts the redline from your preferred language, with the deal scorecard beneath.
- **What next** — the sweep closes with its next moves: copy every drafted redline in one block for the negotiation thread, start the negotiation, open Vera Deal-Room, or ask Vera which redline to push hardest for.

How to use it

1. Select a clause in the playbook rail and read its dossier. Add or edit positions from the rail; each carries a name, a category, your position in plain English, and optionally the precedent wording, the fallback, and the walk-away.
2. If the rail header names an area with no position, add one before comparing.
3. Pick a contract on the right and run the comparison, then click through the graded dots to read each clause's verdict and its redline.
4. See position coverage across all contracts on Coverage.

Related: Coverage · Clause positions (Settings)

COVERAGE

Where: Contracts & documents › Coverage · `/contracts/coverage` · via **More** or Ctrl/⌘ K
Who: everyone

The board says when it is stale: contracts added since the last sweep, or a must-have template edited since, show as a badge on the stage with the run control beside it, and an outdated report tab offers its regenerate button. Nothing reruns without your press.

Your clause library's must-have template, swept across every contract: which agreements carry your required protections, and which do not.

In plain words: your non-negotiables, checked against every contract at once.

When to reach for this

- After setting the clause standard: the sweep is what turns positions into findings.
- Before a renewal: a red Missing cell on a renewing agreement is a negotiation item with a deadline.
- When leadership asks "are we protected?": the coverage report is the written answer.

Before the first sweep, the stage does not draw an empty grid. It states what it is about to do, lists the protections your template requires below it, and offers one action: **Run the first sweep**. Nothing is graded until something has been read.

On the screen

- **The verdict stage** — the page opens on a navy stage, the same language as Second Opinion: one sentence on how many swept contracts are missing a protection you require, four figures counting up beneath it (fully covered, carrying a gap, gaps in total, and posture, the share of checked cells carried), and the page's actions: run the sweep, generate the report, edit the template. The template edit links say where they lead: Settings › Clause positions, where the must-have template lives.
- **The sweep theatre** — **Run coverage check** raises a fullscreen stage where the sweep plays live: position lines tick in as the portfolio is read, a gold arc tracks the real batches, and when the verdicts land the completion line types itself. Esc or the corner button lowers the stage at any time; the sweep keeps running behind it and the stage offers **Watch the sweep live** to return.
- **The board tab** — the protections rail on the left lists every must-have weakest first with a coverage bar; selecting one turns the ledger into that protection's story: one row per agreement with the finding the sweep recorded and its verdict pill, worst first. Expand a row for the evidence in full: the finding, the exact clause quoted from the document, and what your template requires. **All protections** shows every agreement with a verdict square per position (✓ present, ~ deviates, ✕ missing, ? unclear, · not checked) and a gap count; click any square to follow that protection. Verdict filter chips, a search box, and a worst first or vendor A to Z sort work the board.
- **The By protection tab** — the ranked read: the coverage ring, then every protection weakest first with its bar and counts. Opening a row shows the template language it grades against and the agreements missing it, each one press from its contract.
- **The report tab** — the white-paper read of where your protections are strong and where they are missing. If no report exists yet, this tab writes one; regenerating from the stage lands here.
- **What next** — when the sweep finds missing must-have protections, a What next panel targets the worst-exposed vendor: open its renewal, start a case on it, ask Vera about the exposure, or draft the redline from your clause library.

How to use it

1. Run the sweep and watch it play, or lower the stage and keep working. Only contracts with indexed text can be read; the scope line says how many are waiting.

2. Read the stage first, then the protections rail on the board: the top entry is the protection your portfolio is thinnest on.
3. Select that protection and work the ledger from the top, expanding rows to read the clause each verdict rests on. Press the **Missing** chip to isolate the agreements with no responsive clause at all.
4. Open the contract from any row to fix it, or generate the report for the narrative summary.

Tips

- The must-have template lives under Settings › Clause positions, so keep those positions current first; the Clause library workspace is where you compare contracts against them and draft redlines.

Related: Clause library

FAIR RENEWAL RIDER

Where: Renewals & deals › Fair Renewal Rider · `/rider`
Who: everyone

The standard addendum buyers attach to every software renewal: an indexed uplift cap, benchmark repricing, audit conduct rules, no retroactive true-ups, co-termination, and change of control protection. Versioned, public, and free to send.

In plain words: the standard protections, as one attachment, so you never negotiate them clause by clause.

When to reach for this

- Every renewal, as a matter of habit: the rider rides along without spending negotiating capital.
- A contract with a structural gap: the sample Contoso quote carries no uplift cap, and the rider's indexed cap is the ready-made fix.
- Short on time: at two weeks out, one attachment carrying six protections beats six separate asks (the rescue play).

On the screen

- **Adoption counter** — how many real renewals carry the rider through VendorBenchmark. A standard works because everyone asks for the same thing.
- **Open the workroom** — the band above the adoption counter opens the rider as a six document workroom: the cover sheet, the operative language of all six clauses, the rationale brief in three clause families, a risk register pricing each missing protection in the rider's own words, the vendor facing send sheet, and a What now sheet. ⚡ **What now** drafts the cover note that transmits the rider, copies the full text, downloads either edition, and opens your renewals, the negotiation dossier, or the clause library. The standard edition is what goes to the vendor; the negotiator's edition, with the fallback ladders and the push backs already countered, stays on your side. ✦ **Ask Vera** answers from the rider's published text, including what you will hear on each clause and the counter to it.

- **Generator** — pick the contract that is renewing and the rider comes back as a ready-to-send Word document with the parties and renewal date filled in, or, when the vendor's paper on file is a DOCX, as tracked-change insertions inside their own document.
- **Reading room** — the rider itself, every clause annotated with why it is fair, the fallback tiers, and the vendor objection countered.
- **Clause changelog** — what changed in each version of the rider.
- **What next** — after you download the rider, a What next panel points to the renewal it attaches to and copies the send sheet cover note for your own mail.

How to use it

1. Read the rider in the reading room to understand each clause.
2. Pick the renewing contract in the generator and choose the edition.
3. Download the Word document or the tracked-change redline and attach it to your renewal.

Related: Renewals · Buying blocs

VAULT

Where: Contracts & documents › Vault · `/vault` · via **More** or Ctrl/⌘ K
Who: everyone

A room is an access layer, not a second store: every document in one also sits in Your Archive, and renewal terms stay on the agreement ledger.

One workspace over your governed document store, in the same voice as Second Opinion: a navy stage on every page. Projects group documents by matter, for example an acquisition review, and each project is org-wide or limited to named people; documents in a restricted project stay hidden from everyone else.

The front door opens on the stage: your estate in a sentence, documents, rooms, the share read and ready for AI work, encrypted megabytes on file, with a chip for every room. Each chip is a door into that room's own page; a lock marks a restricted room. Three estate-wide tabs sit under the stage: every document across the vault, the coverage matrix across rooms, and every room with its access level.

A room's own page carries its stage: the room's name, its figures, and the drop target. Drag documents or whole folders onto the stage and each file is validated by signature, malware scanned, deduplicated, and indexed automatically; the add existing and cloud import moves (Box, Google Drive, SharePoint) sit beside it as stage buttons. The same three tabs follow below, scoped to the room.

- **Documents** — the room's documents with type, size, state, and per row remove.
- **Coverage** — a vendor by document role matrix built from what the read stage has named: master agreements, order forms, amendments, statements of work. A vendor holding supporting paper with no governing agreement on file is flagged as a gap; papers not yet read for a role sit in the Untyped column rather than being guessed at. The read below the matrix turns the flags into sentences, starting with the missing agreements, which in diligence are usually the first document request.

- **Access & policy** — who can see the room, straight from the database: on a restricted room, the creator and each named member with what their organization role lets them do; on an org-wide room, everyone by role. Under it, the handling policy in plain rows: intake checks, encryption, 60 second signed URL downloads, the audit trail, and the note that Redress analysts read rooms to produce your benchmarks, on the same audit trail.

How to use it

1. Create a project and pick org-wide or restricted access.
2. Open its chip and drop documents on the room's stage, in bulk if needed; each file moves from received to indexed on its own.
3. Read the Coverage tab and request whatever the gap flags say is missing.
4. Run diligence across a project when you want the deal-grade review.

Related: Diligence

DILIGENCE

Where: Contracts & documents › Diligence · `/diligence` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan

End-to-end document review for a deal: every contract in scope is read for change of control, assignment, exclusivity, termination, liability, and data terms, then assembled into a risk register and an analyst report.

On the screen

- **Scope picker** — run the review across all contracts or limit it to one vault project.
- **Risk register** — every finding with its severity and the clause it rests on.
- **Analyst report** — the synthesized read: what needs consent or renegotiation before the deal, and what is clean.
- **Open the workroom** — once a review exists, a navy band above the report binds it into a numbered workroom: the cover sheet, the scope sheet, findings by area, a severity ranked risk atlas, the open questions sheet, and a What now paper. **What now** drafts the internal deal owner briefing and a neutral counterparty information request from this review's own findings and open questions, copies the findings summary or the one line verdict, opens a Vera Deal-Room, and links back to the flagged records. ✦ **Ask Vera** is grounded on this review, so a question about what is still unresolved is answered from the review's own open items.
- **What next** — the review ends with a What next panel built around the top risk: open a case on that counterparty, draft the sends in the workroom, ask Vera about the worst flag, or copy the one line verdict.

How to use it

1. Scope the documents into a vault project first, or run across all contracts.

2. Run a diligence review.
3. Work the risk register, starting with the highest-severity findings.

Related: Vault

REQUESTS

Where: Contracts & documents › Requests · `/requests` · via **More** or Ctrl/⌘ K
Who: everyone

One front door for new spend. Every request is triaged by the AI against what you already own, then routed into sourcing, renewals, or approvals.

In plain words: everyone asks here, the AI checks what you already own, and nothing gets bought twice.

When to reach for this

- Standing up the intake habit: one door, one triage, no side-channel purchases.
- A request arrived: the verdict (net-new, renewal, duplicate, over threshold) decides the route in seconds.
- Someone is about to buy what the estate already owns: the duplicate verdict is this page paying for itself.

On the screen

- **The lobby** — the masthead states what is waiting, how much new spend is under review, how fast the triage desk answers, and how much has been routed.
- **Two columns** — requests needing a decision on the left, everything already routed or dismissed on the right.
- **Triage verdicts** — the AI's read on each request: net-new, a renewal of something you own, a duplicate of existing capability, or over the approval threshold.
- **Routing** — send a triaged request onward as a sourcing project, a renewal, or an approval.
- **Draft the reply** — every triaged request carries a Draft the reply action, composed from the verdict's own facts with the duplicate case leading; copy it or open it in your own mail, and nothing sends automatically.

How to use it

1. Raise a request for the new spend.
2. Read the triage verdict.
3. Route it onward to sourcing, a renewal, or an approval.

Tips

- Owners can set up an intake address under Settings › Integrations so teammates email requests straight in.

Related: Sourcing

SOURCING

Where: Contracts & documents › Sourcing · `/sourcing` · via **More** or Ctrl/⌘ K
Who: everyone

Sourcing opens as a workspace: what the pipeline is carrying, which deals need you today, how far each stage has carried its value, and the project list. The deal desk is one click away and prices every move still open on a live deal.

In plain words: the buying pipeline, from requirements to award, with every open move priced.

When to reach for this

- A net-new purchase big enough to deserve a process (intake to award, the bake-off).
- Mid-process, when "what do I do next?" has a computed answer: the What Next desk.
- The pricing endgame: the deal desk's headroom number is what you walk into the call with.

On the screen

- **Pipeline position** — the masthead states projects in flight, the budget in play, awarded value, and anything stalled past 30 days, with a stage strip under it counting the projects and the money sitting in each stage.
- **Needs you now** — projects at a decision point or stalled past 30 days, each leading with its exact next move, ranked by leverage rather than by date.
- **Stage by stage** — intake to award, with how far the pipeline has carried its value: reached counts every project at or past a stage, here counts the ones sitting in it now, and the money in each is printed alongside.
- **Board and list** — five stage columns or a dense list ranked by leverage, with search, stage chips, and a stalled-only filter.
- **Notes wall** — the desk's shared memory, saved to the workspace, so it reads the same for every teammate on every device: why a supplier was shortlisted, what a vendor said on a call, what changed this week. Pinned entries lead.
- **Project page** — requirements, suppliers, and bids side by side, through to the award decision, led by its own What Next desk (below).

The deal desk

Where: Sourcing › Open the deal desk · `/sourcing/desk`
Who: everyone

One live deal at a time, with every move still open to you priced from the quotes you already hold.

- **What if** — the moves, cheapest first, each with the annual cost, the difference against the quote on the table, the difference against the approved budget, and one line on what it takes: award as quoted, hold the supplier to the top of its own discount band, make it match a competing quote, switch to the cheaper response, or bring the deal back inside the approved budget. The move you are currently on track for is marked.
- **Where each number comes from** — every move carries its basis, and nothing is estimated. A **quoted** figure is the supplier's own annual total. A **price book band** figure is its list price less the discount buyers of your size actually land. A **competing quote** is another response in the same room. A move that cannot be traced to one of those is not shown at all.
- **Headroom** — the gap between the quote on the table and the best move still open, stated in money and as a share of the quote. It is the number to walk into the call with.
- **The ask, drafted** — each priced move drafts a vendor safe ask ready to copy: the target figure and its stated basis only, never your budget, your scores, or another bidder's name.
- **The responses** — what each supplier quoted, scored against your own weighted requirements, with list price, the discount that implies, and the expected band beside it. The leading response is highlighted.
- **Next steps** — the run in order, with everything already true of the deal marked done: the brief, the weighted requirements, the RFP, the shortlist, the bids, the scoring, and the award. A shortlist only counts as done from two suppliers, because one quote is a price and two is a negotiation.

Deals sit on a switcher across the top of the desk, so you can move between them without going back to the list.

The project's What Next desk
Every project record opens with a What Next desk under the stage path: the computed next move as the headline, then the premade moves of the moment, so no step starts from a blank page.

- **Send the word** — Vera drafts the email each stage actually needs, from the project's live numbers: the stakeholder kickoff, the requirements review request, the supplier invitation, the Q&A broadcast (one answer, every bidder, so the process stays defensible), the chase for quiet suppliers with their names already in it, the status update, the score sanity check, the recommendation memo to the budget owner, and after award the winner note, the decline notes, the legal handover, and the CFO outcome note. Drafts are yours to edit and copy; nothing sends itself.
- **Run the AI** — draft or refresh the RFP, find the best suppliers in the field, score every waiting bid against the weighted requirements in one click (each score carries a quote from the bid document that justifies it), and write the evaluation memo. Buttons show what has already run, for example "All 3 bids scored".
- **Go deeper** — the deal desk for this project, Ask Vera, the RFP and evaluation as Word, and the contract decoder once a winner is picked.
- **The supplier field report** — discovery now also renders on the Suppliers tab as a forwardable executive brief: the landscape summary, market context, fit bars across the field, and one honest section per candidate with strengths, gaps, price posture, and why the score, with the Send it on drafter beneath it.

- **Open the workroom** — once a project has a bid, a navy band under the status path opens it as a workroom: the requirements sheet, each bid's stored figures as lettered exhibits, then once two responses are scored a verdict brief that reads the whole field in one paper (the standing, the cross bid price check against bands and the price book, and a what to do first list that prices the band and list rate asks where the stored figures can), the weighted scoring atlas as one toned table, an award brief with a verify before award list, and a What now sheet. ⚡ **What now** drafts the best and final ask to each finalist (priced from the stored figures where the record can price it, in position language that never names the band or another bidder), the award notification, the courteous note to the other bidders (it carries no scores), and the stakeholder update, all from the stored record, then copies the project one liner or the bid table as text and scores the responses on the record. ✦ **Ask Vera** answers from this project's own requirements, bids, and scores.

Send invitations, and bids that file themselves

The RFP tab carries the send flow, so the document you drafted actually travels. Pick the shortlisted suppliers, confirm each contact email, set the response deadline, preview the email, and send. You stay in charge of every send; nothing mails itself.

- Each invitation renders in the document standard with your organization named as sender, summarizes the RFP, states the deadline, and carries **its own private reply address**.
- When the supplier replies with the proposal attached (PDF or DOCX, 8 MB per attachment), the response files straight onto the project: the file is signature checked, its text extracted, and it is auto scored against your weighted requirements exactly as a manual upload would be. The invite flips to Responded, the notes wall gets the update, and the project owner is notified. A reply with no readable attachment is flagged on the wall for a human look instead.
- The **RFQ watchdog** (an agent in the Agent Center) minds the window daily: each quiet bidder gets exactly one reminder, three days before the deadline, and never without a deadline on record. When every invite has answered, or the window closes, it notes the tally on the project wall and nudges you toward evaluation.
- The supplier list on the tab shows each invite's state: Sent, Reminded, Responded, with dates.

Award scenarios

The Evaluation tab ends the guesswork about weights and split awards with a panel of pure math over your stored scores and extracted line items. No AI call is spent on any of it.

- **What-if weights** — drag a category's slider and the scored matrix re-ranks live, with each supplier's movement against the RFP ranking marked. One click resets to the RFP weights; the weights on record never change from here.
- **Where the winner flips** — one line per category stating the exact weight at which rank 1 changes hands, for example "Northline overtakes Meridian if Integration weight rises above 28 percent". No line means the ranking is weight robust.
- **Single award vs split** — where bids carry priced line items, the panel prices every single award against the best two way split by line item group: each contested group goes to the cheaper bidder, groups only one bidder priced stay with it, and every card shows total cost, the delta against the best single award, the combined weighted score (each requirement served by the stronger member), and mandatory coverage.

The math states its own limits: up to 15 line items per bid, the first 8 bids, and no modeling of bundle discounts a vendor might withdraw on a partial award.

- **Adopt a scenario** — saves the chosen scenario on the project and posts the one line summary to the notes wall. For a single award scenario the regular Award button sits right beside it; the award itself stays a deliberate, separate step.

How to use it

1. Walk in and read the marquee; it names the project that needs you first. Click that card.
2. Make the move the focus console states, or run the engine it offers.
3. Press **N** to leave what you learned on the wall, and **U** to log what you did.
4. Before you step away, open the Shift log and compose the **End of shift handoff** for your team.
5. Start a new project from the backstage toolbar, or turn a purchase request into one.

Related: Requests

Next: Chapter 02 · Benchmarks & research

02 · Benchmarks & research

This chapter covers the benchmarking half of the platform: instant self-serve price checks for hundreds of vendors, your whole portfolio scored against the market in one view, the analyst-delivered reports you get when you upload a proposal, the vendor grade book, and the research library. Reach for it whenever you need to know whether a price is fair, where the money is across your stack, or what a vendor is doing in the market.

INSTANT BENCHMARKS

Where: Benchmarks & market › Instant benchmarks · `/benchmarking`
Who: everyone

Self-serve price checks for more than 1,347 vendors and product lines, no upload needed. Answer a few questions about your deal and get a market comparison in about two minutes.

In plain words: type a vendor's name, answer a few questions, and find out whether your price is good.

When to reach for this

- "Am I overpaying for this vendor?" before a renewal call.
- "What discount should I ask for?" the day a quote lands.
- "Which of our vendors deserves attention first?" when triaging the estate (or jump straight to the Portfolio benchmark for the whole answer at once).

Opened from your home screen (app mode), the page walks you straight into the room: the Ask the market toolbar runs edge to edge under the navy band with the search bar ready, and the page header and cross vendor suite row sit below the working surface. Opened from the dashboard, the page keeps the composition described here, header first.

On the screen

- **Ask the market** — the search bar at the top, with the quick asks and the corpus counts beneath it. Vendors suggest instantly as you type, and the preview card beside the list shows the highlighted vendor's benchmark depth, reference deals, your own agreements, and the next renewal. You can also type a plain language question ("am I overpaying for Salesforce?") and the Ask Vera row answers it with grounded, cited figures. Press **/** from anywhere on the page to jump into the bar, then **↑↓** to move, **Tab** to complete the vendor, ↵ to run, and **Esc** to clear.
- **Drop your quote** — the dashed lane under the search field. Drag a vendor quote onto it, or anywhere on the page, and the platform reads the paper instead of asking you to type the deal in. See Drop a quote below.

- **Quick ask pills** — one press starters for the three questions buyers ask most: am I overpaying, what a strong deal looks like, and what to push for.
- **Tool cards** — one tool per vendor or topic: Microsoft, Oracle, Salesforce, SAP, AWS, and more, grouped by category, with your own vendors floated first.
- **Peer comparison** — results show where your deal sits against a cohort of similar companies, with percentile bars.
- **Saved scenarios** — name and save a position and it is kept: you can revisit it, compare it against your other saved positions on the same vendor, and it joins the Portfolio benchmark roll-up. The five flagship tools (Microsoft EA, Oracle ULA, Oracle Technology, Sales and Service Cloud, Oracle Java SE) save positions the same way as the rest of the library.

How to use it

1. Press / (or click the bar) and type a vendor name. Arrow keys move through the matches, Tab completes the name, Enter runs. A vendor with one benchmark opens its tool directly; a vendor with several lists them below the bar.
2. To score your whole stack at once instead of one vendor, open the Portfolio benchmark.
3. For a deeper expert review, upload the proposal for analyst benchmarking.

Tips

- The full set of vendor tools is the benchmark library below. The hub groups them by category so you can find yours fast, or press Ctrl/⌘ K and type the vendor name.
- Typed questions from the bar land in Ask Vera with the question prefilled, so the search and the assistant are the same instrument.

A worked example. The sample portfolio holds 320 Salesforce Enterprise seats. Type Salesforce, pick the Sales Cloud line, enter the seats and the net you pay, and the page answers the three questions in order: the discount your figures imply, the range comparable buyers land at your size, and where that places the deal, with targets to ask for. The whole pass takes about two minutes, and saving it as a scenario means next quarter's check starts prefilled.

If something looks wrong

- The page says there is not enough data to rank you. That is honesty, not failure: too few comparable deals sit in your band. Widen the region or check the deal size; Troubleshooting walks the cases.
- You cannot find the vendor. Search the full benchmark library first; the largest publishers are split by product line, so the tool you want may be under the product's name rather than the company's.

Related: Drop a quote · The benchmark library · Portfolio benchmark

DROP A QUOTE

Where: Benchmarks & market › Instant benchmarks · `/benchmarking`
Who: everyone on a paid plan

The fastest route to a benchmark: hand over the vendor's own paper instead of typing the deal in. Drag a quote, order form, or renewal proposal anywhere onto the workspace and you get back a six page negotiation brief in a deal desk room, usually in about three minutes.

In plain words: drop the vendor's quote and get back the negotiation brief.

When to reach for this

- The renewal quote just arrived and you want your position before you reply.
- You are walking into a pricing call this week with the order form on file.
- Someone asks "is this quote fine?" and you would rather answer with a brief than an opinion.

On the screen

- **The drop lane** — the dashed panel under the search field. Click it to browse, or drag a file onto it.
- **Drop anywhere** — the whole page is the target. The moment a file enters the browser window the workspace dims and a gold outline appears, so you never have to aim.
- **Vendor field** — optional. Naming the vendor helps us match the right cohort when the paper is ambiguous, but we read the vendor off the document either way.
- **The progress ring** — a real percentage, not a spinner. Uploading your bytes drives the first tenth, and the rest moves as each step of the analysis completes. The eight beats are listed underneath and light up as they happen: uploading, reading the paper end to end, pulling out the commercial terms, matching comparable closed deals, ranking this deal inside its cohort, writing the position and the levers, checking every figure against the data, and setting the deal desk room.

How to use it

1. Drag your quote or order form onto the benchmarking page, or click the lane and browse for it. PDF or DOCX, up to 30 MB.
2. Optionally type the vendor name in the field below the lane.
3. Watch the percentage. You can leave the page while it runs: the brief keeps building and stays in your workspace.
4. When it finishes you land in the deal desk room with the finished brief.

Tips

- Upload the document that carries the pricing. A master agreement with no figures gives us nothing to benchmark, and the run will tell you so rather than guessing.
- Your discount is measured against a real list price: the one on your paper, our published price book, or one we find and cite from the vendor. Where no list price can be established from any source, the benchmark says the discount cannot be verified, names exactly what to ask the vendor for, and does not rank the deal rather than showing a made up number. The same deal returns the same discount to everyone.
- We benchmark against real closed deals for that vendor. If the library does not hold a cohort for the vendor on your paper, the run says so plainly instead of ranking you against a near miss.
- Your file is read once and deleted. Nothing about the document is retained.

If something looks wrong

- The upload is refused: the lane takes PDF and DOCX up to 30 MB, verified by contents, and Troubleshooting covers the rejection reasons.
- The run says it cannot benchmark the paper: a master agreement with no figures gives it nothing to price. Drop the document that carries the money, usually the order form or the quote.

Related: The deal desk room · Instant benchmarks

THE DEAL DESK ROOM

Where: opens automatically after a dropped quote finishes · `/benchmarking/room/…`
Who: everyone on a paid plan

Where a finished benchmark lands, and the surface you take into the pricing round.

In plain words: the finished brief, plus the doors that turn it into a negotiation.

When to reach for this

- Reading the result of a dropped quote, which lands here on its own.
- Coming back before the vendor call: the room stays in your workspace, so the brief is where you left it.
- Handing the position to whoever negotiates: the room is the thing to share, not a screenshot of it.

On the screen

- **The command band** — the dark panel at the top carrying the vendor, the deal size and term, and the rank as one word: Poor, Below market, At market, Good, or Top of the market.
- **The three answers** — stated in order before the brief starts. What discount you are being offered. What comparable buyers get, as a range with its midpoint. Where this deal ranks, with the gap in points and in money over the term.
- **The brief** — six pages in the delivered executive brief format: executive summary, the deal snapshot against your vendor's proposal, a twelve row peer table with your own deal placed inside it by rank, cohort statistics, a sensitivity table, the market evidence behind the bands, the discount ladder, the size against rate chart, the rate alignment bridge, four levers with the order to play them, four talking points with lines you can say out loud, and the verdict.
- **Next moves** — three links out of the room: model the deal in the vendor's own tool, open a negotiation case, or ask Vera how to open the conversation.

How to use it

1. Read the three answers at the top first. That is the finding; everything below argues it.
2. Work down the brief to the levers and the sequencing note, which tells you which lever opens and which closes last.
3. Take the talking points into the call. Two of the four carry a script line in quotes.
4. Use **Open a negotiation case** to carry the position into tracked work with a timeline and owners.

Tips

- Every peer row is a real closed transaction, shown by generic sub sector. No counterparty is ever named, and neither are you.
- If the brief says the rate case is closed, that is a genuine result: your rate is already at or above the market and the remaining value is in the structural terms. Play the last three levers and protect the rate.

Related: Drop a quote · Delivered reports

PORTFOLIO BENCHMARK

Where: Benchmarks & market › Portfolio benchmark · `/benchmarking/portfolio`
Who: everyone

Your whole software stack compared with the market in one view: which vendors look expensive, which look fair, and where the money is.

In plain words: every vendor you hold, ranked by how far off the market it sits.

When to reach for this

- "Where should this quarter's negotiation effort go?" answered by money rather than instinct.
- Before budget season, when the savings prize number frames the plan.
- After loading the estate, as the first read of what the paper says in aggregate.

A worked example. The sample portfolio carries four vendors worth $331,200 a year together. The roll-up shows one clearly above market, the Northwind CRM subscription with $28,650 a year in reach across its seat price and discount gaps, and one clearly market-leading, the Fabrikam cloud commit, where the work is defending the position rather than repairing it. That contrast is the point of the page: it separates the deals worth a fight from the deals worth a thank-you.

Your position sits in the masthead: positions placed, how many beat the market median, how many sit above it, and the total annual headroom between what you pay and the median.

On the screen

- **Standing per vendor** — each vendor's position against market pricing. The saved roll-up is built from the **scenarios you save inside the vendor tools**, so a vendor joins this view once you have saved a position for it. Uploading an agreement does not by itself place a vendor here.
- **Savings prize** — the estimated annual amount on the table if above-market deals moved to the median.
- **The seeded start** — with agreements on file but no saved scenario yet, the table is not empty. It opens seeded from your own records: a row for each of your ten largest agreements whose vendor has a self-serve tool, carrying the annual value you already store, each row marked **From your {vendor} agreement** with the renewal month where one is held.
- **The one missing figure** — each seeded row states the single number its tool still needs, and why: the **list value** where the vendor is ranked on discount off list, or the **unit count** (seats, users, cores) where it is

ranked on net fee per unit. A row with no stored annual value asks for that instead.

- **What is not seeded** — agreements beyond the ten largest, and agreements whose vendor has no self-serve tool, are counted and named under the table rather than dropped in silence. Every vendor tool stays open from the benchmark library.
- **Nothing runs until you press Run** — seeded rows are a draft. They are editable and removable, and no row is placed, saved, or benchmarked until you press **Run the benchmark**. Removing them all says so and leaves your workspace untouched.
- **What next** — the roll-up ends with its next moves: rework the largest-headroom position in its vendor tool, start the negotiation, open the savings plan, copy the one line verdict, or take the numbers to Vera or a booked analyst.

How to use it

1. On a fresh workspace, fill in the one missing figure on the rows you care about and press **Run the benchmark** for a quick placement of each.
2. Open a vendor that looks above market to see the detail behind its standing.
3. To keep a position in the roll-up rather than just placing it once, open that vendor's tool and save the scenario there.
4. See ranked savings opportunities in Opportunities.

Tips

- A run from the seeded table is a quick placement against reference data, not a saved position. Saving a scenario inside the vendor tool is what adds a lasting row to this roll-up.
- An Oracle Technology position is ranked on discount off list, and the fee a discount gap moves is a one-time license charge rather than a yearly one, so it contributes no annual headroom to the total above.

Related: Opportunities · Benchmark all my vendors

BENCHMARK ALL MY VENDORS

Where: Benchmarks & market › Delivered reports › Benchmark all my vendors · `/reports/vendor-benchmark`
Who: everyone · requires an AI plan
You get: a three page portfolio benchmark you can download as Word · 1 to 3 minutes

A three page portfolio benchmark generated from every agreement you have uploaded: spend per vendor, market position where our reference data defends one, and an explicit insufficient-data note where it does not.

In plain words: the Portfolio benchmark screen as a document you can forward.

When to reach for this

- The estate read needs to leave the platform: a Word document for finance, not a screen share.
- You want the honest coverage map: page three names exactly which vendors cannot be benchmarked yet and the one step that fixes each.

On the screen

- **Generate button** — scans your agreements, places each vendor in its peer cohort, writes the report, and checks every figure against your stored data. Takes 1 to 3 minutes.
- **Trust strip** — like every AI-generated report, it opens with its grounding state ("Figures verified against your data", or a review notice when a figure needs one), the source count, the generation date, and feedback thumbs.
- **Spend per vendor** — page one: portfolio vitals and an overview of annual spend by vendor, largest first.
- **Market position table** — page two: each vendor's position and the basis for it, a delivered analyst score, a cohort percentile with a ± range, market reference data, or insufficient data.
- **Data coverage** — page three names the vendors we cannot benchmark yet and the one step that closes each gap. We never estimate a position without data.
- **What next** — the report ends with its next moves, built from its own totals: turn the positions into the savings plan, start a negotiation on the largest headroom vendor, copy the one line verdict, ask Vera which vendor to take to the table first, or book the analyst. The Send it on panel drafts the CFO, CIO, vendor, or colleague email from the report itself.

How to use it

1. Generate or regenerate the report.
2. Download as Word; the export matches the on-screen brief.
3. Close a data gap by adding a list price or seat count in the vendor's benchmark tool, which moves it from reference to ranked.

Tips

- This report reads from agreements, not one-off instant checks. The more complete your Agreements list, the more of your stack it can place.

Related: Portfolio benchmark · Delivered reports

DELIVERED REPORTS

Where: Benchmarks & market › Delivered reports · `/studio` (the delivered reports are the Library tab of the Documents app; Team & AI › Deliverables opens the same place)
Where: from the rail · `/reports` forwards to this tab
Who: everyone

Benchmark reports our analysts delivered for your proposals, plus anything still in flight with its estimated delivery date. This page is also your documents library: everything the platform has ever written for your organization lives in the All documents section below the shelf. Report covers on the shelf lead with the annual value found; the overall score sits as a small corner mark.

In plain words: the reports the analysts wrote for you, and the shelf of everything else the platform has ever written.

When to reach for this

- A proposal you uploaded reads Delivered and you want the finding.
- You are hunting for something the platform wrote weeks ago: a memo, a decode, a face-off.
- The renewal is near and the delivered report's levers need to become the ask list.

A worked example. The sample portfolio's delivered report on the Northwind CRM subscription shows the shape every report follows. The finding: the deal is above market, with $28,650 a year in reach. The evidence: the seat price at $571 against a market median of $504, worth $16,750 a year, and a 12 percent discount off list against a 22 percent median, worth $11,900. The structure: a 7 percent uplift cap against a 5 percent median means the premium widens every year it is left alone. The levers follow in order, priced, with the talking points scripted. Read your own reports the same way: finding, evidence, structure, then the play.

On the screen

- **Report page** — a metrics table with percentile bars, market low, median and high, an overall score, and the analyst summary. Reports generated from August 2026 lead with a paged executive brief: a vitals strip, the executive summary, a deal snapshot of your current position against the benchmark supported position, your terms against the market band with a standing verdict per metric, the headline metric drawn against its band, a savings bridge priced from the stored metrics, numbered negotiation levers with a sequencing note, and scripted talking points; earlier reports keep the layout they were delivered in. Long reports show a thin gold reading-progress line at the top of the window, a reading-time estimate in the byline, numbered figures with their sources, and a Top button once you scroll.
- **Open the case** — on a delivered report, "The case file" navy band on the proposal's Report tab opens the benchmark as a full screen case: numbered executive brief documents in a grouped left rail, paged with the arrow keys, with a read meter and Esc to close. Six documents: the cover, the source proposal record, the analyst brief, a metrics atlas with the market ranges and the verdict table, an ask sheet priced from the stored figures, and a What now sheet with a drafted CFO brief, a vendor safe meeting request, chat ready copies, and the PDF. Drafts open in your own inbox, nothing sends from a case, and every figure is read from the report itself.
- **In-flight benchmarks** — proposals still in review, with an estimated delivery date.
- **All documents** — every brief, memo, and report your organization has generated. The library opens as a shelf of document covers, hued by type, with the vendor, title, and date on each spine; a Covers | Table toggle switches to the dense table (sortable, searchable, pageable, CSV export) and remembers your choice. The library also lists your saved drafts, contract decodes, and quote face-offs alongside the AI deliverables. A cover or row opens the document itself, with the owning record one click away in the table's Context column.
- **Origin filter** — beside the type filter, narrow the library by where a document came from: All origins, AI report, Tool run, or Draft.
- **Boardroom Brief, presented** — after writing a Boardroom Brief from this page, Present beside its download actions opens it full screen as a four movement show: the posture, renewals and risk, savings and leverage, then the asks, with the brief's KPI figures on the stage ticker and a printable leave-behind of the screens you walked.

- **Boardroom Brief, as a case file** — once a brief is written, Open as a case file in the dialog and preview footers expands it into the same full screen case reader: a cover sheet with the headline numbers and the bottom line, one exhibit per section of the brief, and a What now sheet carrying the same Word and PowerPoint downloads, a team circulation draft, and the desks to open next.
- **Word downloads** — rows whose type supports it show a clickable Word link in the Formats column, so you can export straight from the list. Most document types export to Word, including vendor dossiers, QBR prep packs, deal handoff memos, audit exposure reports, clause coverage reports, consolidation memos, position trajectory reports, diligence reports, SOW rate reviews, vendor relationship analyses, and saved tool analyses.

How to use it

1. Open a report and walk the metrics.
2. Download the report as a PDF.
3. When you are hunting for something the platform wrote earlier, sort or search All documents instead of revisiting the record that produced it; the row opens the document itself, and the Word link in its Formats column downloads it directly.
4. Bookmark a report for later, and find it again in Saved.

Tips

- A delivered report comes from uploading a proposal for expert review, not from an instant check. Start one on the Proposals screen.

If something looks wrong

- The shelf is empty: no proposal has reached Delivered yet. The in-flight strip shows what is coming and when; an instant benchmark fills the gap today.
- A report you expect is missing from All documents: check the Origin filter first; it may be narrowed to one origin.

Related: Proposals · Saved

SEND IT ON: THE FOLLOW UP EMAIL, DRAFTED

Where: beneath every report and brief, from the delivered benchmark report to dossiers, memos, and tooling reports · AI
Who: everyone

In plain words: the follow-up email, already written, for whichever reader you have in mind.

A report is rarely the last step; someone has to act on it. The Send it on panel at the foot of every report drafts that email for you. Pick the reader and Vera writes a short, friendly note built only from the report above it, never from figures the report does not carry. Nothing sends from the platform: you copy the finished email and paste it into your own mail client.

On the screen

- **Reader chips** — To the vendor (named when the report knows the vendor), To your CFO, To your CIO, To a colleague. Each reader gets its own draft: the vendor note is friendly but firm and careful about what it reveals, the CFO note leads with the money, the CIO note frames the platform decision, the colleague note is a quick heads up.
- **Subject and body** — both fully editable before you copy. Drafts are short on purpose, two paragraphs at most, and sign off with a [Your name] placeholder.
- **Copy email** — copies the subject line and body together, ready to paste into Outlook or Gmail.
- **Redraft** — writes a fresh draft for the selected reader if the first one misses the mark.
- **Confidentiality check** — vendor drafts pass the same scan as negotiation language; if a draft carries an internal figure or a name it should not, a warning lists exactly what to remove.
- **On the dashboard** — the desk rail carries a Send it on card listing your newest documents; Draft email opens the same panel without leaving the dashboard.
- **In theatre mode** — presenting a brief adds an envelope button to the stage controls; it opens the drafting drawer over the show, built from the document on stage, so the follow up is written before the room empties.

How to use it

1. Open any report and scroll past the document to the Send it on panel.
2. Press the chip for the reader you have in mind and give Vera a few seconds.
3. Edit the subject or body until it sounds like you.
4. Press Copy email, open a new message in your mail client, and paste.

Tips

- Drafts only state figures that appear in the report itself. If a number is missing, the email speaks in direction instead; add specifics in your edit if you want them.
- Switching chips keeps each reader's draft, edits included, until you leave the page.

Related: Delivered reports · Renewals

REPORTING STUDIO

Where: Documents › Generate · `/studio/generate` · AI
Where: from the rail · `/reporting` forwards to this tab
Who: everyone on an AI plan
You get: an executive brief of 2 to 5 pages, written from your own stored agreements · a run takes 1 to 3 minutes

56 standardized reports on your own contracts and spend, plus custom commissioned ones. Every run is written from your stored agreements only, addressed to the reader you choose (CIO, CFO, or procurement), at the length you choose (2 to 5 pages), and figure checked against your data after writing. Where the data cannot support a number, the report says not on file instead of estimating.

In plain words: pick a report, pick the reader, and a grounded executive brief writes itself from your own records.

When to reach for this

- Leadership asked for something in writing and the deck-building afternoon is not available.
- The recurring documents (the spend review, the renewal pipeline, the concentration read) should come from live data, not from last quarter's copy.
- You need a document the catalog does not have: the commission desk takes a title and one sentence.

On the screen

- **The commissioning stage** — the dark masthead opens with your estate's own numbers: tracked annual spend, the agreements grounding every figure, and the 56 reports on the desk.
- **Run controls** — Reader, Length, and Scope float beneath the stage as one instrument, applied to whichever report you run next. Report default lets each report address its natural reader; the scope select narrows any run to one vendor.
- **The catalog** — 56 reports across eight groups, each card a miniature executive brief cover: executive and board, spend and cost, renewals and pipeline, vendors and concentration, risk and compliance, market position, negotiation and savings, and data and operations. Search filters by title and topic; the category chips narrow by group; the whole card is the Run button.
- **The Writing Room** — running a report raises a full screen stage where the document writes itself in front of you: the masthead and gold rule land first, the vitals arrive with the fact pack, the prose writes line by line, and the fact check sweeps the figures in gold, beside a five stage checklist, a live progress rail, and an analyst ticker narrating the work.
- **Commission a custom report** — a title and a one sentence brief produce a report the catalog does not have, to the same grounded standard.
- **Recent runs** — your latest reports as chips, one press back to any of them.
- **The run page** — the finished document with its Word export, print to PDF, Reader and Length controls for regenerating, and version history.
- **Present** — runs the report full screen on the shared stage as a working session: the executive summary opens, the vitals get their own screen, each section walks as its own act, and the close carries the recommendations and the bottom line. The stage's envelope button drafts the follow up email without leaving the show.
- **Forward** — emails the report to a colleague with your own note and a 30 day read only link. No account needed on their side.
- **Send it on** — beneath every finished report, the follow up email drafter: to the vendor, your CFO, your CIO, or a colleague (see Send it on).

How to use it

1. Pick the reader and the length in the run controls; leave Reader on Report default to let each report address its natural audience.
2. To report on one vendor, choose it in the Scope control. Reports marked Runs on one vendor require it.
3. Press Run on any catalog card. Generation takes 1 to 3 minutes and lands you on the finished report.

4. On the run page, export to Word, print to PDF, or change the reader or length and press Regenerate.
5. Press Present to walk the report on the stage, or Forward to email it to a colleague as a 30 day read only link.
6. For something the catalog does not have, give the commission desk a title and a brief and press Run.

Tips

- The same report reads genuinely differently per reader: the CFO version leads every section with the number, the CIO version reads the estate as platform strategy, the procurement version is an ordered work list.
- Every run files into All documents on Delivered reports with version history, so a monthly rerun never overwrites last month's edition.
- A report is only as complete as the record: the Data gaps report in the catalog tells you exactly which missing values and dates to fix first.

Related: Delivered reports · Benchmark all my vendors · Proposals

VENDOR SCORECARDS

Where: Benchmarks & market › Vendor scorecards · `/scorecards` · via **More** or Ctrl/⌘ K
Who: everyone

Every calibrated vendor graded as a counterparty: the renewal corridor behind our benchmarks turned into an A-to-F grade, plus how verified buyers rate the vendor's renewal fairness, audit conduct, and flexibility.

In plain words: what this vendor is like to renew with, graded, before you find out the hard way.

When to reach for this

- Before the first call with a vendor you have never negotiated: the grade and the buyer ratings set your expectations.
- When shortlisting: two similar products with different renewal conduct are not similar purchases.

Read against your own estate. The page opens with how many of the graded counterparties you actually hold, what you pay them a year, and what renewing them at their own calibrated medians would add. That last figure, your annual value times the median uplift, is stated per vendor and as a total, and the vendors you hold are listed first in that order. Every corridor is drawn on one shared axis with the median marked and a hairline at a flat renewal, so two vendors compare by eye rather than by reading two sentences.

On the screen

- **The grade** — renewal economics from the calibrated corridor, not a popularity score.
- **Buyer ratings** — one review per organization, shown only as floored aggregates.
- **Public page** — each vendor has a public scorecard the whole market can cite.
- **Put the grade to work** — the vendor page's rail routes onward: run an instant benchmark against the corridor, ask peers on the Buyer Network, build the renewal plan, open the vendor room, or start the

negotiation.

How to use it

1. Rate a vendor you hold paper with.
2. Open the public page for a vendor you are renewing.

Related: Research · Vendors

RESEARCH

Where: Benchmarks & market › Research · `/research` · via **More** or Ctrl/⌘ K
Who: everyone

Our latest white papers and market notes: vendor pricing moves, licensing changes, and negotiation guidance.

In plain words: the analyst library, searchable, with the papers on your own vendors floated first.

When to reach for this

- A renewal is coming and the vendor's Negotiation Toolkit and Renewal Outlook are the homework.
- Something moved in the market (an audit campaign, a licensing change) and you want the considered read, not the headlines.
- Ten spare minutes: the Most liked shelf is reliably worth them.

The series you will find here

- **Negotiation Toolkit** — one per covered vendor: the levers, the counters, and how to run the deal.
- **Negotiation Intelligence Report** — for the major vendors: the seller's calendar, comp plan, and tactics decoded.
- **Renewal Outlook** (quarterly, top 10 vendors) — what your renewal will cost in the next 12 months and what to lock in now: announced price changes with dates, uplift corridors, and a month-by-month timing calendar.
- **Vendor Strategy & Viability Assessment** (annual, top 10 vendors) — where the vendor is taking pricing, packaging, and product, what that does to your cost curve, and a closing posture recommendation.
- Topical papers on specific market events, such as licensing policy changes and audit campaigns.

For covered vendors, the newest Renewal Outlook also appears automatically on that vendor's renewal, contract, and negotiation pages, and the assessment on the vendor's page, so the right paper is at hand where you work the deal.

On the screen

- **The search toolbar** — the page opens on the search field, the topic filters, and a line counting the papers, the vendors, and what arrived in the last month. The field filters the whole library as you type and shows a live match count. Press **/** from anywhere on the page to start typing, and **Esc** to clear. Matched words light up in titles and summaries, and while a search is active the results rank best match first.

- **The library rail** — the dark column down the left of the results, in three blocks. Its **search box** drives the same catalog as the stage above, with two modes: **Keyword** is the strict matcher, where every word you type has to appear, precise and unforgiving of a wrong word; **Meaning** counts any word that lands and weighs topic and summary matches alongside titles, so a question phrased the way you would say it out loud still finds the paper. Both run in your browser over the whole index, so both answer as you type, and the rail says how many papers match. **Your shelf** holds the briefs you have opened in this browser under **Recent** and the ones you starred under **Starred**, one click back to where you were. **Cite what you read** acts on the last paper you opened: **Cite into case** copies the title, the publisher, and the link that opens that exact paper, ready to paste into a case or a brief; **Open the case** takes you to that vendor's benchmark case when the library covers it; **Copy the reader link** copies one link that opens the paper for anyone in your workspace.
- **Topic pills** — every paper is tagged by topic (Negotiation, Renewals, Pricing and inflation, Audit defense, CFO and board, Cloud, AI licensing, Exit and alternatives, Contract terms). One tap narrows the library to that topic; tap again to release it.
- **Vendor chips** — filter to one vendor. The counts on the chips update live with your search and topic, so you always see how the current result set breaks down. Vendors that appear in your own agreements carry a small gold dot.
- **For your vendors** — when nothing is filtered, a shelf above the results keeps the newest research on the vendors you hold contracts with.
- **New badge** — papers published in the last two weeks are marked **New**.
- **Sorts** — recent (or best match while searching), **Most liked · top 25** (the 25 most liked papers platform wide), oldest first, and **By vendor**, which groups the library into a directory with a header per vendor.
- **Cards or Shelf** — a layout toggle beside the sort. Shelf stands every brief as a case spine on a gold rail, vendor on the spine, following whatever filters and sort you have set; a bright gold band is an unread brief, and spines you have read dim to soft gold. Your choice of layout is remembered in this browser.
- **The reader** — papers open in a reading dialog. A note published without an attached file renders as a formatted brief with an estimated reading time at the top. When a paper has an attached file, an **Open in tab** button beside Download opens it in its own browser tab. **Previous** and **Next** (or the left and right arrow keys) walk the whole filtered shelf in order with a position counter, so a series reads cover to cover; read cards carry a green check.
- **After reading** — the reader ends with an After reading strip, in the dialog and in reading mode: benchmark the vendor the paper names, or ask Vera a question about the paper you just read.
- **Reading mode** — the **Reading mode** button in the reader takes the paper full screen on the same stage as the playbook theatre: the room goes truly fullscreen (F toggles it) and the header fades after a few seconds of stillness, returning on any movement. The sun and moon toggle sets the stage light to light or dark independently of your app theme, and the choice is remembered on this browser. The arrow keys still walk the shelf; Escape or the X returns you to the normal reader.

- **Share link and Forward** — both live in the reader header, in the normal reader and in reading mode. **Share link** creates a public link and copies it: anyone with the link can read that one paper for 90 days, no account needed. **Forward** emails the paper to up to five colleagues with an optional personal note.
- **Like button** — one press, one vote, with the live count beside it: what readers across the whole customer base found worth their time. Counts are anonymous; only totals are ever shown.
- **Bookmark icon** — save any paper as a favorite. The bookmark sits on every card and in the reader header, so you can favorite a paper while reading it. Favorites also appear on your Saved page.
- **Favorites chip** — leads the filter row and shows your saved count. Press it and the page becomes your own shelf: vendor chips, search, topics, and sort all work inside your favorites.

How to use it

1. Type what is in front of you, a vendor name, "audit", "briefing the CFO", and open the top match. Results update on every keystroke.
2. Read or download a paper, and press the thumbs up on anything that earned it.
3. Share a paper with a colleague straight from the reader: **Share link** copies a public link, **Forward** emails it with your note. The **Recommend** button on each card opens the same dialog.
4. Short on time? Switch the sort to **Most liked** and read what everyone else rated, or start from the **For your vendors** shelf.

Download limit: reading a paper in the viewer is unlimited, but file downloads are capped at ten per day and thirty per week per person so access stays fair across organizations. The allowance frees up on its own as your recent downloads age out. If you reach the limit and need more before then, ask us and a VendorBenchmark admin can reset it for you right away. All platform limits are collected in Plans, allowances & limits.

Related: Saved · Vendor scorecards

THE BENCHMARK LIBRARY

Where: Benchmarks & market › Instant benchmarks › any vendor · `/benchmarking/[vendor]`
Who: everyone

In plain words: one self-serve price-check page per vendor and product line, roughly 1,286 of them.

When to reach for this

- The vendor you are negotiating has its own page here; this is the deep tool the hub's search lands you on.
- The deal prices on an unusual meter (cores, FUE, endpoints, trader seats): the page names the unit that actually bills you.
- You want the full case treatment: the placed verdict opens as a four document benchmark case with the ask sheet scripted.

The instant benchmarks hub is a front door to roughly 1,347 individual vendor tools, one page per vendor and product line, from Microsoft EA and RISE with SAP to Snowflake, CrowdStrike, and Autodesk, including the AI tool stack (AI coding assistants, model and inference APIs, GPU cloud, data labeling, voice, and video) and the Google estate (Google Ads, YouTube, Marketing Platform, Analytics 360, Maps Platform, Mandiant, Cloud

Support). Each page is a self-serve price check: you enter your own numbers and see where your deal lands against a cohort of comparable transactions, with no upload and no wait.

Real estate and staffing are covered as industries rather than as a handful of names. On the property side, alongside Yardi, RealPage and MRI Software the library prices CRE data and marketplaces (CoStar Group, Cotality, Altus Group), property management platforms (AppFolio, Entrata, Storable), leasing, deal and investment management (VTS, Dealpath, Juniper Square), and worktech, lease accounting and residential brokerage (Eptura, FinQuery, HqO, Lone Wolf). On the staffing side, alongside Bullhorn, Beeline and iCIMS it prices the agency front and back office (Avionté, TempWorks, Crelate, Loxo, Daxtra), talent engagement and experience (Sense, Phenom), recruitment marketing and job advertising (Radancy, Appcast, ZipRecruiter, DHI Group), and the PEO and staffing payroll platform underneath (PrismHR). Neither industry meters per employee, so each page names the unit that actually bills you: per unit under management, per facility, per building, per lease, per recruiter seat, per active job slot, per worksite employee. Four of these pages are built on a published rate card, and two of those tell you the escalator matters more than the discount, because Crelate publishes a 7 percent annual price rise and Loxo is reported at 5 percent.

Retail and commerce has its own category on the hub. It covers the enterprise store estate (Aptos ONE, Cegid Retail, Diebold Nixdorf Vynamic, Toshiba ELERA, Zebra Workcloud), the shelf and the shrink stack (VusionGroup electronic shelf labels, Nedap iD Cloud RFID, Sensormatic, Everseen checkout vision AI, Scandit smart data capture), merchandising, planning and shopper data (SymphonyAI Retail CPG, ToolsGroup, Impact Analytics, dunnhumby, Placer.ai, alongside Blue Yonder and Oracle Retail), digital commerce, marketplace and channel (VTEX, Spryker, Mirakl, Rithum, alongside Shopify Plus, BigCommerce, commercetools and Salesforce Commerce Cloud), product and customer content (Syndigo, Akeneo, Bazaarvoice), and fulfilment, post-purchase and the frontline workforce (Fluent Commerce, Narvar, WorkJam). Retail almost never prices per employee, so each page names the meter that actually decides the bill: per terminal or per store, per lane, per shelf label, per RFID tag, per named planner, per seat, per order, per SKU and recipient, as a share of gross merchandise value, or per frontline employee. Several of these pages say plainly that the discount is not the lever and explain what is, because the money in a retail contract is usually somewhere else: the consumable tag schedule behind an electronic article surveillance install, the year-seven shelf label replacement price, the double-run between an old point of sale maintenance line and its cloud replacement, the take rate rather than the subscription, and the seasonal headcount a frontline licence is struck against.

Energy and utilities is a category of its own, and it is the largest single build in the library. It runs from the operating estate of a utility (ADMS and SCADA, DERMS and virtual power plants, metering and the grid edge, vegetation and asset inspection, crew and field service, retail energy platforms, water data) through oil, gas and industrial operations (subsurface and drilling, production and upstream ERP, pipeline and terminal, industrial analytics, machine health, maintenance management) into energy trading and markets (E/CTRM, trade surveillance, market data and price reporting, power forecasting) and the clean energy transition (storage optimisation and battery analytics, electric vehicle charging, solar and wind design and performance, carbon accounting, utility data access, environment and safety, environmental commodity markets). Several of these are not licence deals and the pages say so rather than forcing them into a discount frame: an efficiency provider that installs equipment at no up-front cost and earns from your energy saving is a financing agreement, where the measurement baseline and the early buyout decide the cost, and a charging network that earns per kilowatt hour turns on who carries the demand charges when a fleet plugs in at once.

ITSM and ops covers the estate around ServiceNow rather than only ServiceNow, because the alternatives are what give a ServiceNow renewal any tension at all. Security is built the same way, covering the shortlist a buyer actually runs a platform against instead of only the platforms themselves. In both categories the page tells you where a credible alternative exists and where one does not, since a renewal with no second quote is a different negotiation from one with three.

Open source vendors that sell commercial support are covered as their own kind of deal, because the software is free and what you are buying is support, indemnification, certified builds, and long term security maintenance. Your genuine alternative is not a competitor, it is running the same software unsupported for nothing, so each page says where the community edition is already sufficient and where the subscription earns its place. The library covers databases and data platforms, Java runtimes, Linux and operating systems, private cloud, middleware and developer tools, content and collaboration, monitoring, and security tooling. Some are unusually clear cut: one monitoring vendor gates no features at all and sells only response time, one hypervisor is free to run in production with published per socket tiers, and one community platform has no paid tier to unlock, which makes a self host threat unusually credible. Two carry a specific threshold worth knowing before you open a quote: MySQL Enterprise does not discount below roughly $100,000 of annual value and then moves from 10 to 60 percent as the deal grows, and Azul runs 30 to 60 percent off list.

Advertising and media has its own category on the hub. Fifty vendors cover the stack a brand, an agency or a publisher actually signs for: the exchanges and supply platforms (Magnite, PubMatic, Index Exchange, OpenX, Equativ, Yieldmo, GumGum), native and content recommendation (Taboola, Teads, MGID), demand side platforms and media buying software (StackAdapt, Viant, Basis Technologies, Zeta Global, Adform, Moloco, MiQ, Smartly, Skai, alongside The Trade Desk and Criteo), verification and fraud (HUMAN Security, Zefr, Pixalate, alongside DoubleVerify and Integral Ad Science), audience measurement and the television currencies (Nielsen, Comscore, Kantar Media, VideoAmp, iSpot.tv, EDO, Samba TV), marketing effectiveness (Measured), mobile and app attribution (Adjust, Branch, Kochava, alongside AppsFlyer), shared advertising identity (ID5), agency ad operations (Mediaocean, which now owns Flashtalking and Innovid), out of home and audio (Vistar Media, Broadsign, AdsWizz, Triton Digital), search and AI visibility (BrightEdge, Conductor), affiliate and partnerships (impact.com, Awin), creators, social and communications (CreatorIQ, Later, Cision), creative production (Celtra, VidMob), supply path curation (Scope3) and consent management (Didomi). Advertising almost never meters per employee, so each page names the unit that actually decides the bill: a percentage of media spend for a demand side platform, a percentage of publisher revenue for an exchange, a cost per thousand impressions for verification and identity, a cost per attributed event for mobile measurement, a percentage of tracked sales for an affiliate network, and a price per screen for out of home. The published anchors sit on the pages: platform fees run about 15 to 20 percent of media and compress toward 7 to 13 percent above roughly $3m of annual spend, sell side take rates average about 14 percent of publisher revenue, and verification runs about $0.05 to $0.10 per thousand impressions. Several pages say plainly that the discount is not the lever and explain what is, because a measurement currency your television deals settle against and the workflow your agency bills through do not negotiate on rate. Ownership matters more here than in any other category, so each page names it: your mobile measurement partner may belong to a company that also sells you media, your audio technology may belong to one of the largest audio sellers, and your television data vendor may be part owned by the broadcasters it measures.

The library also covers professional services under its own category. Twenty services providers, from the strategy houses (McKinsey, BCG, Bain) and the Big 4 (Deloitte, PwC, EY, KPMG) through the global SIs (Accenture, IBM Consulting, Capgemini, NTT DATA), Booz Allen Hamilton, Slalom, Kyndryl, and the offshore majors (Cognizant, Infosys, TCS, Wipro, HCLTech, Tech Mahindra), are benchmarked on blended hourly rate per consultant rather than discount off list: enter your total fees and hours and see where your effective rate lands for that firm and engagement size, with your onshore/offshore delivery mix and commercial leverage as the levers.

Alongside those firms, the library benchmarks the four software vendors that also sell the services to implement their own software, one tool per practice rather than one per firm: ten Microsoft Industry Solutions Delivery practices (Azure infrastructure and migration, data and AI on Fabric, Copilot deployment, Dynamics 365 finance and operations, Dynamics 365 customer engagement, Microsoft 365 and modern work, security, Power Platform, app innovation, and industry cloud solutions), ten Oracle Consulting practices (Fusion Cloud ERP and EPM, HCM, SCM and CX, NetSuite, OCI migration, E-Business Suite, database and Exadata, Oracle Health, and managed services), twenty IBM Consulting practices split across its Strategy and Technology line (strategy, hybrid cloud and Red Hat, application modernization, data and AI, cybersecurity, the SAP, Oracle, Salesforce, Microsoft and AWS practices, IBM iX, finance, supply chain, talent, sustainability, quantum and mainframe) and its Intelligent Operations line (application managed services, business process operations, and cloud platform operations), and ten SAP Services practices (RISE adoption, S/4HANA brownfield conversion and greenfield implementation, premium engagement, SuccessFactors, Ariba, BTP, analytics, Signavio, and custom development). All fifty use the same blended hourly rate metric as the firm-level tools. The point each page makes is the same one: a vendor's own consulting arm can defend a premium for the work only it can credibly do, such as an SAP brownfield conversion or an Oracle database architecture with licensing consequences, and cannot defend one for the work a systems integrator delivers just as well, which is why the rates inside one vendor range so widely.

The largest publishers are covered by product line rather than as one tool, because they do not price as one vendor. Oracle has its own deep set: the Fusion cloud applications (SCM, EPM, and the CX suite of Sales, Service, Marketing, CPQ, and Commerce), Analytics Cloud, the industry clouds (Health, Retail, Hospitality, Financial Services, Utilities, Communications, Transportation Management, Field Service, Primavera Cloud), and the on-prem license and ULA-family estates (E-Business Suite, Siebel, JD Edwards, PeopleSoft, Hyperion, Primavera P6, Pool of Funds, and Perpetual ULA). Salesforce is broken out into Sales, Service, Commerce, Experience, Revenue (CPQ), Field Service, Platform, Agentforce, CRM Analytics, and the Health, Financial Services, and Nonprofit clouds. ServiceNow covers ITSM, ITOM, CSM, HRSD, SPM, Security Operations, IRM, App Engine, Field Service Management, and the Now Assist AI add-on. SAP covers S/4HANA (on-prem and RISE, with RISE Base and RISE Premium benchmarked on separate curves), GROW, Analytics Cloud, Datasphere, Integrated Business Planning, Fieldglass, Commerce Cloud, Sales & Service Cloud, Signavio, LeanIX, SuccessFactors, Ariba, Concur, and BTP (CPEA and BTPEA as separate tools). Pick the exact product line you are negotiating.

The Microsoft EA tool models the November 2025 licensing change directly. Microsoft removed the EA volume discount levels for online services on November 1, 2025, so the tool carries a pricing era control: deals under the new rules benchmark roughly 10 points less discount than the legacy price sheet, unless the deal commits Copilot on an E5 (or E7) estate, which restores pre 2025 discount levels. Set the era to match your agreement; renewals still on a legacy price sheet keep it until they renew.

The SAP tools follow the August 2026 curve guideline: median discount off list with a p05 to p95 range, banded by ACV of the SAP component (under $500k, $500k to 2m, $2 to 10m, $10m and up). RISE Private Cloud has separate Base and Premium curves and separate tools, BTP has separate CPEA and BTPEA curves and tools, GROW carries the shallowest cloud curve, and S/4HANA on premises perpetual is shown as a reference curve only, because cloud and perpetual percentages are not comparable: a 55 percent RISE discount can beat a 70 percent perpetual discount once bundled infrastructure, support and managed services are counted. SAP ECC remains the exception: new ECC purchases price at list, with value delivered only through RISE conversion credits.

The Oracle applications estates (E-Business Suite, JD Edwards, and Siebel) share one discount curve, from the August 2026 curve documents: 50 percent off list is the median at the $100K deal minimum, climbing to 78 percent at $2m and above, so any offer under 50 percent reads as an opening position rather than a market price. The pages say why the discounts run this deep: they are nearly free for Oracle to give, because the commercial objective is the 22 percent support annuity on the net fee, which is why the escalation cap and repricing protection belong in the same negotiation as the rate. Oracle Java SE is the opposite: the published employee tiers already pre discount the rate, so negotiated discounts are thin by design (2 to 4 percent median below $2m of annual spend, 20 percent at $2 to 5m, 40 percent at $5m and above), and the tool carries a documented OpenJDK exit checkbox because the exit alternative, not the rate conversation, is what moves Oracle on Java.

The Oracle ULA benchmark is a deliberate exception to reading the discount at all. It opens with the fee curve by product scope: a focused 2 to 4 product technology ULA lands near $2m for a three year term, 5 to 8 products near $6m, 9 to 14 near $14m, and a 15 plus product estate near $30m, with your own scope and fee placed on the curve live. Beside it sits the signing reason multiplier, the largest hidden pricing variable in the deal: growth driven signings pay the baseline, renewals 1.1 to 1.3x, M&A driven deals 1.1 to 1.4x, and audit settlements 1.3 to 2.0x, because Oracle prices audit driven agreements to the exposure rather than the value. The how Oracle actually prices a ULA panel carries the negotiation guidance: keep growth estimates conservative because they are not contractual, always signal you may not renew, start your clock at least 12 months out (Oracle forecasts your pricing up to 11 months before expiration), sweep cloud and DR instances early because late sweeps came in 1.5 to 2.5x over expectations, and negotiate the terms that outlast the price (the certification clause, public cloud deployment rights, a support fee cap after the ULA ends, coverage for every legal entity, and merger and acquisition rights). A line under the play card links the Oracle ULA certification playbook for the renew or certify decision itself.

On the screen

- **The curve dock** — on vendors covered by an August 2026 curve document (SAP, SAP Ariba, SAP Concur, SAP SuccessFactors, IBM, Microsoft EA, Workday, ServiceNow, the Oracle applications estates of E-Business Suite, JD Edwards, Siebel, PeopleSoft and Hyperion, Oracle Java SE, Oracle Fusion Cloud SaaS across ERP and SCM, HCM, EPM and CX, NetSuite, Salesforce across the core clouds, Marketing and Data

Cloud, Agentforce, Tableau, MuleSoft and Slack, Google Workspace and Google Cloud with BigQuery and Gemini, AWS EDP, Adobe across Creative Cloud, Acrobat and Experience Cloud, Atlassian, Docusign, Zoom, CrowdStrike, Palo Alto Networks, Databricks, Snowflake, Informatica, Okta, Coupa, Veeam and the backup platforms Rubrik, Cohesity and Commvault, Autodesk, Datadog, Twilio, Citrix, Cloudflare, the contact center trio of Genesys, NICE CXone and Five9, Check Point, F5, Akamai, BMC, OpenText, IFS, Infor, SAS, Qlik, Alteryx, OneStream, Anaplan, Celonis, UiPath, Pegasystems, Procore, Epicor, Deltek, Tyler Technologies, Palantir, Freshworks, HubSpot, Sprinklr, Sage, Zendesk, UKG and ADP, Teradata, Nutanix, Wiz, CyberArk, SentinelOne, Zscaler, MongoDB Atlas, Confluent, Elastic, GitLab, JFrog, TIBCO, Boomi, Gong, Box, MuleSoft on its own vCore document, the financial services platforms FIS, Guidewire and Temenos, the market data desks Bloomberg, FactSet, LSEG, S&P Global Capital IQ and Moody's Analytics, the healthcare records vendors Epic, Oracle Health and Veeva, and the experience management pair Medallia and Qualtrics, and the finance and legal platforms Kyriba, BlackLine, Workiva, Vertex, Avalara, Ironclad and Icertis CLM, iManage, Relativity, LexisNexis and Thomson Reuters Westlaw, the collaboration, security and mid market ERP platforms Notion, Miro, Figma, Fortinet, Sophos, Darktrace, KnowBe4, QAD, Acumatica and Unit4, the Cisco estate re benchmarked across hardware and the Enterprise Agreement with Meraki, Webex and Splunk beside them, the information and data platforms Nvidia AI Enterprise, Kaseya, ZoomInfo, Wolters Kluwer, CoStar, IQVIA, Gartner and LinkedIn, the PE owned tool portfolios Quest and One Identity, Flexera, Perforce and ConnectWise, the HR stack Deel, Paycom and Paylocity as one negotiation, iCIMS, Cornerstone and Dayforce, the GRC set AuditBoard, Diligent, Mitratech, Archer and NAVEX, the consolidators Precisely, Rocket Software, Progress, Aptean and the Constellation Software operating businesses, the PAM duopoly Delinea and BeyondTrust on one shared curve, the security mid tier Ivanti, Tenable, Barracuda and Mimecast, the healthcare operations set Meditech, Waystar, R1 RCM and symplr, the education pair PowerSchool and Ellucian on one shared curve, the property manager RealPage with the DOJ settlement repricing line, the AP automation pairs Tipalti and BILL by monthly payment volume and Basware by annual invoice volume, the procurement challenger JAGGAER, the subscription billing vendor Zuora, the legal drafting roll up Litera, the insurance pair Duck Creek and Applied Systems on one shared curve, the ambulatory EHR pair athenahealth and NextGen by practice size, the nonprofit platform Blackbaud, the AEC conversion pair Nemetschek and Bluebeam on one seat banded curve against the December 2026 Revu deadline, Trimble with its published 40 to 80 percent trade in tiers shown as a reference line, and the Hexagon divisions, the insurance utility Verisk negotiated per layer, the data stack Domo, dbt Labs and Fivetran, the Equifax Work Number met with the verification waterfall, the BI challengers Sigma and ThoughtSpot on one shared displacement curve, and the developer tools set Snyk, Dun & Bradstreet, Docker, GitHub with Copilot and Advanced Security priced as one negotiation on both the GitHub and Copilot tools, HashiCorp, Redis, Postman and JetBrains, the library's buyer friendly counter example whose nearly flat curve is the honest reading of a transparent vendor, MSCI with the fund launch basis point line beside the contested Analytics and ESG curve, Clarivate against its documented mandatory uplift posture, Ideagen in the revalidation moat's consolidation window, Conga anchored at its documented 36 to 59 percent list to landed gap, Nintex on the workflow count meter, Hyland against its own public OMNIA pricelist, Morningstar and PitchBook on one shared curve negotiated with the vendor's own investor relations statements, the observability pipeline Cribl by daily data processed, RingCentral by seats in the deflating UCaaS market, Sumo Logic by daily ingest with its printed renewal uplift struck at signature, PagerDuty and its add on ladder, Shopify Plus by annual GMV against the unpublished revenue share conversion, the travel and expense set Navan and Emburse, the interchange pair Brex and Ramp by annual card volume, and the consumption storage pair Pure Storage and NetApp by committed capacity, the supply chain set SPS Commerce by trading partners, E2open, the planning duopoly Kinaxis and o9 on one shared curve, Manhattan Associates and Blue Yonder, six ERP curves re benchmarked on annual contract value, NetSuite, Deltek, IFS, Epicor, Infor and OpenText, the core banking pair Fiserv and Jack Henry on one shared curve, Oracle MySQL by subscribed servers, the proxy advisory duopoly ISS and Glass Lewis on one shared curve, Nasdaq Adenza, Broadridge's unregulated services layer, nCino, and the open source repricing set: SUSE benchmarked by deal posture, Red Hat with its estate engineering line, MongoDB re benchmarked on its documented commit bands, Azul by Java server estate with the metric negotiated before the price, Grafana Labs with the OSS self host path as the structural leverage, GitLab in the Ultimate and Duo era with three meters on one developer, the revenue stack pairs Outreach and Salesloft, Gong and Clari, and Braze and Iterable each on its own documented line, Yardi completing the property duopoly by portfolio units, Jamf by managed devices against the Intune shadow, Intercom and Fin by monthly resolutions on the purest per outcome AI meter, the governance duopoly Collibra and Alation, Workato at its documented 35 to 50 percent list to landed gap, and the superseded Informatica and MuleSoft calibrations re benchmarked onto the IPU and Salesforce estate era, the industrial token meters AspenTech and AVEVA, the password duopoly 1Password and LastPass on one shared curve, and Akamai, Wiz, Medallia and F5 re benchmarked on annual contract value, and the wave twenty seven money meters: Stripe by annual card volume with the Adyen counter quote as the upper line, Adyen by processing volume with the 0.60 percent markup as the only line that negotiates, Experian and TransUnion on one shared bureau curve by annual pull volume with the 2026 FICO repricing and the VantageScore counter in the file, LexisNexis Risk Solutions by transaction volume split between the contested identity lines and the claims data monopoly, the publishing duopoly Springer Nature and Elsevier priced against the published DEAL benchmark and the documented UC, SUNY and DEAL exits, Plaid on the per login item meter with dormancy hygiene as the free discount, MathWorks with the toolbox prerequisite chain audit, Cotality decomposed per dataset under PE ownership, ICE Mortgage Technology by annual loan volume with the FTC mandated alternatives, Addepar by assets on platform against its documented rival gaps, Clearwater Analytics by assets under administration where the breakpoint schedule is the negotiation, SS&C on the per heritage roll up map, and the superseded SAS calibration re benchmarked onto contract value bands with renewal and migration event lines, and the AI coding tools, Cursor, Windsurf, GitHub Copilot with its three purchasing pathways, Replit banded by monthly agent spend, v0 by Vercel, and Lovable banded by monthly credits, and the wave twenty nine consolidation files: Verint and Calabrio on one agent count curve now that Calabrio operates inside Verint's platform, Seismic and Highspot on one seat curve carrying the two paddle era's transacted medians ahead of the February 2026 Permira merger, the enterprise low code tier of Appian, OutSystems and Mendix on one curve with the Power Apps shadow as the structural frame, Everbridge by contact count where the roster is the meter, UserTesting on its credit bundle economics, Grammarly on the shadow IT to enterprise conversion, Canva carrying the reversed 2024 shock repricing and its referenceable Pricing Promise, Airtable with the 84 percent seat architecture case, and the superseded Box calibration re benchmarked onto the AI repricing era's wider seat bands, and the wave thirty AI model provider set on one shared set of contract value bands: Microsoft 365 Copilot against its documented $18 to 22 landed band, Google Gemini and Vertex AI as one Google AI estate with the deepest curve in the provider set (the Gemini tool moves off the Workspace family onto its own curve), xAI Grok as the youngest and most volatile paper in the library, Perplexity with its 8x tier gap read as a negotiation signal, Cohere where the deployment tier decision moves costs 10x before any rate conversation, and Mistral AI whose own open weights cap its managed pricing), a panel above the tool where you choose the commercial model or product (for SAP: RISE Base, RISE Premium, BTP CPEA, BTP BTPEA, GROW, or the on premises reference; for Oracle Java: the standard line or the documented OpenJDK exit line; for Salesforce: the core clouds, the multi cloud commitment, or the flat renewal line) and the deal size band, and type the discount you are offered. The bands follow how each vendor actually prices: contract value for most, seat counts for Adobe Creative Cloud, Zoom, Google Workspace, GitLab, Gong and Box, protected endpoints for CrowdStrike and SentinelOne, protected workloads for Veeam and Wiz, protected capacity in TB for Rubrik, Cohesity and Commvault, agent counts for the contact center platforms, construction volume for Procore, user counts for Epicor and Deltek, platform fee for Palantir, privileged accounts for CyberArk, employee counts for UKG and ADP, production vCores for MuleSoft, hospital bed count for Epic, Oracle Health and Meditech, monthly claim volume for Waystar, net patient revenue for R1 RCM, annual verification volume for the Equifax Work Number, licensed seats for Sigma and ThoughtSpot, trading partner connections for SPS Commerce, subscribed servers for Oracle MySQL, licensed user count for Veeva, transaction volume for Vertex and Avalara, seat counts for Notion, Miro and Figma, licensed seats for ZoomInfo, workers on platform for Deel, employee headcount for Paycom, Paylocity, iCIMS, Cornerstone and Dayforce, monthly invoice and payment volume for Tipalti and BILL, annual invoice volume for Basware, practice size in providers for athenahealth and NextGen, user counts for BlackLine, QAD, Unit4 and Aptean, annual committed spend for AWS, Google Cloud, Databricks, Snowflake and MongoDB Atlas, licensed seats for Nemetschek, Bluebeam and Trimble, daily data volume for Cribl and Sumo Logic, Java server estate for Azul, portfolio units for Yardi, managed devices for Jamf, monthly Fin resolutions for Intercom, annual GMV for Shopify Plus, annual card volume for Brex, Ramp and Stripe, annual processing volume for Adyen, annual pull volume for Experian and TransUnion, annual transaction volume for LexisNexis Risk Solutions, annual loan volume for ICE Mortgage Technology, assets on platform for Addepar, assets under administration for Clearwater Analytics, contact counts for Everbridge, agent counts for Verint and Calabrio, seat counts for Seismic, Highspot, Grammarly, Canva, Airtable and the re banded Box, committed capacity in TiB for Pure Storage and NetApp, seat counts for RingCentral, years since migration for the Atlassian loyalty decay curve, and estate size for Citrix and TIBCO, the program's two uplift curves under Cloud Software Group: their docks read the renewal increase on your quote rather than a discount, lower is better, and the rank inverts accordingly. The athenahealth dock is the library's first rate curve: it reads the negotiated percent of collections by practice size, lower is better like the uplift curves, and prices every point of the rate against your own annual collections. It answers the three questions in order: your discount, the p05 to p95 range comparable customers land with its midpoint, and the rank in words (Poor, Below market, At market, Good, or Top of the market), priced per year when you give the deal value. The full curve for the chosen product renders underneath, and the Microsoft and Google Workspace docks add a pricing era switch, covering the November 2025 removal of programmatic EA discounts and the January 2025 Gemini repricing.

- **Your deal panel** — the inputs on the left: currency, your net fee, unit count, an optional list value for the discount reading, region, industry, and any deal-structure levers such as term length or competitive presence. The verdict updates live as you type. The unit is the one the vendor actually licenses, named on the page throughout: FUE for SAP S/4HANA and RISE, named users for SAP ECC, cores for VMware, consultant hours for a rate benchmark. When your workspace holds exactly one agreement for the tool's vendor, a line above the form offers its stored annual value and renewal date: one click fills the matching fields, everything stays editable, and sample data never prefills.

- **Product** — where a vendor sells several product lines with genuinely different economics, a Product selector sits at the top of the form: OpenAI splits into API committed spend and ChatGPT Enterprise seats, Anthropic into the API and Claude Enterprise seats, Cisco into EA software, hardware off the Global Price List, Meraki, Webex, and Splunk, and the Broadcom estate into VCF, VVF, standalone vSphere subscriptions, and CA mainframe. Switching products opens that product's own calibrated dataset; it is never a filter on the current one.
- **Deal size** — a one-click way to read the market for a size before you have a quote. Picking a band (for example $250k to 1m) fills a typical net value for that size into the form; you can refine the exact figure at any time, and everything downstream updates as if you had typed it.
- **The market discount curve** — on every discount benchmark, a small table under the inputs shows the median discount and the typical range for each deal size band, your band highlighted in gold. Once your deal is entered, your own discount is placed inside your band's row, so the first read is always the three answers in order: what you are offered, what comparable customers get at your size, and where that ranks.
- **Open the case** — the navy band at the top of the results. It opens your placed verdict as a full screen benchmark case of four executive brief documents, every figure computed from your live position. Your deal sheet is two pages: the order form facsimile with Vera's markings and margin pins on every placed term (risk pins in brick), then a reading ledger (field, what the paper says, what it means for you, a verdict pill) and the deep dive panels on the clauses that decide the deal. The verdict brief runs three pages: the position summary, a commercial journey chart from list paper to your paper to the close target, a peer scatter drawn from the real cohort with your gold ringed dot called out, section by section verdict words (Ahead, At market, Behind, Exposed), a numbered How the target gets achieved sequence, a pull quote, and the bottom line. The field atlas shows all four instruments in sequence: the range bars, the percentile ladder with your YOU cell placed among the cohort's quantiles, the captured versus defended dollar stack whose totals reconcile with the verdict line, and the full verdict table. The ask sheet carries eight to ten sequenced asks (anchors first, protections second, the rate once the table is set, timing last), each with its clause anchor, its priced value, a Say this script line, and an If they push back counter. Press any figure, tag, or chip inside the case and it jumps to the exact clause the number was read from; Play the read walks the deal sheet one lit term at a time. The case recomposes from your live position every time you open it, levers included.
- **The workroom dock and rail** — the icon strip on the case's right edge and the finding blocks under the document list. Press ⚡ and the paper slides left to make room for the What now panel beside it, so you can read the verdict brief and fire the moves at the same time; press ✦ and the same split opens Ask Vera, grounded on the whole case, with one chip per document. Click a chip and the cited document opens on the stage. Questions asked here land in the same saved conversation every time you reopen this case. Esc unwinds one layer at a time: the open dock panel, then the open room, then the case. Below the document list the rail carries Find a case (search every vendor benchmark by name, recent cases first; a

recent reopens with its case on the stage and your numbers prefilled, and a vendor you have not benchmarked yet lands on its page so you can enter your deal, because a case is only ever built on figures you supplied), a Research search over your own uploaded documents, and your shelf: the last benchmark cases you worked, under Recent and Starred tabs, each reopening its page with your numbers prefilled, plus a Star this workroom button for the open case.

- **The case rooms** — the pills in the top band beside the case title: Benchmark case, Compare room, Deal room, Second opinion. They swap what is on the stage without closing the case. The chrome, the document rail, the dock and every figure stay where they are, so nothing is lost and the way back is the Benchmark case pill, any document in the rail, the back button at the foot of the room, or Esc. On a narrow window the same pills ride the end of the document strip.
- **Compare room (in the case)** — your saved positions for this vendor, listed newest first. Tick up to three and press Replay them together: each one is replayed alongside the position on the page now, through the same peer-cohort math against today's data rather than the standing stored at save time. You get the scenario verdict (which position is strongest, the gap in discount points or fee per unit, and what that gap is worth per year at your volume), a standing band with every position dotted on it and the strongest in gold, and the table underneath: metric value, standing, cohort median, top-quartile target, annual headroom to median, and the save date. With no saved positions yet, the room offers to save the one on the page so the next quote has something to land against.

- **Deal room (in the case)** — what the Vera Deal-Room adds to this benchmark: it reads the contract itself and assembles the decode, the vendor intel, a credible alternative, the brief and the playbook around it. Below that sit your existing deal rooms, this vendor's first and flagged, each opening at its own stage, and one control to hand in a contract and start a new one.
- **Second opinion (in the case)** — the same shape for the four page verdict brief: what a second opinion adds to this benchmark, your existing opinions with their verdict lines and this vendor's first, and one control to run a new one. Both of these read a contract and take a few minutes, so they run as their own jobs; this case stays on your shelf, one click from where you left it.
- **What now** — the case file's sixth document, and the answer to "what do I do next with this benchmark?". The same moves also ride the ⚡ dock panel beside every page of the case. Twenty plus moves in five families, every one running on the case's own numbers. Send the word: Vera drafts six emails you edit and send from your own inbox, to the vendor account team (opening the negotiation, or putting your asks on the table), your CFO, procurement, your reseller, and your own team; the vendor facing drafts never reveal your percentile, your cohort, or where the numbers come from. Every draft is written against what is known about how that particular publisher negotiates, where it concedes, and where its fiscal year sits, so an Oracle email and a Salesforce email do not read the same. Put it in front of people: copy a chat ready summary or the one line verdict, mint a 30 day share link that recomputes as the market moves, or copy a prefilled page link for a workspace colleague. Turn it into paper: generate the three negotiation toolkit briefs as executive brief PDFs, or download the one page position snapshot. A brief is written from scratch against your position and takes about three minutes, so the move shows a progress meter with the

seconds elapsed and the usual duration while it works; the estimate tightens against your own past runs, and you can keep reading the case while it finishes. Keep it working: save the position as a named scenario, with or without market alerts, set the renewal date, or jump back to the levers. Go deeper: Ask Vera with the case loaded as the question, compare vendors side by side, open a full negotiation workup, or book an analyst review with the agenda prefilled.

- **The email desk** — what opens when you pick one of the six drafts. Three controls sit around the draft. **The angle** is a row of buttons across the top: the reason the email is making its case. The vendor facing drafts offer competitive pressure (the default, which names an alternative you could credibly move to, says the price is above market for a customer your size, and asks for a competitive offer before you go further), budget pressure (the line is capped and today's number does not fit, no alternative named and no target stated), and market price (you know what this goes for at your size and the next term has to reflect it). The internal drafts have their own three: your CFO email can ask for the decision, lead with the money, or lead with the risk of renewing untouched; procurement can take the frame, the money, or the timing; your team note can route contact through you, be a light heads up, or ask for the usage evidence. Pressing an angle rewrites the whole email around that premise. **Ask Vera for a rewrite** is the field under the draft: say what to change in your own words ("shorter and firmer", "name our March renewal", "drop the last paragraph") and she rewrites what is on screen, so edits you already made by hand survive; leave it empty for a fresh take on the same angle. The alternative named in a competitive draft is researched rather than assumed, from your own alternatives analysis when you have one and a live search when you do not, and when nothing can be verified the email refers to your alternatives in general terms rather than naming a company. Nothing sends from here; you copy it or open it in your own email app.
- **The conversation log** — under every vendor facing draft: save this email and it becomes the opening move on a negotiation track for that vendor, which is where their reply will land. If a track is already open the panel says how many messages are on it and the draft files onto the same thread rather than starting a second one. Closing the window with an unsaved draft asks once before it lets the ask go. When they answer, paste their reply into the counter desk (**The vendor replied? Paste it, get your response**): Vera reads the message against your position *and* against everything already on the log, so a number they are quietly repeating or a concession they are walking back is visible, then tells you what the account team is doing, what to hold, and drafts the response, which never reveals your percentile, your cohort, or where the numbers come from. **Save it to the log** keeps the whole round: their message, your response, and her read of it. **The conversation log** move opens the tracks themselves. If you would rather work the thread from your own inbox, the desk links to the ghost writer, where you forward vendor mail to a private agent address instead of pasting.
- **Where this lands** — a percentile band pinned under the inputs showing the middle 50% of your peer set with your deal marked, and the verdict pill (favorable, mid, or above market).
- **The brief** — a placed deal reads as a six section decision brief. 01 The verdict: the document masthead opens with the plain call in one colored line (a good deal, a fair deal, or you're overpaying) and a one-sentence reading, then the judgment sentence, the stat band, and the market field with your deal as the gold dot. Beneath it, the recommendation: the play in one bold line (what to open at, what to settle at, the floor you will not sign above) followed by the moves as a short numbered list, all written from the same targets and clock every instrument runs on. 02 Your market, your numbers: the region and industry

selects on the left re-cohort every figure for real, and on discount benchmarks a list-value dial lets you test how every discount figure reads if the list anchor moves, your net spend staying fixed, with your standing recomputed live; below it, every comparable deal plotted by deal size. That chart reveals itself the first time it scrolls into view: the comparable deals light up one by one, a gold trend line sweeps through the field, and your dot lands with a gold pulse; after that, every change renders instantly. 03 The deal snapshot: your position against the benchmark supported one, dimension by dimension. 04 The ask: the three targets (defensible floor, strong close, stretch anchor), each priced, with the savings bridge printing what you pay today against the target and the open. 05 Timing and the plan: the engagement roadmap, one card that opens with the four numbers on the clock (the vendor's quarter close, the open window, your notice deadline, and the renewal, each with its days remaining), then three phase lanes (prepare, engage, paper it) on a month ruled time axis with every deadline dropped through the lanes, and closes with the three moves in the same phase colors, each composed from your own figures. Add your renewal date on the card and the notice deadline joins the roadmap. 06 What next: Ask Vera on one side, and the Take it forward console on the other, the page's one dark panel, with the recommended next step first and every further path numbered beneath it. Every instrument carries a short "what this is, how to read it" note, and hovering any figure offers an enlarged view with the same reading.

- **The brief and the working file** — the toggle beside the verdict chapter. The brief (the default) is the ninety second read above. The working file adds every remaining instrument at the end of the page: the peer comparison suite, the placement curve, the deal played forward, region and signing-year medians, the deal-structure levers, the should-cost waterfall, your leverage scorecard, the pushback counters, and the scenario theatre where you drag seats, term, and price and the totals reflow. The working file only adds; nothing in the brief changes.
- **Saved scenarios** — name a position and save it to track it over time, feed the Portfolio benchmark roll-up, and get alerted when that vendor's market moves. Two or more saved positions on the same tool compare side by side on one standing band. Comparison is per tool: a Microsoft EA position and an Oracle ULA position are measured on different metrics and are never placed on the same band.
- **The five flagship tools save the same way** — Microsoft EA, Oracle ULA, Oracle Technology, Sales and Service Cloud, and Oracle Java SE each carry their own dataset and their own math, and each now carries the same **Saved scenarios** panel. Name the position, press **Save**, and it joins the roll-up like every other vendor. The figures are re-run on the server when you save, so a saved position is the tool's own math rather than a number a browser sent. Saving is only offered once the tool has placed your deal against a large enough peer cohort, and a workspace can hold up to 200 saved scenarios.
- **The money still on the table** — the flagship tools close on the annual gap between your position and the market median, then the moves that go after it: start the negotiation carrying these numbers, compare your saved positions, open the roll-up, open the vendor room, get a second opinion on the actual order form, copy the one line verdict, ask Vera what she would do, or book the analyst walkthrough.
- **Action bar** — download the placement as a PDF, or create a share link (valid 30 days, optionally kept live so it recomputes against current data each time it is opened).

How to use it

1. Find your vendor from the hub cards, the search box, or the Ctrl/⌘ K palette, and open its page.

2. Enter your deal on the left. The verdict, targets, and play light up as soon as your deal is placed against a deep enough cohort.
3. Read the play and take the target and counters into the negotiation. Save the scenario or share the result with your team.

Tips

- When too few comparable deals match your position to rank it honestly, the page says so instead of showing a noisy percentile. Widen the region to all regions, or check the deal size.
- Pick your currency at the top of the deal panel and the whole page follows it: you enter euros or pounds and every figure on the page, in the case documents, and in the PDF export reads back in the same currency. The reference deals are held in US dollars and converted at a fixed rate, printed under the selector with the date it was set, so a saved position replays identically rather than drifting with a live feed.
- For vendors whose discount scales with deal size, such as Oracle cloud, your deal is ranked inside its own value tier, so a small deal is not judged against whale discounts and a large deal is held to the deeper discount its size should command.
- An instant check is a fast read you run yourself. For a full, hand-verified analysis of a specific quote or contract, upload it for analyst benchmarking; the vendor page links straight into the upload wizard with your vendor and deal value prefilled.

Related: Instant benchmarks · Proposals · Portfolio benchmark

---

Next: Chapter 03 · Tooling & playbooks

03 · Tooling & playbooks

Instant benchmarks answer a few questions about one deal. The tools in this chapter go further: they read your own exports and documents, a usage report, a cost export, a signed contract, a photographed quote, and build a plan from your real numbers. The chapter closes with the vendor playbooks, full self-serve negotiation and migration engagements for ten major vendors.

Nearly every tool is free to explore with sample data on any plan; the few that only make sense on your own inputs, such as Quote Face-Off and Negotiation self play, say so on their pages instead of loading a sample. Uploading your own data, and running the AI over it, requires a Professional or Enterprise plan, and each AI run counts against your plan allowance; the tool tells you when a step needs it, and Plans, allowances & limits defines every gate. Every contract-family document you hand an intake tool, the three decoders, the Verdict Bar, Second opinion, the Vera Deal-Room, Proposal comparison, and the vendor playbook drop zones, is kept: it saves to your workspace as a contract, stored privately and visible only to your organization, and joins your benchmarking queue automatically. Benchmarking extracts only high level data points, the vendor, the price, the term, and the key commercial terms; the document itself never leaves your organization, every saved document can be deleted from its contract record at any time, and an owner can turn the saving off entirely under Settings, Security & compliance, Data controls. Three more tools now say plainly which side of that line each of their file types falls on, because they take more than one kind: the Time machine keeps **both** documents; Quote Face-Off keeps **PDF quotes** and never stores a photo capture; the Rate Card Benchmark keeps **PDF and DOCX** documents and never stores an **XLSX** rate card, because a spreadsheet of rates is a data file rather than a contract. The same owner switch under Settings, Security & compliance, Data controls governs all of them, and turning it off means those tools read in memory and keep nothing. Phone captures, invoices, audit letters, and spend exports keep the old rule, read in memory and never stored, and the browser-only calculators parse your export on the page so your usage data never leaves it.

Each tool section below carries a **You get** line: the artifact the tool produces and roughly how long a run takes, so you know what you are waiting for before you start. In the product, every tool page opens with the same strip under its title: when to use it ("Use this when …"), the typical finding it surfaces, its data-privacy posture, and, where the tool ships a sample, a "Runs on sample data before any upload" note, so you can judge a tool before you feed it anything. Every tool page also carries an **All tools** link at its top left, one click back to the tool menu from anywhere in this chapter.

Every chart enlarges, and so do the playbook cards. This is platform wide, not just in the tools: hover any chart on any screen and an **Enlarge** button appears in its corner; click it, or the chart itself, and the figure opens in a near full screen window with its title, its description, and a "How to read this" note. Hovering a bar, ribbon, dot, or segment, on the page or in the enlarged view, shows the exact figure behind it in a tooltip the moment the pointer touches it. In the playbooks the text windows enlarge too: click anywhere in a card or KPI tile that is not a button or a field and it opens in the same window with the text scaled up for reading; selecting text never triggers it. Press Esc or click outside to close the enlarged view. Charts and tiles that are themselves links, such as a dashboard tile, keep navigating on click.

ALL TOOLS

Where: Tooling › All tools · `/tooling`
Who: everyone

The tooling hub: every upload-driven optimizer and calculator on one wall, grouped by job. Start here when you have your own data and want a plan built from it rather than a quick market check.

On the screen

- **What do you want to do?** The hub opens with a band of six doors: Portfolio optimization, Contract intelligence, Renewal playbooks, License optimization, Cloud spend, and AI spend. Click a door and the wall filters to that group, with filter pills (All instruments plus the groups) appearing so you can switch without scrolling back up.
- **Your team used recently**: a chip row of the tools your team has opened before, shown once the team has some history, so the instruments you actually use stay one click away.
- **Verdict Bar hero**: the dropper fronting the Contract intelligence group. Not sure which tool fits? Drop the document here and it routes itself (see below).
- **Instrument cards**: one card per tool, grouped under the same six groups as the doors. Each card names the situation it is for and the typical finding it surfaces.
- **Pointers between tools**: overlapping instruments point at each other: Vera Deal-Room carries a Most complete badge, Snap a quote points at Quote Face-Off, and each optimizer cross-links its vendor playbook. Follow the quiet line on a card when it fits your situation better.
- **Start here**: the spend optimization tool is featured as the first stop for a budget owner.
- **Narrated AI reports**: generating an AI report anywhere on the wall shows what the analyst engine is doing, step by step with a timer.

How to use it

1. Start from the doors: pick the job you came to do and the wall narrows to the tools that do it.
2. If you know the tool you need, open its card, or pick it from the recently-used chips.
3. If you only have a document, use the Verdict Bar dropper and let it pick the engine.
4. Try any tool on sample data first to see the output before you bring your own.

TOOL FINDER

Where: Tooling › Tool finder · `/tool-finder`
Who: everyone

Say what you are trying to do and open the tool that does it. The finder lists every tool on the platform under the job it does, with the situations it answers phrased the way you would say them, and every result ends in a Start button.

On the screen

- **Job dropdown**: the fifteen jobs the platform does, from checking a price to handling an audit. Picking one narrows the list to that job's tools.
- **Search box**: type your situation in your own words, for example audit letter or renewal quote. The list narrows as you type, and the situations that match your words are bolded on each card.
- **Result cards**: each card names the tool, says what it produces, lists the situations it answers, and ends in a **Start** button that opens the tool directly.
- **Count line**: how many of the tools currently match your job and search.

How to use it

1. Pick the job from the dropdown, or type your situation in your own words.
2. Read the matching cards and find the one whose situations sound like yours.
3. Press **Start** to open the tool.
4. Not sure how to phrase it? Ask Vera (Ctrl or Cmd J) instead; she names the one tool that fits and links it.

Tips

- The same index behind this screen also powers Vera's routing and Chapter 14, the printed "Which tool do I use?" reference, so all three always agree.

VERDICT BAR

Where: Tooling › Verdict Bar · `/verdict`
Who: requires an AI plan
You get: the document classified, routed to the right engine, and judged, with red flags and the next step · about a minute

In plain words: drop any document; the right engine finds it and judges it.

When to reach for this

- A document landed and you do not know, or care, which of the thirty tools fits it.
- You want the one-minute read before deciding whether the deal deserves a deeper workup.
- A teammate asks where to start: this is always a correct answer.

The one front door when you have a document and no idea which tool fits. Drop any quote, contract, order form, or invoice; it is classified, routed to the matching engine, and judged against the market in about a minute. The screen follows the same shape as Contract decode: a drop stage, a full screen read, and the settled verdict as a working room.

On the screen

- **The drop stage**: the navy panel at the top. Drop the file anywhere on it or press Get the verdict to pick one. The optional vendor field tells the engine who the counterparty is when the document does not say so clearly. Under the panel: the accepted types (PDF and DOCX to 30 MB, photos and screenshots to 3.5 MB) and how many verdicts your organization has left today.

- **Samples**: two stored runs, a Microsoft EA quote and a ServiceNow renewal contract, that play the whole show without an upload. Useful before you bring a real document. Sample runs are labelled on screen and never touch your allowance.
- **The read**: once a file lands, the stage takes over the screen. A percentage and a gold bar track progress, the document turns under a light beam, and the engine feed prints what it is doing line by line: what it detected, what it classified the document as and how confident it is, which engine it routed to, what it extracted, and every flag with the money attached.
- **The reveal**: when the verdict is in, the stage reveals it as an executive brief: the one line judgment, the scorecard, the evidence the engine read, what the paper is doing to you, the clauses already in your favor, and where the deal sits in the market. Press Skip to jump to the end, or Work the verdict to step onto the page.
- **The verdict page**: the run bar (vendor, file, date, with Drop another document and Replay the reveal), then the verdict itself as one executive paper, the same format a Second Opinion delivers: the one line verdict as the headline, an executive summary, the stat tiles, and numbered sections carrying everything the engine read. On a quote that is the deal on the table with every line item checked against the published price book, your placement against the peer set with the counter priced, and the flags on the paper with the exact clause quoted; on a contract, where you stand, the changes worth asking for with paste ready language, the clauses worth protecting, and the terms against the market; on an invoice, the reconciliation and where the dispute drafts live. The paper's fine print states how the document was classified, at what confidence, and what the judgment was grounded in.
- **The docked verdict bar**: fixed at the foot of a settled verdict, the Second Opinion pattern: the one line verdict always in reach, Copy the one-line verdict, and Download PDF, which prints the clean brief alone.
- **The verdict room**: the navy band opens the whole verdict as a full screen case file: numbered executive brief papers in a grouped left rail, the arrow keys page, a gold meter tracks how much you have read, and Esc closes. A contract verdict holds the cover sheet, the verdict brief, the ask sheet and the benchmark table; an invoice verdict holds the cover sheet and the handoff into Invoices; a quote verdict opens its own case with up to six papers including the quote as parsed, the dollar bridge to the computed counter, and a priced counter sheet. The right hand dock carries What now, the moves drafted from this run's own findings (the redline reply, the vendor question list, the legal briefing, the copyable ask list), and Ask Vera grounded on the whole verdict. On a quote case those moves are the requote draft with the counter and every marked passage, the chat ready summary, the one line verdict, the ask list as a meeting agenda, and the handoffs into the deal room, the vendor room and the renewal calendar; under the document list sit a search over your own uploaded documents and your shelf of recent workrooms. Drafts open for you to copy or edit; nothing sends from a room.
- **What next**: every verdict ends with its next moves. A decoded contract carries its landmines onward: copy the full ask list with the paste ready replacement language, open Vera Deal-Room on the same file, or start the negotiation; a judged quote leads into the deal room and an invoice into the dispute drafts. Ask Vera and the analyst booking ride along, preloaded with the verdict.

- **Saved to your workspace**: when the document routed as a contract, a note under the verdict says where it went: it was filed under Contracts, private to your organization, and joined your benchmarking queue. The note links to the contract record and carries a delete that removes the stored document on the spot; quotes, photos, and invoices are read in memory and never stored. Owners can turn document saving off under Data controls.
- **Your verdicts**: past runs stand as spines on a shelf, gold banded once the verdict is in. Switch to the table view for the document, what it was classified as, its headline figure, status, and date.

How to use it

1. If you have no document handy, press one of the samples and watch a full verdict play out.
2. Drop your own document on the stage, or press Get the verdict and pick the file.
3. Watch the read: it classifies, routes, and flags as it goes.
4. When the reveal ends, press Work the verdict.
5. Press Open the room to work the papers, or take a move from What next.

Related: Contract decode · Snap a quote · Your first verdict

SPEND OPTIMIZATION

Where: Tooling › Spend optimization · `/tooling/spend-optimization`
Who: requires an AI plan
You get: a five to ten page savings plan of about twenty prioritized moves · two to four minutes

In plain words: the whole budget's savings plan, about twenty prioritized moves, from your own agreements.

When to reach for this

- Budget season: the run-rate reduction plan, written and prioritized, before finance asks for one.
- After the estate lands: the first "where is the money" pass across every vendor at once.
- When single-vendor wins stop compounding: this is the portfolio view of the same work.

The plan's Theatre, Share, Forward, Word, and print controls sit in one toolbar row above the plan, alongside its version history; the generate card holds only Generate and Regenerate. Exporting or sharing a plan already on file keeps working even when AI generation is unavailable.

The savings plan for your whole software budget. Rank your top vendors by annual spend, see every renewal date on one calendar, and get an analyst-grade plan of about twenty prioritized moves to lower the run rate, written by Claude Opus 5 and checked against your own agreements. Once a plan is on file, the tool opens on stage: the plan presents itself as a full screen executive brief, and Escape steps down to the classic desk described below.

The brief on stage (the front door once a plan exists)

- **Seven sections, walked like a briefing**: the executive brief (a counted opening with the verdict, then a key findings page where every line is computed from your own agreements), the portfolio (the spend map with the five largest vendors beside it, the concentration pareto), the renewal calendar (the eighteen month horizon with flags on every vendor renewing blind, then a countdown in days to the next open window), the savings opportunity, the action plan (the move deck, the priority board, the 90 day sprint, the roadmap lanes), the decision (the targets, the bottom line beside the sequence), and execution (the what next cards and the project plan). Arrow keys walk the brief; G opens the overview grid.
- **The agenda adapts to your data**: when benchmarks have priced headroom, the savings section carries the run rate math with its per vendor table, the headroom board seating every analyzed vendor, and value in play with the vendor by vendor status; when they have not, one honest benchmark coverage screen shows where the signal stands per vendor and what opens the money up, and the header ticker tracks the spend renewing inside 12 months instead of zeros.
- **The priority board**: every move in one of three effort lanes (quick wins, standard plays, structural), the money first inside each lane, with its triage state on the card; click a move to open it in the deck.
- **The spotlight**: click any treemap tile, headroom bar, or target card and that vendor's dossier slides in: spend, share, renewal window, benchmark signal, the computed reasons it is worth a fight, and its moves with triage buttons.
- **Triage from the stage**: the deck shows one move at a time with accept, done, and dismiss; verdicts persist to the plan and the header ticker reprices headroom in play live, with delta chips on every change.
- **The workspace on stage**: Share mints the 30 day read only link, Forward emails the plan with the verdict figures, Word downloads the export, the envelope has Vera draft the follow up email, and the leave behind prints exactly the screens you walked, all without leaving the stage.
- **Autoplay**: press A or the header pill and the show walks itself one screen at a time; press again to take the wheel back.
- **Presenter kit**: presenter notes (N), gold ink over any screen (I), fullscreen (F), and light or dark stage lighting, remembered per user.

On the screen (the desk, under the theatre)

- **Vendor picker**: take your top 10, 25, 50, or 100 vendors by annual spend, or choose them yourself.
- **The board**: spend per vendor, every renewal on one eighteen month calendar, and the market position wherever the data defends one.
- **The portfolio scan**: generating opens a full screen scan: your analyzed spend as a board of vendor blocks coloring in by renewal urgency, an honest progress figure over the two to four minute composition, and a live ticker narrating each vendor's spend, agreements, and renewal date. When the plan lands, an executive readout counts up the verdict: spend analyzed, benchmark headroom, renewals inside 12 months, and the opening moves. Close it any time; the plan still lands on the page.
- **The plan**: a five to ten page optimization plan of about twenty recommendations across right-sizing, renegotiation, consolidation, terms, and timing, each anchored to a renewal date. When benchmarks place headroom, page one leads with the savings waterfall from analyzed spend down to the optimized run rate.

- **The treemap**: the vendor board opens as tiles: size is annual spend, color is renewal urgency. In the pick-vendors scope, clicking a tile toggles that vendor in or out of the analysis.
- **The briefing case**: the finished plan splits into five documents drawn as first-page miniatures: the C-level brief on where the money is, the vendor targeting memo with the evidence per target, the full move list, the renewal calendar, and the phased roadmap. Click one and it zooms open into a reading room with a document rail, arrow-key paging, read state, and gold chips that jump to the evidence behind a figure.
- **Move triage**: every move in the walkthrough takes a verdict: accept, mark done, or dismiss. Verdicts save to the plan itself, so the whole team sees them and they survive reloads; press a verdict again to clear it.
- **The tracker**: a live band above the plan recomputes as you triage: headroom in play, headroom engaged (an accepted or done move on the vendor), and headroom dropped (every move on the vendor dismissed). Every dollar is the vendor's benchmark headroom from your own deals.
- **Open the workroom**: a navy band above the tracker opens the plan as a six document workroom on a dark stage: the cover sheet, the spend ledger as the exhibit, the optimization brief, every move priced by family in the action atlas, the sequenced plan sheet with a first step per move, and the What now sheet. Under the document list sit a search over your own uploaded documents and your shelf of recent workrooms. Two panels dock on the right: ⚡ **What now** drafts the note to the budget owner (analyzed spend, identified headroom, where the savings sit, and the first three moves with their deadlines, figures included) and the opener to the lead vendor (the entitlement, the usage, and a line by line renewal quote, stating plainly that pricing is under review and never quoting your headroom), copies a chat ready summary, the one line plan and the first moves, and opens the renewal calendar, the Deal Room, your two highest priority vendor rooms, and the analyst booking. ✦ **Ask Vera** answers from this plan alone: the ledger, the measured headroom, the renewal clock, and the moves in plan order. The sample plan carries the same band, so the room is visible before you upload anything.
- **Groundedness**: every figure is checked against your stored agreements before the plan is shown.
- **Share and Forward**: Share mints a 30 day read only link that always shows the current plan (copy it, or revoke every link in one click). Forward sends the plan to a colleague as an executive brief email with the verdict figures, the opening moves, and the same kind of link. No account needed on their side.

How to use it

1. Pick your vendor set by spend, or hand-pick it (the treemap tiles toggle vendors in the pick-vendors scope).
2. Generate the plan and watch the portfolio scan while Opus composes it. This takes two to four minutes.
3. Return with a plan on file and the brief opens by itself: walk the sections with the arrow keys, put vendors under the spotlight, and triage moves from the deck. Escape lands on the desk; the Theatre button raises the stage again and offers to resume where you left off.
4. On the desk, walk the plan move by move and give each a verdict: accept, mark done, or dismiss. The tracker above the plan shows the value in play as you work.
5. Press **Open the workroom** and work the What now panel: draft the note to the budget owner, open the lead renegotiation, and put the renewal dates on the calendar.
6. Take any move deeper with Vera, or book the analyst walkthrough; a dot on the renewal calendar jumps to that vendor's move.

7. Share or forward the plan to the budget owner, or export it to Word or PDF, from the stage or the desk.

Related: Portfolio benchmark · Opportunities

SPEND INTELLIGENCE

Where: Tooling › Spend Intelligence · `/tooling/spend-intel`
Who: everyone · metered by plan
You get: a classified spend report with duplicate and price variance findings, a six month forecast, and an analyst narrative · two to four minutes

In plain words: your raw spend export, classified, checked for duplicates and price creep, and forecast.

When to reach for this

- The AP export exists and nobody has ever read it as a whole.
- Duplicate payments and quiet unit-price steps hide in exactly this data.
- Finance wants a six month forward view with buying windows marked.

Upload a raw AP or ERP spend export (CSV) and get the whole ledger read back. Every supplier lands in a 20 category spend taxonomy, two deterministic detectors work your own invoice history for duplicate payments and unit price steps, the monthly trend and seasonality project the next six months with buying windows marked, and the analyst narrative says what to chase first. The file is analyzed in memory and never stored; the report saves to your workspace.

On the screen

- **Upload stage** — drop the CSV or browse for it. Vendor, amount, and date columns are required (common column names map automatically); invoice references, memos, GL accounts, SKUs, and unit prices sharpen the detectors when present. Up to 8 MB and 20,000 rows.
- **Sample ledger** — the **See it on the sample ledger** button renders the full report instantly on a fictional 18 month, 41 vendor ledger, no upload and no allowance used. The **CSV template** link downloads a miniature export in the same format, with a planted duplicate pair and a planted price step, so you can watch the detectors fire on it.
- **The ledger scan** — running an analysis raises a full screen scan with an honest progress figure and a live board: the moment the rapid first pass lands, its real numbers stream in (rows read, spend totaled, categories forming, findings ticking) while the deep pass finishes behind it. Close it any time; the report still lands in your workspace.
- **Verdict band** — the flagged dollars, the spend analyzed with its row and month counts, the projected next quarter against the last three months, and the classified share as a ring. A note confirms every headline figure was recomputed from the ledger.
- **Where the money goes** — the category composition as a dense bar list: spend, share, vendor count, each category's own yearly trend, and its top vendors inline.

- **The findings ledger** — duplicate invoice suspects (identical amounts inside 45 day windows, re-keyed references) and price variance outliers (unit price steps above 15% of your own trailing median, months beyond 2.5 sigma of a vendor's norm), each with severity, the amount at stake, and the evidence rows (reference, date, amount). Mark each finding **Worth chasing** or **Dismiss**; dismissals leave the flagged total live and your triage saves to the run.
- **The next six months** — actual monthly spend as a solid line, the projection dashed with a one standard deviation band, values printed at the data, and deterministic buying window notes per category beneath it.
- **The analyst read** — the narrative over the computed numbers: where the money concentrates, which findings to chase first, and three to five prioritized moves with the stake quantified.
- **Classifications worth a human glance** — suppliers the model placed below 60% confidence, with the guessed category and the reason it was unsure.
- **Your ledger reviews** — every run saved to the workspace as a shelf or a dense table with the flagged total, finding count, and status; reopen any run, rerun on a fresh export, or delete it.
- **Open the workroom**: a navy band above a saved review opens the analysis as a case file on a dark stage: the cover sheet, the analyst brief, the category composition, the findings ledger with every evidence table and your triage state, the six month forecast with its buying windows, and the What now sheet. The ⚡ dock beside the paper carries the moves: the hold list for the payment run (the duplicate suspects with their evidence, figures included), the statement request to the flagged vendor (your own postings only, never the analysis behind the flags), the finance readout, chat ready copies of the verdict, the findings, and the composition, and handoffs into invoice intelligence and the renewal calendar. ✦ **Ask Vera** answers from this ledger alone: the flagged total with your triage applied, the composition, the projection, and the analyst moves. The rail searches your uploaded documents and keeps your shelf.

How to use it

1. Press **See it on the sample ledger** first if you want to see the output shape before exporting anything.
2. Export the ledger from your AP or ERP system as CSV (vendor, amount, and date at minimum) and drop it on the stage.
3. Watch the scan or leave the page; the analysis takes two to four minutes and saves either way.
4. Triage the findings: mark the real ones **Worth chasing** and dismiss the rest; the flagged total follows your verdicts.
5. Take the top category into the spend optimizer, reconcile flagged invoices in invoice intelligence, or put a buying window on the renewal calendar from the What next panel.

Tips

- Negative rows (credits) count toward totals but are never flagged as duplicates.
- Mixed currency exports are summed as printed and the report says so; filter to one currency for exact totals.

Related: Spend optimization · Invoices

CONTRACT DECODE

Where: Tooling › Contract decode · `/tooling/contract-decode`
Who: everyone · metered by plan
You get: a clause-by-clause plain-English decode, up to ten paste-ready asks, and a term-by-term benchmark · the first read streams at once, the deep read follows within a few minutes

Upload one vendor contract and watch the AI decode it live: what every consequential clause means in plain English, up to ten changes worth asking for with paste-ready language, the terms already in your favor, and a term-by-term market benchmark.

In plain words: the contract, translated, with the changes worth asking for written out.

When to reach for this

- Before signing anything: what does this paper actually say, and what would a buyer change?
- Before a renewal: the current contract's terms are your baseline and your ammunition.
- When legal is a bottleneck: the decode is the read that decides whether legal needs to look at all.

A worked example. Decode the sample Northwind master agreement and the read surfaces exactly the terms a buyer negotiates: the auto-renewal with its 60 day non-renewal notice, the 8 percent renewal price cap written into the paper, the 99.5 percent uptime commitment with service credits capped at 30 percent of monthly fees, and the 60 day data retention after termination. Each lands with what it means for you and, where it falls short of market, the paste-ready ask.

Opened from your home screen (app mode), the decoder walks you straight into the room: the upload stage runs edge to edge under the navy band with the dropper ready, and the page header and meta row sit below the working surface. Opened from the dashboard, the page keeps the composition described here, header first. The SOW Decoder and MSA Decoder share the same behavior.

On the screen

- **Upload**: drop the contract or quote. The document saves privately to your workspace, files itself as a contract with the vendor filled in by the decode, and the benchmark is requested automatically under your plan's allowance, so a decoded document is on the analyst path the moment the reveal ends.
- **Run bar**: a finished decode opens with a bar across the top: the run's inputs, a **Decode another** button, and the run actions: a **Favorite** star (starred decodes stand under Starred on the workroom shelf), replay the reveal, **Hide the document** to collapse the split view, open the original document, **PDF** and **Word** exports of the decode in the executive-brief theme, **Redline the Word file** on Word decodes, and delete.
- **Redline the document and send it out**: the platform writes your chosen redlines into the real document, ready to go back across the table. On a .docx, **Redline the Word file** produces true Word tracked changes: each quoted clause struck through with the replacement language inserted. Every change and comment carries your own name as the reviewer, and the document carries no VendorBenchmark markings, so it can go straight back to the vendor as your redline. On a PDF, **Redline the PDF** marks each clause as a highlight whose margin note carries the replacement language and the reasoning (a PDF cannot hold

Word revisions; the marks open in any PDF viewer). The findings in your asks are applied (before any triage, every open finding is), plus your own redlines; the other open findings can ride as margin comments, so nothing flagged travels silently. Download the file, or send it to any address: the vendor's rep, a colleague, or yourself. A send is completely plain: your message (drafted from your chosen changes as short bullets, fully editable) and the attachment, nothing else, your name on the From line, replies landing in your own inbox, and no VendorBenchmark branding anywhere on the email or the document. The stored original is never modified.

- **Make the redlines your own**: every finding card is editable. **Edit** on the ask box rewrites the replacement language in your own words (Reset restores the AI's); **Add your note** pins your comment to the finding for your team and for the margin comments of the outgoing document; and **Add your own redline** under the findings takes a clause you quote verbatim from the document plus your replacement language, and from then on it travels exactly like the AI's asks: into the copy list, the ask sheet, and the redlined file. Everything saves to the decode the moment you press Save.
- **The cover sheet**: a compact brief-style header over the report: your market standing, the money on the table, and the standing if the asks land, each figure with a small chip linking to the report section that argues it; a table of contents with reading times that jumps to the clauses, the keepers, or the term benchmarks; your first moves lifted from the top asks; and a coverage note stating what was read and whether every figure passed verification. The money figure recomputes live as you accept or dismiss suggestions. Every jump scrolls the page to the right section of the report.
- **The working report**: the full review reads directly on the page under the cover sheet. The verdict opens as a scan, not a wall: one headline sentence, the figures, **The read in brief** as bullets, and a two-column read of **Working for you** (the terms already in your favor, topped up with any term benchmarked above the 60th percentile; when nothing qualifies the column says so honestly) against **What can be better** (the three strongest open findings, priced). Then every finding renders as a card with its severity, its page, the verbatim clause it responds to, the plain-English read, why it is risky, and the paste-ready replacement language with a copy button. **Add to my asks** and **Dismiss** sit on each card, and the verdict figures, the triage strip, and the drafts recompute as you work the list. **Open the original document** on the run bar opens the stored PDF in its own tab, on any machine, visible only to your organization.
- **The document, redlined, beside the report**: on a wide screen a decoded PDF opens as a split review: the original document on the left with every flagged clause highlighted in place (red for high risk, amber for caution, green for the terms in your favor), margin pins carrying each finding's dollar value, a gold thread linking the three highest-leverage clauses, and a minimap of the whole document; the report on the right. The two sides read together: press a page reference on any finding card and the document scrolls to that exact clause; press a highlight or a margin pin in the document and the report scrolls to its card. Dismissing a finding fades its highlight. Zoom controls sit on the viewer's corner (75 to 300 percent; press the percentage to snap back to fit width), and **Expand** opens the document at full size over the whole screen with the working report beside it, every control live: triage, edited language, notes, and your own redlines, with Escape stepping back to the page. **Hide the document** on the run bar collapses the split back to the classic layout, and the choice is remembered on your browser. Decodes of Word documents and older decodes without a stored copy read in the classic layout.

- **Two expert reads**: underline tabs above the report switch between the **Negotiation expert** (the door's own decode on the contract and SOW doors) and a **Legal review** written by a buyer-side technology lawyer; on the MSA Decoder the legal review leads and the negotiation expert is the second view. The second read is generated on demand from the same stored document with one press, saves onto the decode automatically, keeps its own accept-and-dismiss triage, and redlines the document with its own findings when it is open.

- **Live decode**: clauses explained as the pass runs, with a preliminary read first and the deep read behind it.
- **Ask list**: the changes to request, ranked, each with language you can paste into a redline.
- **Market benchmark**: where the deal stands term by term.
- **The next-steps rail**: four cards carry the decode forward; in the split review they sit in a row directly under the cover sheet, before the document and the report, and in the classic layout they run as the rail beside the report. **Next steps** hands the run into negotiation, the vendor room, Second opinion, the Deal Room, the renewal calendar, the linked contract record, or the vendor's research in the library when coverage exists. **Draft emails** opens three drafts composed from the findings themselves (the redline reply, the vendor question list, the legal brief); edit them and send from your own inbox, nothing sends from the platform. **Share this decode** invites a teammate into the decode by email (members work it live; anyone else gets a 30-day read-only copy), mints the read-only link, or forwards the decode; when an invited teammate first opens the invitation, the decode files itself under Starred on their own shelf. **Your ask sheet** builds from the findings you add to your asks: the accepted asks with their language, ready to copy, or emailed to your own address with one press.
- **Team discussion**: a comments thread under the report, shared with your organization. @mention a teammate and they are notified with a link straight to this decode; threads support one level of replies, and roots can be resolved and reopened. A half-typed comment autosaves as a draft on your browser and is restored when you come back.
- **Vendor research**: when the research library covers the decoded vendor, the matching analyst briefs (the negotiation prep brief, the renewal outlook, the strategy assessment) appear as cards below the report.
- **The decode workroom**: the navy band under the run bar opens the decode as a case file of executive-brief papers: the verdict brief (the narrative read with the executive summary, the clause read, and what to do first), the ask sheet with the paste-ready language on each ask, and the term-by-term benchmark table. An icon dock on the right edge splits the stage: ⚡ puts the decode's moves beside the paper (draft the redline reply from the flagged clauses, draft the vendor question list, brief legal with the clause table, mint a share link, carry it into Second opinion or the Deal Room), and ✦ opens Ask Vera grounded on the whole decode, with one chip per paper that jumps the stage. The rail carries your other decodes, a keyword and semantic search over your uploaded documents, and your shelf of recent and starred decodes. Dismiss a finding on the decode page and the papers, the money, and the drafts all follow it.

- **The shelf**: your saved decodes stand as case spines with the vendor on each: a gold band means decoded, a dark red spine means the run failed. A toggle switches to the dense table with standing and suggestion counts; each decode reopens from either view and can be forwarded with a share link.

- **Sample decodes**: under the upload panel, open a finished sample and work the whole thing before committing your own contract. Three are hand written by our analysts: a Salesforce order form, a Microsoft EA quote and a ServiceNow renewal. Each plays the full reveal and then the complete report beside its illustrative document, redlined exactly as your own would be, and every control works: add findings to your asks or dismiss them, edit the language, pin your notes, write your own redlines, and download or email the redlined sample PDF. Sample work autosaves on your browser only (a sample has no workspace record), none of your allowance is used, and every sample has its own link you can send to a colleague. Inviting teammates, the @mention discussion, and workspace saving unlock on a decode of your own document.

How to use it

1. Upload the contract, or open one of the three sample decodes first to see exactly what comes back.
2. Read the split review: the redlined document on the left, the report on the right. Press a finding's page reference to see the clause in place, and press Add to my asks on the findings you want to carry.
3. Press the second tab above the report to run the other expert read (the legal review here, the negotiation expert on the MSA Decoder); it saves onto the decode automatically.
4. Email yourself the ask sheet, open a draft from the rail, or share the read-only link with a colleague.

Related: Vera Deal-Room · Clause library

SECOND OPINION

Where: Tooling › Second opinion · `/tooling/second-opinion`
Who: everyone can preview the sample · running your own contract needs a paid plan (trial plans are limited to the sample)
You get: a four page benchmarked executive brief: your placement against the peer set, the money on the table priced, five arguments that work, ready email drafts, and three priced paths · minutes

Drop a software contract and Vera returns a benchmarked second opinion in minutes. Where Contract decode reads every clause, Second opinion answers the first question: where does this deal sit against the market, what is it worth to fix, and what do I say. Use it when a renewal quote or a fresh contract lands and you need the market read before you engage.

In plain words: is this deal fine, what is it worth to fix, and what do I say.

When to reach for this

- The agreed deal is drafted and one hour of checking is still cheap (the pre-signature recipe).
- A stakeholder wants independent reassurance before they countersign.
- You suspect a deal is fine and want that confirmed in writing rather than assumed.

Opened from your home screen (app mode), the tool opens on the room: the drop stage runs edge to edge under the navy band, and the page header and meta row sit below it. Opened from the dashboard, the page keeps the composition described here, header first.

On the screen

- **The drop zone**: drop the contract as PDF or DOCX, or browse for it. The document is read for the terms that price the deal.
- **The reading**: the dark stage that shows the extracted terms live as the read runs: vendor, product, price, seats, and term land as they are found, so you see what the brief will be built on before it is written. Beside the document, a progress tracker states the phase in words (parsing the order form, extracting commercial terms, placing the deal in the cohort), the percentage in figures, how many terms have been extracted so far, and the elapsed clock. The bar reaches 100% only when the run is finished, never before. Press **Skip to the brief** and the same percentage and clock follow you onto the page, so a run you leave and come back to still says where it is.
- **The masthead**: a finished opinion opens under one navy masthead. The verdict is the headline, with the one line under it and the deal's meta (vendor, product, event, file, date). Beside it sit the three figures every benchmark on the platform leads with: the discount this paper offers as a percentage off list, what comparable customers get at the same size as a range with its midpoint, and where the deal ranks in one of five words (Poor, Below market, At market, Good, Top of the market) with the gap in points and the money on the table. The masthead also carries the opinion's actions (replay the reveal, email a copy, download the PDF, start another opinion, delete) and the numbered tab band that runs the whole opinion. The number keys 1 to 6 switch views, and each view keeps its own link, so a forwarded URL lands on the exact view you were reading.
- **Verdict** (view 1): the executive read on one screen: the summary, the bottom line, the deal ledger (population, net ACV, discount, uplift, payment, flex), and where you sit in the band drawn as a meter, with the peer figures under it.
- **The brief** (view 2): the four page paper in the executive brief format: the same three figures across the masthead, your placement against the peer set on a scatter, the money on the table priced, the five arguments, the talking points, and the three priced paths. Where the vendor paper prints no list price, and most enterprise cloud quotes do not, the brief says the discount cannot be verified from the document, still shows you what comparable customers get, and makes asking for the list price line by line one of the five arguments. The discount is always taken on the whole deal, total list against total net over the term, never line by line and never per unit. The docked verdict bar rides with the paper, and Download PDF prints this document.
- **Arguments** (view 3): the five arguments and the talking points as one working view; every talking point carries a Draft this email that lands on that draft, open.
- **Draft emails** (view 4): the five ready drafts, one per argument: read the full body in place, then copy the draft or open it in your own mail app. Nothing sends from the platform.
- **Negotiation** (view 5): the brief for the table. The three priced paths stand as the high level strategy ladder with the note on how to choose, followed by the tactical plays written from your own verdict figures (cap the uplift before price, ask for the band you are in, take the flex band, park the signature against the vendor's fiscal year end), how to open the conversation, and a budget posture line. Beside them sits the desk's curated view of the vendor at the table, covering the eight major publishers: fiscal year end and quarter closes, the best signing window, what the vendor's sellers are paid to attach right now, the main competitors worth naming with one line on when each is credible, and the table notes a buyer should

know before the first call. Vendors without a curated desk keep the strategy and plays, with a pointer into the Vera Deal-Room for live intel.

- **Next moves** (view 6): the saved contract note, the door into the opinion workroom, and the handoffs.
- **The opinion workroom**: Open the workroom, on the masthead from any view, puts the brief on one dark stage with the five drafts, the copy moves, and the handoffs docked beside the paper (the dock opens on the moves), Ask Vera grounded on the whole verdict, and your past opinions, document search, and shelf on the rail.
- **The handoffs**: take the same contract deeper in one click: Contract decode for the clause by clause read, the Vera Deal-Room for the full negotiation package, or Proposal comparison when a revision comes back.
- **Email a copy**: send the brief to yourself or a colleague; enter the address and the brief arrives as an executive brief email.
- **The sample**: a sample Workday renewal shows the full application instantly, no upload and no AI allowance used, so you can judge the output before you bring your own paper. Trial plans use the sample to see exactly what the tool produces; upgrade to a paid plan to run a second opinion on your own contract.
- **Saved to your workspace**: a note on Next moves says where the contract went: it was filed under Contracts, private to your organization, and joined your benchmarking queue. The note links to the contract record and carries a delete that removes the stored document on the spot. Owners can turn document saving off under Data controls.

How to use it

1. Drop the contract, or press Try the sample Workday renewal first.
2. Watch the reading extract the terms, then land on the Verdict view: the summary, the bottom line, and your standing on one screen.
3. Walk the band with the number keys: the full paper on 2, the arguments on 3, the drafts on 4, the negotiation brief on 5.
4. Open the draft that fits, copy it into your own inbox, or email a copy of the brief to a colleague.
5. When the fast read earns a deeper one, follow Next moves into Contract decode or the Vera Deal-Room.

Related: Contract decode · Vera Deal-Room · Verdict Bar

PROPOSAL COMPARISON

Where: Tooling › Proposal comparison · `/tooling/proposal-comparison`
Who: everyone can preview the sample · running your own proposals needs a paid plan (trial plans are limited to the sample)
You get: a benchmarked executive brief that surfaces every change between two proposals, each pricing and term change marked improved worse or unchanged, with the arguments and ready email drafts · minutes

Drop two vendor proposals and Vera shows exactly what changed. Where Second opinion reads one contract against the market, Proposal comparison reads one proposal against another and answers a different question: what moved, did it move your way, and what is the movement worth. It handles both situations. A same vendor revision, the original against the revised proposal, tells you whether the vendor's new draft actually helped. Two competing vendors, side by side, tells you which bid is stronger. Use it when a vendor sends back a revised proposal, or when you are choosing between two.

On the screen

- **The two drop slots**: Proposal A original and Proposal B revised, each dropped as PDF or DOCX or browsed for. Slot A is the baseline and slot B is what changed against it, whether that is the same vendor's revised paper or a competing bid. Both are read for the terms that price the deal.
- **The reading**: the dark stage that shows the changes as they are found: the terms extracted from both proposals and the differences between them land live, so you see what the brief will be built on before it is written.
- **The change ledger**: every change grouped into pricing changes and term changes, each one marked improved worse or unchanged against the baseline, so the movement in the deal reads at a glance.
- **Pricing changes**: the price lines that moved, old value against new, with the difference priced, so a lower unit price, a higher uplift cap, or a shifted payment term each shows what it is worth.
- **Term changes**: the contract terms that moved, from renewal notice and price protection to liability and termination, each read in plain English and marked whether the change helps or hurts your side.
- **Open the case**: a finished comparison shows a navy band above the brief that opens it as a seven document case file: both proposals as exhibits, a delta atlas over every compared term, the ask sheet that reverses the takebacks and banks the wins, and a What now sheet with the vendor reply and the internal note drafted from your own numbers. The case now carries the workroom dock: press ⚡ and the same moves work beside whichever paper is open; press ✦ and Ask Vera opens in the split, grounded on the whole comparison, with one chip per document that jumps the stage. The rail lists your past comparisons, a keyword and semantic search over your uploaded documents, and your shelf of recent and starred comparisons. The sample opens the same case, marked Sample.
- **Email drafts**: five ready drafts behind buttons: edit the subject and body in place, then copy the email or open it in your own mail app. Nothing sends from the platform.
- **The handoffs**: take either proposal deeper in one click: Contract decode for the clause by clause read, or Second opinion for the benchmarked market read on a single proposal.
- **Email a copy**: send the brief to yourself or a colleague; enter the address and the brief arrives as an executive brief email.
- **The sample**: a sample revision shows the full comparison instantly, no upload and no AI allowance used, so you can judge the output before you bring your own paper. Trial plans use the sample to see exactly what the tool produces; upgrade to a paid plan to compare your own proposals.
- **Saved to your workspace**: a note under the brief says where both proposals went: each was filed under Contracts as a vendor proposal, private to your organization, and joined your benchmarking queue. The note links to both records and carries a delete for each. Owners can turn document saving off under Data controls.

How to use it

1. Drop the original in slot A and the revised or competing proposal in slot B, or press Try the sample revision first.
2. Watch the reading extract the terms and surface the changes, then read the change ledger and the brief.
3. Check each pricing and term change against improved worse or unchanged to see whether the deal moved your way.
4. Open the email drafts, edit the one that fits, and copy it into your own inbox, or email a copy of the brief to a colleague.
5. When a single proposal earns a deeper read, follow the handoff into Contract decode or Second opinion.

Related: Second opinion · Contract decode · Vera Deal-Room

BUILD COMPETITIVE SCENARIOS

Where: Tooling › Build competitive scenarios · `/tooling/competitive-scenarios`
Who: everyone can preview the sample · running your own review needs a paid plan (trial plans are limited to the sample)
You get: a competitive options review with three to five cost scenarios priced against your renewal, a vendor-ready page, and two what-next emails · minutes

Name the product you are renewing and Vera researches the credible alternatives, prices three to five cost scenarios against the renewal over your term, and writes a competitive options review you can put in front of the vendor. Where Second opinion reads one contract against the market, this tool answers a different question: what are my alternatives actually worth, and what does a credible option on the table give me. Each scenario carries the license, migration, and transition cost, so renewing as proposed, repricing to market, and switching to each alternative are priced side by side. Use it when you want a believable alternative in the room before you renew, so the vendor knows you are looking.

On the screen

- **The brief form**: name the product you are renewing, the current annual spend, and the seats and term. That is all the model needs. Nothing you enter is stored beyond the finished review, and a sample Workday review shows the full brief before you run your own.
- **Start from your estate**: a row of chips above the form, read from the agreements you already store. A chip fills the vendor, the stored annual value, and the term, and everything stays editable before the run. A value is offered only where it is stored in a currency the tool reads, a term only where the stored dates give one between one and five years, and sample data never appears.
- **The market scan**: the dark stage that shows the scan live: the credible alternatives found, each priced at your scale with its migration cost and timeline, and the scenarios being modeled against your renewal.
- **The options review**: the vendor-safe executive brief: the alternatives evaluated with their fit and pricing posture, the cost scenarios priced over your term as a chart and a table, the modeling assumptions stated plainly, and a measured bottom line. This is the page you hand to the vendor.
- **Use this in the room**: the leverage layer that stays with you and is never printed and never in the shared copy: the moves to make with the alternatives, so you use them without bluffing a walk-away.

- **The what-next emails**: two ready drafts behind a button, one to the vendor rep signaling you have completed an options review and are evaluating alternatives, one to your leadership with the recommendation and the review attached. Edit in place, then copy or open in your own mail app. Nothing sends from the platform.
- **Share it**: email a copy to the whole team (separate several addresses with commas), or copy the workspace link that opens for members of your organization.
- **The handoffs**: take the deal deeper in one click: the Vera Deal-Room to run the negotiation, Second opinion for the benchmarked read on the incumbent contract, or Proposal comparison once the vendor comes back.
- **The sample**: a sample Workday HCM review shows the full options review instantly, no input and no AI allowance used. Trial plans use the sample to see exactly what the tool produces; upgrade to a paid plan to run your own.
- **Open the workroom**: a navy band above a saved review opens it as one case file on a dark stage: the vendor safe options review as the paper, the leverage sheet that stays with you, and the scenario ledger year by year. The ⚡ dock beside the paper carries both what next drafts, the copies, and the handoffs, and the ✦ dock is Ask Vera grounded on the whole review. The rail searches your uploaded documents and keeps your shelf.

How to use it

1. Name the product you are renewing and enter the current annual spend, or press Use the sample Workday review first.
2. Add the seats, the unit, the term, and the proposed uplift if you have them, then press Build the scenarios.
3. Watch the market scan find the alternatives and price the scenarios, then read the options review.
4. Open Use this in the room for the leverage moves, and the what-next emails to send from your own inbox with the review attached.
5. Email a copy to the team or copy the workspace link, and follow a handoff into the Vera Deal-Room to run the negotiation.

Related: Vera Deal-Room · Second opinion · Proposal comparison

SOW DECODER

Where: Tooling › SOW Decoder · `/tooling/sow-decoder`
Who: everyone · metered by plan (shares the decode allowance with Contract decode)
You get: a buyer-side review of a consulting statement of work, up to ten paste-ready fixes, and the rate and fee terms benchmarked · the first read streams at once, the deep read follows within a few minutes

In plain words: the SOW read like a buyer-side sourcing advisor, hunting the clauses that become change orders.

When to reach for this

- A statement of work is on the table and nobody on your side prices SOWs for a living.

- The engagement is time and materials and you want the caps and acceptance criteria checked before, not after.
- The rates deserve their own look too: pair it with the Rate Card Benchmark.

Upload a consulting or systems-integration statement of work and watch the AI review it live like a buyer-side sourcing advisor. It hunts the failure modes that cost SOW buyers money: scope written as activities instead of outcomes, assumptions that quietly become billable change orders, deliverables without acceptance criteria, deemed-acceptance windows, junior-heavy staffing pyramids, uncapped time and materials, expense policies without caps, and payment tied to calendar dates instead of accepted work. Missing protections (no termination for convenience, no key-personnel clause, no expense cap) are flagged too, anchored to the nearest passage so you land in the right place.

On the screen

- **Upload**: drop the statement of work as PDF or DOCX, with an optional consultancy name. The document saves privately to your workspace and files itself as a contract, with the benchmark requested under your plan's allowance.
- **Live review**: the same full screen reveal as Contract decode: a scan ticker while the passes run, then the executive brief with your market standing, the money on the table, and the ask ledger.
- **Ask list**: up to ten changes, leverage ranked, each with paste-ready replacement language reflecting the positions procurement teams actually win: holdbacks of 10 to 20 percent released at final acceptance, a signed change order before any out-of-scope work, objective acceptance criteria with rework at the firm's cost, termination for convenience on 30 days notice, and named key personnel with approval rights over substitutions.
- **The working report**: the full review reads on the page: every finding with the verbatim passage it responds to, the plain-English read, and the paste-ready language, with Add to my asks and Dismiss on each card. The original SOW opens in its own tab from the run bar, restored automatically from the stored workspace copy, visible only to your organization. The next-steps rail beside it carries the drafts, the invite and share controls, and your ask sheet; the favorite star, PDF and Word exports, and the team discussion work exactly as on Contract decode.
- **Market benchmark**: the commercial terms the SOW prints (blended rates, discount off standard rates, escalation, expense cap, holdback, termination notice) placed against the market for engagements of that firm tier and size.
- **The decode workroom**: a finished review opens as the same workroom as Contract decode: the verdict brief, the ask sheet, and the benchmark table as executive-brief papers, with the redline reply, the consultancy question list, and the legal brief drafts docked beside the paper, and Ask Vera grounded on the whole review.
- **The shelf**: your saved SOW reviews stand as case spines on their own shelf, separate from contract decodes; each reopens, replays, forwards, and shares like any decode.
- **Sample review**: under the upload panel, open a finished sample and read the whole thing before committing your own document. The SOW sample is a tier one firm's fourteen month ERP implementation

on uncapped time and materials, with five day deemed acceptance and an unnamed team. It plays the full reveal and then the complete report. Nothing is uploaded, nothing is saved, and none of your allowance is used; the sample has its own link you can send to a colleague.

How to use it

1. Not ready to upload? Open the sample review first and see exactly what comes back.
2. Upload the statement of work before you sign it.
3. Read the ask list on the page; accept or dismiss each suggestion as you go.
4. Copy the ask list into your response to the firm, or forward the review to the deal owner with a share link.

Related: Contract decode · MSA Decoder · Rate Card Benchmark

MSA DECODER

Where: Tooling › MSA Decoder · `/tooling/msa-decoder`
Who: everyone · metered by plan (shares the decode allowance with Contract decode)
You get: the buyer-side legal review of a master services agreement, clause by clause, with paste-ready redlines and the key legal terms benchmarked against standard negotiated positions

In plain words: the MSA reviewed the way a technology-transactions lawyer would, exposure first.

When to reach for this

- A master agreement arrived and outside counsel is expensive, slow, or both.
- Before the redline round: the deviations from standard negotiated positions, with replacement language ready.
- When the missing clauses matter as much as the present ones: absent protections are flagged as changes too.

Upload a master services agreement (or master subscription, cloud, or framework agreement) and watch the AI run the review a technology-transactions lawyer would, triaged by exposure: liability, indemnities, and data first, then IP, termination and exit, commercial terms, and the boilerplate with teeth. Every clause that deviates from the standard negotiated position becomes a change with replacement language; protections that are simply missing (no data-incident carve-out, no breach-notice window, no transition assistance) are flagged as changes too.

On the screen

- **Upload**: drop the MSA as PDF or DOCX. The document saves privately to your workspace and files itself as a contract, with the benchmark requested under your plan's allowance.
- **Live review**: the full screen reveal: overall risk as written, your standing against market positions, and the clause ledger as it lands.
- **The clause checklist**: the review walks the buyer-side checklist: the liability cap measured against annual fees and its carve-outs, the data-incident super-cap, the IP infringement indemnity and its exclusions, work product ownership, the DPA and breach notification window, warranties and their disclaimers,

termination and transition assistance, payment terms and the right to withhold disputed amounts, insurance, assignment and change of control, non-solicits, publicity, and the order of precedence between the MSA and its SOWs.

- **Redlines**: each change carries paste-ready language at the standard negotiated middle ground: a cap of 1 to 2 times annual fees with the five standard carve-outs, breach notice within 48 to 72 hours, a 90 day deliverables warranty from acceptance, termination for convenience on 30 to 60 days notice, and transition assistance at contracted rates.
- **Keepers**: the clauses already at or better than market, listed so nobody negotiates them away.
- **The working report**: every finding on the page with the verbatim passage it responds to and the paste-ready language, triage on each card; the original agreement opens in its own tab from the run bar, restored automatically from the stored workspace copy. The next-steps rail beside it carries the drafts, the invite and share controls, and your ask sheet; the favorite star, PDF and Word exports, and the team discussion work exactly as on Contract decode.
- **The decode workroom**: a finished review opens as the same workroom as Contract decode: the verdict brief, the redline sheet, and the benchmark table as executive-brief papers, with the redline reply and the legal brief drafts docked beside the paper, and Ask Vera grounded on the whole review.
- **Sample review**: under the upload panel, open a finished sample and read the whole legal read-through before committing your own agreement. The MSA sample is vendor paper with a three month liability cap, no carve outs, a one way indemnity and no breach notification clock. Nothing is uploaded, nothing is saved, and none of your allowance is used; the sample has its own link you can send to a colleague.
- **The shelf**: saved MSA reviews live on their own shelf and forward with a share link.

How to use it

1. Not ready to upload? Open the sample review first and read the whole legal read-through.
2. Upload the MSA before it goes to signature (or before you spend outside counsel hours on the first pass).
3. Walk the changes on the page; accept the ones you will fight for and dismiss the ones you will concede.
4. Copy the redline list for your legal team, or forward the review with a share link.

Related: Contract decode · SOW Decoder · Clause library

VERA DEAL-ROOM

Where: Tooling › Vera Deal-Room · `/tooling/dossier`
Who: everyone · your own data requires a Professional or Enterprise plan
You get: six exportable documents, from the market benchmark to the priced walk-away · pieces stream in over several minutes

In plain words: one upload, the whole negotiation pack: benchmark, decode, brief, playbook, intel, and the priced walk-away.

When to reach for this

- A deal big enough to deserve the full pack, and a meeting soon enough to need it (the meeting recipe).
- You want the credible alternative priced with crossover math, not asserted.
- The room becomes the working surface for the negotiation, not a one-time report.

Drop one contract and the whole negotiation deal room assembles itself with no further input: the deal placed against peer reference data, the decoded contract with its leverage-ranked ask list, a combined five-page analyst brief, the strategy and tactics playbook, a live web-researched vendor intel briefing, and the credible walk-away alternative priced with crossover math. The theatre is the front door: opening Vera Deal-Room with a finished case on file raises the curtain on the Deal-Room Theatre, the case presented full screen in five acts and run as a complete workspace. Escape steps backstage into the room itself, which runs as one application: a navy masthead carrying the counted verdict and the room's actions, and a numbered tab band over everything the room can do, with every document opening on the case file's stage.

Opened from your home screen (app mode), the Deal-Room opens on the room: the drop stage runs edge to edge under the navy band, and the page header and meta row sit below it. Opened from the dashboard, the page keeps the composition described here, header first.

On the screen

- **The drop stage**: the navy hero shows the six sample documents fanned like an opened case. Drop your contract anywhere on the page, or use the button; the file saves to your workspace as a contract and joins your benchmarking queue.
- **The assembly desk**: while the pipeline runs, each of the six documents is a paper sheet on a desk: a ghost outline until its stage starts, a typing skeleton while it generates, then its real first page the moment it lands, ready to read while the rest assemble. A ticker narrates each stage, and the money at stake surfaces in a gold band as soon as the decode finds it. When the last document lands, the sheets stack and a gold seal stamps the delivery.
- **The masthead**: a finished deal room runs as one application under a single navy masthead: the deal's title and source line on the left, the room's actions and the counted verdict on the right (market standing today, money at stake across the asks, where the deal lands if every ask succeeds, and days to the signing deadline once the room has one). At its foot sits the tab band that runs the whole room.
- **The tab band**: numbered tabs over the room's sections: Overview, Documents, Next moves, Five minute brief, Negotiation desk, and Cover to cover. The number keys switch tabs, each tab can be linked to directly, and the band's right end carries the read meter (documents read against documents delivered) and the Open the case file door.
- **The case file**: clicking any document, on any tab, opens it as a paper on the case file's dark stage, the whole room bound as one numbered workroom: the cover sheet with the mandate and the verdict figures, every document of the room as an exhibit with read tracking, and a closing What now sheet whose forward and share moves land on the room's own Next moves board. The rail lists every document, the arrow keys page from paper to paper, and Escape returns to the tab you left. Under the document list sit a search over your own uploaded documents and your shelf of recent deal rooms. On the right, ⚡ **What now** carries the same eleven moves as the Next moves rail, so the four Vera drafts (the meeting request, the opening asks, the leadership brief, the working session invite), the workspace link and the forward ride beside every exhibit rather than living on a separate board. On your own rooms, ✦ **Ask Vera**

docks next to it, grounded on this deal's asks, plan, and documents; a worked example carries the moves but not Vera, because she never grounds on data that is not yours.

- **Reporting modes**: three ways to read the same case. Documents shows the cover sheet (verdict numbers with proof chips, the table of contents with read state, your first three moves) and the document wall of real first-page miniatures. Cover to cover reads every landed document as one continuous paper; the arrow keys move between documents. Five minute brief compresses the whole case into a one page executive readout, assembled verbatim from the deal room's own documents, for the reader who has five minutes before the call.
- **The Deal-Room Theatre**: the full screen show, and the default entrance: arriving at Vera Deal-Room with a finished deal room opens it on stage behind a curtain rise (arriving at a specific document link, or through the All deal rooms door, does not). The curtain rises on the lobby first: continue the latest case (the card carries its standing and money at stake and says where it picks up), open any other finished deal room from the shelf, or drop a new contract right there, drag and drop or browse. Choosing continue lands on the workroom's What next board; your saved mark wins when you left one, and the overture sits at the top of the bill. The show runs in five acts: the verdict (the counted overture, then the five minute brief), the instruments (the ask ladder), the documents (the cover sheet with read state, then one paper per screen, each with its own download, forward, copy link, and print toolbar), the workroom (the What next board and the war plan), and the curtain call. Presenter notes, the overview grid on G, light or dark stage, and a printable leave-behind of the screens you walked all ride along, and the show remembers where you left off. The View the negotiation strategy button backstage re-enters it any time.
- **The ask ladder**: every ask is a switch. Turn one off in front of the room and the value selected, the projected standing, and the header ticker recompute live. The scenario dial prices three postures, conservative, expected, and every ask, so leadership sees the honest range.
- **Stage tools**: three presenter tools float over every scene: the rehearsal clock on R (a five minute target that turns red when you run over), pushback cards on Q (the likely vendor objection on each of your strongest asks, with the counter drawn from the deal room's own clause read), and the spotlight on S (follows your pointer and dims the rest of the stage). The envelope in the header drafts the follow up email without leaving the stage.
- **The workroom**: the act that answers "what do I do next". The What next board is three columns of operable cards: reach out (Vera drafts the vendor meeting request, the negotiation opener, the CFO brief, or the colleague working session invite; you edit and send from your own inbox), report it (the five minute brief, the PDF pack, the Reporting studio), and prepare (the war plan, the call rehearsal, Self play, the analyst handoff).
- **The war plan**: the negotiation as a four phase project plan with Gantt bars on a real week timeline ending at your renewal date (or a sixteen week runway when the contract states none): read and align, open the channel, negotiate, close and sign. Tasks come from the playbook's own sequencing and your top asks; checking one off fills its phase bar, the ticker counts moves done, and the printed leave-behind includes the plan. Point at any task and the line shows you where to click: its checkbox pulses a blue ring, a click marker appears at the end of the text, and the whole line is the switch, so clicking the words checks the task off.

- **The desk saves itself**: the ask ladder's switches, the scenario dial, the war plan's check offs, and your place on stage persist per deal room. Once a project room exists (below), the desk also saves to the room itself, so progress follows you across devices and shows on your dashboard's Your deal room card.
- **One tab at a time**: the active tab gets the whole width under the masthead; there is no side menu to manage. Where you stand is always visible on the band, and any document you open from a tab returns you to that tab when you close the case file.
- **The overview**: the room's control panel, and the deal on charts. Press Set up the deal room and answer five questions: your goal for the deal, what a win looks like, the signing deadline (prefilled from the contract's renewal date when it states one), who works the deal, and what matters most. The overview then shows the goal band, five vitals (progress, days to deadline, standing, money at stake, team size), three charts (the percentile ladder with the deal as written and where it lands if every ask succeeds, the top asks as value bars, and the market band with your own price marked on it), and the plan as a Gantt: every task in the four phase plan on its own row with a checkbox, its bar on the shared timeline, and the Today line through it. Tick tasks off on the chart, and assign any task to teammates with their initial chips (an optional RACI over the plan Vera populated from the playbook and your top asks); assignments and check offs save to the room. Before the room exists, the overview shows the charts and the two minute setup invitation.
- **Vera in the room**: an Ask Vera box sits above the charts on the overview, and Vera reads the room itself while you are in it: the ask list with what each ask is worth, the six documents and which ones you have read, the goal and the deadline, and every planned move still unticked. Ask "What should I do next?" and the answer names the next move on your own plan and the ask worth the most, not general negotiation advice. The room keeps its own saved conversation, so a question asked here is waiting for you next week when you open the room again, and Ctrl/⌘ J, the floating panel, and the voice assistant all answer from the same reading while you stand in the room. See Vera AI.
- **The team**: invite colleagues by email address; they get an executive brief invite with the room link (a workspace sign in is required to open it) and they join the status report list. Remove anyone with one press.
- **Room chat**: comments live in the room and can anchor to any document, the next moves, or the war plan; filter by anchor with the chips. Everyone in your workspace sees the same chat, and new comments ride into the next status report. Type @name (first name, full name, or the start of their email) and that member gets an email with the comment and a button into the room; mentions render highlighted in the chat.
- **Notes and updates**: a notes card on the overview for meeting notes, vendor call summaries, and anything the team should know. Notes are visible to the whole room, ride into the next status report, and @name notifies there too.
- **The status report**: Vera writes a status report from the room itself, progress, the deadline clock, moves completed since last time, and new comments, and emails it to everyone in the room. Pick every day, weekly on one day, or exactly the days you want, plus the hour; pause or change the schedule any time.
- **The negotiation desk**: under the Five minute brief sits the toolkit for whoever actually sits at the table: the strategy in one screen, talking points with the exact paste ready words for each ask, how the seller across

the table is compensated (quota, accelerators, quarter end), what drives this deal, and the six pitfalls that cost buyers money.

- **Market position by instant benchmark**: when the stated terms cannot be placed against the reference cohort, Vera runs her instant benchmark from the deal facts themselves, the stated quantity and price against the calibrated market bands, so Market position arrives with a percentile, the market range, and a verdict instead of Not available. When the terms cannot be unit priced at all, the document states the clause read’s own placement instead. Either way the basis is spelled out, and deals assembled before this existed are filled in the next time you open them.
- **The curtain call desk**: the show ends on your first three moves with the whole desk under them: Vera drafts the meeting request, the negotiation opener, or the leadership brief (you edit and send from your own inbox; nothing sends from the stage), and forwarding, the workspace link, the full PDF pack, the analyst handoff, and the call rehearsal sit beside them. Nothing requires stepping backstage.
- **Next moves**: its own tab, the three move families laid side by side, and it answers "now what". Vera drafts the email for the move you pick: the note to your vendor requesting the renewal conversation, the opener that puts your strongest asks on the table, or the internal brief for your CFO. Every draft is short, friendly, and free of dash punctuation; edit the subject and body in place, then copy the email or open it in your own mail app. Nothing sends from the platform.
- **Forward**: sends the whole deal room, or any single document, to a colleague as an executive brief email carrying the verdict figures and a link that opens right here in the workspace. The recipient needs a VendorBenchmark login with access to your workspace; the deal room never travels on a public link.
- **Copy link**: copies the workspace link to the deal room, or to the open document; it opens for members of your workspace after they sign in.
- **Download the pack**: all six documents as one PDF behind a generated cover page; each document also exports on its own.
- **Have an analyst price this**: on the next-move rail: opens the advisory concierge desk with the booking agenda pre-filled from this deal room: the vendor, the money on the table, the percentile as written, and the number of asks.
- **The shelf**: saved deal rooms stand as case spines: a gold band means delivered, an assembling spine fills as its documents land. A toggle switches to the dense table.

How to use it

1. Drop the vendor contract anywhere on the page.
2. Read the documents as they land on the desk; the money numbers arrive within the first couple of minutes.
3. When the seal stamps, read the case your way from the tab band: any document you click opens on the case file's stage with the rest of the pack one arrow key away, Documents holds the cover and the grid, Cover to cover the continuous read, Five minute brief the executive readout, and the Negotiation desk the table toolkit. Present, in the masthead's actions, raises the full screen Deal-Room Theatre when you want to walk the case in front of a room; it only opens when you press it.
4. Press Set up the deal room on the overview, answer the five questions, and run the deal from there: invite the team, comment on anything, check moves off the runway, and schedule the status report.

5. From the next-move rail, let Vera draft the vendor or leadership email, edit it, and copy it into your own inbox; forward the deal room or one document to a colleague when they need the full paper.
6. When you want a second read, use Have an analyst price this to take the deal room to a senior analyst with the agenda already written.

Related: Contract decode · Negotiations

SNAP A QUOTE

Where: Tooling › Snap a quote · `/tooling/snap-a-quote`
Who: everyone · metered by plan
You get: every line item checked against list, the market verdict, and the counter to send · about a minute

Photograph, screenshot, or drop any vendor quote and get an instant verdict: every line item read off the paper, each SKU checked against published list prices, the deal placed against comparable deals, and the counter to send back.

In plain words: photograph the quote, get the line-by-line check and the counter.

When to reach for this

- The quote arrived as a PDF, a screenshot, or a photo of someone's screen; all three work.
- You want the counter in a minute, not a workup: this is the fastest paper-to-reply path on the platform.
- Several quotes for the same need deserve Quote Face-Off instead.

On the screen

- **Capture**: a phone photo, a screenshot, or the PDF. The capture is read in memory and never stored.
- **Or see a finished verdict first**: a worked example under the capture panel, a five line renewal quote for 5,400 workers whose 16.3 percent blended discount hides one line at 25 percent and another at nothing. Pressing it replays the example as an already finished read, verdict and workroom included. Nothing uploads, nothing saves, and no snap is spent.
- **Line items**: every line read off the quote, with each SKU checked against list.
- **Verdict line**: where the whole deal sits against comparable deals.
- **The counter**: the response to send back.
- **The snap workroom**: the navy band above the verdict opens the read as two executive-brief papers, the verdict brief and the quote sheet, with the requote request draft, the copy moves, and the handoffs docked beside the paper, Ask Vera grounded on the whole quote, and your recent and starred snaps on the rail.
- **Saved snaps**: only the verdict saves to your workspace.

How to use it

1. Press the worked example if you want to see a finished verdict before snapping your own quote.
2. Snap or drop the quote.

3. Read the line-by-line check and the verdict.
4. Send the counter.

Related: Verdict Bar · Quote Face-Off

QUOTE FACE-OFF

Where: Tooling › Quote Face-Off · `/tooling/quote-faceoff` · via **More** or Ctrl/⌘ K
Who: everyone · metered by plan
You get: the aligned tale of the tape, the best-mix "Franken-quote" priced, and an analyst memo with the counter · a few minutes

Drop two or three competing quotes for the same need, vendor direct, reseller, or CSP. The board aligns equivalent SKUs into one tale of the tape, benchmarks every line, marks who wins where, and builds the Franken-quote: the best mix across all papers, priced.

In plain words: the quotes lined up SKU by SKU, and the best mix across all of them priced.

When to reach for this

- Two or three papers for the same need, and the decision deserves a written why (the recipe).
- Direct against reseller against CSP: the alignment is exactly where those comparisons go wrong by hand.
- The endgame of a bake-off, when the bids are finally on the table.

On the screen

- **The ring**: one corner pad per quote with the VS mark between them: Quote A and Quote B first, a third corner appearing once two are staged. Drop files anywhere on the stage, click a corner to browse, or paste screenshots (Ctrl+V).
- **What is kept**: a **PDF quote** saves to your workspace under Contracts, private to your organization, and joins your benchmarking queue; a note under the result names it, links it, and offers to delete it, which removes the record and the stored file and writes a deletion certificate. A **photo capture or screenshot** is read in memory and never stored, so a phone snapshot of a quote leaves nothing behind. The result footer repeats the rule. Owners can turn the saving off under Data controls, and with it off nothing you hand the tool is kept.
- **The sample bout**: **Or watch a finished bout first** under the ring shows a judged face-off before any upload; nothing uploads and no run is spent.
- **The bout card**: while the deep read runs: the papers facing off across the VS mark, with the judging steps beneath.
- **Run bar**: a finished face-off opens with a bar across the top: the run's inputs, a **Run another face-off** button, and the run actions (replay, forward, share, delete).
- **Tale of the tape**: equivalent SKUs aligned across the quotes, each line benchmarked, the winner marked per line. When you watch a run finish, the rows score themselves clause by clause, the Franken-quote last.

- **The Franken-quote**: the best mix across all papers, priced, so you can see what the ideal combination would cost.
- **Analyst memo**: which quote to take, which terms to pull from the losers, and the counter to send.
- **The finish row**: a judged bout closes with one-press copies of the priced counter and the top asks, and Ask Vera answers from the whole board.
- **Save and export**: file the whole result (tale of the tape, Franken-quote, and memo) to Deliverables, download it as a Word brief, or print it, from the actions under the verdict.
- **The fight cards**: every saved face-off as a card: the bout title, the corners, the one-line verdict, and the judgment status. Click one to reopen its board.
- **Open the workroom**: a navy band above the judged board opens it as one case file on a dark stage: the cover sheet with the verdict tiles, the tale of the tape as an executive brief paper with the Franken-quote row, the analyst memo with its citations intact, and the padding ledger as an exhibit. The ⚡ What now dock beside the paper carries the counter to the winning column, a last call per losing column, the internal decision note, the copy moves, the save and Word export, and the Deal-Room handoff; the ✦ dock is Ask Vera grounded on the whole bout. The rail searches your uploaded documents and keeps your shelf of past face-offs.

How to use it

1. Stage the competing quotes for the same need on the corner pads.
2. Read the tale of the tape line by line, then the Franken-quote.
3. Take the memo's counter back to your preferred supplier, and save the result to Deliverables so the comparison is on file.

Related: Snap a quote · Sourcing

AUDIT LETTER DECODER

Where: Tooling › Audit letter decoder · `/tooling/audit-decoder` · via **More** or Ctrl/⌘ K
Who: everyone · metered by plan
You get: who is really auditing you, a business-day countdown, demanded versus implied, and the first-hour response plan · about a minute

In plain words: the audit letter decoded before you reply to it: who, when, what is actually demanded, and what not to do.

When to reach for this

- The day the letter lands, before anyone replies or runs a vendor tool (the audit recipe).
- Distinguishing a formal audit from a voluntary SAM engagement dressed as one.
- Turning a threatening deadline into a business-day countdown with a plan against it.

The moment a software audit letter lands, drop it here, PDF or a phone photo, or paste a screenshot with Ctrl+V. The decoder identifies who is actually auditing you (Oracle GLAS, a Big Four firm on IBM's behalf, a voluntary SAM engagement), turns the stated deadline into a business-day countdown over a calendar that runs from today to the due date, separates what is demanded from what is merely implied, and hands you the three things not to do first. Not sure what you would get? **See a decoded sample** shows a fully decoded Oracle GLAS notice before you upload anything.

On the screen

- **Dropper**: the audit letter as a PDF, photo, or pasted screenshot, on the navy upload stage. It is read in memory and never stored.
- **Countdown wall**: the business days left as a large number, the deadline date, the letter's own wording, and a calendar strip from today to the due date with weekends dimmed.
- **Who is auditing**: the real party behind the letterhead, and a pressure meter for how hard the letter pushes.
- **The three do-nots**: numbered cards with the first-hour mistakes to avoid, grounded in how this vendor's audits typically run.
- **Scope with a verdict**: each demand marked demanded or implied, with our read of whether complying is contractually required, summarized in one strip above the table.
- **The response plan, on your calendar**: each step mapped onto a real date inside the window the letter leaves you, linked to your contracts on file and the full defense playbook pre-filled.
- **The safe first reply**: a copy-ready holding reply that acknowledges receipt, concedes nothing, and routes everything through one coordinator.
- **Exports**: **Print or save as PDF**, and **Send it on**, which drafts the covering email for you to send from your own mail. The contract-match rail is deliberately left off the printed pack, so an over-deployment figure of your own is never handed out by accident with the decode.
- **Open the workroom**: a navy band above a stored decode opens it as a case file on a dark stage: the cover sheet with the business day countdown and the three do-nots, the decoded letter with our read on every scope demand, and the response plan beside what your records show. The ⚡ What now moves run on the stored decode: brief the deal owner and legal with the full figures, send the auditor a neutral acknowledgment that concedes nothing and carries no figures, copy the one line summary and the do-nots, and draft the defense playbook on the matched contract. The ✦ dock is Ask Vera grounded on the letter: the auditor, the deadline, the scope reads, the plan, and the license position on record. The rail searches your uploaded documents and keeps your shelf.

How to use it

1. Drop the letter the day it arrives, before anyone replies to the auditor.
2. Read the countdown wall, then the three do-nots, then who is auditing and the demanded-versus-implied split.
3. Copy the safe first reply, then follow the response plan into the defense playbook.

Related: Verdict Bar · Diligence

RATE CARD BENCHMARK

Where: Tooling › Rate card benchmark · `/tooling/rate-benchmark` · via **More** or Ctrl/⌘ K
Who: everyone · metered by plan
You get: every day rate placed against its market band, the annualized overpay per line, and an analyst memo · a few minutes

In plain words: every day rate on the card placed against its market band, with the overpay priced per line.

When to reach for this

- A consulting or SI rate card is on the table, or buried inside an SOW.
- The negotiated rates deserve enforcement: the vault remembers them so every later review checks against them.
- Services spend has never had the scrutiny the software line gets; it usually pays for the look.

Upload a services rate card or SOW and see every day rate on the ladder: the market band per role, delivery region, and vendor tier, your rate as the dot, and the annualized overpay per line. Services hide 30 to 40% of software spend; this is where it surfaces. The page also holds your MSA rate card vault: record the rates you actually negotiated once and every SOW review, benchmark board, and invoice check enforces them from then on.

On the screen

- **Upload**: the rate card or SOW, dropped anywhere on the panel.
- **What is kept**: a **PDF or Word** document saves to your workspace under Contracts and joins your benchmarking queue, with a note that links it and offers to delete it. An **XLSX rate card** is a data file rather than a contract, so it is read in memory and never stored, and the result footer says so. Owners can turn the saving off under Data controls.
- **Exports**: the memo carries the shared deliverable toolbar: **Print or save as PDF**, **Word**, and **Send it on**, which drafts the covering email for you to send from your own mail. A sample run keeps print and Word.
- **The ladder**: one row per role with the market band and your rate marked against it.
- **Role mapping**: correct the AI's role mapping inline if a title was misread; the ladder reprices as you fix it.
- **Your contracted rates**: when a saved MSA rate card matches the read vendor, a strip above the mapping states how many lines price above what you already negotiated, and a **Your MSA** column appears per line: the contracted day rate, the delta, and the annualized overage, all computed in code. Contract compliance, not a market ask.
- **Blended rate deconstruction**: any blended or offshore-blended line gets a **Deconstruct this rate** panel: the plausible senior/standard/junior mixes at that price, the implied vendor cost per day against the billed rate, and the margin range each read implies, from a curated delivery cost basis (five roles by four regions).
- **Your rate cards**: the vault below the tool. List, create, edit, and delete the organization's negotiated rate cards, one dense grid per card: role as the MSA prints it, benchmark role, region, level, unit, and the contracted rate. Cards can also arrive in one click from a finished SOW rate review (**Save as rate card**).

- **Analyst memo**: which lines to contest, the blended-rate trap, and the counter-structure to propose.
- **Theatre**: once lines are benchmarked, the Theatre button on the verdict band presents the findings to procurement full screen: the opening position, each role family against its market band, then priced negotiation targets and the analyst memo. Presenter notes guide the walkthrough, and the leave-behind prints the screens you covered.
- **Open the workroom**: a navy band above a benchmarked card (real runs, not the sample) opens it as one case file on a dark stage: the position cover sheet, the full ladder worst gap first with the annual figure per line, the analyst memo once one exists, and the unmapped lines as an exhibit with the reason each could not be judged. The ⚡ dock carries the rate challenge to the vendor, stating only the day rates we need per role and their own printed rates, never the band behind them, plus the internal position note with the full placement and the annualized overpay, the copies, the Word brief, and the Deal-Room and vendor record handoffs; the ✦ dock is Ask Vera grounded on the stored extraction and your saved corrections. The rail searches your uploaded documents and keeps your shelf.

How to use it

1. Record your MSA rate card once, by hand in **Your rate cards** or with **Save as rate card** on a SOW rate review.
2. Upload the rate card or SOW to benchmark.
3. Fix any role mappings the AI got wrong.
4. Read the Your MSA column first (already agreed, just enforce it), then the market gaps, then take the memo into the negotiation.

Related: Contract decode · The benchmark library

TIME MACHINE

Where: Tooling › Time machine · `/tooling/time-machine` · via **More** or Ctrl/⌘ K
Who: everyone · metered by plan
You get: a narrated diff of the two papers and the real price walk line by line · a minute or two

Drop last year's paper and this year's for the same vendor. The machine narrates what actually changed: the terms you won, the protections that quietly disappeared, and the real price walk line by line. Renewals smuggle in term degradation, and the machine reads for it so you do not have to.

In plain words: last year's paper against this year's, with every quiet change named.

When to reach for this

- The renewal draft arrived and "nothing else changed" deserves verification, not trust.
- Post-signature archaeology: what did we actually give up over the last three renewals?
- The history you kept in Your Archive is exactly what this machine eats.

On the screen

- **Two uploads**: last year's document and this year's.
- **Start from your estate**: a row of chips above the uploads, read from the agreements you already store. Where a vendor holds two readable documents, the **pair** itself is offered and the comparison runs on the stored files with no re-upload and nothing new saved. Other chips fill the vendor name only, and the line beneath says why: the workspace stores fewer than two readable documents for those vendors.
- **What is kept**: on the upload path, **both** documents save to your workspace under Contracts, private to your organization, and join your benchmarking queue, with a note that names each one, links it, and offers to delete it. On the stored-pair path nothing is uploaded and nothing is saved again. Owners can turn the saving off under Data controls.
- **Or see a finished read first**: a worked example under the two upload slots, two order forms for the same vendor a year apart. It walks a 19.7 percent price walk that decomposes into unit price on products you already had, added volume and one new product, and a retired line that takes money back out, plus three protections that quietly left the paper. Pressing it replays the example as an already finished read, workroom included. Nothing uploads, nothing saves, and no read is spent.
- **The narrative**: what you gained, what you lost, and where.
- **Price walk**: the real increase line by line, not just the headline number.
- **Saved reads**: the walk delta is kept with each saved read.
- **Theatre**: a finished read carries a Theatre button that plays the comparison full screen for a renewal review: the verdict, what you gained and lost, the protections that quietly left the paper, then the price walk a stretch of lines at a time while the cumulative figure counts up in the header. The show ends on your position and the watchouts to raise before signing.
- **The timeline workroom**: the navy band above the read opens it as two executive-brief papers: the timeline brief, with the price walk drawn as a journey (last year, the movers, this year, every value printed at the data), what moved each way, and the quiet losses as a verdict table; and the line ledger, every priced line both years with the movement per line. The dock beside the paper drafts the renewal reply that asks every changed term and quiet removal back to last year's language, briefs the deal owner with the full figures, and hands off to the negotiation, the vendor room, and Second opinion; Ask Vera rides the same dock, grounded on the whole read.

How to use it

1. Press the worked example if you want to see a finished read before uploading your own paper.
2. Upload both versions of the same vendor's contract.
3. Read the narrated diff, watching for dropped protections.
4. Take the findings into the renewal.

Related: Compare · Renewal playbook

SELF PLAY

Where: Tooling › Self play · `/tooling/self-play` · via **More** or Ctrl/⌘ K
Who: requires an AI plan

You get: eight full negotiation transcripts, the outcome distribution, and the swing move priced · several minutes; leave it running

Name the vendor and the annual price, and the machine plays your renewal to completion eight times: your negotiator on one side, the vendor's rep in persona on the other, four strategies against two rep temperaments. It is a simulation against a modeled counterparty, not a prediction.

In plain words: your renewal played out eight times before you play it once.

When to reach for this

- The day before the real thing (the rehearsal recipe).
- Choosing a strategy: the outcome distribution shows which approach holds up across temperaments.
- Finding the swing move: the one choice that moves the outcome most is named and priced.

On the screen

- **Setup**: the vendor and the current annual price.
- **Start from your estate**: a row of chips above the setup, read from the agreements you already store. A chip fills the vendor and today's annual price where your estate holds one; the quote and the target stay yours. An annual value is offered only where it is stored in a currency the tool reads, and sample data never appears.
- **Outcome distribution**: the eight results drawn as a thousand strands.
- **Transcripts**: the full script of each of the eight matches.
- **The swing move**: the sequencing decision that separated the good closings from the bad ones, with its dollar figure.
- **Save and export**: file the distilled result (verdict figures, strategy comparison, and key excerpts) to Deliverables, download it as a Word brief, or print it, from the actions above the strands.
- **The exercise workroom**: the navy band above the strands opens the distilled brief as a workroom: the same document that saves and exports, on one dark stage, with the team plan email drafted from the simulation's own numbers, the save and Word export, and the handoffs docked beside the paper, and Ask Vera grounded on the whole exercise.

How to use it

1. Enter the vendor and your annual price.
2. Run the eight matches; this takes a while.
3. Read the swing move, then read the transcripts where it played out.
4. Save the run to Deliverables so the read survives the session.

Related: Negotiations · Vera Deal-Room

REPORT EXTRACT

Where: Tooling › Report extract · `/tooling/report-extract` · via **More** or Ctrl/⌘ K
Who: everyone · the custom-field scan needs an AI plan
You get: an Excel workbook of every contract in your chosen layout · seconds, or a few minutes when a scan runs first

Everything document search has read across your contracts, exported as one workbook: one row per contract, columns you choose. Use it when you need the contract register in Excel, when your SAM tool (Xensam, Flexera One, ServiceNow, Snow, USU) wants an agreement import, or when your team's own report template should fill itself.

On the screen

- **Layout** — six starting points: the full extract with every field, or layouts shaped to the agreement imports of Xensam, Flexera One, ServiceNow SAM, Snow License Manager, and USU. Below them, upload your own XLSX template.
- **Fields** — every column the workbook can carry, grouped: identity, term dates, renewal mechanics, values in the original currency and in USD at today's rate, documents, detected clauses, evidence quotes, and your custom fields. Tick what you need; the preview follows.
- **Column mapping** — with a template loaded, each of your template's columns and the field that fills it. Columns are auto-mapped by header name; fix any of them with the dropdown. Unmapped columns are left exactly as they are.
- **Custom fields** — anything the standard fields miss: payment terms, governing law, data residency. Add a field with a name and what to look for, then **Scan contracts** re-reads every indexed contract and records the value with the verbatim quote it was read from. A blank cell honestly means the paper does not state it.
- **The preview** — the first rows exactly as the workbook will carry them, with the contract and indexed counts. The download carries every row.
- **Send it to your SAM tool** — the SAM layouts drop into each tool's agreement import, and the Xensam connector under Settings › Integrations syncs its software inventory into License intelligence the other way.

How to use it

1. Pick a layout, or upload your template and check the column mapping it proposes.
2. Adjust the fields until the preview shows the columns you want.
3. Optionally add custom fields and press **Scan contracts**; the scan runs in the background and the coverage counts fill in as it reads.

4. Press **Download Excel**. A fresh workbook carries a Contracts sheet and an About sheet; a filled template comes back as your own file with the data under its headers.

Tips

- Review the file in your SAM tool's import preview before committing the import; every tool maps a few columns its own way.
- Money exports in the original contract currency and in USD at today's rate, each as its own column, so pick whichever your report needs.
- A record joins the scan once document search has read its files; the Indexed column shows exactly which rows that is.

Related: Document search · Vera AI with your SAM tool · The integration catalog

CLOUD COST OPTIMIZER

Where: Tooling › Cloud cost optimizer · `/tooling/cloud-cost` · via **More** or Ctrl/⌘ K
Who: everyone · your own exports require a Professional or Enterprise plan
You get: an executive brief, a full FinOps report, and a phased action plan · a minute or two after the export parses

In plain words: your cloud bill read for waste, with the findings staged and the plan phased.

When to reach for this

- The monthly cloud bill grew again and nobody can say exactly why.
- Before a committed-spend negotiation: know your optimized baseline before committing to the unoptimized one.
- The export parses in your browser, so the first pass costs nothing but the drop.

Pick your hyperscaler on the navy stage and drop a daily cost export from OCI, AWS, Azure, or GCP. Every run plays the FinOps scan: a full screen stage where a beam sweeps a miniature of your estate while a live ticker narrates real findings parsed from your file in the browser, then an executive reveal with the figures counting up. Behind it sit a savings dashboard, an executive brief, a full FinOps report, and a phased action plan you can track.

On the screen

- **Before you commit**: a card above the stage pointing at the Commitment Playbook, the ten part research series on sizing and running a committed spend agreement. It opens the document the series anchors on your vendor, and links the whole series, so you can read the method before the tool hands you a number.
- **The stage**: choose the cloud from the four provider marks and drop the cost export on the zone that lights up. The file is parsed in your browser and never uploaded.
- **The FinOps scan**: while the AI writes (one to three minutes on a real upload), the estate miniature colors in dot by dot (green healthy, gold rising, brick always on, blue new arrivals) and the ticker calls out your run rate, 24/7 bands, weekend deltas, and spike days. The reveal shows the verdict and the top findings, then opens the report.

- **The dashboard**: count up figures, the money bridge waterfall from today's run rate to the optimized number, a weekday by week spend rhythm heatmap with spike days ringed in gold, per service sparklines, and your recoverable share placed against the market's 18 to 25 percent band. Hover the daily spend chart for a crosshair with every service's cost on that exact day.
- **Findings triage**: mark each top finding Pursue, Later, or Dismiss; the choice is saved for your whole team and the tally shows what the pursued set is worth per year.
- **Dial the ambition**: drag between the Conservative, Realistic, and Aggressive scenarios and watch the annual savings, ROI, and payback track the thumb.
- **Executive brief and full report**: the short read and the deep read.
- **Share and forward**: mint a 30 day read-only link to the brief, or send it by email as an executive brief with the headline figures and top findings. Both carry aggregates and the written assessment only; the cost export itself never left your browser. Links can be revoked at any time.
- **Action plan**: a phased plan with a momentum header: the overall completion ring, per phase progress, and a savings activated tally that counts up as steps are checked off.
- **What next**: the result ends with a What next module seeded from the run's own numbers, so the follow-through starts from this estate rather than a generic list.
- **Upload-over-upload tracking**: a resource-level export unlocks a per-instance explorer, and your next upload adds a since your last export ribbon with the run-rate delta and the biggest movers.
- **Open the workroom**: a navy band above the result opens all six views as one case file on a dark stage: the executive brief and the full report as papers, the dashboard and the resource ledger as exhibits, and the action plan and the commitment sizing as action papers. The ⚡ dock beside the paper splits the stage and carries the moves, and the ✦ dock is Ask Vera grounded on the whole analysis. The rail lists your other runs, searches your uploaded documents, and keeps your shelf.
- **The moves in the dock**: draft the commitment ask to your provider (the level your run rate implies, with the discount schedule, rollover, true-down right, and overage rate asked for by name), draft the CFO savings memo (the measured savings range with the evidence under each spend pattern), hand the action plan to its owners split by the owner each step names, book the change into your renewal calendar, jump to the commitment levers to rerun the sizing, watch drift against the plan, mint the share link, copy the plan, and export the brief to Word. Nothing sends from the workroom and no alert is armed: drift is the run-rate delta between your last export and the next one, and it says whether the plan is landing.

How to use it

1. Select the provider on the stage and upload the daily cost export (or press Try with sample data to watch the whole ceremony first).
2. Watch the scan, read the reveal, then open the full report.
3. Read the dashboard: the money bridge says what the number is, the heatmap says where the rhythm is wrong. Triage the top findings and dial the scenario that matches your appetite.
4. Press **Open the workroom** and read the brief with the ⚡ dock open beside it; draft the CFO memo and the commitment ask from there.

5. Work the phased plan and watch the savings activated tally climb; re-upload later to track the run rate.
6. Share the brief with the budget owner by link or email when the case is ready.

Related: Workload comparator

WORKLOAD COMPARATOR

Where: Tooling › Workload comparator · `/tooling/workload-compare` · via **More** or Ctrl/⌘ K
Who: everyone · your own data requires a Professional or Enterprise plan
You get: a placement verdict over your whole estate, every workload priced across the four clouds, and the moves ranked by saving and effort · the math is instant and runs in your browser, the AI report takes a minute or two

In plain words: every workload priced on all four clouds, and the ones on the wrong platform named.

When to reach for this

- A cloud negotiation needs a credible alternative: the cross-cloud price is the evidence.
- Platform strategy season: where the estate would sit if placed by price rather than history.
- A consolidation question (the recipe) where workloads, not licenses, are the unit.

Save analysis files the placement verdict brief and its cost atlas to your Documents library, where it reopens at `/tooling/analyses/workload-compare` without a re-run. Your working session also stays in this browser per organization.

Describe your workloads or import an inventory CSV and the comparator prices each one on AWS, Azure, GCP, and OCI at list, then opens with the verdict: the optimal multi-cloud mix, the annualized saving, how many workloads sit on the wrong platform today, and a three sentence executive read written live from your numbers. A sample portfolio loads on arrival so every panel is visible before you upload anything.

On the screen

- **Placement verdict**: the optimal mix per month, the saving annualized, the misplaced count, and the widest cross-platform spread, with an estate shape line (how much spend runs 24/7, what is licensing, egress, and GPU).
- **What if switches**: four scenario toggles that re-price the whole estate live: owned Windows/SQL licenses (Azure Hybrid Benefit), a negotiated 20% discount, doubled egress, and a one year commitment cap. A chip shows what the scenario saves or adds against plain list.
- **Executive read**: three analyst sentences over your numbers: what the mix is worth, the single biggest lever, and what to do first.
- **Open the workroom**: once results exist, a band under the executive read opens the comparison as a six document case file, recomposed live from the current comparison: all four contenders' cost lines as exhibits, the cost atlas with per line deltas, the verdict brief (with the four single platform totals and the optimal mix drawn as a labelled bar chart), the switch sheet with the one time cost, the breakeven month, and the cumulative net position at twelve, twenty four, and thirty six months, and a What now sheet with a full figures stakeholder draft and links to the four clouds' commitment benchmarks. The same moves

ride the ⚡ dock beside every paper, the ✦ dock is Ask Vera grounded on the whole comparison, and the rail searches your uploaded documents and keeps your shelf. Moves copy to your clipboard or open in your own inbox; nothing sends.

- **Estate map**: every workload as a block sized by monthly cost, colored by where it runs today or where it should run, with the movers outlined in gold. Click a block to open its breakdown.
- **Cheapest platform per workload**: the comparison matrix: cells shade with cost, non-winners show the premium over the winner, each workload carries its cost driver chip (compute, licensing, GPU, storage, or egress led), and every column sorts.
- **The full picture**: the selected workload's cost anatomy on each cloud, the on-demand and commitment prices, a negotiation angle with a one to five leverage meter, and a migration effort read.
- **Where the money should move**: flow ribbons from today's platforms to the optimal placement, sized by monthly cost.
- **Where the saving comes from**: a waterfall splitting the saving between pricing discipline (commitments and spot, no migration needed) and replatforming.
- **36-month projection**: cumulative spend on today's path against the optimal mix, with an adjustable one-time migration estimate and the breakeven month marked.
- **Which moves to make first**: each move plotted by monthly saving against migration effort; start in the quick wins corner.
- **Commitment plan, data gravity, and the price of one cloud**: which pricing vehicle each cloud's share rides on, the portfolio's egress bill on each platform, and the premium every single-cloud consolidation pays over the mix.
- **Exports**: the full quote set as CSV, a copyable executive summary, and the AI FinOps placement report.
- **What next**: the result ends with a What next module seeded from the run's own numbers.
- **Your session keeps**: the portfolio, switches, and results stay in this browser per organization, so leaving or reloading does not lose the run; **Start over** clears it.

How to use it

1. Describe or import your workloads, honest sizes, egress, and licensing give an honest answer.
2. Read the verdict and the executive read, then click through the estate map and the matrix.
3. Flip the what if switches that match your reality, owned licenses or a negotiated discount change winners.
4. Press **Open the workroom** and read the verdict brief and the switch sheet with the ⚡ dock open beside them; draft the stakeholder brief from there.
5. Work the quick wins corner first, and take the flows, waterfall, and CSV into your architecture and pricing talks.

Related: Cloud cost optimizer

M365 OPTIMIZATION

Where: Tooling › M365 optimization · `/tooling/m365-optimization` · via **More** or Ctrl/⌘ K
Who: everyone · your own data requires a Professional or Enterprise plan
You get: inactive seats, tier right-sizing, Copilot sizing, and a per-user reassignment plan · instant, analyzed in your browser

Save analysis files the playbook to your Documents library, where it reopens at `/tooling/analyses/m365-optimization` without a re-run. The per-user action list stays in your browser and is never part of the saved copy.

Upload your Microsoft 365 usage export to find inactive accounts, right-size E7/E5/E3/E1 tiers, and size Copilot. Generating the playbook opens the estate scan, a full screen reveal where every licensed user is a dot that colors in as its verdict lands, ending in an executive readout with the money bridge from today's spend to the optimized mix.

In plain words: your Microsoft seats against actual usage, with the surplus priced.

When to reach for this

- The EA renewal or true-up is inside a year: walk in with the surplus already found (the recipe).
- Copilot sizing is on the table and deserves usage data rather than the vendor's enthusiasm.
- Quarterly hygiene: inactive seats accrue quietly between renewals.

A worked example. The sample estate's Microsoft position is 1,800 E3, 350 E5, and 900 F3 seats, and the flagged exposure on the E3 pool alone runs $64,800 a year. That is the shape of finding this tool exists for: not a dramatic discovery, just seats nobody uses and tiers nobody right-sized, priced and turned into a per-user reassignment plan you can actually execute.

On the screen

- **Before you commit**: a card pointing at the Commitment Playbook, the ten part research series on sizing and running a committed spend agreement. On this screen it opens the Microsoft document, on ramps, co-terms and the multi-year shape, and links the rest of the series.
- **Usage upload**: your M365 usage export, analyzed in your browser so your user data never leaves the page. A sample 500 user estate is one click away if you want to see the result first. Your session stays in this browser per organization, so leaving or reloading does not lose the analysis; **Start over** clears it.
- **The estate scan**: the reveal that plays when you generate the playbook: the scan beam, the live findings ticker, the headline numbers with your waste share placed against peer estates, the money bridge (stamped when the math closes), the seat-flow map of where every seat goes, and the largest individual moves. A renewal inside 120 days shows a countdown chip. Skip it with one click or replay it any time with **Replay the reveal**.
- **Verdict band**: the four numbers that matter (annual spend, identified savings, net position, accounts flagged) pinned above the playbook. Drag the inactivity threshold anywhere from 14 to 180 days, or use the 30/60/90 presets, and flip the Copilot appetite; every number in the band and the report recomputes in place.

- **Scenarios A and B**: freeze the current assumptions into a slot, change them, and read the two positions side by side with deltas. Load either scenario to put the whole playbook on its numbers.
- **The playbook**: the print-ready report, opening on an executive cover page: estate inventory, the silence curve, the what-Microsoft-sees table (assigned seats against the seats actually working, per tier), the per-department waste rollup with the worst offender named (when your export carries a Department column), tier right-sizing with evidence, Copilot sizing with a per-appetite return model, the savings plan line by line, and the 12-month realization curve of what lands when.
- **Share readout**: a 30-day read-only link to the executive readout for a CFO or a colleague. It carries aggregates only, never user names; the per-user plan stays in your workspace.
- **Send to EA playbook**: carries the optimized seat counts into the Microsoft EA renewal playbook, where they appear as your opening position for the quantity conversation.
- **Reassignment plan (CSV)**: the per-user action list, ready for your admin team.
- **Open the workroom**: a navy band above the result opens the playbook and the seat commitment sheet as one case file on a dark stage, with the ⚡ dock beside the paper and the ✦ dock for Ask Vera grounded on the whole estate. The rail searches your uploaded documents and keeps your shelf.
- **The moves in the dock**: draft the seat true-down ask to Microsoft (your optimized seat count against today's assignment, asking for the true-down right, price protection, and the published tier rates by name), draft the CFO savings memo, hand the three reassignment waves to IT with their counts, book the renewal into your calendar, mint the share link, export the reassignment CSV, copy the savings plan, and run the EA playbook handoff without leaving the paper. No draft names a person; the user list never leaves your browser.

How to use it

1. Upload the usage export (or start from the sample data).
2. Answer the questionnaire and generate the playbook; watch the estate scan land the verdict.
3. Scrub the threshold and Copilot settings in the verdict band until the assumptions match your reality; freeze scenarios A and B if two positions are on the table.
4. Press **Open the workroom** and draft the seat position for Microsoft and the CFO memo with the playbook still on screen.
5. Share the readout with the budget owner, send the optimized counts to the EA playbook, and work the reassignment list into your next true-up or renewal.

Related: EA renewal playbook · Spend optimization

SPEND STUDIO

Where: Tooling › Spend Studio · `/tooling/spend-studio` · via **More** or Ctrl/⌘ K
Who: everyone · sizing on your own usage export requires a Professional or Enterprise plan
You get: one money workspace with four AI vendor books behind a picker, plus launch desks for the portfolio spend plan, the cloud tools, and the license right-sizers · a minute or two per book

Spend Studio is the one place to size a committed-spend deal with an AI vendor and to work the rest of your savings tools. It opens on four tabs.

On the screen

- **Before you price the AI**: a card above the tabs pointing at the AI Pricing Playbook, the ten part research series on buying AI when the meter is new: translating seats, tokens, credits and outcomes into one comparable number, what counts as a billable outcome, credit expiry and revaluation, the AI clause set, and what to do when the model behind the endpoint changes. It opens the document anchored on whichever AI vendor you have selected, so picking Claude, OpenAI or Copilot changes what it shows.
- **Before you commit**: a card above the tabs pointing at the Commitment Playbook, the ten part research series on sizing and running a committed spend agreement: how to size the number off your own forecast, what a bigger commitment is actually worth, what happens to what you do not spend, and how to run the drawdown review afterwards.
- **Commit sizing**: pick the vendor, Claude, OpenAI, Copilot Cowork, or Agentforce, and that vendor's book opens with its own price book. Each book reads your usage export (or sizes workloads from prompt counts), adds the seat side of the contract, runs the commit-readiness gate, and prices the minimum, medium, and maximum commitment before you sign. The four books are documented in full below. Your choice stays in the link, so `/tooling/spend-studio?vendor=claude` opens straight into the Claude book and can be shared.
- **Plan**: opens Spend optimization, the all-vendor plan that ranks your whole stack by annual spend and writes the prioritized savings moves.
- **Cloud**: the Cloud cost optimizer and the Workload comparator, for cloud placement and FinOps savings.
- **Licenses**: the M365 optimizer and the VCF licensing right-sizer.

How to use it

1. Open Spend Studio and, under **Commit sizing**, pick the vendor you are negotiating with.
2. Drop in that vendor's usage export (or start on the sample data) and work the book to a sized commitment.
3. Use the Plan, Cloud, and Licenses tabs to reach the portfolio plan and the cloud and license tools from the same desk.

Related: Spend optimization · Cloud cost optimizer

AGENTFORCE OPTIMIZER

Where: Spend Studio › Commit sizing › Agentforce · `/tooling/spend-studio?vendor=agentforce` (old `/tooling/agentforce` links forward here)
Who: everyone · your own data requires a Professional or Enterprise plan
You get: a commit verdict with a robustness score, three priced commitment levels, a quote check, a drafted negotiation ask, and the commit workroom · the model updates as you type, the AI report takes a minute or two

The Agentforce book inside Spend Studio sizes a Salesforce Flex Credit commitment from your own consumption. Feed it your Salesforce Digital Wallet export (from a POC or production), or type your monthly credit burn if agents are not live yet, then set the term, growth, and planned use cases, and the tool prices the minimum, medium, and growth commitments against the forecast. The page opens on sample data, clearly marked, so every panel is visible before any upload; nothing you upload leaves your browser.

On the screen

- **Five numbered steps**: the model's inputs down the left. **Consumption data** takes the Digital Wallet CSV or a typed monthly credit burn; **Term & growth plans** sets the term, organic growth, and the planned new use cases (case deflection, a voice agent, and more), each with its monthly actions, start month, and ramp; **Risk posture** sets which forecasting mistake stings more; **Negotiation levers** holds the negotiated discount, credit rollover, and the scenario odds; and **Check Salesforce's quote** prices any quoted commitment and discount against your own forecast, flagging shelfware when the quote sits well above your expected burn and true-up overage when it sits below.
- **Commit verdict**: the recommended level, the annual commitment, expected savings versus pay as you go, expected utilization against its break-even point, and a plan robustness score: how many of roughly thirty perturbed futures still pick the same level.
- **Commitment levels**: minimum, medium, and growth cards, each with expected utilization against its break-even tick and where the expected money goes (committed spend, unused spend, overage).
- **Forecast chart**: credit burn drawn as an uncertainty cone from conservative to aggressive against the commit levels, with overage shading and a monthly headroom strip.
- **Discount staircase**: the three levels placed on the volume-discount ladder, with a nudge when committing slightly more costs less.
- **What drives the burn**: annual credits and yearly cost by source, baseline versus each planned use case.
- **Total cost by scenario**: every level priced across the conservative, expected, and aggressive paths, with a pay-as-you-go reference.
- **Your negotiation ask**: a copy-ready note with the sized number and the terms worth asking for.
- **Save analysis and AI report**: file the sized result into your Documents library, or generate the full commitment-sizing white paper.
- **Open the workroom**: opens the sizing report as a case file on a dark stage, with the ⚡ dock beside it and the ✦ dock for Ask Vera grounded on the sizing, including how each level lands under each demand scenario. The dock drafts the Flex Credit commitment ask to Salesforce (the credits you will commit on your term, with rollover, the overage rate, a true-down right, and the discount as a rate per credit asked for by name), drafts the CFO commitment memo, copies the whole commitment position, books the term into your renewal calendar, and opens the Salesforce room or the Deal Room. When a commitment does not beat pay as you go, the memo says so plainly.

How to use it

1. Drop in your POC or production Digital Wallet export, or type your monthly credit burn; the sample stays on until you do.
2. Work down the steps: term, growth, and planned use cases, then the risk posture and the negotiation levers. The verdict, the three levels, and the charts update as you type.
3. When Salesforce quotes you a number, paste it into **Check Salesforce's quote** to see how it prices against your own forecast.
4. Press **Open the workroom** and draft the commitment ask and the CFO memo with the sizing report on screen.
5. File the result with **Save analysis**, or generate the AI report for the full white paper.

Related: Salesforce SELA

VCF LICENSING

Where: Tooling › VCF licensing · `/tooling/vcf-licensing` · via **More** or Ctrl/⌘ K
Who: everyone
You get: your billable core count, three priced scenarios, and the negotiation playbook · counts are instant in your browser, the AI playbook takes a minute or two

In plain words: your real billable VCF core count under Broadcom's rules, and the renewal scenarios priced.

When to reach for this

- The VMware renewal is coming and the 16-core minimum makes intuition wrong.
- You need the number before the meeting: counts are instant, in your browser, from the RVTools export you already have.

Upload your RVTools export to see exactly how many VCF cores Broadcom will bill you for under the 16-core minimum, build three pricing scenarios against real-world discount bands, and walk into the renewal with a verdict: the market curve, your leverage scored, the signing window dated, and an AI negotiation playbook.

On the screen

- **Upload stage**: drop the RVTools export on the dark stage; it is analyzed in your browser so your inventory never leaves the page. **Try with sample data** loads a full 48-host estate with spend and deadline so every screen renders before you upload anything.
- **Estate scan**: the parse narrated line by line: hosts found, sockets counted, the 16-core minimum applied, and the phantom cores it adds (cores billed that your hardware does not have) before the billable total counts up.
- **Savings odometer**: the money on the table at your target scenario, reacting live to every change in the builder.
- **Three scenarios**: conservative, target, and best case, each with a committed-core count, term, and a discount slider showing the benchmarked typical range; a verdict chip says whether your discount sits below, inside, or above that band.

- **Leverage scorecard**: six negotiation levers (exit alternative, benchmark anchor, full-estate commit, multi-year term, quarter-end timing, executive sponsor) checked off into a strength score, with one click to apply the earned lift to your best case.
- **When to sign**: Broadcom's fiscal quarter closes plotted with your deadline as a T-minus chip and a dated four-phase plan ending at the right quarter close.
- **Market curve and tier ladder**: realized discounts by committed cores with your estate marked, all five volume tiers listed, and a flag when you are close enough to price the deal one band deeper.
- **Cost heatmap**: annual cost at every term and discount step, the typical band outlined, your scenarios plotted.
- **Verdict band**: the analysis opens on the four numbers that matter (billable cores, list exposure, target, savings) over a deal journey from list price down to best case.
- **Estate evidence**: a five-year cumulative cost chart, the 16-core minimum drawn host by host with phantom cores as dashed cells, a cluster strip, and a sortable host explorer with the worst phantom offenders first.
- **Consolidation lab**: toggle keep, consolidate, or retire per cluster and watch the core count, band, and annual cost reprice.
- **The walk-away, priced**: your target rate against directional per-core costs of Nutanix, OpenShift, Azure Local, and Proxmox, plus the vSAN entitlement and the VVF crossover if you run compute only. The card links the third-party support playbook for pricing that exit in full.
- **What next**: the result ends with a What next module seeded from the run's own numbers.
- **Verdict card and exec summary**: export a navy verdict card as a PNG or copy a one-paragraph exec summary for email.
- **Playbook**: the AI negotiation guidance for the reprice; only aggregate totals are sent, never host names.
- **Save analysis**: file the result into your Documents library as a durable document; host names are stripped before anything is stored.
- **Open the workroom**: a navy band above the verdict opens the licensing analysis as a case file on a dark stage. The ✦ dock is Ask Vera grounded on your parsed estate: the hosts, the clusters, the billable cores, the phantom cores, the volume band, and the three priced landing zones. The rail searches your uploaded documents and keeps your shelf.
- **The moves in the dock**: draft the true-down ask to Broadcom (your own core count, the phantom cores priced, and relief on the hosts under the minimum, a true-down right at each anniversary, the vSAN entitlement in writing, and price protection for the term asked for by name), draft the CFO memo, book the renewal into your calendar, export the host worksheet, copy the three landing zones, and open the Deal Room or the Broadcom room.

How to use it

1. Upload the RVTools export (or load the sample) and watch the estate scan land on your billable core count.

2. Fill in the renewal facts, check off the leverage you actually have, and tune the three scenarios against the shaded benchmark bands.
3. Generate the analysis, test consolidation moves and the walk-away, then take the verdict card, the exec summary, and the playbook into the negotiation.
4. Press **Open the workroom** and draft the true-down ask with the analysis on screen, or put a question to Vera about the estate from the ✦ dock.

Related: Playbooks

ORACLE ERP CLOUD PRICING CALCULATOR

Where: Tooling › Oracle ERP pricing · `/tooling/oracle-erp-pricing` · via **More** or Ctrl/⌘ K
Who: everyone
You get: the Fusion ERP order priced, the term cost under an uplift, and the deal's tier on the Fusion discount ladder · instant, in your browser

Build the Oracle Fusion ERP order the way Oracle quotes it: pick the cloud services on your proposal, enter the quantities, and apply a discount deal wide or line by line. The calculator prices the subscription, shows what the term costs once the annual uplift compounds, and places the net annual fee on the Fusion discount ladder so you can see whether the discount is good for a deal your size.

The catalog carries Oracle's own published figures for every component: the list price, the metric it is sold on, the stated minimum quantity, and the Oracle part number, read from the Oracle Fusion Cloud Service Global Price List. A weekly job re-reads that list from Oracle and updates the prices here when Oracle republishes it, so the date printed under the catalog is the list you are actually pricing against. List is a ceiling anchor, not a price anyone pays, which is what the discount and the ladder are for. Every price stays editable, so where your own proposal differs your figure wins. No order leaves your browser unless you choose to send it.

On the screen

- **Pick the cloud services**: the Fusion ERP catalog grouped by family (enterprise resource planning, tax, procurement, enterprise performance management, supply chain, warehouse and logistics, subscription management, analytics and AI, environments and options). Each tile shows Oracle's published list price, the metric it is sold on, and the stated minimum. Click a tile to put it on the order at its minimum quantity.
- **Your order**: one row per component with the Oracle part number, the metric, the quantity, the list price per unit, an optional per line discount that overrides the deal discount, and the resulting annual net. A quantity below Oracle's published minimum is flagged on the line, because Oracle writes the order at the minimum. Lines with no quantity yet are excluded from every total.
- **Apply the discount and the term**: the deal wide discount off list, the subscription term in years, and the annual uplift. Oracle's standard opening position is 8 percent uncapped, so the uplift starts there.
- **The four tiles**: annual list, net fees to Oracle per year, the term total with the share of it that is uplift, and the ladder tier with its median and 90th percentile.

- **Market position**: the full Fusion discount ladder with your tier highlighted and your blended discount marked on it, plus a plain reading of where you stand. Tiers band on net fees payable to Oracle per year, the annual subscription value after discount, not on the list value of the components.
- **The priced report**: the executive brief version of the order, the line table, the ladder, the year by year term cost, and the four moves to make first. It prints and forwards as one document.
- **Send the report**: email the priced order to up to ten colleagues with an optional note. The email carries the figures on the page.
- **Invite others to this calculator**: email colleagues a link to the calculator itself. The invitation carries no figures from your order and grants no access to your workspace.

How to use it

1. Open the calculator and press **Try with a sample order** to see every panel, or start picking components straight away.
2. Add each cloud service on the Oracle proposal, then set the quantity. The list price fills itself from Oracle's published list; overwrite it where your proposal quotes a different figure.
3. Set the discount, the term, and the uplift. Give any line its own discount where Oracle discounted it differently, which is common on non-production environments.
4. Read the market position: your tier, the median and 90th percentile for it, and the gap in dollars a year.
5. Send the report to the people who need it, or invite them to price their own view.

Related: Oracle HCM Cloud pricing calculator · Oracle ULA

ORACLE HCM CLOUD PRICING CALCULATOR

Where: Tooling › Oracle HCM pricing · `/tooling/oracle-hcm-pricing` · via **More** or Ctrl/⌘ K
Who: everyone
You get: the Fusion HCM order priced, the term cost under an uplift, and the deal's tier on the Fusion discount ladder · instant, in your browser

The same instrument for the HCM side of Fusion, on the same published price list. Pick the cloud services on your proposal, enter the quantities, apply a discount deal wide or line by line, and the calculator prices the subscription and places the net annual fee on the Fusion discount ladder. ERP Cloud and HCM Cloud ride the same discount curve, so a combined subscription negotiated as one deal reaches tiers that neither pillar reaches alone.

On the screen

- **Pick the cloud services**: the Fusion HCM catalog grouped by family (core HR, talent, workforce management, payroll and compensation, analytics and AI, environments and options), every line carrying Oracle's published list price, metric, minimum, and part number. The nine country payroll services are listed separately, as Oracle prices them.

- **Your order**: one row per component with the part number, the metric, the quantity in that metric, the list price per unit, an optional per line discount, and the resulting annual net.
- **Apply the discount and the term**: the deal wide discount off list, the subscription term, and the annual uplift.
- **The four tiles**: annual list, net fees to Oracle per year, the term total, and the ladder tier.
- **Market position**: the Fusion discount ladder with your tier highlighted and your blended discount marked on it.
- **The priced report**, **Send the report**, and **Invite others to this calculator** work exactly as they do on the ERP calculator.

How to use it

1. Open the calculator and press **Try with a sample order**, or start picking components.
2. Add each HCM cloud service on the proposal and set the quantity in that line's metric, usually employees. The published list price fills itself; overwrite it where your proposal differs.
3. Set the discount, the term, and the uplift, then read your tier and the distance to its 90th percentile.
4. Send the report, or invite colleagues to the calculator.

Related: Oracle ERP Cloud pricing calculator · Oracle ULA

CLAUDE OPTIMIZER

Where: Spend Studio › Commit sizing › Claude · `/tooling/spend-studio?vendor=claude`
Who: everyone · your own data requires a Professional or Enterprise plan
You get: a spend scan reveal, a commit-readiness check, and a committed-spend deal sized at three levels · a minute or two

Upload your Claude Console usage export or size each workload from prompt counts, add the seat side of the contract, run the commit-readiness check, and size the committed-spend deal, minimum, medium, or maximum, before you sign. Everything is analyzed in your browser; the export never leaves the page.

On the screen

- **The spend scan**: when an export (or the sample) lands, a full-screen reveal scans your bill: one dot per slice of monthly spend colors in under the beam (brick means a premium tier doing commodity work, gold means a cacheable prompt prefix, blue means batchable traffic, green means right-priced) while a ticker calls out findings from your file. It closes on an executive readout: the verdict, a strip placing your optimization headroom against the Claude estates we size, the money bridge from today's run-rate to the optimized mix, and the commit podium with the level worth signing stamped. Press Escape or Skip to jump ahead; **Replay the reveal** brings it back.
- **Live verdict band**: the four numbers that matter (run-rate today, optimized mix, the commit to sign, the saving vs pay-as-you-go) with a verdict pill, plus two **what-if switches** that flip prompt caching or the Batch API off so every number on the page snaps back to the unoptimized bill.
- **Usage or prompt profiles**: upload the Console export, or size workloads from prompt counts.

- **Model right-sizing**: each workload on the cheapest model that clears its quality bar, with a lever strip attributing the saving to re-tiering, prompt caching, and the Batch API, a Sonnet 5 introductory-pricing note while the promotion runs, and a spend-flow map tracing where the traffic should move.
- **The commit dial**: drag across commitment sizes and watch expected cost, unused risk, and overage recompute on a cost curve with the sweet spot marked; a growth slider in step 4 re-flexes the whole forecast the same way.
- **Take these to the table**: the negotiation angles your numbers support (band anchoring, rollover, overage at the committed rate, a ramp), ranked by what each is worth.
- **When the savings land**: the realization curve: caching first, re-tiering next, batch after, and what year one captures.
- **What one unit costs**: per-workload cards pricing one request, message, or run today versus on the optimized mix.
- **Seats**: add Team or Enterprise seats to the model.
- **Commit-readiness check**: whether your run rate is stable enough to commit, scored on a ring gauge with the open steps listed.
- **Share, hand off, export**: **Share readout** mints a 30-day read-only link for a CFO (aggregates only, the export never leaves your browser), **Send to EA playbook** opens the Anthropic EA playbook from your sized commit, **Board slide** downloads a one-page PPTX readout, and **Save analysis** files the result into your Documents library.

How to use it

1. Upload the usage export (or load the sample) and watch the spend scan land on your verdict.
2. Classify each workload in step 2 so the right-sizing and the levers price correctly.
3. Add the seat side, scrub the growth and commit dials until the numbers hold, and run the readiness check.
4. Forward the readout link or the board slide, then take the sized commitment into the agreement, or send it straight into the Anthropic EA playbook.

Related: Anthropic EA playbook · OpenAI optimizer

OPENAI OPTIMIZER

Where: Spend Studio › Commit sizing › OpenAI · `/tooling/spend-studio?vendor=openai`
Who: everyone · your own data requires a Professional or Enterprise plan
You get: a commit-readiness check and a committed-spend deal sized at three levels · a minute or two

Upload your OpenAI usage export or size each workload from prompt counts, add the ChatGPT seat side of the contract, run the commit-readiness check, and size the committed-spend deal, minimum, medium, or maximum, before you sign.

On the screen

- **Usage or prompt profiles**: upload the usage export, or size workloads from prompt counts.
- **Seats**: add ChatGPT Team or Enterprise seats.
- **Commit-readiness check**: whether your run rate is ready to commit.
- **Deal sizes**: the committed-spend deal at minimum, medium, and maximum.
- **Save analysis**: file the sized result into your Documents library as a durable document you can reopen any time.

How to use it

1. Upload the usage export or build workloads from prompt profiles.
2. Add the seat side and run the readiness check.
3. Take the sized commitment into the deal.

Related: Claude optimizer

COPILOT COWORK OPTIMIZER

Where: Spend Studio › Commit sizing › Copilot Cowork · `/tooling/spend-studio?vendor=copilot`
Who: everyone · the AI report on your own inputs requires a Professional or Enterprise plan
You get: a twelve-month credit forecast, three sized commit options, a readiness verdict, and an AI cost report · a minute or two

Copilot Cowork, Microsoft's agent layer inside Microsoft 365 Copilot, bills metered Copilot Credits on top of every Copilot seat: the first Microsoft 365 line that behaves like cloud consumption instead of a seat. This tool models what it will actually cost before you enable it or prepay for it. Facts current as of July 2026; the live tier sheet in Azure reservations governs at purchase.

On the screen

- **Your Cowork plan**: enabled users, how many already hold Copilot seats, weekly active share, tasks per active user, adoption growth, and the task mix (light, medium, heavy, each with editable credits per task). Everything recalculates as you type.
- **The forecast KPIs**: month-one and year-one credits, the pay-as-you-go year cost at $0.01 per credit, and the all-in figure including any new Copilot seats the rollout forces.
- **The twelve-month burn**: monthly credit consumption with your adoption growth compounding each quarter.
- **Commit options**: Floor, Balanced, and Ceiling P3 pre-purchase sizes against the forecast, each with its tier discount, upfront cost, expected pay-as-you-go overflow, and the expiry risk if usage lands under forecast. Prepaid units expire at term end and purchases are final, which is why the tool sizes commits to measured burn, never ambition.
- **Commit readiness, five checks**: a measured pilot, spending caps configured, a named budget owner, named workloads, and the EA's Reserved Instances policy enabled. The score decides which commit option the tool recommends, including recommending none.

- **The seat prerequisite and the E7 question**: the incremental Copilot seats a rollout forces, and the E7 math ($99 bundles E5, Copilot, Entra Suite, and Agent 365, about $18 under the parts, with no Cowork credits included).
- **The asks**: rate protection, rollover or credit-back, auto-renew off, an adoption-gated second purchase, and the effective discount stated against your existing discounts.

How to use it

1. Enter the pilot you actually intend to run, not the org chart, and tune the task mix.
2. Read the forecast and the commit options; answer the five readiness checks honestly.
3. Follow the recommendation: pay as you go until the checks pass, then commit at the floor.
4. Generate the AI report and take the asks into the agreement conversation.

Tips

- Type `/cost` inside Cowork during the pilot to see real per-task credit costs, then replace the band defaults with your measured numbers.
- The Microsoft EA playbook's Copilot Decision screen carries a compact version of this model, so the renewal case and the Cowork sizing reconcile.

Related: Microsoft EA playbook · M365 optimization

VENDOR PLAYBOOKS

Where: Tooling › All tools › Renewal playbooks · `/tooling#wall-renewal-playbooks`
Who: everyone · your own data requires a Professional or Enterprise plan
You get: a guided engagement you can pause and resume, with live models and an executive report at the end · runs over days, not minutes

In plain words: a full renewal engagement for one vendor, run yourself over days, ending in an executive report.

When to reach for this

- The big-vendor renewal (Microsoft EA, SAP RISE, Oracle ULA, and their peers) where the engagement deserves weeks, not an afternoon.
- You would have hired help for this deal: the playbook is the methodology, self-run.
- Learning a vendor's game before playing it: each playbook teaches while it works.

Full renewal and migration engagements, run yourself. The playbooks live as cards in the tool wall's Renewal playbooks group; the old separate Playbooks page redirects there. Each playbook is a guided Prepare, Analyze, Execute, Deliver journey for one vendor, with live recalculation as you change inputs, work-back timelines, a negotiation desk, and an executive report at the end. Everything runs in your browser and progress saves as you go; every playbook works fully on its sample estate on any plan, and uploading your own data is the Professional or Enterprise step. Every playbook ends on a What Next desk rather than the report.

On the screen

- **Playbook cards**: one card per vendor engagement, ten in all.
- **The stages**: inside each, the Prepare, Analyze, Execute, Deliver path with its live models.
- **Executive report**: the CFO business case and the exec report the journey produces. The report screen closes with a **Take this forward** row: start a negotiation, open the vendor room, or attach the report to a contract, with a copyable target line for the deal.
- **Your own records at the end**: on the Microsoft EA, RISE with SAP, Oracle ULA certification, and IBM ELA playbooks, the closing desk reads your workspace rather than offering generic links. If a case or negotiation is already live on that publisher, it says so and links it instead of offering to start a duplicate. If a matching agreement is on file, **Open the {vendor} renewal record** goes to your own renewal, with the notice clock and the paper of this cycle.
- **File this target on the renewal**: on those same four playbooks, one button writes the target the playbook just computed, and the savings behind it, into that vendor's discussion, dated and attributed to you, so the number survives the browser tab. The note lands on the shared vendor discussion, so it shows on the renewal record and the vendor page alike. It appears only where a real agreement in your workspace matches the publisher, and never on the sample walkthrough; there is no confirmation step and no undo, so read the figure before you press it.

Because every playbook shares that shape, each entry below gives its route, what the engagement is for, the key moves it walks you through, and when to reach for it.

Microsoft EA

`/tooling/ms-ea-renewal` · Run a Microsoft Enterprise Agreement renewal from scoping to the CFO business case, solo in your browser or as a team project the whole deal team works on together, with a built-in consultant that adapts every screen to your answers.

The Microsoft playbook runs like a real engagement, and now it opens like one too. **The first open is a guided kickoff in three steps:** the Discovery interview first, then Opus 5 drafts your project plan from your answers and your numbers, then a confirm step where you check the four fields every date derives from (client, agreement, anniversary, where in the term you are), adjust anything, regenerate if the anniversary moved, and step onto Mission Control. You can skip the kickoff with one click and run it later from the Discovery screen; either way the playbook remembers and opens on Mission Control from then on. Team projects, shared view links, and screen deep links never see the kickoff.

The interview itself works the same wherever you run it (goals, the Microsoft relationship, constraints, Copilot intent, Azure posture), and you choose how: answer the structured questions yourself, or **talk it through with Vera**, by voice (tap to talk) or typing, and she extracts your answers into the questionnaire live. Both modes fill the same answers, so you can start in one and finish in the other, and everything Vera captures stays reviewable in the form. Opt in and Vera saves a digest of the interview to her long-term memory (manageable on the Ask page), so she recalls your renewal anywhere in the app; the playbook state itself stays in your browser. As you answer, a live **What I'm hearing** panel plays the forming strategy back, then commits to a named posture such as Structured cost reduction or Leverage-building exit track. From that point every screen opens with **Your consultant's read**: what this page is for, what to look at first, and how your answers changed it, plus a leverage score out of 100 with the gaps to fix. The journey runs 32 focused screens across Prepare, Analyze, Execute, and Deliver, one idea per page, carried in a **left hand menu** grouped by phase: every screen always visible, each phase with its completion percentage, the active screen a filled blue pill (small windows keep the compact stepper instead). Everything explains itself on hover, from KPI tiles to table columns. The results ticker also carries a **Theatre** button: pick the light or the dark stage, or resume where the last show left off, and the whole playbook runs full screen, one screen at a time, everything still editable. The stage goes truly fullscreen, its chrome fades after a few seconds of stillness and returns on any movement, and arrow keys move between screens; crossing into a new phase raises a short title card with the act number and the running savings figure. Press G for a grid of all 32 screens and jump anywhere in one click, N for the presenter note (a one line analyst takeaway under each screen), and I for the ink pen: gold strokes drawn over the screen on a live call that fade on their own and clear when the screen changes. The live position numbers ride in the header and briefly show the delta each edit contributed, Ask Vera opens right over the stage, and the show lives in the page address, so a shared link drops a colleague onto the same screen of the same show. The leave-behind button prints exactly the screens you walked as a navy and gold packet. Escape steps back out one layer at a time to exactly where you were, and the stage choice is remembered. It opens with What's Changed, the research brief on everything that moved since your last signature (the July 2026 price increase, Cowork and metered credits, E7, the MCA-E push). The Analyze phase adds the Suite Ladder (E1 to E7 with six step-down questions), Substitutions (the components you may be paying twice for, with the cheaper Microsoft shape and the third-party alternative per row), and the BoM Benchmark (your bill of materials priced line by line against the peer unit-price corridor); the Azure screen teaches how to structure the commitment across reservations, savings plans, and the MACC. At the table, Opus 5 prepares the negotiation strategy pack (internal talking points, scripted lines for Microsoft, the trade list, the watchouts) and the two-to-four page brief for your CFO and CIO. After discovery, your project plan is one click away and hard to miss: Mission Control carries a **Your project plan** card where you generate it (and read the headline, the steps in order, and the key dates without leaving the overview), and the Renewal Plan screen renders the full dashboard (key dates on your calendar, workstreams with owners, internal questions, risks, and the three-number position), rule-based instantly and personalized by Opus 5 on one click. The Intake Analysis screen then has Opus judge your loaded evidence before you brief the advisor, and the Commit Scenarios screen prices four ways to commit side by side with a live recommendation.

The renewal cost corridor. Mission Control opens with the three numbers that frame the renewal, drawn as a waterfall: the worst case if the agreement is left to Microsoft (current spend plus the expected list uplift, plus Cowork credits drifting at pay as you go when your discovery says agents are in play), the current annual spend, and the best case with every identified saving landed, the concessions target negotiated, and any Cowork spend committed at the discounted ladder instead of drifting. The corridor recomputes live from the same state as every other screen, so each seat count you correct and each lever you set moves it, and the headline shows what the corridor is worth per year, the distance between doing nothing and running the playbook.

The Cowork Cost Lab (Analyze phase, after the Copilot Decision) prices the agent layer before Microsoft prices it for you. Set the rollout (users, seat coverage, weekly actives, tasks per user, quarterly growth), the task mix on Microsoft's published planning bands (light 70 to 200 credits, medium 400 to 600, heavy 1,500 and up), and the **model mix per LLM** at Microsoft's published rate tiers: mini models at 0.1 credits per 1,000 tokens, general models (GPT-4.1, GPT-5 chat, Claude Sonnet 4.6) at 1.5, deep reasoning models (GPT-5 reasoning, Claude Opus 4.6) at 10. The lab shows the model-mix multiplier, the cost per weighted task, the year-one burn, and the prerequisite Copilot seats the rollout quietly adds, then prices floor, balanced, and ceiling commitments against the published nine-tier P3 discount ladder (5% at 300K credits to 20% at 300M), with the expiry risk of every prepaid option stated next to its discount and a recommendation gated on your readiness score. Every rate is pinned to the Microsoft price documents listed at the top of the screen, verified on a stated date and re-fetched daily by Policy Watch: if Microsoft changes the price list, the screen flags it and links the change analysis on the Policy Watch page.

The Feature Heatmap (Analyze phase, after the Cowork Cost Lab) draws the whole suite argument in one picture once the ladder, the E5 profile, and the Copilot decision are worked: features down the side, the E1 to E7 rungs across the top, every included feature colored by how deeply your estate actually uses it. Levels derive from the same telemetry as the rest of the playbook (SKU utilization, the E5 segmentation, the Copilot adoption bands) and every row can be corrected by hand; corrections mark themselves and persist. The argument card composes the lines to say at the table, we only use this many of the features that separate E5 from E3, with the right-profiling dollars as the evidence value, and the send desk makes the grid travel: Vera drafts the feature-use email to Microsoft, the CFO brief, or a colleague note from the grid, and a monospaced text version copies straight into Teams or Slack.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone on the team works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, the four stages with progress, a you-are-here marker, the next step, and the anniversary countdown, expanding into a full Gantt of the 16 work-back activities with a today line. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox (they can test numbers, nothing they change is saved). Every screen carries a notes and corrections box, comments and corrections sync to the whole team, and the project log holds meeting notes and decisions as the renewal runs. Any page can be shared with a revocable view link from its notes box. The playbook's questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone, on the platform or not: the tokenized page opens pre-filled, and their answers merge back into the project under their name. When a phase completes, the project mails everyone an exec-brief formatted update (KPIs, phase table, what happens next, what is coming due); a weekly digest goes out Monday mornings, and every recipient chooses what they receive through a personal preferences link in each email footer, with the same toggles editable per person on the Project Team screen.

Starting from your contracts. The Data Intake screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the enrollment pack in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, evidenced quantities, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Apply what looks right with one click and start the interview from the contract instead of a blank form. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen: what to tell Microsoft from the first call. Five message windows from T minus 9 months to signature, each with a goal, lines to deliver roughly verbatim, and what must not leave the room; the five standing principles (one voice, share process not position, everything on the record); six lines that cost money with their replacements; and a scripted first meeting agenda. The read at the top tailors itself to your discovery answers.

The assessments. The Azure Commit, Copilot Decision, and Cowork panels each carry a weighted assessment bank (ten, ten, and eight questions) that scores your readiness out of 100 and ranks the fixes by what they are worth at the table; every question explains on hover why a veteran asks it.

The Deliver phase runs five screens: the Outcome Dashboard (KPI wall, savings mix, the waterfall from spend to target, preparation and assessment verdicts, the Azure landing, and what remains open, each tile clicking through), the Business Case, the Executive Report, What Next, and Updates & Sharing, where you send a project update on demand and see exactly who receives what.

The What Next desk ends the journey the way the Benchmark Case does: an open desk that puts the computed position to work, every move running on the playbook's live numbers. Six emails Vera drafts for you to edit and send from your own inbox: the Microsoft meeting request, the asks letter, the feature-use case, the CFO brief, the procurement frame, and the internal team alignment note; the Microsoft-facing drafts carry your asks and your own usage facts, never your savings target, the corridor, or the strategy. Around them: the copy-ready position digest and the heatmap text block for chat, the management brief generated right on the desk, one-click routes to the Executive Report and the CFO Business Case, Ask Vera with the position preloaded, and exit ramps into a tracked negotiation, the Microsoft vendor room, benchmarking, and an analyst session.

Key moves

- Run the guided kickoff: the discovery interview sets the strategy with you, Opus 5 drafts the project plan from it, and you confirm the timeline before Mission Control opens. Or start from your contracts and let Opus pre-fill the interview.
- Read the renewal cost corridor first: worst case, current, best case, and what the gap is worth, then watch it close as you work.
- Create a team project so the deal team, and invited outsiders, work one live state instead of a screenshot chain.
- Load the license estate, then work E5 right-profiling, the Copilot decision, the Cowork Cost Lab (the task mix, the per-LLM model mix, and the P3 commit ladder with its expiry risk), hygiene and SA benefits, and the Azure MACC, each on its own page, and run their assessments until the fixes list is empty.
- Build the Feature Heatmap once the suite screens are worked, correct any level you know better, and send the feature-use case to Microsoft from the grid.
- Follow the Comms Plan: set the frame at T minus 9 months, land the anchors before the first quote, and never name a budget.

- Apply the recommended lever settings derived from your interview, then stress-test them in the scenario studio.
- Take the peer benchmarks, program split, and a table plan tailored to Microsoft's past behavior with you into the rounds.
- Finish on the Outcome Dashboard, the CFO business case, and an executive report framed to the win you defined, then work the What Next desk: brief your own side first, book the Microsoft session, and let the project send the closing update itself.

When to use it: your Microsoft EA renewal is inside 12 months.

RISE with SAP

`/tooling/rise-with-sap` · Run the move from ECC or on-premise S/4HANA to RISE from scoping to the CFO business case, solo in your browser or as a team project the whole deal team works on together, with a built-in consultant that adapts every screen to your answers.

Like the Microsoft playbook, RISE runs as an engagement: the discovery interview (answer the questions, or talk it through with Vera by voice) commits to a named posture, every screen opens with **Your consultant's read** and a leverage score, the levers arrive pre-set from your answers with one-click apply, and the Negotiation Desk adds questions and counterplays for the tactics you report. The strategy layer is current to July 2026, including the EU's binding SAP commitments (no reinstatement fees, capped back-maintenance, landscape segmentation), which reshape the 2027 deadline conversation. The same planning layer as the Microsoft playbook applies: a Conversion Plan dashboard after discovery (Opus 5 personalizes it on one click), an Intake Analysis verdict on your evidence, and Commit Scenarios pricing RISE now, the third-party support bridge, the on-premise conversion, and staying put side by side with five-year totals; the alternatives card links the third-party support playbook it names, so the bridge option is one press from its full runbook. The project plan also lives on Mission Control: generate or regenerate it right on the overview and read the headline, the steps, and the key dates without leaving the front page; the discovery interview now also asks your BTP intent, SAP's platform attach, so the credits in the proposal get priced against a real roadmap rather than accepted as padding. Mission Control also leads with a **cost corridor**: a single waterfall showing what the move costs on SAP's opening proposal left unmanaged (worst case), what you pay on the current platform today (current run rate), and where the playbook lands the subscription once the FUE count is rebuilt, the discount is anchored to the peer corridor, BTP is right-sized, and the conversion credit is applied (best case). It recomputes as you work the screens, and the swing between the two ends is the number the playbook exists to close.

It opens like an engagement too. The first open is a guided kickoff in three steps: the Discovery interview (answer the questions yourself or talk it through with Vera), then Opus 5 drafts the project plan from your answers and your numbers, then a confirm step on the fields every date derives from (client, current platform, target signature, and the maintenance cliff) before Mission Control opens. Skip it with one click and the playbook opens on Mission Control from then on; team projects, shared view links, and screen deep links never see the kickoff. The screens are carried in a **left hand menu** grouped by phase, each phase with its completion percentage and the active screen a filled blue pill (small windows keep the compact stepper), and the results ticker carries a **Theatre** button: pick the light or the dark stage and the whole playbook runs full screen, one screen at a time, arrow keys to move, everything still editable, Escape to step back out exactly where you were. The plan stays live on the Conversion Plan screen.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, the four stages with progress, a you-are-here marker, the next step, and the countdown to your target signature, expanding into a full Gantt of the 16 work-back activities with a today line. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox (they can test numbers, nothing they change is saved). Every screen carries a notes and corrections box, comments and corrections sync to the whole team, and the project log holds meeting notes and decisions as the conversion runs. Any page can be shared with a revocable view link from its notes box. The playbook's questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone, on the platform or not: the tokenized page opens pre-filled, and their answers merge back into the project under their name. When a phase completes, the project mails everyone an exec-brief formatted update (KPIs, phase table, what happens next, what is coming due); a weekly digest goes out Monday mornings, and every recipient chooses what they receive through a personal preferences link in each email footer, with the same toggles editable per person on the Project Team screen.

Starting from your contracts. The Data Intake screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the license orders, the maintenance base letter, or the RISE proposal in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, the user counts the paper evidences applied as SAP's proposed classification per band, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Apply what looks right with one click and start the interview from the contract instead of a blank form; your evidence counts stay yours to set on the Users & FUE screen. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen: what to tell SAP from the first call. Five message windows counted back from your target signature, each with a goal, lines to deliver roughly verbatim, and what must not leave the room; the standing principles, including the SAP-specific one, keeping the audit and measurement track strictly separate from the commercial track; six lines that cost money with their replacements (never concede the 2027 deadline as binding); and a scripted first meeting agenda. The read at the top tailors itself to your discovery answers on cliff pressure, classification evidence, and alternatives.

The assessments. The Users & FUE, RISE Offer, and Negotiation Desk screens each carry a weighted assessment bank (FUE classification, BTP commitment, and contract terms; ten, eight, and eight questions) that scores your readiness out of 100 and ranks the fixes by what they are worth at the table; every question explains on hover why a veteran asks it.

The FUE Heatmap. The Analyze phase ends with a heatmap that makes the classification argument in one picture. RISE has no suite ladder to grade the way a Microsoft agreement does; its one great price lever is how users are classified, so the grid puts the four FUE bands SAP prices against (Advanced, Core, Self-service, Developer) across the top and your user populations down the side, and colors each population by how intensely it actually works in SAP. It shows what a blanket Advanced count hides: the approvers, requisitioners, and self-service users who belong a band, or two, below where SAP proposes, with the annual value of the reclassification printed above the grid and the argument to make at the table composed live beneath it. Correct any level by hand where you know better than the derived read, and the argument, the copyable text grid, and the emails all recompose around your number.

The What Next desk. The Deliver phase carries an open desk that puts the whole position to work. Vera drafts six emails from your live numbers: the meeting request, your asks, and the FUE-activity case to SAP, plus briefs for your CFO, procurement, and your own team. The SAP-facing drafts carry only your own asks and your own usage facts, which are the argument, and never your savings target or the corridor; the internal drafts carry the full position. You edit each one and send it from your own inbox, nothing sends from the platform. Alongside the emails sit copy-ready digests (the position in five lines, or the FUE heatmap as a text grid), the management brief written on one click, and the exit ramps into a tracked negotiation, the SAP vendor room, a market comparison, and an analyst pressure test.

The Deliver phase runs five screens: the Outcome Dashboard (KPI wall, the savings mix across FUE right-classification, discount, BTP right-size, and the conversion credit, the waterfall from SAP's proposal to your target, preparation and assessment verdicts, the five-year TCO read with the payback year, and what remains open, each tile clicking through), the Business Case, the Executive Report, the What Next desk, and Updates & Sharing, where you send a project update on demand and see exactly who receives what.

Key moves

- Answer the discovery interview and let the playbook set the strategy with you, or start from your contracts and let Opus pre-fill it.
- Create a team project so the deal team, and invited outsiders, work one live state instead of a screenshot chain.
- Rebuild the FUE count from usage evidence, baseline your current run rate, and run the three assessments until the fixes list is empty.
- Follow the Comms Plan: set the frame at T minus 9 months, keep the audit track separate, and never treat the 2027 cliff as binding.
- Decompose SAP's offer, price the alternatives against it, and work the migration timeline and the commercial asks.
- Read the cost corridor on Mission Control, then open the FUE Heatmap to build the classification argument and send the FUE-activity case to SAP from the What Next desk.
- Finish on the Outcome Dashboard, the five-year TCO business case, and an executive report, and let the project send the closing update itself.

When to use it: SAP is pushing you from ECC or on-premise toward RISE.

Oracle ULA

`/tooling/oracle-ula-certification` · Decide renew versus certify on the numbers, then exit clean, solo in your browser or as a team project the whole certification team works on together.

The Oracle playbook runs the ULA endgame like a real engagement: the discovery interview commits the strategy, the License Assessment holds the deployment inventory and every data source, the Renew vs Certify screen decides the path on five-year numbers, and the Maximize and Certification Desk screens execute it. Everything the other deep playbooks have applies here too, with certification semantics throughout.

It opens like an engagement too. The first open is a guided kickoff in three steps: the Discovery interview (answer the questions yourself or talk it through with Vera), then Opus 5 drafts the project plan from your answers and your numbers, then a confirm step on the fields every date derives from (client, agreement, the ULA end date, and the notice window) before Mission Control opens. Skip it with one click and the playbook opens on Mission Control from then on; team projects, shared view links, and screen deep links never see the kickoff. The screens are carried in a **left hand menu** grouped by phase, each phase with its completion percentage and the active screen a filled blue pill (small windows keep the compact stepper), and the results ticker carries a **Theatre** button: pick the light or the dark stage and the whole playbook runs full screen, one screen at a time, arrow keys to move, everything still editable, Escape to step back out exactly where you were. The plan stays live on the Timeline & Activities screen.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, the four stages with progress, a you-are-here marker, the next step, and the countdown to the ULA end date, expanding into a full Gantt of the 16 work-back activities that runs one month past the end date to carry the post-cert VMware exit. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox. Every screen carries a notes and corrections box that syncs to the team, the project log holds meeting notes and decisions, and any page can be shared with a revocable view link. The questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone: the tokenized page opens pre-filled and their answers merge back under their name. Phase completions mail an exec-brief update; a weekly digest goes out Monday mornings; every recipient controls what they receive.

Starting from your contracts. The License Assessment screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the ULA ordering document, master agreement, amendments, and support renewals in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Processor counts the paper evidences apply with one click; the deployment inventory ultimately owns those numbers. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen built around the ULA-specific discipline: your renew-or-certify intent is the single most valuable fact you hold, and it is spent last. Five message windows counted back from the ULA end date, from the quiet evaluation frame through absorbing the renewal push, deciding inside while holding outside, the notice and GDR window, and the close; six lines that cost money with their replacements (never volunteer deployment talk, never invite Oracle's measurement scripts, never share the VMware topology); and a scripted first meeting agenda. The read at the top tailors itself to your VMware position, audit history, and GDR readiness.

The assessments. The License Assessment, Maximize, and Certification Desk screens each carry a weighted assessment bank (certification evidence, VMware maximization, and certification desk discipline; ten, eight, and eight questions) that scores your readiness out of 100 and ranks the fixes by what they are worth; every question explains on hover why a veteran asks it.

The Deliver phase runs five screens: the Outcome Dashboard (the certified quantity composition across physical, VMware maximized, and the cloud 365-day average, the waterfall from the renewal's five-year cost down to certifying, the buffer years, preparation and assessment verdicts, and the renew-signal verdict, each tile clicking through), the Business Case, the Executive Report, the What Next desk, and Updates & Sharing.

The discovery interview runs in either mode: answer the structured questions (VMware footprint, cloud evidence, audit history, and now your Oracle Java exposure, the cross-sell lever Oracle raises mid-negotiation), or **talk it through with Vera** by voice or typing and she fills the questionnaire live. After the interview, Mission Control carries **Your project plan**: one click and Opus 5 turns your answers and your numbers into the engagement plan, the steps in order, key dates against the certification clock, workstreams with owners, and the risks with mitigations, all on the overview.

A built-in consultant. Once the interview is answered, the playbook consults. Every working screen opens with **Your consultant's read**: a named posture drawn from your goal and situation, a leverage score out of 100 with the strengths and the gaps behind it, and a per-screen headline that tells you what the page is for and what to do first, tailored to your growth outlook, your VMware footprint, and how defensible your deployment evidence is. It is rule-based and instant, not an AI call, so it updates the moment you change an answer.

The planning screens. After Discovery, a **Certification Plan** screen turns your answers into the project plan a firm would charge for: the default path, key dates counted back from your ULA end date, workstreams with owners, the internal questions, and the three-number position, generated rule-based instantly and rebuilt by Opus 5 on one click. In the Execute phase, a **Commit Scenarios** screen prices the ways to commit side by side (maximize-certify-exit, certify as-is, or renew) with the five-year totals and a recommendation that updates live as your answers change.

The cost corridor. Mission Control leads with a corridor: what renewing the ULA costs over five years (the new license fee plus the support that compounds on it forever), what five years of support at today's rate costs as a neutral baseline, and where certifying and exiting lands it (support unchanged, no new fee, plus the one-time certification project), as one waterfall that recomputes as you work the screens. The swing between renewing and certifying is the number the playbook exists to capture.

The Certification Heatmap. The Analyze phase ends with a heatmap that makes the certified position visible product by product. For each option the ULA covers it shows what you have deployed today, the count you can defensibly certify to, the banking headroom in between, and the perpetual license value that headroom is worth. It flags the products carrying real headroom (certify each at its maximized count, since support is fixed and the headroom is then free forever) against the drop candidates with none (stop paying 22% support on them rather than folding them into a renewal). Java and cloud sit below the products as context rows, kept off the ULA paper. Correct any level by hand and the argument and the emails recompose.

The What Next desk. The Deliver phase carries an open desk that puts the whole position to work. Vera drafts six emails from your live numbers: the meeting request, your asks, and the certified-position case to

Oracle, plus briefs for your CFO, procurement, and your own team. The Oracle-facing drafts carry only your own deployment facts, which are the argument, and never your savings target or the corridor; the internal drafts carry the full position. You edit each and send it from your own inbox, nothing sends from the platform. Alongside the emails sit copy-ready digests (the position in five lines, or the certification heatmap as a text grid), the management brief written on one click, and the exit ramps into a tracked negotiation, the Oracle vendor room, a market comparison, and an analyst pressure test.

Key moves

- Answer the discovery interview (form or with Vera), or start from your contracts and let Opus pre-fill it, then generate the project plan on Mission Control.
- Create a team project so SAM, DBAs, infrastructure, legal, and procurement work one live state.
- Assess the license position, review the contracts and entities in scope, and run the three assessments until the fixes list is empty.
- Follow the Comms Plan: hold the evaluation posture, keep the measurement track separate, and let the notice letter carry the decision.
- Maximize deployments on VMware before the certification count is taken, and work the cloud and certification rules.
- Read the cost corridor and the consultant's read on Mission Control, then open the Certification Heatmap to see what to certify at its maximized count and what to drop rather than renew.
- Finish on the Outcome Dashboard and the CFO case, send the certified-position case to Oracle from the What Next desk, and let the project send the closing update itself.

When to use it: your Oracle ULA is approaching its certification date.

IBM ELA

`/tooling/ibm-ela` · Decide whether to renew the bundle or right-size out of it, solo in your browser or as a team project the whole renewal team works on together.

The IBM playbook runs the ELA renewal like a real engagement: the discovery interview commits the strategy, the License Position holds entitled versus deployed versus S&S paid for every product with the data sources beside it, the Renew vs Right-size screen decomposes IBM's quote and decides the path on five-year numbers, and the Shelfware and ILMT screens execute it. Everything the other deep playbooks have applies here too, with renewal semantics throughout.

It opens like an engagement too. The first open is a guided kickoff in three steps: the Discovery interview (answer the questions yourself or talk it through with Vera), then Opus 5 drafts the project plan from your answers and your numbers, then a confirm step on the fields every date derives from (client, agreement, the ELA end date, and where in the term you are) before Mission Control opens. Skip it with one click and the playbook opens on Mission Control from then on; team projects, shared view links, and screen deep links never see the kickoff. The screens are carried in a **left hand menu** grouped by phase, each phase with its completion percentage and the active screen a filled blue pill (small windows keep the compact stepper), and the results ticker carries a **Theatre** button: pick the light or the dark stage and the whole playbook runs full screen, one screen at a time, arrow keys to move, everything still editable, Escape to step back out exactly where you were. The plan stays live on the Timeline & Activities screen.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, the four stages with progress, a you-are-here marker, the next step, and the countdown to the ELA end date, expanding into a full Gantt of the 16 work-back activities. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox. Every screen carries a notes and corrections box that syncs to the team, the project log holds meeting notes and decisions, and any page can be shared with a revocable view link. The questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone: the tokenized page opens pre-filled and their answers merge back under their name. Phase completions mail an exec-brief update; a weekly digest goes out Monday mornings; every recipient controls what they receive.

Starting from your contracts. The License Position screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the ELA/ESSO schedule, the Passport Advantage agreement, amendments, and IBM's renewal quote in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Entitlements the paper evidences apply with one click; deployment counts stay yours to set from your own discovery. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen built around the two IBM disciplines: your compliance posture is discussed with nobody while ILMT is remediated, and nothing un-itemized gets discussed. Five message windows counted back from the ELA end date, which is also IBM's December 31 fiscal close, from the quiet frame through absorbing the Cloud Pak and watsonx pitch, landing the decomposed anchors, contesting into the year-end rounds, and closing on the Passport Advantage record; six lines that cost money with their replacements; and a scripted first meeting agenda. The read at the top tailors itself to your ILMT posture, shelfware stance, and Cloud Pak intent.

The assessments. The License Position, ILMT & Audit Exposure, and Negotiation Desk screens each carry a weighted assessment bank (license position, ILMT compliance, and renewal terms; ten, eight, and eight questions) that scores your readiness out of 100 and ranks the fixes by what they are worth; every question explains on hover why a veteran asks it.

The Deliver phase runs five screens: the Outcome Dashboard (the S&S base split across kept, shelfware, and bundled lines, the waterfall from IBM's ask down to the right-sized target, the audit exposure until remediated, preparation and assessment verdicts, and the renew-signal verdict, each tile clicking through), the Business Case, the Executive Report, the What Next desk, and Updates & Sharing.

The discovery interview runs in either mode: answer the structured questions (Cloud Pak and watsonx intent, ILMT coverage, the shelfware stance, and now your recent audit posture, because an audit letter near a renewal is leverage on IBM's side of the table), or **talk it through with Vera** by voice or typing and she fills the questionnaire live. After the interview, Mission Control carries **Your project plan**: one click and Opus 5 turns your answers and your numbers into the engagement plan, the steps in order, key dates against the year-end close, workstreams with owners, and the risks with mitigations, all on the overview.

A built-in consultant. Once the interview is answered, the playbook consults. Every working screen opens with **Your consultant's read**: a named posture drawn from your goal and situation, a leverage score out of 100 with the strengths and the gaps behind it, and a per-screen headline that tells you what the page is for and what to do first, tailored to your ILMT position, how much shelfware you are willing to cut, and whether you can decompose the Cloud Pak bundle. It is rule-based and instant, not an AI call, so it updates the moment you change an answer.

The planning screens. After Discovery, a **Renewal Plan** screen turns your answers into the project plan a firm would charge for: the default path, key dates counted back from your ELA end date, workstreams with owners, the internal questions, and the three-number position, generated rule-based instantly and rebuilt by Opus 5 on one click. In the Execute phase, a **Commit Scenarios** screen prices the ways to commit side by side (right-size and rebuy, renew the bundle, or settle-then-right-size) with the five-year totals and a recommendation that updates live as your answers change.

The cost corridor. Mission Control leads with a corridor: what IBM's bundled renewal ask costs (last year's S&S base, the uplift, and the Cloud Pak padding), what your current S&S base costs as a baseline, and where right-sizing lands it (the bundle declined, the uplift capped, shelfware S&S dropped, and IBM's selective repricing absorbed), as one waterfall that recomputes as you work the screens. The swing between the bundle and the right-sized run rate is the number the playbook exists to capture.

The Right-Sizing Heatmap. The Analyze phase ends with a heatmap that makes the position visible product by product. For each product it shows what you are entitled to, what you actually deploy, the utilization, and the S&S sitting on the gap. It flags the heavy-shelfware products to right-size (drop the maintenance on undeployed entitlements while keeping the perpetual licenses) against the healthy ones to keep, with the total shelfware S&S printed above the grid. Correct any level by hand and the argument and the emails recompose.

The What Next desk. The Deliver phase carries an open desk that puts the whole position to work. Vera drafts six emails from your live numbers: the meeting request, your asks, and the right-sizing case to IBM, plus briefs for your CFO, procurement, and your own team. The IBM-facing drafts carry only your own deployment facts, which are the argument, and never your savings target or the corridor; the internal drafts carry the full position. You edit each and send it from your own inbox, nothing sends from the platform. Alongside the emails sit copy-ready digests (the position in five lines, or the right-sizing heatmap as a text grid), the management brief written on one click, and the exit ramps into a tracked negotiation, the IBM vendor room, a market comparison, and an analyst pressure test.

Key moves

- Answer the discovery interview (form or with Vera), or start from your contracts and let Opus pre-fill it, then generate the project plan on Mission Control.
- Create a team project so SAM, procurement, legal, and the platform teams work one live state.
- Build the license position, decompose IBM's renewal quote line by line, and run the three assessments until the fixes list is empty.

- Follow the Comms Plan: hold the line-by-line posture, keep compliance off every commercial call, and let the year-end close work for you.
- Drop shelfware subscription and support while keeping the perpetual licenses, and remediate ILMT before it becomes audit leverage.
- Read the cost corridor and the consultant's read on Mission Control, then open the Right-Sizing Heatmap to see what to right-size and what to keep.
- Finish on the Outcome Dashboard and the CFO case, send the right-sizing case to IBM from the What Next desk, sign inside IBM's year-end close, and let the project send the closing update itself.

When to use it: your IBM ELA renewal quote just landed.

Adobe ETLA

`/tooling/adobe-etla` · Run an Adobe ETLA renewal inside the one window the agreement allows, solo in your browser or as a team project the whole renewal team works on together.

The Adobe playbook runs the renewal like a real engagement: the discovery interview commits the strategy (answer the questions or talk it through with Vera), the License Position holds seats contracted versus assigned versus active for every product with the data sources beside it, the Renew vs Right-size screen decomposes Adobe's proposal and decides the path, and the seat and AI screens execute it. Everything the other deep playbooks have applies here too, with ETLA semantics throughout.

It opens like an engagement too. The first open is a guided kickoff in three steps: the Discovery interview (answer the questions yourself or talk it through with Vera), then Opus 5 drafts the project plan from your answers and your numbers, then a confirm step on the fields every date derives from (client, agreement, the ETLA end date which is also Adobe's fiscal close, and where in the term you are) before Mission Control opens. Skip it with one click and the playbook opens on Mission Control from then on; team projects, shared view links, and screen deep links never see the kickoff. The screens are carried in a **left hand menu** grouped by phase, each phase with its completion percentage and the active screen a filled blue pill (small windows keep the compact stepper), and the results ticker carries a **Theatre** button: pick the light or the dark stage and the whole playbook runs full screen, one screen at a time, arrow keys to move, everything still editable, Escape to step back out exactly where you were. The plan stays live on the timeline screen.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, expanding into a full Gantt of the work-back activities. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox. Every screen carries a notes and corrections box that syncs to the team, the project log holds meeting notes and decisions, and any page can be shared with a revocable view link. The questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone: the tokenized page opens pre-filled and their answers merge back under their name. Phase completions mail an exec-brief update; a weekly digest goes out Monday mornings; every recipient controls what they receive.

Starting from your contracts. The License Position screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the ETLA order form, license schedule, amendments, and Adobe's renewal proposal in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, the seats the paper evidences applied as the contracted counts, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Active-usage counts stay yours to set from the Admin Console evidence. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen built around the Adobe disciplines: your seat inactivity data is released once, as an anchor, never volunteered early; every generative AI line goes to separate paper; and the VIP alternative is mentioned once, factually. Five message windows counted back from the ETLA end date, which is also Adobe's late-November fiscal close; six lines that cost money with their replacements; and a scripted first meeting agenda. The read at the top tailors itself to your AI intent, profile drivers, and exposures.

The assessments. Three weighted assessment banks score the seat position, the AI and premium lines, and the renewal terms out of 100 (ten, eight, and eight questions), each with a verdict and the fixes ranked by what they are worth; every question explains on hover why a veteran asks it.

The Deliver phase runs four screens: the Outcome Dashboard (Adobe's ask decomposed, the right-sized target, the inactive seat waste, and every assessment verdict, each tile clicking through), the Business Case, the Executive Report, and Updates & Sharing.

Key moves

- Run the guided kickoff, or start from your contracts and let Opus pre-fill the interview.
- Create a team project so IT, procurement, and the creative leads work one live state.
- Pull the Admin Console evidence on real seat activity, cut inactive seats, right-profile All Apps and Acrobat, and run the three assessments until the fixes list is empty.
- Follow the Comms Plan: release the usage evidence once as an anchor, decompose the generative-AI padding onto separate paper.
- Price the VIP alternative and sign inside Adobe's fiscal year close.
- Finish on the Outcome Dashboard and the CFO case, and let the project send the closing update itself.

When to use it: your Adobe ETLA renewal window is opening.

Salesforce SELA

`/tooling/salesforce-sela` · Renew the Salesforce wrap on your measured usage, not the quote, solo in your browser or as a team project the whole renewal team works on together.

The Salesforce playbook runs the SELA renewal like a real engagement: the discovery interview commits the strategy (answer the questions or talk it through with Vera), the License Estate holds provisioned versus active seats for every cloud with the data sources beside it, the renewal screens decompose the quote against your measured usage, and the Agentforce screen right-sizes the consumption commitment. Everything the other deep playbooks have applies here too, with SELA semantics throughout.

It opens like an engagement too. The first open is a guided kickoff in three steps: the Discovery interview (answer the questions yourself or talk it through with Vera), then Opus 5 drafts the project plan from your answers and your numbers, then a confirm step on the fields every date derives from (client, agreement, the anniversary the plan works back from, and where in the term you are) before Mission Control opens. Skip it with one click and the playbook opens on Mission Control from then on; team projects, shared view links, and screen deep links never see the kickoff. The screens are carried in a **left hand menu** grouped by phase, each phase with its completion percentage and the active screen a filled blue pill (small windows keep the compact stepper), and the results ticker carries a **Theatre** button: pick the light or the dark stage and the whole playbook runs full screen, one screen at a time, arrow keys to move, everything still editable, Escape to step back out exactly where you were. The plan stays live on the timeline screen.

Running it as a team. The Project Team screen (Prepare phase) turns the playbook into a shared project saved to your workspace: everyone works on one live state with autosave, and a project trackboard rides under the phase stepper on every page, expanding into a full Gantt of the work-back activities. Editors need a workspace account; viewers can be anyone, invited by email with a read-only link that opens the live project as a sandbox. Every screen carries a notes and corrections box that syncs to the team, the project log holds meeting notes and decisions, and any page can be shared with a revocable view link. The questionnaires, the discovery interview, the readiness check, and the three assessments, can be emailed to anyone: the tokenized page opens pre-filled and their answers merge back under their name. Phase completions mail an exec-brief update; a weekly digest goes out Monday mornings; every recipient controls what they receive.

Starting from your contracts. The estate screen opens with the contract starting point: pick agreements already uploaded to the portal or drag the SELA master agreement, order forms, amendments, and Salesforce's renewal quote in (PDF or DOCX), and Opus 5 reads the paper into the playbook, agreement facts, the seats the paper evidences applied as the provisioned counts, the terms that will matter with a favors-you, watch, or risk verdict on each, risks, and interview answers the paper supports with a citation for each. Thirty-day active counts stay yours to set from your own telemetry. Dropped documents are kept: each saves to your workspace as a contract and joins your benchmarking queue, and the saved records are listed under the analysis for one-click open or delete.

The comms plan. The Prepare phase includes a Comms Plan screen built around the Salesforce disciplines: the account team reads internal champions for renewal certainty, so one voice matters double; adoption weakness is never volunteered; un-itemized bundles are returned, not discussed; and the close sequences into the January 31 fiscal year end. Five message windows counted back from the anniversary; six lines that cost money with their replacements; and a scripted first meeting agenda. The read at the top tailors itself to your edition drivers, add-on shelfware, and Agentforce intent.

The assessments. Three weighted assessment banks score the license position, the platform and AI lines, and the renewal terms out of 100 (ten, eight, and eight questions), each with a verdict and the fixes ranked by what they are worth; every question explains on hover why a veteran asks it.

The Deliver phase runs four screens: the Outcome Dashboard (current spend, the quote as uplifted, the right-sized target, the savings mix across edition right-profiling, hygiene, and the Agentforce right-size, and every assessment verdict, each tile clicking through), the Business Case, the Executive Report, and Updates & Sharing.

Key moves

- Run the guided kickoff, or start from your contracts and let Opus pre-fill the interview.

- Create a team project so admins, procurement, and the business owners work one live state.
- Segment the Unlimited base against real usage, clear the add-on shelfware, and run the three assessments until the fixes list is empty.
- Follow the Comms Plan: one voice, adoption weakness never volunteered, un-itemized bundles returned.
- Right-size the Agentforce commitment to the measured conversation burn.
- Sign inside the January 31 year-end close with the CFO case, and let the project send the closing update itself.

When to use it: your Salesforce SELA renewal is on the horizon.

Anthropic EA

`/tooling/anthropic-ea` · Size an Anthropic enterprise agreement off your own forecast.

Key moves

- Load the Console usage and right-size the model mix behind your evals.
- Wire caching and the Batch API, and trim inactive seats.
- Size the commitment against the true-up risk curve.
- Sign with rollover, price protection, and the CFO case in hand.

When to use it: Anthropic put an enterprise agreement on the table.

Workday

`/tooling/workday-renewal` · Negotiate a Workday renewal against the market, not the bundle.

Key moves

- Benchmark your discount off list against the peer corridor.
- Decompose the ask and the expansion bundled into it.
- True the worker bands down to reality in the only window they reopen.
- Keep the Workday-loyal implementation partner out of the negotiation room, and sign inside the January 31 close.

When to use it: your Workday renewal quote bundles more than you asked for.

ServiceNow

`/tooling/servicenow-renewal` · Renew ServiceNow on your numbers, against a corridor that runs past 45%.

Key moves

- Benchmark your discount off list.
- True down idle seats and modules in the only window one-way pricing allows.
- Cap the uplift at 5% or lower, and refuse every product without a funded owner.
- Keep the ServiceNow-loyal implementor out of the commercial room, and start a year early so a credible no earns the one-time license credits.

When to use it: your ServiceNow renewal is coming and the uplift looks steep.

Third-party support

`/tooling/third-party-support` · Move stable OEM support to a third-party provider on your terms.

Key moves

- Decide eligibility product by product.
- Run a blind tender across every provider without disclosing what you pay the OEM.
- Sign far below the standard OEM rate.
- Use the incumbent re-tender path for estates already on third-party support.

When to use it: you pay OEM support on stable products you barely call.

The What Next desk

`?screen=whatnext` on Adobe ETLA, Salesforce SELA, Anthropic EA, Workday, ServiceNow, and third-party support · reached from the end of the playbook or the last item in the phase rail.

Where the four largest playbooks (Microsoft EA, Oracle ULA certification, RISE with SAP, IBM ELA) have always closed on a full desk, the other six used to stop at the executive report. They now close the same way, in three parts.

The desk carries each playbook's own ask, so the six read differently even though the shape is shared.

On the screen

- **The drafted vendor email**: the ask that playbook exists to land, composed from the figures your own run produced, in an editable subject and body. Copy it or open it in your mail app; nothing sends from the platform, and the draft carries no greeting or sign off so you add your own. Editing takes over from the computed draft, and one link restores it.
- **The exit ramps**: two or three moves that carry the vendor with them, the matching benchmark tool, a negotiation case seeded on that vendor, and the costed alternative or the BATNA playbook that prices your walk away.
- **Ask Vera and the position digest**: a question carrying your run's headline numbers, and a copyable summary of where the run landed.

Vendor-facing drafts state your asks and your own usage facts only. Your target, your corridor, and your savings estimate stay in the digest and the Ask Vera question, never in the email. The third-party support draft is a tender invitation by design: it publishes the scope and the timetable and never what you currently pay the OEM.

Next: Chapter 04 · Insight

04 · Insight

The intelligence layer over everything you have uploaded. It turns your proposals, agreements, and invoices into spend and concentration analytics, ranks the savings still on the table, proves negotiated savings against the invoices that followed, reads your documents for search and grouping, and connects you to an anonymous community of other verified buyers. Where a screen leans on AI, it is marked below and needs an AI plan.

YOUR ORGANIZATION

Where: Insight › Your organization · `/org`
Who: everyone in the organization

The whole estate on one page: what your organization actually spends per year, vendor by vendor, with every agreement consolidated. This is the page you forward to finance, so it is open to every member rather than owners only.

In plain words: the one-page answer to "what do we hold and what does it cost."

When to reach for this

- Finance asks for the estate: this page is built to be forwarded, and every member can see it.
- Sanity-checking the consolidation: agreements rolled up per vendor, one line each.

The number it leads with is what you pay per year under the agreements in force **today**. Not a contract total, and not anything whose term has ended. A vendor you have renewed three times counts once, because three sequential three year deals are nine years of paper and one live agreement.

On the screen

- **Open the workroom** — the band at the top opens the estate as a workroom. Up to six documents: the cover sheet, the vendor ledger, the flagged overlap atlas, the renewal horizon, a consolidation sheet with three options per overlap (the overlap papers appear only when two live agreements overlap by more than 180 days), and a What now sheet. ⚡ **What now** drafts the estate report for a stakeholder, copies the estate summary, and opens renewals, The Desk, or your biggest agreements. Only live annual figures are counted, and an overlap is flagged rather than silently merged, so the number is the estate's own. ✦ **Ask Vera** answers from your estate's own vendors, years, and renewals. Several pages in this chapter open the same way; each section says what its workroom holds.
- **Annual spend, live** — the headline. Underneath it, what adding every contract on file would have claimed instead, so the difference is explained before anyone asks.
- **Why this is lower than your spreadsheet** — the gap in money, and the reason for it.
- **Renewal windows** — what comes up this year, next year, and the year after, by the annual value at stake, with the largest vendor in each named.

- **Spend per year** — each agreement spread across the years its term actually covers, so a three year deal is a steady line rather than a spike in one year. Agreements with no dates are left out of the chart rather than guessed into a year.
- **Vendors by annual spend** — one row per company, consolidated across every agreement and legal entity you hold with them. Different spellings of the same vendor (Salesforce, Inc., Salesforce UK Limited, SFDC Ireland Limited) merge into one row, and the merged entities are listed under the name so nothing is hidden. Vendor aliases your organization has confirmed are applied first. A vendor with earlier renewals says so. Two live agreements overlapping by more than six months under the same recorded name are flagged as overlapping terms; a company holding paper through two entities at once is normal contracting and is not flagged.
- **The move on a flag** — a flag now settles itself rather than just warning you. An **overlapping terms** flag links the two agreements open side by side in the comparison view, worst overlap first, so one look tells you whether it is two real products or the same contract loaded twice. An **uplift ahead** flag drafts the objection: the percentage, the years it steps between, your current annual spend with that vendor, and a request for the calculation in writing before the new pricing takes effect. Copy it, or open it in your own mail client. Nothing sends from the platform. The uplift marker reads the signed fee schedule, meaning a rise the contract already prices, rather than an uplift-cap clause.
- **Show history** — the ended agreements and what they were worth, visible on request and never added into the live figure.
- **Your five biggest agreements** — the largest live agreements, one per vendor, with one button to send them all for a benchmarked second opinion.
- **Second opinions** — every verdict the organization has run, with its status, linking to the full report.
- **What to do next** — a note to your CFO, a board report, a renewal project, a Vera Deal Room, market research on your vendors, and the renewal watchdogs.

How to use it

1. Read the headline, then the line explaining the difference from your own spreadsheet. Take that difference to whoever owns the spreadsheet.
2. Check the three renewal windows. The one with the most money in it is where your next quarter goes.
3. Scan the vendor table for **overlapping terms**. Each one is either a product nobody told procurement about or the same contract twice, and both are worth ten minutes. Press **Open the two terms side by side** on a flagged row to settle it in one look.
4. Press **Analyse these 5** to get a benchmarked verdict on your largest agreements. Results appear on this page as they finish.
5. Pick a next move from the bottom of the page.

Tips

- Two live agreements from one vendor are counted, not merged, until you say otherwise. We will not quietly halve your spend on a guess.
- A renewal signed a few months before the old term ends is normal and is never flagged as an overlap.
- Contracts still inside a migration that has not gone live do not appear here at all.

Related: Portfolio · Spend · Migration Studio

THE VERA WIRE

Where: Benchmarks & market › The Vera Wire · `/wire` (the Market Desk: Today, Policy watch, Weekly report, and Your watchtower as tabs; the last two open their own rooms as a pill row beneath)
Who: everyone

The always-on market news desk. Every licensing, pricing, and vendor move the platform tracks lands here in one feed, newest first, each item read for what it means for your estate. It writes itself: the wire aggregates Terms Watch price changes, the Price Index, newly published research, the weekly vendor report, and Policy Watch document changes every day, so it stays current without anyone curating it. A condensed strip also rides on your dashboard under the Ask bar and on the welcome screen.

On the screen

- **The cover** — the lead story of the day, with Vera's read on what to do about it and a chip counting the items that touch vendors you hold.
- **Lenses** — one tap filters the wire to Licensing, Pricing, Vendor news, or Research.
- **My vendors only** — a toggle that narrows the feed to vendors you hold a contract with.
- **The feed** — a dated timeline of every signal; a colored dot marks the category, a pill carries the verdict, Favorable, Caution, or Risk, and a chip marks your vendors.
- **Vera's read** — tap any item to expand Vera's one line take: the move to make, grounded in your numbers.
- **Your move** — an item touching a vendor you hold ends with a Your move line: the vendor room, that vendor's benchmark, and Ask Vera seeded with the headline.
- **The rail** — price movers, Terms Watch changes, and newly published research, each linking to the source.

How to use it

1. Read the cover story and Vera's read for the headline move of the day.
2. Tap a lens to focus on Licensing, Pricing, Vendor news, or Research.
3. Turn on My vendors only to narrow the wire to your own estate.
4. Tap any item to expand Vera's read, then follow the source link for the detail.

Tips

- The strip on your dashboard and welcome screen shows the same lead and the next signals, so the headlines meet you before you open the full wire.

Related: Weekly report · Research · Policy Watch · Platform pulse

POLICY WATCH

Where: The Vera Wire › Policy watch · `/wire/policy`
Who: everyone

Vendors reprice you through their documents, not just their price lists. Policy Watch fetches the public licensing and cloud policy documents of the top-20 vendor set every day, 32 documents in all, and diffs each one against yesterday's copy: the Microsoft Product Terms and Universal License Terms, the Oracle partitioning, cloud licensing, support, and Java SE policies, AWS service terms, the Google Cloud terms, and the standing agreements of SAP, Salesforce, IBM, Adobe, ServiceNow, Workday, and their peers. Vendors publish on their own rhythm, Microsoft monthly on the first business day, AWS continuously, but the sweep runs daily because vendors do not always keep to their own calendar. When the fine print changes, the change is logged with the exact lines that moved and Vera's read on what it means, and material changes publish to The Vera Wire under the Licensing lens.

What you see

- **Vitals** — documents watched, vendors covered, material changes in the last 30 days, and when the last daily sweep ran.
- **Signals** — the subscription panel (see below).
- **Change ledger** — every caught change, newest first: the headline, what changed in the document, a verdict pill, Favorable, Caution, or Risk, and Vera's read on why it matters and the move to make. Minor wording edits are kept in the ledger but marked and never sent to the Wire. Each change row ends with the same Your move trio the Wire carries, scoped to the change: the vendor room, that vendor's benchmark, and Ask Vera seeded with it.
- **Show the changed lines** — expand any change to see the exact lines added and removed, the evidence behind the read.
- **What we watch** — the full registry by vendor: each document with its type, why it matters to buyers, its known publication rhythm, when it was last checked, and a link to the live document at the vendor's site.

Signals: alerts per vendor or all

- Follow **every tracked vendor** with one switch, or pick vendors one by one with their bells.
- Choose your channels independently: **Email me** sends one exec brief email per sweep with the changes on your vendors and a one click unsubscribe that touches nothing else; **Notify my dashboard** lands each material change on your notification bell, tunable like every other alert kind in settings.
- Both channels are opt in and off by default, carry material changes only, and share one vendor list. Minor wording edits never alert.

How to use it

1. Read the change ledger newest first; the vitals tell you in seconds whether anything moved since your last visit.
2. On a change, read Vera's read, then expand Show the changed lines for the primary evidence, for example before raising it with your vendor rep.
3. Follow your vendors with their bells and turn on the channel you want in Signals; from then on the fine print comes to you.

Tips

- Material changes also land on The Vera Wire, so you do not need to visit this page daily; with a subscription on, you do not need to visit it at all.
- A document shown as unreachable means the vendor moved or blocked the page on the last check; the platform keeps checking and the history stays intact.

Related: The Vera Wire · Watch · Research

PORTFOLIO

Where: Spend & savings › Portfolio · `/portfolio`
Who: everyone

In plain words: where you stand and where the risk sits, computed from the paper you already filed.

When to reach for this

- The quarter's "where should effort go" question, answered by position rather than noise.
- Before leadership conversations: this page and the Estate map are the picture they want.
- After data-in: the first proof the filing paid off.

The position and risk lens on your estate, built automatically from proposals and agreements that carry an annual value. Portfolio, Spend, and Renewals share a tab row at the top with a one line note on what each covers, and every chart lives on exactly one of them: this page argues where you stand and where the risk sits, and links to the others for values and dates.

On the screen

- **Vitals** — tracked annual spend, benchmarked proposals, renewals in notice, and realized savings.
- **Data confidence meter** — how complete your estate records are, blended from pricing, dates, benchmarks, categories, and owners. It names the three highest-impact fields to add next, each a deep link straight to the exact records that need them.
- **Spend mix** — the category donut beside the concentration read: how dependent you are on your biggest vendors, top three share included.
- **Treemap** — every vendor sized by spend and colored by standing.
- **Market position** — where your deals sit against the market, with the upcoming renewals beside it.
- **Vendor risk** — the vendors carrying elevated risk right now.

How to use it

1. Scan the vitals and the spend mix to see where the money and the dependence sit.
2. Read market position for the leverage, then the vendor risk strip for what needs attention.
3. Add annual values to agreements to enrich this view; renewal outlook lives on Renewals insights and exports on Spend, one click away.

Tips

- The richer your annual values, the sharper every chart here. Missing values leave gaps.

Related: Spend · Opportunities · Agreements

MONTHLY BRIEFING

Where: Spend & savings › Monthly briefing · `/briefing` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan
You get: a five minute spoken brief, chaptered so you can skip · a minute or two to generate

Your portfolio as a five minute spoken brief, one per month: what renewed, what is landing, where the leverage is, and the one number that matters. Chaptered, skimmable, and grounded in the same data as every page in the app.

In plain words: the month, spoken, in five minutes.

When to reach for this

- The commute, the treadmill, the five minutes before the leadership call.
- The monthly cadence between quarterly board packs.

On the screen

- **Player** — the current month's brief, split into chapters so you can jump to the part you want.
- **Open the workroom** — a band above the player lays the current briefing out as a five paper workroom: the month at a glance, every chapter as a timed ledger, the full script verbatim, a provenance sheet, and a What now sheet. ⚡ **What now** plays the recording, re-records it with today's numbers when AI is enabled (reporting the real outcome rather than a generic line), forwards the briefing to a colleague as a drafted note, copies a chat ready summary or the full transcript, and follows where the script points into your renewals, the market wire, or your dashboard. ✦ **Ask Vera** answers from the briefing's own script. The room is rebuilt from the live briefing on every open.
- **Generate** — builds this month's briefing from your live portfolio when one is not already waiting.

How to use it

1. Open Monthly briefing and press play, or generate the current month if the player is empty.
2. Skip between chapters to reach renewals, spend, or leverage directly.

Related: Portfolio · Vera AI

ESTATE MAP

Where: Spend & savings › Estate map · `/portfolio/map` · via **More** or Ctrl/⌘ K
Who: everyone

In plain words: the estate as one picture: size is spend, color is market position.

When to reach for this

- The slide leadership actually remembers: one canvas, exported, no deck-building.
- Spotting concentration and outliers faster than any table shows them.

Your whole software estate on one canvas: every vendor sized by spend and colored by market position, for the quarterly leadership conversation. It runs the same deterministic portfolio scan as Portfolio, promoted to a full canvas with a matching slide export.

On the screen

- **Vitals** — tracked annual spend with the share of it that is benchmarked, the spend sitting over market, and the headroom to the median. The share counts only vendors the canvas actually colors, so the number and the picture always agree.
- **Verdict lens** — chips under the masthead, each carrying the spend behind it. Press one and the canvas isolates that verdict; press it again to bring the whole estate back.
- **The map** — one tile per vendor, area sized by annual spend, color set by benchmark position, taking the full width of the screen. Vendors we cannot benchmark yet read as "Not benchmarked yet". Hover or tab a tile and the inspector under the canvas spells that vendor out in full.
- **Off-map note** — how many vendors carry no annual value and so do not appear on the canvas.

How to use it

1. Read the biggest tiles first: that is where the money and the leverage are.
2. Press the **Over market** chip to see only what is priced above the cohort, which is the version of this canvas worth putting on a slide.
3. Click a tile to open that vendor's page.
4. Export the canvas as a slide for the board deck.

Related: Portfolio · Strategy

ANOMALIES

Where: Spend & savings › Anomalies · `/anomalies` · via **More** or Ctrl/⌘ K
Who: everyone

Every finding can now be closed out: **Resolve** files an optional recovered annual amount, moves the finding to a Resolved section, and closes any linked task, and the verdict band totals recovered dollars against dollars at risk. Billed-above-contract, unused-seat, and no-uplift-cap findings also offer **Draft the fix**, a copy-only letter built from the finding's stored facts; nothing sends from the platform.

What the weekly watchdog found while scanning your portfolio: unusual uplifts, missed notice windows about to bite, duplicated spend, and other things worth a look.

In plain words: the weekly sweep's findings: what looks wrong in your portfolio, worth a human look.

When to reach for this

- The weekly glance: findings arrive scored, not alphabetized.

- An uplift breach like the sample Northwind case (a 14 percent ask against a 7 percent cap) is exactly what lands here.
- Before renewing anything: check whether the watchdog already flagged the vendor.

On the screen

- **Vitals** — what the open findings put at risk each year, how many are open, and how many carry an AI verification.
- **Findings list** — each anomaly with its severity and the evidence behind it.
- **What a finding looks like** — before your first scan, three worked examples of the kinds of finding the watchdog raises, marked as illustrative. Your first scan replaces them.
- **Triage controls** — mark a finding resolved or not relevant; the watchdog will not raise it again.

How to use it

1. Open a finding and check the evidence.
2. Act on it, then resolve or dismiss findings you have handled.

Related: Billing watch · Invoices

RIGHTS REGISTER

Where: Spend & savings › Rights register · `/obligations` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan

In plain words: the rights your contracts already grant you, harvested, with their expiry windows watched.

When to reach for this

- Before spending negotiating capital on something the paper already grants: true-downs, swap windows, price holds.
- Planning an exit: transition assistance and data-return rights are leverage you own (the sunset recipe).
- When a right is about to expire unused: the watchdog's alert is this page working.

Everyone benchmarks the price; this page harvests the contract rights you already paid for. A quote-anchored pass reads each contract and records your buyer-side rights (true-downs, swap windows, price holds, SLA credits, benchmark and termination clauses) with their exercise windows, then the weekly watchdog alerts you when a right is about to expire unused.

On the screen

- **Register** — every harvested right, its status, and the exact clause it was read from, so you can act with the contract in hand.
- **Harvest stats** — active rights, contracts scanned, and how many are still pending a harvest.

- **Right families** — before anything is harvested, the six kinds of right the harvester looks for, plus the one step that unlocks the register.
- **What next** — under the register, the next moves. The featured move is a pre-drafted, clause-quoting exercise letter for the next closing right, sent from your own mail or copied. The rest follow: open the clause behind the window, start the negotiation with that vendor, ask Vera which expiring right is worth most, or book the analyst to sequence the windows.

How to use it

1. Upload contracts and index them on Document search, then harvest here.
2. Work the rights with a near expiry so nothing you paid for lapses unused.

Related: Document search · Clause library

OPPORTUNITIES

Where: Spend & savings › Opportunities · `/opportunities` · via **More** or Ctrl/⌘ K
Who: everyone

Everything worth doing about your vendors, ranked in one list by value and urgency together, with the reasoning for each and the reading that prepares you for it.

In plain words: the savings, ranked, so the next negotiation picks itself.

When to reach for this

- Choosing the quarter's targets by money: the sample portfolio's list is led by the Northwind CRM deal's $28,650 a year.
- Feeding the pipeline: each opportunity opens into the vendor's benchmark and the negotiation that captures it.
- The list is only as complete as the benchmarking behind it; a thin list means run more checks, not that the estate is fine.

On the screen

- **Vitals** — how much of your spend sits above market, what you pay for licenses nobody uses, and the prize: the two added together, which is what the whole queue is playing for.
- **Next best actions** — the ranked queue. Every renewal window, benchmark gap, shelfware finding, price anomaly and stack overlap in one order, scored on what it is worth and how soon it closes, which is why a smaller number with a deadline can outrank a larger one without. Every ranked lever ends in one action that carries the money onward: into the renewal playbook, or into a new case.
- **The live tier** — an action whose deadline falls inside thirty days, or has passed while the cycle is still open, is marked live: a gold rule down the side, the amount at stake set large, and the deadline as a red pill. The header counts them for you, so you can see at a glance how much of the list is urgent.

- **Why now** — one sentence per action explaining what makes it worth doing today, written against the figures beside it: the annual dollars, the deadline, or the vendor's own fiscal close.
- **Read this before the meeting** — where a negotiation prep brief exists for the vendor, it appears under the action as a button. Three pages on what the vendor is selling, the facts about your own estate to hold before they quote, the talking points, and what they will say with the counter to each.
- **The window workbench** — the what-if. Two sliders, the annual uplift you hold to and the discount you win, applied to every renewal in the window at once. It shows the three year cost against the baseline, what you keep, and which vendors it comes from. The baseline is what happens if nobody acts: each agreement uplifted by its own contractual cap, or by a typical renewal uplift where the agreement has no cap, which is why an uncapped agreement costs money before anyone negotiates.
- **The window workroom** — **Open the workroom** turns the queue into a full case file: a cover sheet with what is live and what is at stake, the ranked queue with its reasoning, the renewal window with the uplift baseline each agreement carries, and a What now sheet of moves. Ask Vera from inside the room and she answers from this queue, with your deadlines and exposures in front of her.
- **Ranked vendors** — the vendors above market, with the estimated annual saving per vendor if the deal moved to the market median.
- **The portfolio strategy** — the generated portfolio strategy persists with its saved date, and a regenerate action refreshes it when the queue moves.

How to use it

1. Read the count in the header first. It tells you how many actions are inside thirty days, which decides whether the list needs attention today.
2. Work the live tier top down. Where a prep brief is offered, read it before the first call rather than after.
3. Move the workbench sliders to see what holding the line across the whole window is worth, and note the agreements with no uplift cap: capping those is usually the cheapest win on the page.
4. Open the workroom when you want the whole position in one file, or something to forward.
5. Turn an opportunity into a tracked negotiation, and book an analyst to pressure-test the plan if the deal is large.

Related: Portfolio benchmark · Negotiations · Renewals · Research

GUARANTEES

Where: Spend & savings › Guarantees · `/guarantees` · via **More** or Ctrl/⌘ K
Who: everyone

Engagements where the fee is at risk: if invoice-verified savings land under the guaranteed floor by the end of the measurement window, the fee is waived. Every guarantee settles on invoice evidence, never on claims. There is nothing to create here; guarantees are underwritten by VendorBenchmark on engagement start, only where the dataset is deep enough to stand behind a floor.

On the screen

- **Vitals** — guaranteed engagements, how many measurement windows are still live, and how many have settled.
- **Guarantee rows** — per vendor: the guaranteed annual floor, the savings verified so far, the measurement window, and the status (Open, or Settled with the floor met or the fee waived).
- **Certificate** — open a guarantee for the pinned projection, confidence, model version, and settlement.
- **Before your first guarantee** — the page names the candidates rather than leaving you to work them out: your largest renewals still ahead, with their dates and annual values, and the two ways forward. Whether any of them qualifies depends on how deep the dataset is for that vendor, which is the question to put to us.

How to use it

1. Track open guarantees against their window and the savings verified to date.
2. Open the certificate for the full underwriting detail.
3. With no guarantee yet, read the candidate renewals and press **Ask whether a renewal qualifies** to put the question to the team.

Related: Managed renewals · Savings proof

EXPERT BENCH

Where: Benchmarks & market › Expert Bench · `/experts` · via **More** or Ctrl/⌘ K
Who: everyone

Vetted negotiators and ex-vendor insiders you can bring into a deal: former EA desk leads, LMS auditors, and commercial architects, presented by verified track record under bench codenames.

On the screen

- **Vitals** — how many negotiators are on the bench, the specialties they cover, your open requests, and the intro turnaround.
- **For your renewal** — experts who work a vendor you have a live renewal with sort to the top of the bench, nearest renewal first. Each of those cards leads with a panel naming that vendor, the renewal date, the annual value, and the notice deadline, all read from your own agreement, so you are matching a person to a real deadline rather than to a specialty.
- **Track record** — verified from engagement records when the expert joins, re-checked annually.
- **Engagement shapes** — a working session, a fixed-scope review, or a full embedded engagement.
- **The request topic** — on a matched profile the topic arrives already written from your renewal record: the vendor, the date, the value, and the notice deadline. A line above the form says where it came from, and every word stays editable before you file it.
- **The intro** — the advisory desk confirms the quote and you meet the expert by name on the first call.

How to use it

1. Read the top of the bench: the experts there are matched to your nearest renewals, not ranked by popularity.
2. Open the profile and check the prefilled topic against what you actually want from the session.
3. Edit the topic if it needs it, then file the session request.

Related: Advisory

SAVINGS PROOF

Where: Spend & savings › Savings proof · `/savings`
Who: everyone

Every claimed saving proven or disproven against the invoices that followed: the ledger of what your negotiations actually returned against your benchmarked baselines.

In plain words: savings that survived contact with billing, and only those.

When to reach for this

- Quarter end (the reporting recipe): the proven set is the number finance accepts.
- After any close: record the outcome while the baseline is fresh.
- When a claim is challenged: the row names the invoices that confirmed it.

On the screen

- **Vitals** — annual savings recorded, what the same outcomes are worth over their terms, and how many outcomes are on the ledger.
- **From your closed deals** — the card above the ledger, listing closed cases whose certificate is not on it yet, each with the vendor, the closing date, the baseline, the accepted figure, and the annual saving. **File this as a saving** opens the record form with the vendor, the contract, the currency, the baseline, the negotiated figure, and a note naming the certificate already filled in, so filing is a review of numbers you already have rather than data entry. Everything stays editable, and a filed deal drops off the list.
- **Savings ledger** — each outcome with the baseline, the negotiated price, the annual saving, and the term total.
- **Cumulative curve** — annual savings stacked by the month each outcome was recorded.
- **Open the workroom** — the band above the proof hero opens the ledger as a five document workroom: the evidence ledger row by row, the bridge from claimed to invoice verified savings with every accepted, awaiting, and contested outcome narrated on its stored verdict (omitted while nothing is claimed), a forwardable CFO brief, and a What now sheet. ⚡ **What now** drafts the CFO note from the ledger's own verified figures, copies the savings one pager, opens the proof statement, mints a verified win badge, shelves the record on the board, opens The Desk, and puts the next win on the ledger. ✦ **Ask Vera** reads

the ledger, so a question about what is actually verified is answered from what the ledger verified rather than from what was claimed.

How to use it

1. Record a negotiated outcome as each deal closes.
2. Export the ledger as CSV for finance.
3. Find the next saving to go after in Opportunities.

Related: Verified wins · Opportunities

VERIFIED WINS

Where: Spend & savings › Savings proof › Verified wins tab · `/savings/wins`
Who: everyone

A tab on Savings proof, sharing its underline tab row. The bragging layer on top of the ledger: mint a public, revocable badge from any invoice-verified saving, follow your Deal Score against the calibrated corridors, and (owners only) list your organization on the opt-in board.

On the screen

- **Ready to mint** — only savings confirmed against a real invoice appear here; claims never mint.
- **Deal Score** — where your confirmed deals land inside the corridors, 100 being the strong end.
- **The board** — opt-in, owners only, ranked under whatever name you choose; it opens at five listed orgs.

How to use it

1. Mint a badge from a verified saving and copy the public link.
2. Check your Deal Score against the corridors.

Related: Savings proof · Vendor scorecards

SPEND

Where: Spend & savings › Spend · `/spend`
Who: everyone

In plain words: what the estate costs, in figures finance recognizes.

When to reach for this

- "What do we actually spend?" answered from the ledger, not from estimates.
- Concentration questions: on the sample ledger, one $410,000 cloud agreement dwarfs the four core vendors' $331,200 combined, and that shape is the finding.
- The denominator for every savings story you tell.

The values and finance lens: spend analytics over time, built from the annual values on your proposals and agreements, with invoice actuals beside them. Headline numbers paint immediately and the heavier charts stream in.

On the screen

- **Vitals** — total annual spend, how much of it is up for renewal inside ninety days, and the vendors tracked.
- **Category breakdown and trend** — spend by category and over time, including what renews when.
- **Forecast and actuals** — the forward forecast, and invoice actuals when you have imported them.
- **Vendor chart** — spend by vendor, with a CSV export.
- **Exports** — every estate export in one card: the Excel and PDF reports and the raw CSVs. Portfolio links here.

How to use it

1. Read the trend to see where spend is growing and what renews next.
2. Add missing annual values on agreements to fill the picture.
3. Concentration lives on Portfolio; the link under the category chart takes you there.

Related: Portfolio · Invoices

INVOICES

Where: Spend & savings › Invoices · `/invoices` · via **More** or Ctrl/⌘ K
Who: everyone

Invoice line-item intelligence: what the vendor actually billed, reconciled line by line against what you contracted. Every line gets a verdict with the dollars at stake.

In plain words: every invoice line checked against the contract, with the overbilling priced.

When to reach for this

- Standing hygiene: set the forward-in address once and every invoice checks itself (the dispute recipe).
- A bill feels wrong: the line verdicts replace the feeling with a figure.
- After a price-list move: whether your invoices already reflect it is a fact, not a worry.

On the screen

- **Vitals** — what has been billed above contract, what you have recovered, what is in dispute, and the worst offender by dollars.
- **Line verdicts** — per line: ok, price above contract, billed over entitlement, off-contract item, uplift-cap breach, or above MSA rate card (a consulting line billed higher than the rate on your saved MSA rate card; see Rate Card Benchmark).

- **Open the workroom** — once findings exist, a band above the invoice list opens the analyzed invoice with the most at stake as a six document workroom: the parsed invoice with every flag marked, an overbilling brief, the recovery bridge, the itemized recovery sheet, and a What now sheet. ⚡ **What now** drafts the vendor dispute letter and the AP payment hold note from the invoice's own lines, files the dispute on this page's own desk in one click (reporting the real outcome rather than a generic line), copies the dispute summary, opens the agreement the invoice bills against, and opens a negotiation dossier. ✦ **Ask Vera** answers from this invoice's own line items and the contract terms they should have matched. Every figure is read from the stored invoice lines.
- **File as verified saving** — on a dispute resolved with money recovered, one press writes the recovered amount to the savings ledger with the invoice number in the note. It is counted once, as a one-time recovery, never multiplied across a term, so a recovery does not inflate your annual savings figure. The card then reads **On the savings ledger** with a link to it.
- **Forward-in address** — connect the invoice rail to get a private email address; invoices forwarded there are processed automatically.

How to use it

1. Upload an invoice, or set up email-in so they process automatically.
2. Work the lines flagged above contract or over entitlement, dollars first.
3. When a dispute is resolved with money back, file it as a verified saving so the recovery joins the proof ledger.

Tips

- Setting up the forward-in address is an owner action under Settings › Integrations.

Related: Billing watch · Anomalies

STRATEGY

Where: Spend & savings › Strategy · `/strategy` · via **More** or Ctrl/⌘ K
Who: everyone

Once a strategy is generated, both the state-of-the-stack panel and each category strategy offer **Export for the board** (Word) and **Board deck** (PowerPoint), each rendering the stored generation in the executive brief theme.

An annual, board-ready read on your software stack: category strategies, vendor posture, and the year's negotiation calendar in one document.

In plain words: the year's plan for the stack, as one document you can defend.

When to reach for this

- The annual planning sit-down (the recipe) and the board pack it feeds.
- When posture questions arrive ("are we consolidating or diversifying?") and deserve a written answer.

On the screen

- **Vitals** — annual spend covered, categories in the stack, renewals in the next twelve months, and how many categories already carry a strategy.
- **Category sections** — per software category: current position, market movement, and the recommended play.
- **Export for the board** — downloads the current generation as a Word brief in the executive theme, ready for the board pack unedited.

How to use it

1. Generate or refresh the strategy.
2. Work the category sections in spend order: the recommended play on the biggest category is the year's first move.

Related: Portfolio · Estate map

VENDOR CALENDAR

Where: Renewals & deals › Vendor calendar · `/vendor-calendar` · via **More** or Ctrl/⌘ K
Who: everyone

The timing desk: the fiscal year ends and quarter closes of the top 40 software vendors, read against your own notice deadlines. Vendors discount hardest when their quarter is closing, so this is when to time your asks. The navy masthead states the position in figures: how many vendors close within 90 days, how many inside 30, how many of your deadlines meet a close window and what they are worth a year, and how many vendors you track.

In plain words: when each vendor's quarter closes, so your ask lands when their discount authority peaks.

When to reach for this

- Booking the vendor conversation: their calendar, not yours, decides the good weeks.
- The annual plan (the strategy recipe): quarter ends layered against your notice windows.
- A stalled deal: sometimes the fix is waiting three weeks for their quarter to squeeze them.

On the screen

- **Four views on one band** — the masthead's band switches between **Timeline**, **Agenda**, **Close pressure**, and **Your renewals**. Press **1** to **4** to switch views and **/** to jump to the vendor search.
- **Timeline** — every vendor's forward closes on one axis anchored on today, the next 30 days shaded as the pressure band, your notice deadlines drawn on the matching vendor rows. Search, the horizon control (30 days to 12 months), and the Tracked filter ride a pinned toolbar.
- **Agenda** — every quarter and year close inside the horizon with your own notice deadlines threaded in, one chronological ledger grouped by month. Deadline rows are tinted so your dates stand out from theirs.

- **Close pressure** — the twelve months ahead as a grid, one card per month with its close count and the vendors closing in it; the busiest window is flagged in gold.
- **Your renewals** — the bench where one of your notice deadlines meets a vendor's close, soonest first, each with the one line timing read and an alignment strip reading left to right from today: your notice bar, their close diamond, and the shaded gap between them is your negotiating window.
- **The vendor dossier** — click any vendor anywhere (a timeline row, an agenda line, a pressure chip, the bench) and their dossier opens on the right: their next four closes, the timing read, and your agreements on that vendor. Each agreement carries its four moves: draft the notice letter, open the vendor in a Benchmark case, book the negotiation window as calendar entries, and watch the deadline, which stays off unless you turn it on. The dossier has an address, so a shared link lands in it.
- **The rail** — the twelve month close density strip, your tracked vendors with each one's next close, and the roads onward: your renewal calendar, your agreements, and the watchtower.
- **The star** — track a vendor anywhere on the desk and it lands on your dashboard and in the rail watchlist.

How to use it

1. Read the Agenda for the window ahead: your deadlines and their closes in one list.
2. Line a renewal up against the vendor's quarter end for the deepest discount window, and land the ask inside the shaded 30 day band.
3. When a deadline of yours meets their close, open **Your renewals** and work it from the bench: the letter, the case, the calendar file.

Related: Renewals · Agreements

MEETING WAR ROOM

Where: Renewals & deals › Meeting war room · `/meetings` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan · the app joins on the Enterprise plan (workspaces from before July 29, 2026 and trials keep access; see Plans, allowances & limits)
You get: a pre-call pack with a printable one-page agenda · assembles in about a minute

In plain words: the pre-call pack, assembled from your own numbers, before every vendor meeting.

When to reach for this

- Any vendor call worth walking into prepared, which is all of them.
- Connected to your calendar, it preps without being asked; Troubleshooting covers the connection when packs stop appearing.
- The debrief slot after: what happened stays with the meeting, not in someone's head.

One meeting is one room. Name the vendor and the time, or drop the calendar invite, and the pre-call pack assembles from your stored numbers: where you stand, what is open, the figures to cite at the table, and the three questions to ask. After the call, the debrief slot keeps what happened with the meeting.

On the screen

- **The intake** — pick the vendor (your tracked vendors autocomplete) and when the meeting happens, or drop an .ics invite anywhere on the panel. Meetings found in the invite are matched against your vendors and become one-click packs.
- **The room** — the posture headline, a countdown to the meeting, and three fact panels: Where we stand, Open items, and Ammunition. Every line comes from facts already on file (spend, notice deadlines, benchmark standing, open disputes, vendor pressure).
- **The three questions, in order** — cue cards you can read out loud on the call.
- **The agenda paper** — a one-page agenda in the executive brief format. **Print the agenda** prints just the paper.
- **After the meeting** — the debrief slot: how it went (advanced, concession, stalled, information only) and two lines on what happened. Saved with the meeting, waiting whenever the room reopens. Saving a debrief adds a What next panel: file the commitments in the vendor room, ask Vera about the vendor's stated position, or assemble the follow-up pack.
- **The shelf** — upcoming meetings stand as tent cards (weekday, vendor, time, headline); past meetings sit in a dimmer ledger below. Click any of them to enter its room.
- **Connect the calendar feed** — add your private ICS feed (Outlook: Publish calendar; Google: Secret address in iCal format) and the agent assembles packs on its own, about a day ahead. It reads only titles and times.

How to use it

1. Name your next vendor meeting (vendor plus time) and press **Assemble the pack**, or drop the calendar invite on the page.
2. Read the room before you dial in, and print the agenda if you want it on paper.
3. After the call, write two lines in the debrief slot so the outcome stays with the meeting.
4. Optionally connect a read-only calendar feed under Settings › Integrations; the agent then preps a pack for every vendor meeting it spots, and its notification opens the room directly.

Tips

- **Refresh the pack** rebuilds the room from today's facts, useful when a dispute closed or a new benchmark landed since it was assembled.
- Packs are grounded only in stored data. If a panel says nothing is on file, upload the agreement first so the room has numbers to stand on.

Related: Vera AI · Vendor calendar

DOCUMENT SEARCH

Where: Contracts & documents › Document search · `/doc-search` · via **More** or Ctrl/⌘ K
Who: everyone

The Records Hall: ask a question in plain English across every contract and proposal you have uploaded and get the exact clause back, cited to the contract it came from, then compare, save, restore, and present what you found without leaving the room.

In plain words: ask your contracts a question, get the clause back with its citation.

When to reach for this

- "What did we agree about termination with this vendor?" answered by the paper itself.
- A dispute needs the exact language: the clause, cited, beats a recollection.
- One question across many contracts at once belongs in Review tables instead.

Opened from your home screen (app mode), the hall leads with the question bar ready and the page header sits below it. Opened from the dashboard, the page keeps the composition described here, header first.

On the screen

- **The desk rail** — down the left side, the room's fixtures: the question shelf, premade briefs, and the shared work list.
- **The question shelf** — the questions buyers run most, grouped by what you are hunting (exposure, money, renewal, and what to improve), including list the least favorable terms and which terms to improve at the next renewal. One click runs any of them.
- **Premade briefs** — one click runs the search and then writes the executive findings brief from whatever comes back, unattended: **What to fix first**, **Renewal exposure**, and **Liability exposure**.
- **The work list** — what this room has flagged, shared with your whole workspace. On any exhibit, press the bookmark to flag its contract as **Improve the terms**, **Renewal negotiation**, or **Watch**; the flag becomes a workspace task that waits in the rail, on the Tasks page, and in My work until someone checks it off. Each item links back to its contract.
- **The index observatory** — a cluster at the top of the stage states your archive at a glance: a coverage ring showing how much of the archive is searchable, plus documents indexed, contracts, and vendors, and how many are still waiting to index. Vendors are counted as companies: Salesforce UK Limited, SFDC Ireland, and Salesforce, Inc. count as one Salesforce.
- **The library wall** — your contracts stand as a wall of spines, one per document, grouped and labeled by vendor company, so every legal entity of the same company shelves together. A spine is lit gold where its text is searchable and dim while it waits. Hover a spine for the contract name, click it to open the contract. When an answer cites a vendor, that company's spines glow gold, and hovering an exhibit or a citation chip flares its spines brighter.
- **The ask bar** — type a question across your whole portfolio, or press **/** anywhere in the room to focus it. Type ahead surfaces your standing questions, recent asks, vendors on the wall, and proven clause patterns as you type. The scope chip counts how many documents are indexed, Try chips offer starting questions, your starred questions wait as gold Standing chips, and your last three questions wait as Recent chips.
- **Index button** — appears when documents are waiting. Indexing runs as a progress bar that names each document as it becomes searchable and lights its spine.
- **Sample documents** — when your archive is empty, a **Try with sample documents** button (also offered in the empty state below the ask box) loads four realistic sample contracts, already indexed, so you can try a

search before uploading anything. The sample is clearly labeled and removable from the dashboard at any time.

- **The staged search** — while a search runs you see each stage: your question turned into legal search terms, the scan across your indexed documents, passage ranking, and the drafted answer.
- **The finding header** — when the answer lands the room turns dark and the documents light up. The header states the question, then the method and the counts as a plain line (a literal scan always, a literal and semantic scan when meaning-based retrieval contributed, the clauses cited, the documents read), a star that keeps the question standing on the stage, and the finding's moves: **Write the brief** in gold because it is the thing you forward, then **Send the exhibits**, **Present**, and **Copy**.
- **The bottom line** — the cross-contract synthesis in a sentence or two on cream paper, with gold numbered citation chips inline. Hover a chip to light the matching exhibit and its spines on the wall; click it to open the exact clause on the workbench. **Read the full analysis** unfolds the rest of the written answer when you want it.
- **The evidence and the dispositions** — the results are split by what they are. A **view** is a way of looking at what was found and lives in the docket on the left: **Evidence** and **Coverage**. A **disposition** is a thing the finding becomes and waits permanently on the right, each showing whether it exists yet: the findings brief, Send it on, the project plan, and What next. Nothing is hidden behind a tab.
- **Evidence** — every cited clause arrives stacked in order, so you read the whole bundle by scrolling instead of pressing Next once per exhibit. **Show one at a time** switches to reading them singly, with previous and next and the **J** and **K** keys; **Show all** puts them back. Each clause is a lit sheet: the exhibit number and its topic, the vendor, the document it came from, the exact quote with your question's key terms flared, why it answers, and the sheet's own actions: **Reading room**, **Copy as citation**, **Work list**, **Pin**, and **Open the contract**.
- **The docket** — the left column lists every cited clause with its number, vendor and contract. Select one to jump to its sheet; when the whole bundle is on screen it scrolls to that sheet rather than replacing it. The same column switches between the Evidence and Coverage views.
- **Coverage** — every vendor in your portfolio accounted for: matched vendors glow gold, silent ones stay gray, the strip counts them, and the contracts the answer cited link straight through.
- **The project plan** — ask Vera to draft it from the dispositions and the finding becomes a 6 to 12 week action plan on a Gantt board: phases group the bars, weeks run across the top, and the owner (You, Procurement, Legal, CFO, IT) rides each bar. The board opens over the finding rather than replacing the evidence. The plan is saved with the finding, states what done looks like, copies as text for your project tracker, and redrafts on demand.
- **Send it on** — one door, not two. **Send the exhibits** in the header, or **Read and send** in the dispositions, opens the note with the finding, the cited clauses and your optional message, addressed to anyone you name. Nothing sends until you press send.
- **Pin and compare** — pin exhibits to the bench tray at the bottom right, then press **Compare** to see the pinned clauses side by side on paper columns.
- **The reading room** — any exhibit opens full focus: the clause in large type on cream paper, why it answers, previous and next a keystroke away, and copy as citation one click away.

- **Saved findings** — every answer saves to your browser with its pins and its project plan. The strip above the answer restores any saved finding instantly, with no re run and no second AI charge, and coming back to Document search later opens on a **Pick up where you left off** shelf listing your recent findings. Remove a finding from the strip with its ×.
- **What next** — the finding ends in moves rather than in more questions: open a negotiation with the vendor the answer leads with, flag the weakest clause for the next renewal, or press Vera on what the exposure is worth.
- **Press further** — at the foot of the finding, chips offer three drafted follow ups built from what the answer found, beside the **Ask a follow-up** button.
- **Forward** — sends the finding to a colleague as an exec brief email directly from the platform: the answer, the exact clauses quoted, your optional note, and a workspace link that re-runs the search live. The recipient needs a VendorBenchmark login with access to this workspace to use the link; the quoted evidence travels in the email itself.
- **Generate report** — turns the findings into a short executive brief, about four pages: an executive summary, where the exposure and the leverage sit, the cited clauses verbatim, and a prioritized list of next moves. The brief is built only from the clauses your search cited, so every line is grounded in your own contracts. Print or save it as a PDF from the report toolbar. The report carries its own Send it on desk below the paper.
- **Present** — runs the finding full screen on the shared stage in three acts: the answer with a coverage map across your portfolio, one screen per cited clause with the clause enlarged as the evidence, and a closing bottom line that recaps who matches and where each quote lives. Arrow keys move between exhibits; the presenter notes, ink pen, and printable leave-behind work as on every other show.
- **The keyboard bench** — **/** focuses the ask bar, **J** and **K** walk the exhibits, **Enter** opens the reading room on the active exhibit, **P** pins it, and **G** takes the stage. Typing in a field is never hijacked.

How to use it

1. If you have nothing uploaded yet, press **Load sample documents** to try the search on four sample contracts.
2. Press **Index documents** if any are waiting, and watch the wall light up as each one becomes searchable.
3. Ask a question like which agreements auto renew, tap a Try chip, or press **/** and let type ahead finish the thought.
4. Read the bottom line, hover a gold citation chip to see which exhibit and which spines it lights, then click it to open the exact clause on the workbench. Walk the exhibit list with **J** and **K** or the arrows on the sheet.
5. Pin the two or three clauses that matter with **P** and press **Compare** to put them side by side.
6. Star the question to keep it standing on the stage, then follow a Press further chip to run the next question in the line.
7. Press **Generate report** to turn the findings into a four page brief, or **Present** to walk the finding full screen, ending on the bottom line.
8. Send it on: open the **Send it on** tab and let Vera draft the follow up email for the reader you pick, or press **Forward** to email the finding, clauses quoted, straight to a colleague.

9. Open the **Project plan** tab and press **Draft the project plan** to turn the finding into a dated Gantt board.
10. Flag what needs fixing: press the bookmark on an exhibit and file its contract on the work list as improve the terms, a renewal negotiation item, or one to watch. The whole team sees it until it is checked off.
11. Leave and come back whenever: your findings wait on the **Pick up where you left off** shelf with their pins and plans intact, and the work list keeps the follow ups.

Related: Document clusters · Rights register

DOCUMENT CLUSTERS

Where: Contracts & documents › Document clusters · `/doc-clusters` · via **More** or Ctrl/⌘ K
Who: everyone

The survey bench now says when the map is stale: documents added since the last grouping show as a badge with the count, and the newer documents are named as not yet in any group, with the Re-group control beside it. Nothing regroups without your press.

Your contracts grouped automatically by what the documents actually say, subject matter rather than folder names. The grouping is deterministic text math; the AI only names the groups.

On the screen

- **Cluster cards** — each group with its AI-written name and the documents inside it.
- **What next** — under the map, the next moves: price the overlap in the IT spend optimizer, open the ranked opportunities, or ask Vera where the biggest cluster can collapse first.

How to use it

1. Run or refresh the clustering.
2. Open a document from its group.

Related: Document search

REVIEW TABLES

Where: Contracts & documents › Review tables · `/review-tables` · via **More** or Ctrl/⌘ K
Who: everyone · requires an AI plan
You get: a quote-backed answer grid across all your contracts, exportable to Excel · fills a few documents at a time

Ask the same questions across many contracts at once. Each column is one question (termination notice, renewal cap, liability cap) and the AI fills in a quote-backed answer for every document.

In plain words: one question per column, a quote-backed answer per contract, exportable.

When to reach for this

- The estate-wide questions ("who has an uplift cap?", "what are our notice periods?") that would otherwise take a week of rereading.

- Diligence and audits: the answer grid is the evidence table, to Excel in one click.
- Anything the Coverage sweep does not already ask; coverage is your standard, this is your bespoke question.

On the screen

- **The drawer** — the masthead states how many tables you hold, how many documents they cover, how many answers are extracted, and what share of every cell is filled.
- **In the drawer** — each table as a card with a live mini grid of its answers and a fill ring, opened with one click.
- **Columns** — your questions. Start from a template or add your own, up to twelve.
- **Fill table** — extracts every missing answer, a few documents at a time, with the supporting clause behind each cell.
- **Export** — the whole grid as an Excel workbook, quotes included.
- **What next** — once answers land, the grid ends with its next moves: export it for the deal file, fix a weak answer with preferred language from the clause library, start a negotiation, or ask Vera which answers should worry you most.

How to use it

1. Create a table from the Termination & renewal template, or add your own columns.
2. Fill the table, then click a filled cell to see the exact clause it came from.

Related: Document search · Clause library

DATA HEALTH

Where: Spend & savings › Data health · `/data-health` · via **More** or Ctrl/⌘ K
Who: everyone

What the platform cannot see for your organization and what filling each gap unlocks. The scan runs on every load; a prioritized plan tells you which gap to close first.

In plain words: what the platform cannot see yet, ranked by what seeing it would unlock.

When to reach for this

- The migration finish line (Chapter 18): "done" is a score here, not a feeling.
- Before an audit: the exposure estimate says how defensible the entitlement picture is today.
- Monthly: one closed gap a month compounds faster than any single feature.

On the screen

- **Vitals** — how much of your estate the platform can see, how many gaps are open, how many records are visible, and your audit exposure.
- **Completeness score** — the share of data checks you have covered. Each open gap names the feature it would unlock.
- **Audit exposure** — an estimate of your license compliance exposure, built from the entitlement data connected so far.
- **AI plan** — the gaps ranked by value, refreshed whenever the underlying gaps change.

How to use it

1. Close the highest-value gap first.
2. Connect a data source to fill the entitlement gaps (an owner action under Settings › Integrations).

Related: Integrations

COMMUNITY POLLS

Where: Benchmarks & market › Community polls · `/polls` · via **More** or Ctrl/⌘ K
Who: everyone

Four standing polls answered by other software buyers on the platform: the most difficult vendors, audits in the last 12 months, coverage requests, and renewal price increases. Vote to see how your peers answered; individual ballots stay private.

On the screen

- **Vitals** — ballots on the floor, boards open to you, the vendors in play, and when the next poll opens.
- **Board index** — a docket of the four boards that follows you as you scroll: a rail on wide screens, a chip strip on small ones. Boards still waiting on your ballot carry a blue dot; voted boards a green check. Click a board to jump straight to it.
- **Results bars** — unlock after you vote, so results are not biased by peeking. A live board names its leader in the header and badges the top three ranks.
- **Ballot search** — on long vendor lists (the coverage board runs the full curated list), type to filter instead of scrolling; your picks stay pinned above the list so a search never hides them.
- **Suggestion box** — what peers vote on next. The strongest suggestion becomes the board that opens on the first of the month.

How to use it

1. Walk the index top to bottom and cast a ballot on every board with a blue dot; each vote reveals that live board.
2. On a long ballot, search the vendor name; remove a pick from the pinned row if you change your mind.
3. After a vote, follow the **Next open board** link at the foot of the results to the next board waiting on you.

Related: Platform pulse · Buyer Network

PLATFORM PULSE

Where: Benchmarks & market › Platform pulse · `/pulse` · via **More** or Ctrl/⌘ K
Who: everyone

What other teams on the platform are using, published from real usage. Aggregate counts only; no team is ever named.

On the screen

- **The pulse report** — which features teams across the platform use most, drawn from real usage.

How to use it

1. Read the published pulse to see what peers lean on, and try a feature you have not.

Related: Community polls

BUYER NETWORK

Where: Benchmarks & market › Buyer Network · `/buyer-network` · via **More** or Ctrl/⌘ K
Who: everyone

Vendor-anchored discussion with other verified software buyers, composed as a trading floor: the market corridors are always on the table, even when the conversation is quiet. Everyone posts under a stable desk handle with banded, platform-verified credential chips; nobody sees a name, an email, or a company. This is one of two rooms under the Buyer Network umbrella: a shared tab strip links these threads with the one-to-one Buyer intros.

On the screen

- **Vitals** — open threads, peer replies on the board, vendor rooms, and how many carry a calibrated corridor.
- **The rooms** — vendors with live discussion plus flagship quiet rooms, each carrying the house corridor: the calibrated band where that vendor's renewals settle, drawn with its median tick and the low, median, and high figures printed at the data. Click a room to filter the board to that vendor.
- **Your desk** — the pseudonym you post as, with your verified chips and your contribution record (threads started, replies posted). It stays yours, so a reputation can build on it.
- **Credential chips** — verified buyer and estate bands computed from platform data, never self-claimed.
- **The house read** — the calibrated renewal corridor, shown live in the composer the moment you type a recognized vendor, and charted beside every thread so the conversation starts from the same market facts the benchmarks stand on.
- **Sort lenses** — Latest, Most active, or Unanswered. Threads still waiting on a first reply carry a gold flag so help finds them.

- **Verified position chip** — opening a thread can attach your organization's anonymized market position on that vendor, a percentile band verified from your stored standings and rounded to 5 points so it can never fingerprint a deal. It renders as a gold chip on the thread; your identity never rides along, and the option only attaches when your org actually holds a standing on the vendor.
- **Report** — every thread and reply carries a Report control. One click flags it for the platform moderation team; nothing is removed until a moderator reviews it.

How to use it

1. Open a thread on a vendor you are renewing; the composer shows the house read as you type the vendor, so you can cite the corridor in your opening post. On a quiet board, one-click openers prefill a question peers argue about.
2. Click a room, or use the vendor filter, to read everything the network holds on one vendor.
3. Sort to Unanswered and give a peer their first reply; your desk's record grows with every contribution.
4. Switch to Buyer intros from the tab strip when you want a private one-to-one.

Related: Buyer intros · Vendor scorecards

BUYER INTROS

Where: Benchmarks & market › Buyer intros · `/network` · via **More** or Ctrl/⌘ K
Who: everyone

The Buyer Network's intro room: talk to another buyer who runs the same vendor, one to one. Anonymous until you both accept, with contact details shared only on a mutual yes. The page is a launcher; the work happens in the full-screen **intros workspace**, five pages on one dark stage with Ask Vera docked beside every one of them. The shared tab strip links back to the Buyer Network threads.

On the landing

- **The network pulse** — the page opens on the numbers: organizations active, requests waiting for a peer, calls connected, and the share rated a good call, plus the most asked-about vendors of the last 60 days. Press a trending vendor and the workspace opens on the ask composer with it filled in.
- **Enter the workspace** — the navy band opens the workspace on The floor; beneath it, one card per workspace page shows its live figure (open requests, decisions waiting on you, organizations in the network, your profile strength) and opens that page directly.
- **Sample view** — one press opens the workspace over a realistic example week, clearly marked with actions disabled, so you can feel how the network works before yours is warm.

In the workspace (pages on the left rail, ← → to move between them, Esc to close)

- **The floor** — the board at a glance: the pulse verdict and figures, the requests that fit you best, trending vendors, recent connections, and the most helpful desks, with your next move one press away.
- **Talk requests** — the ask composer and the open board. One-click challenge templates (renewal uplift, audit posture, packaging change, price sanity check, switching away), a live reach line that says how many

organizations get notified the moment you post, and a live preview of exactly what peers see: your alias, country, and industry, never your organization. The composer warns if your text contains an email or phone number; contacts are shared automatically on a mutual yes, so they never belong in the post. Requests on vendors from your talk list (and anything sent directly to you) surface first in a For you lane.

- **Requests & matches** — the reciprocity ledger (your helper credits translated into the planned renewal discount, 1% per credit toward the 10% cap), then every match on four stages: Offered, Mutual yes, Connected, Rated. Connected matches carry the other person's contact details and a suggested 30 minute agenda built from the vendor and the challenge, ready to copy into the calendar invite.
- **Our profile** — a profile strength meter that names exactly what is missing, watched vendors as add-and-remove chips, and the visibility switches (accept direct requests, appear in the network).
- **Community network** — the directory with industry facets and vendors-in-common badges, plus the most helpful desks, recent connections, and the network's reach by industry and country. Helpers rank first.
- **The dock** — three panels beside every page: **What now** with your next moves on the board, **Ask Vera** grounded on the live board (the open requests, your matches, your profile), and **Who sees what**, the three-stage anonymity ladder.

How to use it

1. Enter the workspace and read The floor; the requests that fit your talk list lead.
2. Post that you want to talk about a vendor (a template gets you moving), or volunteer on an open request.
3. Mutually accept, reveal, and only then share contact details; use the suggested agenda for the call, then rate it. Good calls earn the helper a credit toward the renewal discount.

Related: Buyer Network

Next: Chapter 05 · Team, AI & assistant

05 · Team, AI & assistant

This chapter covers the shared-work queues where your team hands work back and forth (tasks, approvals, and team coverage), Vera the AI analyst grounded in your own portfolio and documents, the automations and deliverables built on top of her, and the ways your team talks to each other and to our analysts. Everything here is about getting work assigned, answered, and moving.

TASKS

Where: Team & AI › Tasks · `/tasks`
Who: everyone

Shared to-dos for your team as one runway with four views: follow-ups on renewals, negotiation actions, and anything else worth not forgetting.

In plain words: the team's shared to-do list, attached to the work it belongs to.

When to reach for this

- Splitting a renewal's actions across people, with owners and dates.
- Handing work over cleanly: reassign the open items when someone leaves or a portfolio changes hands.
- Your own queue lives on My work; this page is the team-wide view.

On the screen

- **The runway masthead** — states where the work stands: open, due today with the overdue count called out, in progress, and cleared this week. These four move as you work, so the figures are always the board you are looking at. The New task action lives here too.
- **Four views, one board** — the band under the masthead switches views: **Board** (the drag and drop lanes), **My day** (every active task in landing order, grouped past their date, due today, tomorrow, this week, further out, and no date), **Ledger** (the whole runway as a dense table, with a toggle to include Done), and **The bench** (who is carrying what, with a load bar per teammate and their next due task). Press 1 to 4 to switch, / to search, N for a new task.
- **The task file** — click any task in any view to open its file, a right hand panel where the status, due date, owner, and notes all edit in place. The file is in the page address, so a forwarded link opens straight on that task.
- **The pick up tray** — the bench view leads with tasks that have no owner. **Take it** assigns one to you with a click; unowned work is how a date gets missed.
- **One click complete and defer** — in My day, the circle at the left of a row marks it done, and the Today and +1w moves push its date without opening anything.

- **The pace instrument** — the rail answers when the runway clears: drag the slider to test a weekly clearing pace against the active count.
- **Delegate to Vera** — every open task carries a sparkle control (on the card, the row, and inside the file). Delegating hands the task to Vera, the AI teammate: she reads the task and its linked record, writes the finished work product into the task notes, moves it to Done, and notifies you. Requires an AI plan.
- **What next** — the runway ends with a What next block built from the board's own numbers, including handing the most overdue task to Vera.

How to use it

1. Create a task with N or the New task button, assign it to the teammate who owns it, link it to the vendor or renewal it belongs to, and set a due date.
2. Run the week from My day: work the list top to bottom, complete with the circle, and push dates with the defer moves when the plan changes.
3. Once a week, open The bench: empty the pick up tray and balance the load bars before anything slips.
4. To see just the items assigned to you rather than the whole team, use the Assigned to me filter or open My work.

Related: My work · Team

APPROVALS

Where: Team & AI › Approvals · `/approvals` · via **More** or Ctrl/⌘ K
Who: everyone

Sign-offs your team routes through the platform, for example a negotiation mandate waiting for a budget owner. The page is a four tab workspace: the queue leads, the setup sits behind its own tabs.

In plain words: the sign-offs waiting on people, and the chain they run through.

When to reach for this

- You are an approver: the queue is where deals wait on you.
- A deal needs its documented sign-off before signature; the chain itself is configured by owners under Deal approvals.

On the screen

- **The gate** — the masthead states what waits and what it is worth: requests needing a human decision, the combined annual value on the line, how many have run past their SLA, and the median time from request to decision over the last 90 days.
- **Auto-cleared** — under the masthead, the share of the last 90 days' decided requests that your policies cleared before anyone read them. The rest reached a person because a rule said they should.
- **Queue** — every request waiting on a human, oldest wait first. Each row shows the wait or SLA chip, who it is routed to, and the AI risk summary; approvers decide right on the row. Each pending request also

carries a Vera's read line: the ask, the vendor's delivered benchmark position, and the routing rule that queued it. The tab count turns red when a request is past its SLA.

- **Decided** — the decision ledger: cycle time analytics (median days from request to decision, decisions per month), plus a search box and an outcome filter (approved, rejected, auto approved, human decided).
- **Automation** — your auto approve policies and the rule proposals learned from your own decision record. Every auto approval lists the guards it passed.
- **Vera's review** — on a deal sign-off request, Vera reviews the snapshot before the chain votes: the latest quote against the mandate target, the concession balance, and the price movement. Her Flagged or Clear verdict, the checks behind it, and a short analyst note render beside the approval panel. She advises the chain and never casts a vote.
- **Routing** — the authority matrix (which type and amount band routes to whom, with optional second steps and SLAs) beside the approval threshold.
- **New request** — the button on the masthead opens a side sheet to raise a renewal or spend item for sign off.
- **The record behind the ask** — every row links the thing being approved, named for what it is: open the contract, the negotiation, the case, the renewal, the sourcing project. It is there on pending rows, so you can read the record before you decide rather than after, and it is offered again the moment you approve or reject.
- **What next** — the page ends with a What next block built from the queue's own numbers.

How to use it

1. Open the Queue tab, open the record behind anything you are not sure about, then approve or reject each item. A rejection asks for a reason; it becomes the audit trail.
2. To raise something for sign off, use **New request** on the masthead; routing rules pick the decider.
3. Track past sign offs and your approval cycle time on the Decided tab.

Related: Deal approvals setup · My work

TEAM

Where: Team & AI › Team · `/team` · via **More** or Ctrl/⌘ K
Who: everyone

Who owns what across your organization, and more importantly what nobody owns: renewals without an owner, open tasks, and assigned findings.

In plain words: the ownership map, with the unowned things surfaced first.

When to reach for this

- The weekly "is anything falling between people?" check: an unowned renewal inside its window is the finding that matters.
- Rebalancing load when someone is drowning or departing.

On the screen

- **The bench** — the masthead states the coverage verdict: people in the workspace, owned renewals and open tasks in flight, what sits in the pick up tray with no owner and what that is worth a year, and the renewal cycles closed in the last eight weeks.
- **Team pulse** — renewal cycles closed per week over the last eight weeks, once anything has closed.
- **Member rows** — each teammate with the renewals, tasks, and findings on their plate.
- **Unowned renewals** — agreements with no renewal owner. Assign these before a notice window gets missed.
- **What next** — the page ends with a What next block built from the bench's own numbers, the unowned work leading.

How to use it

1. Scan the unowned renewals first and assign an owner to each one.
2. Use the member rows to balance who is carrying what.

Tips

- Owners can invite a teammate from here when there is more work than people to cover it.

Related: Renewals · Members

BOARD

Where: Team & AI › Board · `/board`
Who: everyone

The team's durable forum beside the chat stream: decisions, vendor intel, questions, and wins, shelved so they can be found again months later.

On the screen

- **Category tabs** — Decision, Vendor intel, Question, Win, and General, with pinned posts leading every list.
- **Posts** — a title, the write-up, and threaded replies. Record links in a post render as grounded cards (the same cards chat shows), and @mentions notify the teammate.
- **Link a record** — the composer's control for inserting a proposal, renewal, deal, or vendor reference; the record's page lists the post under **Discussed in**.
- **Templates** — the composer's Start from pills prefill the three shapes teams shelve most: a Decision memo, a Vendor review, and a Renewal retro.
- **Accepted answers** — on a question, the author (or an owner) marks the reply that settled it; the list shows a check so nobody re-litigates.
- **Pins** — owners pin the posts the whole team should keep seeing.
- **Open the workroom** — on the unfiltered view, once a post exists, a navy band opens the board as a workroom. It carries a cover sheet with live counts, an exhibit per category, the question docket, and a

What now sheet. ⚡ **What now** drafts the board digest for circulation, emails it to your team, copies the digest or the one line summary, proves the wins as savings, opens the estate the records govern, takes the next decision to the Desk, shelves the next record, and opens the newest one. ✦ **Ask Vera** answers from the records themselves rather than from a summary of them. It reads live from your newest posts; nothing is generated or stored, and nothing sends from a workroom.

How to use it

1. When a decision lands (walking away from an uplift, standardizing on a vendor), write it up as a **Decision** post and link the deal it came from.
2. Ask the team with a **Question** post; accept the reply that answers it.
3. File what a rep told you as **Vendor intel** so the next negotiator starts warmer.

Tips

- Chat is for the conversation; the Board is for what should outlive it. When a thread reaches a conclusion, shelve the conclusion here.

Related: Chat · Tasks · Record pages

VERA AI

Where: Team & AI › Vera AI · `/vera`
Who: everyone · requires an AI plan
You get: a purpose built brief in the executive house style, opened on screen and filed under your reports · a run takes a few minutes

One page, one job. Drop a contract, quote, renewal or SOW on it, or browse, or pick one already on file. Then say what you want in your own words, or choose one of eleven requests. Vera reads the document as a specialist in that vendor and writes the brief that answers you. **You do not need to know which tool does what**: you say the outcome, Vera picks.

On the screen

- **The masthead** — the door's vitals on the navy stage: briefs written, documents on file, the eleven requests on the menu, and the date of your latest read, with the three steps of every run (reads the paper, places you on the curve, writes the brief) beneath them.
- **The drop zone** — drag a PDF anywhere onto it, press **browse** to choose one, or use the dropdown to work a document already in your workspace.
- **The one line** — type what you want in plain words. Vera matches it to the nearest of the eleven requests and tells you which one it chose before anything starts, so a wrong guess costs you nothing. Press **Pick a different one** to override it.
- **Or pick one** — the eleven requests: benchmark this against other customers; I want to know my options; tell me what is good and bad with this agreement; how should I prepare for the negotiation; make me a renewal plan; review the contract terms and pricing; create me a negotiation strategy; create talking

points for my negotiation; benchmark this and arm me for the negotiation; review the terms clause by clause, pricing aside; plan my final negotiation call. Each names the question a buyer is really asking underneath it, and hovering any of them explains what Vera does with the document and what the finished report contains, so you can pick the right read without a test run.

- **The button** — names what you are about to get, for example "Have Vera write the pricing benchmark".
- **Try it on a sample deal** — beside the button when you have not added a document. Opens a finished example of the read you have selected, written on an illustrative deal and rendered as the real document, so you can see what it produces with no wait and no run before dropping your own paper.
- **See an example of each** — beside the eleven requests. Opens the examples gallery, a finished example of every request written on an illustrative deal and rendered as the real document, so you can pick the request that fits before dropping your own paper.
- **Recent reads** — the numbered dark panel on the right, your most recent Vera reads, up to twenty, each named by the deal it covers with the read type and date beneath it. One click reopens a finished brief, so going back to something you ran last week does not mean hunting through the reports library.

How to use it

1. Drop the PDF on the zone, press **browse** to choose one (up to 25 MB), or pick a document already on file from the dropdown.
2. Say what you want in your own words, or pick one of the eleven requests. Hover a tile first if you want to see what that read does and what the finished document contains.
3. Press the button. It names the document you are about to get.
4. Stay and watch the reading card, or leave the page: the run keeps going, and the finished brief opens itself when you are here or arrives as a notice when you are not.
5. On the finished brief, use the rail on the right: download the PDF, email it to a colleague, or take the suggested next read on the same deal.

While it runs the page stays put: the drop zone itself becomes the reading card, with your file named at the top and a live percentage beside it, and the rest of the page waits dimmed behind it. Inside the card a five step checklist reports the stage the run is actually on, each step ticking off as it completes: reading the paper, pulling the commercial facts, placing it against the market, writing the brief, filing it. These are the real steps, not a timer, so a slow stage looks slow. Under the checklist, a **Read off the paper** list fills in as the read establishes facts: the vendor, the product, the term, the commitment totals in US dollars and any stated discount, up to eight rows, each ticking in as the extraction lands. Every row comes straight from your document; a fact the paper does not state never appears. When the brief is ready the page opens it for you. A run takes a few minutes, and you do not have to wait on the page: leave it and a small marker in the corner tracks the brief, then a notice arrives bottom right with a link straight to it when it is ready. The brief lands in your reports either way, so a closed tab loses nothing.

What comes back is shaped to the request, in one consistent house style. The benchmark request returns the benchmark brief: four headline figures, the executive summary, the deal snapshot, the sensitivity if the baseline is wrong, the market evidence, the levers in the order to play them, the talking points, and the verdict. Peer figures are drawn from the reference set rather than written by the AI, so if there is no comparable cohort for your vendor the brief says so plainly and names what to ask the vendor for, instead of inventing a comparison.

The other ten requests each return a purpose built document. **Tell me what is good and bad with this agreement** prices every clause as an asset working for you or a liability working against you, shows them on one diverging ledger netted into a single position, records a right you never used at nil, and lists anything it cannot price yet as unpriceable with the reason. **How should I prepare for the negotiation** decomposes the proposal into a rate per unit per month for every line, prices each way the deal can be restructured, and lays out the information to demand in writing, the opening, target and walk away for each item, and the order to run the conversation. **Make me a renewal plan** works backward from the renewal date to a dated runway: the clauses to secure before you sign, the two windows where you can act on your own, the renewal cost under each scenario, and the gate at every milestone with the notice date marked as the one deadline that cannot be recovered. **Create me a negotiation strategy** reads your leverage honestly, ranks your objectives from must win to sacrificial, plots every ask on a trade matrix by what it costs the vendor against what it is worth to you, and sets out the currency to trade, the rounds, the plays, the closing sequence and a response to every predictable vendor move. **Review the contract terms and pricing** answers both questions at once: it reconciles every line to its actual unit rate and flags which prices were negotiated round rather than derived, reads the volume schedule, and works through each protection and the condition that removes its value at renewal, with a break-even curve wherever a conditional renewal cap has to be weighed against flexibility. **I want to know my options** sets out every real path open on the deal, from a like for like renewal to a re-mix of the licensing to a combined programme, prices each against the same baseline, ranks them by what they are worth against the weeks and the risk each carries, and plots them on a value against effort chart so the efficient paths stand apart from the ones that cost more time for less. **Create talking points for my negotiation** returns a script written to be spoken: the position said the same way by everyone on your side, the questions to ask in order and what each answer settles, the asks with the exact words and how to frame them, the objections with an answer to each, the phrases that lose money and what to say instead, and a pocket card to carry in. **Benchmark this and arm me for the negotiation** returns a three part pack: page one places your deal against the vendor's own discount curve where one exists or an illustrative peer trend otherwise, with your as proposed, target and stretch positions and the headroom between them; page two is ten tactics specific to that vendor, each with what to think about; page three is the arguments that move it, built on its own incentives and ranked by what they cost it to grant, closing with the one to lead with. **Review the terms clause by clause, pricing aside** returns a terms only assessment with pricing deliberately left out: every clause is graded against a fixed five value key (Protects, One sided, Ambiguous, Weak or Absent), each with the drafting change to seek beside it, the headline counts how many clauses protect you and how many protections are absent altogether, and it closes with the terms the order form does not address and the three fixes to make first in dependency order. **Plan my final negotiation call** returns a two page battle plan for the last session: every item forced into one of five statuses (Must win, Won, Open, At risk, Lost) with the conditions of signature first and the sunk items last, the session minute by minute with one hard rule per block, a concession ladder that releases each rung only against a stated precondition, the counterparty's best closing moves with the pre agreed answer to each, and win, loss and walk defined before the call so none is decided in the room. Every figure and date is computed from your document rather than written by the AI, and every report now argues in US dollars, converted from the contract's own currency at today's live exchange rate.

Forwarding and what next. Every finished report has a rail down the right side. Copy a link to send a colleague in your workspace, press **Download PDF** to take the document exactly as shown, or **Email it to a colleague** and Vera drafts the note for whichever reader you pick (a colleague, your CFO, your CIO or the vendor), ready to copy into your own mail client. Under the rail, what to do next: Vera suggests one or two of the other ten reads that follow naturally from the one in front of you, so a benchmark points you to a preparation brief and a strategy points you to talking points, each opening Vera with that request already chosen for the same deal; or take the numbers into a negotiation.

Related: Vera command desk

VERA COMMAND DESK

Where: Team & AI › Vera command desk · `/ask`
Who: everyone · requires an AI plan

In plain words: the analyst on call, answering from your own records with citations.

When to reach for this

- Any question that starts "what do we…" or "am I…": your portfolio is her context.
- One keystroke from anywhere: Ctrl/⌘ J opens her over whatever you are looking at.
- Asking well is a skill worth two minutes: the recipe covers it.

A worked example. Ask "am I overpaying for our CRM?" on the sample portfolio and the answer is grounded: the Northwind agreement's $142,800 a year and its delivered benchmark ($571 a seat against a $504 median) are cited inline, and the answer states the $28,650 in reach rather than gesturing at "potential savings". That is the standard to hold her to: numbers with sources, or a plain statement that the data is not on file.

Vera is an AI analyst grounded in your actual portfolio and uploaded documents, and her page is a command desk: a fixed application frame rather than a scrolling page. The desks stand in a dark rail on the left, the work happens on one white paper at a time in the middle, and a dock on the right edge holds what to do next and where you stand. Four desks: **The conversation**, **Read a document**, **Vendor playbook**, and **Specialist agents**. Switching desks never loses work; a report that is generating keeps generating while you talk. The full conversation history with memory and saved threads is **Full chat** in the command bar, not a second entry point.

On the screen

- **The command bar** — the desk switcher, the workspace vitals (records on file, how many renew inside 90 days and what they are worth a year, reports written), and **Full chat**. Every figure is read from your stored records before you type anything.
- **The desk rail** — the four desks, numbered, with the one you are on as a white pill. Under them: **Pick back up**, your recent conversations with the last question asked, **Reports Vera has written**, the newest four

with a link to the full library, and the allowance meter showing the agent reports you have left this period.

- **The start options** — the conversation desk opens with your own choices, not a vendor Vera picked for you. **Work a contract on file**: pick any agreement and the Read a document desk opens with it already chosen. **Drop a contract here**: drag a PDF onto the tile (or press it to choose one) and it arrives on the reader already staged. **Pick a conversation back up**: your recent threads with the last question asked. Each option shows only when your workspace can honour it, so a fresh workspace simply gets the drop tile.
- **Where you stand** — the standing position lives in the dock, behind the crosshair icon. The headline names the renewal worth working and how many days out it is, with the money at the window and the time to decide under it, and one button into the wall those figures came from. Every figure is read from your stored records. A workspace with nothing on file gets the honest version instead: what Vera can answer today, and what one upload adds.
- **The composer** — docked at the foot of the conversation paper, caret already in it, so you can type without clicking anything. A scope chip says what the answer will be read from, **Attach** hands Vera a document to read, and **Tools** (or typing a forward slash in an empty box) opens the three generator desks. Under it, a suggested question or three built from your own estate: they name your vendors, your contract count, and your nearest renewal rather than offering generic prompts.
- **The conversation** — your question is set behind a gold rule, Vera's answer behind her mark beside it, with the reasoning trail above and the verification line below. Asking a follow up adds to the thread rather than replacing the answer you are reading, and the thread scrolls under the docked composer. Follow up suggestions appear above the composer, read from the thread you are actually in. Every exchange saves to your Ask history.
- **The What now dock** — the lightning icon on the right edge, open by default on a wide screen. It reads your workspace and recommends the one move worth making first: with nothing on file it tells you to give Vera something to read, with agreements but no report it tells you to run your first one, and once a renewal comes inside 180 days it names that vendor as the one worth working now. **Start here** carries the first move with a gold edge; **Take it to the vendor** opens the vendor room, a negotiation, or the Deal-Room with those numbers; **When a report is not enough** books the analyst or hands a near renewal to the team. A move that names a desk opens it in place. Escape closes the drawer, and Escape on a generator desk returns to the conversation. Seeded sample agreements never count towards the figures or the advice.
- **Read a document** — its own desk, reached from the rail, **Attach**, or the tools menu. Pick where to start from (a contract you already have, a PDF you upload, or no contract at all), say in a line or two what you are trying to achieve, and press **Analyze and build the report**. You can also drag a PDF anywhere onto the panel. The file is read once for the analysis and then deleted.
- **Watching the read, and where it lands** — the read takes the whole screen while it runs, the same stage a contract decode plays on: your document sketched at the left with a sweep running down it, the percentage complete, and the passes named as they happen, from the pricing tables and the discount structure through liability and termination rights to scoring leverage on every finding. The percentage is an estimate against how long a read usually takes rather than a countdown, and while your file is still uploading the bar shows the real upload figure instead. When the analysis lands it opens as a room: full screen on the exec brief paper, with the document's own page underneath when you close it. The room is

part of the address, so forwarding that link opens a colleague straight into the analysis, and **Back** closes the room onto the page. Closing the stage while it is still working only stops you watching it; the analysis keeps running and is waiting under **Deliverables** when it is done.

- **Vendor playbook** — its own desk: pick a major vendor and Vera writes a ready negotiation playbook with no document needed.
- **Specialist agents** — its own desk, the named specialists as compact rows, grouped into analysis specialists, negotiation deliverables, and your own team's agents, each showing what it delivers and whether it reads your PDF. The whole set fits in about one screen. Open one, choose the contract, and generate. The run keeps going if you switch back to the conversation.
- **Chat** — answers cite whether a figure comes from your contracts or from market experience. Your conversation is saved and private to you, and the assistant keeps a long-term memory of durable facts about you (reviewable and deletable under Assistant memory).
- **Market questions without your figures** — ask which vendors the platform covers and Vera reads the benchmark library itself: the total, the shape by category, and whether a specific vendor is in it. Ask what a good discount is on a vendor and she answers from that vendor's reference set: the typical (median) discount off list and what top-quartile buyers get, or for per-unit vendors the typical and best-quartile net fee. Once your own figures are on file she runs the exact standing on them instead.
- **Grounded figures** — when an answer states a figure taken from your portfolio or from a live data query, that figure is grounded in your stored data. The prose itself stays clean of source tags; the provenance line under the answer records where the numbers came from, the same way generated reports list their sources below the document.
- **Verification line** — an answer that states figures is checked after it finishes streaming, by the same groundedness check that guards reports. A line under the answer then says **Figures verified against your data**, or flags the specific figures to review. No line means the answer had no figures to check.
- **Staged changes** — when you ask Vera to change something (for example move a renewal stage), she stages it as a card with **Confirm** and **Reject** buttons. Nothing runs until you confirm. The same cards appear in the Ctrl/⌘ J dialog, and deciding in either place applies everywhere.
- **What the answer leaves as** — an answer worth keeping ends with the moves that keep it, on the Vera desk and in full chat alike. **Save as a document** writes the question and the answer into your Deliverables as a formatted draft you can share, email, or print, carrying a note that its figures reflect the stored data at the moment it was written, so whoever you forward it to knows to re-ask before acting on a number that may have moved. **File as a task** opens a task titled *Follow up*, with the full answer in the notes and the conversation linked, unassigned so nobody is emailed. **Open the case** appears only when the answer names a vendor whose renewal falls inside 365 days, and it states the days remaining before it opens the case. The row shows on answers that state figures, or that run long; a one-line answer gets none of it.
- **"Can I help you?"**: every Vera surface opens with the same question and the same six intents, phrased the way you would say them: I have a renewal coming up, I want to check if a price is fair, I need to review a contract, I want to cut our software spend, I am preparing for a vendor negotiation, show me what this platform can do. Answer in your own words, in a sentence or a paragraph, and Vera names the one tool that does it, links straight to it, and says what it will do and what it needs from you. If what you want could go two ways she asks the one question that decides it, then recommends. She routes from the platform's

own registry of every screen, so the link always opens a real page, and if you are already on the right screen she tells you which control to press instead of sending you in a circle.

- **Vera on every page**: the button at the bottom right of every screen carries her name. Pressing it (or **Ctrl/⌘ /**) opens her with the same question, and the written guide for that screen sits underneath. **Ctrl/⌘ J** opens the same conversation as a dialog from anywhere.
- **Ask Vera about this**: every screen that produces something to read (a deal room, a generated report, a tool result, a benchmark) carries an Ask Vera box beside it. Vera reads the artifact itself, not just the page name: the figures it states, the section it argues, the asks it lists, and, on a deal room, which planned moves are still unticked. That is what makes "What should I do next?" answerable with your next actual move rather than general advice. The same reading reaches Ctrl/⌘ J, the floating panel, and the voice assistant while you stand on that page, so wherever you ask, Vera is looking at the same screen you are.
- **A conversation per thing**: each artifact keeps its own saved conversation. Ask four questions about your Salesforce deal room, come back next week, open the room, and Vera picks up exactly where you left off, with the room's current numbers in hand. The Ctrl/⌘ J header says which conversation you are in and carries a one click switch back to your main portfolio chat. Everything Vera has learned about you spans all of them; only the transcripts are separate.
- **Pick back up**: the desk rail lists your most recent conversations, and the chat page lists every one you have open with Vera, newest first, with what you last asked and how many exchanges are saved. Opening one takes you back to the report it belongs to.
- **Vera as a teammate** — beyond answering questions, Vera holds a named seat on the team. You can delegate a task to her from Tasks and she does the work; she automatically reviews every deal sign-off request before the chain votes (see Approvals); on a negotiation she preps the room before a vendor call with a pre-call brief; and Competitive Quotes lets her negotiate with up to ten vendors at once. Everything she does is grounded in your stored deal data and appears on the record's timeline under her own name.

How to use it

1. Start from the options on the desk: pick a contract to work, drop a PDF on the tile, or pick a conversation back up. Or ask about any vendor, price, or clause in plain language, or press one of the suggested questions.
2. Not sure where to start? Open **Where you stand** in the dock (the crosshair icon), then take the first move in the **What now** dock, the one with the gold edge.
3. To commission a document instead of an answer, switch desks: press **Attach** if you have the contract, or pick **Read a document**, **Vendor playbook**, or **Specialist agents** in the rail, the desk pills, or the tools menu (a forward slash in an empty composer opens it).
4. Keep asking. A follow up keeps the answer you are reading, and the suggestions above the composer are read from the thread you are in.
5. Finished reports land under **Reports Vera has written** in the desk rail; press **All reports** for the full library.
6. Open Vera anywhere in the app with **Ctrl/⌘ J**, so you never have to leave the page you are on. On an ordinary page it is the same conversation as this one; standing on a report or a deal room it opens that artifact's own conversation instead, grounded in what is on screen, with the same verification and confirm cards.

7. On any report, use the Ask Vera box under it and start with **What should I do next?**; the answer names the specific move, ask, or clause from that document.
8. Prefer talking? Spoken answers name their sources in words: "your contract shows" for figures read from your documents, "from market data" for market figures.

Tips

- Vera answers only from your organization's own contracts and our market data, and every figure is tagged with its source so you can tell a stored number from a market estimate.
- The floating **Ask Vera** button (bottom right on every screen, Ctrl/⌘ /) is the same Vera: it opens on her question with the screen's written guide underneath, and a question asked there continues here and in Ctrl/⌘ J, grounded in the screen you were on when you asked. The bottom right corner itself is shared: when the Talk to Vera panel, an upload progress chip, or a walkthrough is active, the widgets stack side by side instead of covering one another, with the help button stepping into its quiet round form.
- Vera's long-term memory is fully reviewable and deletable under Settings › Assistant memory.

Related: Assistant memory · Deliverables · Workflow agents

WORKFLOW AGENTS

Where: Team & AI › Workflow agents · `/workflows` · via **More** or Ctrl/⌘ K
Who: everyone

An agent is a job you describe once and never do again: an AI analysis chained into an email or a task, running on a schedule, on every proposal upload, or when a renewal enters its notice window. The page is a workforce desk with four views on the masthead band.

In plain words: the routine you keep doing by hand, chained once and run on a schedule.

When to reach for this

- Anything you caught yourself doing weekly: the "Do this every week" control on record pages builds the entry for you.
- Routing platform output to where the team lives: an analysis into an email, a Slack post, a task.
- Supervising what you built: runs, failures, and pauses show in the Control room and the Agent Center.

On the screen

- **The masthead** — how many agents are on staff of the number you have built, runs in the last 30 days, how many runs are paused waiting on someone, and last of all **Estimated hours back**, labeled as the estimate it is: 20 minutes of manual work a completed run, not a measured figure. Where real money is linked to the work, that money takes its place. The Hire an agent and Build one from scratch actions live here.
- **The roster is the front door** — with nothing on staff yet, the page opens straight onto the role gallery: the fifteen roles of the negotiation desk on big cards first (each stating what it watches, what it writes, and

how many documents a run produces), the rest of the catalogue beneath. A search box and a **"Everything they can do…"** dropdown filter the cards by capability (research a vendor, read a contract like an expert, defend an audit, watch the market weekly, and so on). Picking a card opens the hire wizard already pointed at that role.

- **Four views, one desk** — the band switches views: **The staff** (every agent with its live heartbeat, its step chain drawn as a strip, and its Run now, Pause, edit, and delete controls), **Runs** (the console, with a badge on the tab the moment a run pauses for a decision), **Builder** (a standing bench: pick a trigger, then stack steps in order; details from whatever set the run off, the vendor, the agreement, the deadline, fill in by themselves), and **Wiring** (the agent inbox, the ghost writer, and a plain reference of every trigger). Press 1 to 4 to switch, H to hire an agent.
- **The negotiation desk** — the fifteen prebuilt roles cover the whole arc of a deal: prepare a negotiation from a vendor name or from your uploaded contract, a call script aimed at a goal you state in your own words, the full deal check (the benchmark verdict on how good the deal is, the contract decoded clause by clause, then an independent second opinion that reads both, ending on the Deal Room question), the market monitor, the renewal war room, a quote teardown, a pre-signature red flag scan, audit defense, price increase pushback, the walk-away builder, rightsizing before renewal, the board and CFO approval pack, the first-call rehearsal, and the deal close handoff. Each role writes several full documents per run, every one its own model call, and a chain that ends in an email mails the whole pack.
- **The rail** — beside every view: the runs-per-week pulse over the last eight weeks, any run paused for a sign off with a link straight to it, and the payoff instrument: drag a runs-per-week pace and it states the hours given back a quarter, at 20 minutes of manual work saved a run.
- **Editing an agent** — the pencil on any staff card opens its draft on the Builder view; saving returns you to the staff.
- **Runs** — every run keeps a per-step record; open one to read the documents it produced.
- **The run as a live desk** — launching an agent opens the full-screen desk immediately: a gold progress ring fills step by step around a page writing itself, the narration line streams what the agent is doing right now, and every document in the left rail moves from queued to writing to filed, readable the moment it lands. Back to the app sits top left, your other recent runs and a Hire another agent button ride the left rail, and the Put it to work rail carries the real next moves: email me the pack, rerun with today's data, a chat-ready summary, three suggested next hires, and the schedule card. The agent is always saved automatically; the schedule card asks when it should run by itself (weekly, on every upload, on renewal windows, or only on demand), and the choice stays yours. Arrow keys page the filed documents, Esc closes. Behind the desk, the run page keeps the briefing-case cover, the document miniatures with Word exports, and the step by step run log.
- **Turning the emails off** — every email an agent sends closes with a link to Settings › Notifications, so turning them down or off is always one press away.

Hire an agent, the wizard

Four steps, and every one arrives already answered so you are correcting rather than composing.

1. **What should happen without you?** The full role gallery: the negotiation desk first, then every other role, with search and the capability dropdown. If you hold a renewal inside its notice window the renewal roles rise, and the page says why.
2. **Point it at your vendors.** Pick one of your own vendors and the role fills in around it; the line under the picker rewrites itself with your real numbers as you choose. Vendor lists show one card per company, legal entities like Oracle America or SFDC Ireland are folded into one vendor with the yearly spend summed. The market monitor asks differently: tick up to 30 vendors to watch, from your own records and from the market's major publishers, and its weekly brief reports dated vendor news in plain English, with no advice attached. Roles that work toward an outcome (the contract prep and the call script) also ask **what you want to achieve**, in your own words, because an agent that knows the goal writes a sharper document.
3. **Here is the chain we built** — the trigger, then every step in order with the document it writes. Nothing here has to be right the first time; the builder can change all of it later.
4. **Watch it run** — the rehearsal. It runs the whole chain on your own records, and every email goes only to you, nothing is created for anyone else, and nobody else is notified. You read what it wrote before the agent is ever live.

How to use it

1. Press **Hire an agent** (or pick a card on the landing roster), pick the vendor or tick the watch list, state the goal where the role asks, and run the rehearsal.
2. Read what came back. If it is right, the agent is already saved and you turn it on; if it is not, open it in the builder and change the step that was wrong.
3. Set a workflow to fire on every proposal upload when you want it to run itself; the market monitor mails its brief weekly on its own.

Related: Vera AI · Proposals & upload

VENDOR AGENTS & AGENT ROOMS

Where: Team & AI › Workflow agents › **Run a vendor agent** · `/workflows/agents`
Who: everyone · requires an AI plan

One agent, one vendor, one room. Pick a vendor and an agent role, press **Run the agent**, and the agent writes its documents one at a time, every figure drawn from your stored records or labeled as the desk's market read, and files all of them into a single Agent Room you keep, reread, and rerun whenever your records change.

On the screen

- **The vendor** — two groups in one list. **Your vendors** comes from your stored agreements, largest first, with sample paper excluded, and names the specialist desk where one covers the vendor. **The specialist desks** lists all 30 desks, every one pickable even before you hold an agreement with that vendor; a desk holds the vendor's fiscal calendar, contract watch points, sales motion, and negotiation plays. Any other vendor can be typed in by name; the room then runs on your stored records alone, and the page says so plainly before anything starts.

- **The agent** — six roles plus **Build your own**, each card stating how many documents it writes; the selected card fills navy with a gold check. **Vendor Watchdog** reads how the vendor's business is moving and what their field is selling. **AI Research** is the homework in one brief. **Negotiation Prep** writes the vendor card, the talking points, and the strategy. **AI Negotiation Assistant** is the full desk in six documents. **Vendor Deep Research** is the business behind the logo, including the plays that fit. **Renewal War Room** works the renewal end to end.
- **Build your own** — the seventh card. Start from a situation (prepare for a vendor call, keep watch, work a renewal, do the homework, check before signing, brief the leadership, or the full desk) and the document chain pre-fills; then tick exactly the documents you want. The position page always opens the room, built from your stored records with no AI writing.
- **What is this negotiation, and what do you want?** — one box, your own words, required: "renewal due in March; the quote landed with an 8% uplift; we want the cap at 3% and a shorter term." Every document in the run is written toward what you wrote, in the voice of a specialist who has worked this vendor's agreements for 25 years, and the position page opens with your assignment quoted back. Anything the stored data cannot answer gets named as missing, never invented.
- **The run** — pressing **Run the agent** takes the whole screen instantly: a full-viewport opening stage where the position page writes itself while the room is prepared. The reader then opens on the position, your stored data on page one, the moment it is written; you read while the rest of the case file writes in the background. Pending documents sit in the rail with their live states (writing, queued), a quiet progress band shows what the agent is doing with an honest count that never claims a page before it lands, and each document becomes readable the moment it files. A document takes a few minutes; leaving the page pauses the run where it stands, and reopening the room resumes it from the last filed page. Nothing is lost to a closed laptop.

- **The room** — a finished run reads as a case file: the documents on a numbered rail, one page at a time on the report paper, with previous and next at the foot. The first page is always your position, built only from your stored agreements.
- **The byline on every page** — each document says where it stands: figures checked against your stored data, or exactly how many figures could not be verified and should be read as directional. Market context from the desk is always named as the desk's read, never presented as your data.
- **Rerun with today's data** — starts a fresh room for the same vendor and role against your current records. The old room stays on your list, so runs become a record of how your position moved.
- **Put it in front of people** — the rail's hand-off block. **Copy the direct link** copies the room's address; it opens for people in your organization only. **Send to this teammate** picks a colleague and puts a review task on their desk with a note in their inbox, so the room lands as work rather than a link. And on a

complete room, **Email me the case file (PDF)** sends the whole room to your own inbox as one document: a cover page, then every document in reading order, each page carrying its own provenance line.

- **Team discussion** — a thread under the reader, the same discussion the rest of the platform uses. Type @ and a name to notify a teammate; replies, resolving, and edits work exactly as they do on a renewal or an approval.
- **Your agent rooms** — every room you have run, newest first, with its status. A room that hit an error keeps the documents it finished and offers to resume.

How to use it

1. Open Workflow agents and press **Run a vendor agent**, or go to `/workflows/agents` directly.
2. Pick the vendor. If it carries a specialist desk the page names it; if not, the run still works from your stored records.
3. Pick the role that matches the job in front of you, or **Build your own** and tick the documents. Then say what the negotiation is and what you want; the run does not start without it.
4. Press **Run the agent**. Each document counts toward your AI allowance, and the page says so before you start.
5. Watch the stage as the room fills, or come back later; the run resumes when you reopen it.
6. When your records change, open the room and press **Rerun with today's data**.

Good to know

- The position page is deterministic: it is assembled from your stored agreements without any AI writing, so it is always exact.
- A vendor with no specialist desk still gets a room, built from your records alone, and every page says that is what it is.
- Rooms belong to your organization. A link to a room only opens for people in your org.

Related: Workflow agents · Vera AI · The Agent Center

DELIVERABLES

Where: Team & AI › Deliverables · `/studio` · via **More** or Ctrl/⌘ K (Library, Generate, Drafts, and Lessons tabs)
Who: everyone · requires an AI plan

Every delivered benchmark, AI brief, tool run, and draft in one place, laid out as an inbox: folders on the left, the list in the middle, the document on the right. Reading one never leaves the page.

Your position sits in the masthead: reports delivered with the latest date, how many are in production, and how many vendors they cover, on the same band as the Library, Generate, Drafts and Lessons tabs.

On the screen

- **Folders** — **Inbox** holds everything that is not archived and carries the unread count. **Flagged** holds the ones you marked to come back to. **Delivered to you** is the analyst team's work, **Drafts** your unfinished ones, and **Archive** everything you have filed away. **In production** underneath is not a folder but a count of the benchmarks still being written; they arrive in the inbox when they are published.
- **The list** — newest first, one row per document: the vendor or counterparty, the document title, and a line saying what kind of document it is. An unread document carries a blue dot and sits in bold; once you open it the row goes quiet. **Search this folder** filters what you are looking at.
- **The reading pane** — the document's type, title, vendor, and the day it landed, then **Read it** to open the full document, **Word** to download it where the type supports one, and **Open the record** to go to the contract or deal it belongs to.
- **Triage** — **Flag** marks a document to come back to, **Unread** puts the dot back so it stays on your list, and **Archive** takes it out of the inbox without deleting anything. Archived documents stay readable in the Archive folder and can be restored.
- **Reads across the whole library** — under the inbox, the AI narratives that read across every document rather than one: the board report and the weakest metrics card.
- **Commission a deliverable** — the full catalog of document types, folded away once your library has something in it and open by default when it does not.

How to use it

1. Work the bold rows first. Those are the documents nobody on your side has opened.
2. Select one to read it in the pane beside the list, then press **Read it** for the full document.
3. Flag what you need to come back to, and archive the rest so the inbox stays the things that still need you.
4. Create a new one from Vera AI, or open **Commission a deliverable** at the foot of the page.

Tips

- Read, flagged, and archived state is held in your own browser, so your unread count is yours and not your colleagues'. Clearing site data resets it, and nothing is ever deleted by archiving.

Related: Vera AI · Deliverable exports, sharing & history

DELIVERABLE EXPORTS, SHARING & HISTORY

Where: the toolbar above any AI deliverable, on `/studio` and on the record pages that generate briefs, memos, and reports
Who: everyone · requires an AI plan

Every AI deliverable carries the same three companions wherever it renders: a download that matches the on-screen document, a public share link guarded by a confidentiality scan, and a version history you can diff and revert. As of August 2026 the flagship report views (benchmark narratives, the portfolio board report, the all-vendors benchmark, Reporting studio runs, and the scenario and comparison memos) assemble these through one shared toolbar, so exports, share, present, print, the Send it on drafter, and history sit in the same place on every report.

In plain words: every document can leave the building safely: exported, shared, and versioned.

When to reach for this

- Forwarding a brief outside the workspace: the share link's confidentiality scan is the safety net.
- "What changed since last month's edition?": the version history diffs it.

On the screen

- **Download** — exports the deliverable as Word (and, for the flagship reports, PowerPoint), rendered in the same executive-brief theme as the on-screen document. Most document types support it; the Formats column in Delivered reports › All documents shows which.
- **Share** — mints a public link and copies it to your clipboard, so a colleague without an account can read the document. Before the link exists, the deliverable is scanned for content your organization would not want external: named employees, walk-away numbers, internal targets, and internal notes.
- **Check before you share** — when the scan flags something, the findings render with quoted excerpts and no link is minted. Sharing anyway is an explicit, audited choice; overriding a walk-away number or internal target takes a typed confirmation, not one click.
- **History** — every regeneration is kept as a version. Open the history to see a word-level diff between any version and the current one, and revert to an earlier version when a regeneration went the wrong way.

How to use it

1. Download the Word export when the document leaves the platform; it matches the on-screen brief exactly.
2. Press **Share** to send a read-only link instead of an attachment. If the confidentiality check raises findings, read the excerpts before deciding.
3. Regenerated a report and prefer the old one? Open **History**, compare the diff, and revert.

Tips

- A share link shows the document as it stands; revoking and reissuing is the way to cut off an old copy.

Related: Deliverables · Delivered reports

THE COMMISSION RAIL

Where: under the verdict on record and insight pages, the gold sparkle chips labeled "Commission the desk"
Who: everyone · requires an AI plan
You get: a house executive brief rendered in the white-paper format, saved to Deliverables · about a minute

In plain words: one-click orders for the brief this page is begging for, labeled with its own numbers.

When to reach for this

- The page's verdict deserves to become a document someone else reads: the chip is the shortest path.
- You are not sure what to commission: the chips are already scoped to the page's own situation.

Two to four gold sparkle chips under a page's verdict, each a one-click order for an executive brief about the page you are on. The chips are labeled with the page's own numbers ("Renewal brief: Salesforce, 41 days out"), never generic, and the rail always ends with a blue question chip that opens Vera pre-filled with a question scoped to the page.

On the screen

- **Gold chips** — each one commissions a specific document: press it and the brief generates in a reading dialog with stage-by-stage narration.
- **The finished brief** — an executive summary, numbered sections, a bottom line, and a sources list; every figure is grounded in your stored data, traced in the sources list at the foot of the brief.
- **The blue question chip** — closes the rail: it opens Vera with a page-scoped question already typed, for when you want an answer rather than a document.

How to use it

1. Press the chip whose label matches what you need; the document writes itself in the dialog.
2. Read it there, or find it again later on the Deliverables shelf, where it lands automatically.
3. Want a conversation instead of a brief? Press the blue chip and ask Vera.

Related: Deliverables · Vera AI

ADVISORY

Where: Team & AI › Advisory · `/advisory` · via **More** or Ctrl/⌘ K
Who: everyone

Book time with a buyer-side analyst: a human expert who has negotiated with your vendor before, for the moments a report is not enough.

In plain words: a human who has sat across from your vendor, bookable.

When to reach for this

- The stakes outgrow self-serve: a large audit, a hostile renewal, a deal above your comfort line.
- Before signing something unusual: an hour of expert read on an odd structure is cheap.
- Vetted deal specialists with named track records live on the Expert Bench; this desk is our own analysts.

The desk is an application frame, the same command-desk grammar as Vera AI: a command bar with the live allowance, a dark rail with the desk facts and your bookings, the work on white papers in the middle, and a dock on the right edge.

On the screen

- **The command bar** — sessions left (monthly plus purchased), booked this month against your allowance, and bookings awaiting confirmation. Every figure is your stored usage.
- **The desk rail** — the desk facts, your bookings as one-line jumps, and the privacy note. The desk runs Monday to Friday, 4am to 11pm Eastern, shown in your local timezone, and the earliest booking is 72 hours out.
- **The booking paper** — the headline states how many sessions are still open, then the desk's availability for the next two weeks, the buy-extra buttons for when the month runs short, and the form: your agenda, then either three proposed times or, where instant booking is enabled, a direct calendar slot. The right side of the card explains the three steps of a session.
- **Your bookings** — every request as a ticket on its own paper: confirmed sessions carry the date stub and an Add to calendar link, pending ones list your proposed times, and each can be cancelled.
- **Start from a topic** — in the dock, behind the sparkle icon: four briefs. The first two are composed from your own two nearest notice-window renewals and name the vendor, the renewal date, and the annual value, so the agenda arrives written about a real deal; on an empty workspace they fall back to the standard renewal strategy and pricing benchmark briefs. The last two (vendor consolidation, contract terms) are always the standard ones. One press opens the desk with the agenda pre-filled and leaves you one field from a request.
- **The exact insider** — a dock group that appears when an Expert Bench member works the vendor of your nearest renewal: it links their profile, and your renewal is already on the request when you open it.
- **Present a session** — every booking carries a Present button that runs the meeting full screen during the call: the ticket of record, your brief in your own words, the agenda point by point with a covered counter, then agreed actions you capture live. The printed leave-behind carries exactly the screens you walked.

How to use it

1. Write your agenda, or open the dock's **Start from a topic** panel for a pre-filled brief.
2. Propose three times that work for you; we confirm one and send a calendar invite. Where instant booking is on, pick a slot directly instead.
3. Track the request in Your bookings; add a confirmed session to your calendar from its ticket.
4. If it can wait, message the team instead and they will pick it up in Chat.

Related: Chat · Expert Bench · Your engagements

YOUR ENGAGEMENTS

Where: direct link, shared when an engagement starts for your organization · `/delivery`
Who: everyone (the page lists only engagements run for your organization)

Every advisory engagement we are running for your organization: where it stands, what happens next, and what we need from you. When your organization commissions consulting work (a Microsoft EA renewal run, a Salesforce SELA review, an audit defense), the engagement lives here as a workspace you share with our team.

On the screen

- **Engagement cards** — one per engagement: its status (Setup, Active, On hold, Closing, or Closed), the phases completed out of the total with a progress bar, and the next milestone by name.
- **The engagement workspace** — open a card for the full workspace, organized in tabs: **Overview** (the phase plan and its activities), **Timeline**, **Team** (who is on the engagement, both sides), **Discussion** (a comment thread on the engagement record), **Documents**, and **Deliverables** (published work products). EA and SELA renewal engagements open on a dedicated analysis workspace tab.
- **Join card** — if you have a seat on the engagement team but have not joined yet, a card at the top asks you to confirm your name and role.
- **Data requests** — activities our team marks as data requests appear in the plan; providing the data checks them off and keeps the phases moving.

How to use it

1. Open the engagement and read the progress bar and next milestone; that is the one-line status.
2. Answer anything waiting on you: join the team from the join card, provide open data requests, and reply in Discussion.
3. Read published deliverables in the Deliverables tab as the team ships them.
4. Anything urgent between milestones goes through Chat; the thread and the workspace reach the same team.

Tips

- Internal working activities on our side are not shown to you; the plan you see is the client-facing plan, and every view you take of the record is written to the audit trail like any other document access.

Related: Advisory · Chat · Managed renewals

PARTNER PORTAL

Where: Team & AI › Partner portal · `/partner` · via **More** or Ctrl/⌘ K
Who: members of a partner-program firm (the rail entry appears only when VendorBenchmark has enabled partner access for your organization)

Your client book. Every client organization linked to your firm on one table: their contract dates, their next renewal with its notice deadline, and how engaged the account is. Visibility is deliberately narrow, dates and activity only; each client's files, prices, benchmark results, and spend stay inside their own workspace.

The portal is an application frame, the same command-desk grammar as Vera AI: three desks as pills on one command bar (**Client book**, **Prospect proposals**, **Marketing**), the book's vitals beside them, a dark rail that jumps to any section of the desk you are on, the work on a white paper in the middle, and a dock on the right edge carrying **What now** (invite a client, build a proposal, import your list, send a note, the partner guide) and **The boundary** (the trust rule, one press away on every desk).

A client's page, and the quarterly business review it composes once that client has turned on outcome sharing:

The public directory buyers browse, and the branded landing page an invited client sees:

On the screen

- **The command bar** — clients, renewals inside 90 days across the whole book, items needing attention today, and invitations accepted out of invitations sent.
- **The desk rail** — every section of the book one press away (Your desk today, Clients, Managed renewal requests, the horizon, Grow your book, White label, Program, Meetings), your invitation funnel as four figures, and the data-boundary note.
- **Your desk today** — a ranked attention queue: service requests, notice deadlines that passed or land inside 30 days, clients with no activity in 30 days, workspaces that never signed in, meeting proposals a week without an answer, and invitations about to expire. Each row links to the client or invitation it is about. When nothing qualifies, the queue says so.
- **Draft it with Vera** — a chip on every queue item except the expiring invitation. It writes the note from that item's own client, vendor, and dates: under 120 words, in your voice, addressed to a name you fill in, and never naming VendorBenchmark or Vera. Edit the subject and the body in place, press **Redraft** for another take, then **Copy email**. **There is deliberately no send path.** You send it from your own mail client, under your own name, so the client relationship stays yours and the platform never appears in it.
- **Client book** — one row per linked client: adoption stage, the account owner inside your firm, contracts, next renewal (vendor and date), renewals inside 90 days, active members this week, and last activity. The adoption stage reads Active, Getting started, Quiet, Dormant, or Never signed in, derived from real sign-in activity. Rows with a renewal inside 90 days are tinted. Assign the account owner from each client's page; the morning digest names them on items about their accounts.
- **Outcomes on client pages** — when a client's owner turns on outcome sharing (their switch lives on their Your advisory partner page, off by default and revocable), that client's page shows the results of the work: realized savings with their invoice-verified share, delivered benchmark verdict scores, and closed renewal outcomes. Files, prices, and the analyst narrative are never included, and every look you take is audit logged. With sharing on, **Build the quarterly business review** composes a client-ready executive brief under your brand, savings, verdicts, renewals closed and coming, adoption, and prioritized recommendations, ready to print to PDF.

- **Managed renewal requests** — when a client presses Get it managed on one of their renewals, the request lands here and at the top of Your desk today: client, vendor, renewal date, and their note, with Mark contacted and Close. The platform only tracks the state; scope, terms, and fees are agreed between your firms directly.
- **Renewal horizon** — the next twelve months of tracked renewals across your whole book, counted month by month.
- **Grow your book** — the invitation funnel (sent, awaiting reply, accepted, active now, with your conversion rate), then the invite form: client company, contact, email, and an optional personal note. The email and its landing page go out under your firm's brand. Open invitations show their expiry date with resend and revoke; accepted ones link to the client workspace they became and show how alive it is. **Prospect proposals** (`/partner/proposals`), the second tab (pictured below): enter a prospect's vendor, unit count, spend, and term, and the calibrated market engine builds a one-page savings estimate under your brand, verdict, percentile, cohort range, and gap to median, shareable by an expiring link you can withdraw at any time. Your tier caps proposals per 30 days.
- **White label** — your brand on the client experience: display name, brand color, and tagline, with a live preview. The logo comes from Settings › Branding. Only workspace owners can edit it. On the white-label tier (a Premier benefit, granted by VendorBenchmark), the section adds **Your own sending domain**: register a domain you control, publish the DNS records shown, press Check verification, and from then on invitations and campaigns send from advisor@your domain under your name alone, and your logo and name take over the rail lockup in the client workspaces you serve.
- **Partner guide** — the Program section header links to `/partner/guide`, which is a document rather than a room and so stays off the tab row: the whole offering in one printable executive brief, from the tier ladder and the agreement to outcome sharing, white label, and the data boundary. Read it once, or hand it to a new hire on their first day.
- **Program** — where the partnership itself lives. The tier card shows your rung (Registered, Select, or Premier), the caps it carries, and what the next rung unlocks. The agreement card carries the rate on record: offers from the VendorBenchmark team appear here with their terms, your workspace owner accepts in product, and the accepted version and rate stay visible from then on. The directory card drafts your public listing (one paragraph, focus areas, regions); firms on Select and above can switch the listing live in the public directory at `/partners`, where buyers request introductions that land straight in your contacts. The morning digest switch is here too: weekday mornings, firm owners get the attention queue by email, and a clear desk sends nothing.
- **Your clients' home screen** — curate what client workspaces open on: pick the apps that lead and add a one line headline under your brand, with a live preview of what clients see. The platform standard stays the default until you choose otherwise, the full catalog always stays one press away, and each client's own edits win. Owner only.
- **Messages** — every client page carries a direct message thread between your firm and that client. The client reads and replies on their Your advisory partner page, their owners get notified, and the thread is private to your two firms; partners cannot message anyone who is not their client. The client's side of that page also carries their outcome sharing switch, so what you can see is always their call.

- **The boundary** — in the dock, behind the shield icon on every partner desk: what you see (dates, renewal stages, org level activity, meetings) and what always stays the client's (files, prices, benchmark results, member identities). Show it to a hesitant client; it is the same list we show them.
- **Meetings** — every meeting between your firm and your clients, with its status.

How to use it

1. Start with **Your desk today**; it is the call list, ranked. A passed notice deadline outranks everything else.
2. Open a client row for the full picture: renewal dates with notice deadlines and stages, contract term dates, the 30 day activity chart, and your engagement record (linked since, meetings held).
3. Watch the **Adoption** column: a Quiet or Dormant client two months before a renewal is the account to call first, and a Never signed in workspace needs a walkthrough, not a renewal pitch.
4. Invite a new client from **Grow your book**. When they accept, their workspace opens on a free trial, already linked to your book, and it carries a "served by" mark for your firm. Invitations are single use, live 30 days, and can be revoked or re-sent with a fresh link; the funnel shows what each one became.
5. Set your brand once under **White label**; every invitation and every client workspace mark uses it from then on.
6. Propose a meeting from the client page: up to three times and an agenda. The client's workspace owners get a notification, confirm one time or decline, and both sides see the answer.
7. Need the full record for an engagement? Ask the client to invite you into their workspace as a member; the portal never opens their files.

Clients join your book two ways: your own invitations, or a link made by the VendorBenchmark team after a client agrees in the engagement paperwork. If a client is missing, contact your partner manager.

Related: Partner marketing · Notifications · Members

PARTNER MARKETING

Where: Team & AI › Partner marketing · `/partner/marketing` · via **More** or Ctrl/⌘ K
Who: members of a partner-program firm

Your firm's own marketing desk: a contact list you build yourself, newsletter sends that go out under your white-label brand, and a read log showing who opened what. Nothing here comes from client workspaces; every contact was added by your own team.

On the screen

- **The command bar** — your audience (subscribed and reachable, with how many have unsubscribed), campaigns sent, emails delivered, and the share read.
- **Three steps to your first send** — while your list is empty, the page states the whole path: paste your book, write the note and test it, send and read the ledger. It disappears once you have contacts.
- **Import your list** — one paste, up to two thousand rows. A bare email works, so does the `Name <email>` form every mail client copies, and so does a comma or tab separated row straight out of a spreadsheet.

Anyone already on file is skipped rather than duplicated, so pasting the same list twice adds only the new names.

- **Contacts** — add one at a time by email, name, company, and role. Unsubscribe or remove any contact; unsubscribed contacts never receive a send.
- **Write a note** — subject and body (blank lines start new paragraphs), a live preview of the branded email, a test send to your own inbox, and the real send to every subscribed contact. Emails go from our verified address under your brand's name, with your email as the reply-to.
- **Sent campaigns** — each send with delivered, opened, and clicked counts; expand one for the recipient-by-recipient log: clicked (and how many times), read (and when), delivered but unread, or not delivered.
- **The desk rail** — jumps to each section of the desk, what your tier allows (contacts against the cap, sending name, reply to), and how reads are counted. The boundary the portal runs on stands in the dock behind the shield: these contacts are yours, nothing is harvested from a client workspace, and no client's prices, documents, or benchmark narrative can be pulled into a marketing email.

How to use it

1. Set your brand once on the Partner portal White label card; every send uses it.
2. Paste your whole contact list into **Import your list**, or add people one at a time below it. The list is yours alone; clients' member lists are never visible to you.
3. Write the send, or start from one of the three templates (a renewal season note, a quarterly market note, an introduction) and replace the bracketed slots. Press **Send me a test**, check your inbox, then send for real. The caps rise with your program tier, from 5 campaigns per day and 500 recipients per send on Registered up to 20 and 5,000 on Premier.
4. Watch the read log. Opens are measured by an email pixel, so a recipient whose mail client blocks images can read without registering; treat opens as a floor. Links in the body are click tracked per recipient, and a click also counts the open, so the log distinguishes read from acted even when the pixel is blocked.
5. Every email carries one-click unsubscribe. We honor it automatically and mark the contact unsubscribed for you.

Related: Partner portal · Notifications

SAVED

Where: Team & AI › Saved · /saved · via **More** or Ctrl/⌘ K
Who: everyone

A comparison spine shows a stale badge and a **Compare again** link in place of its saved date when either compared document has picked up a newer version.

Reports and research you bookmarked with the bookmark icon, plus every side-by-side comparison you ran, kept on shelves so you can find them again.

In plain words: everything you bookmarked, in one place.

When to reach for this

- Rebuilding context on a deal: the papers and reports you starred while working it.
- The bookmark is everywhere (cards, readers, report headers); this page is where they all land, and your home screen's Saved & shared panel shows the newest.

On the screen

- **Saved items** — everything you bookmarked, newest first.
- **Comparisons** — every side-by-side scan you ran on Compare, saved to this shelf automatically, plus any a teammate added you to (marked **Shared with you**). A spine reopens the pair on the compare desk.

How to use it

1. Open a saved report or paper straight from the list.
2. Add to the list by clicking the bookmark icon on any report or research paper.
3. Reopen a comparison from its spine. The **X** on a spine clears it from your shelf only; the analysis itself stays on both contract records.

Related: Delivered reports · Research · Compare

SAVED ANALYSES

Where: Team & AI › Deliverables › the **Drafts** tab · `/studio/drafts` (old `/drafts` links forward here)
Who: everyone (writing a new analysis happens in Vera AI, which requires an AI plan)

Negotiation playbooks the AI has written for you, kept in one list as the Drafts tab of Deliverables. When you ask Vera to analyze a document, the playbook she writes is saved here automatically, so you can revisit it, print it, or share a link with colleagues.

On the screen

- **The list** — every analysis with its playbook title, vendor, and creation date; search and sort across all of them.
- **Sharing column** — **Link active** means a share link exists for that draft (with its expiry date); **Private** means only your organization can see it.
- **The draft page** — opens the full playbook with a toolbar: **Share link** mints a public read-only URL and copies it, **New link** replaces it (the previously sent URL stops working), **Email** drafts a message carrying the link, **Revoke** kills the link, **Print** produces a clean paper copy, and **Delete** removes the draft for good.

How to use it

1. Create one from Vera AI: choose "Ask Vera to analyze a document for you", and the result lands here.
2. Open a draft to reread it before a call, or press **Print** for the meeting copy.
3. Share it by minting a link; revoke the link when the conversation is over. A minted link expires on its own, and **New link** invalidates whatever you sent before.

Tips

- Your saved analyses also appear in the All documents library on Delivered reports, so there is never a second place to hunt.

Related: Vera AI · Deliverables

PLAYBOOK LIBRARY

Where: Team & AI › Playbook library · `/playbooks` · via **More** or Ctrl/⌘ K
Who: everyone (a run against your own contract requires an Enterprise plan)
You get: a tailored edition of any playbook, rewritten around your contract's numbers · a run takes a few minutes

Every playbook the AI has written for you, shelved as one library. The move that makes it live: drop a contract onto any volume and the playbook reruns around your actual numbers, every figure and recommendation recomputed from the paper you dropped.

On the screen

- **The shelf** — one bound volume per playbook: the agent that wrote it, the title, the vendor, and the date. **Open** reads it; **Run your contract** picks a PDF or DOCX to rerun it against.
- **Drop targets** — drag a contract over the page and every volume lights up; drop it on one and a confirm bar stages the run.
- **The run room** — while the run tailors: the two stages (reading the commercial terms off your contract, then the rewrite), with your extracted numbers appearing as chips the moment they land.
- **Your numbers band** — the facts read off your contract: annual value, licensed quantity, term, renewal date, and uplift cap. Every figure in the tailored document stands on these.
- **The tailored edition** — the rewritten playbook, rendered as the same executive document, with **Share** and **Open in the library**. It joins the shelf next to its source.

How to use it

1. Ask Vera for a playbook (or open one an agent already wrote); it lands on this shelf as a volume.
2. Drag your contract over the shelf and drop it on the volume you want rerun, then press **Run it**.
3. Watch your numbers land, then read the tailored edition; it is saved to the shelf and ready to share.

Tips

- The dropped contract is read once for its commercial terms and never stored; only the tailored playbook is saved.
- A run is grounded in what the document states. If the contract does not print a value or a quantity, the rewrite says so instead of inventing one.

Related: Vera AI · Saved analyses · Vera Deal-Room

REFERRALS

Where: Team & AI › Referrals · `/referrals` · via **More** or Ctrl/⌘ K
Who: everyone

Your personal referral code. When a company you refer becomes a paying customer, you earn ten percent of their spend for their first twelve months. We email you when a referral signs up and again when they subscribe.

On the screen

- **Your code and link** — share these; signups through them are credited to you automatically.
- **How it works** — the four steps up front: share, they sign up, they subscribe, you get paid, with the email you receive at each step.
- **Referral status** — who signed up, who is now paying, and what you have earned.
- **Payout terms** — payouts clear thirty days after we receive the customer's payment, and we contact you for your payout details once a referral starts earning.

How to use it

1. Copy your referral link and share it with a company that would benefit.
2. Watch the referral status list to see who signed up and what you have earned.
3. Watch your inbox: when you first get your code we email you a short "how it works" explainer, then we email you again the moment a referral signs up and when they become a paying subscriber, with what the referral is worth and how the payout works.

Related: Refer & earn

REFER & EARN (THE INVITE DIALOG)

Where: top bar › the gold **Refer & earn** pill (also in the account menu's help actions)
Who: everyone

The fast way to act on your referral code: a dialog that emails up to five colleagues at other companies a personal invitation, each carrying your code so anyone who joins from it is tracked as your referral automatically. This is deliberately not the team invite: no access to your workspace is granted, and your friends land on the public request-access page.

On the screen

- **Your referral code** — shown on a ticket panel with a **Copy invite link** button, and the terms in one line: when a company you invited becomes a customer, you earn 10% of what they spend for their first 12 months.
- **Friends' emails** — up to five addresses per send.

- **Personal message** — the note carried in the invitation email, up to 1,000 characters. It opens pre-filled with a ready-made note, and the dropdown beside the label offers five alternatives: quick and casual, before your next renewal, how it works, check a quote, and short and direct. Whatever you pick stays fully editable, and you can clear it to write your own.

How to use it

1. Press the gold **Refer & earn** pill in the top bar.
2. Enter up to five email addresses. Keep the pre-filled note, pick one of the five templates from the dropdown, or write your own.
3. Press **Send invites**. Track who signs up and what you have earned on Referrals.

Tips

- Inviting a teammate into your own workspace is a different action: that is Members, owner only, and it grants access. This dialog never does.

Related: Referrals · Members

CHAT

Where: top bar (inbox bell › Team chat) · `/chat`
Who: everyone

Your team's real-time chat: channels for the whole organization, a private support line to our analyst team, and the AI assistant on call in any channel. Chat lives inside the inbox bell: unread messages count into the bell's badge, and the pinned **Team chat** row at the top of the inbox panel opens this screen.

On the screen

- **Channels** — team-wide rooms; #general is always there, and anyone can add more.
- **Vera AI team** — the private support thread; only you and our team see it.
- **@assistant** — mention the assistant in any channel to get an AI answer grounded in your portfolio, visible to the whole team.
- **# links a record** — type **#** in a channel and pick a proposal, renewal, deal, or vendor; the reference renders as a grounded card with the record's live status, the record's timeline notes where it was discussed, and its watchers hear about it. DMs stay out of this; they are private.
- **Deal channels** — a channel opened from a negotiation's **Deal channel** button is bound to that record: the header links back to it, and the deal's activity streams in as quiet centered lines between the conversation.
- **Escalations** — hover any channel message for two carries: **File as a task** puts it on the Tasks runway deep-linked back to the message, and **Promote to a Board decision** shelves it on the Board, quoted with its author and a source link.
- **Transcript export** — owners see a download control in the channel header that exports the full channel as a plain-text transcript for retention or legal review. Every export is recorded in the audit log.

How to use it

1. Message your team in a channel, or start a new one for a specific deal.
2. Ask the AI with **@assistant** when you want an answer the whole channel can see.
3. Message the **Vera AI team** thread to reach our analysts directly.

Related: Vera AI · Advisory · Support

SUPPORT

Where: Team & AI › Support · `/support` · via **More** or Ctrl/⌘ K
Who: everyone

Support tickets: one tracked case per issue, with a number, a status, and the full thread kept on the record. Chat stays the fast lane for quick questions; a ticket is for anything you want followed to resolution, and everyone in your organization sees the same list.

On the screen

- **Vitals** — how many tickets are with our team, how many are waiting on a reply from your side, and your workspace's own median time to a first reply, measured from the tickets we have actually answered rather than promised.
- **Open a ticket** — the masthead action, opening as a dialog: subject, category, priority, and what happened; one issue per ticket. Picking a priority states the first reply target that comes with it, so the commitment is visible before you send.
- **First reply targets** — stated in the rail in full: two hours for urgent, four for high, eight for normal, twenty four for low, measured from when you open the ticket, Monday to Friday. Your measured median sits in the masthead beside them so you can check the target against what actually happens.
- **Open** — every ticket still in flight, most urgent first, each carrying its own clock. A ticket nobody has replied to yet shows how long it has been waiting against its target, and one that has run past that target is marked with a red edge and red text. An answered ticket shows how long the first reply took, against the same target.
- **Status** — **Open** is with our team, **Waiting on you** means we replied and need something from you, **Resolved** means we believe it is fixed and asks you to confirm, **Closed** is done.
- **Settled** — closed tickets stay readable with the thread and the resolution intact.
- **Ticket thread** — the full conversation on one page; our replies carry a blue edge, and replying always hands the ticket back to our team.
- **Faster than a ticket** — the rail names the two things that answer sooner: Vera, who answers most platform questions instantly and reads your own documents while she does it, and the support chat, which reaches the same team without opening a ticket.

How to use it

1. Press **Open a ticket** in the masthead, give it a one-line subject, pick a category and priority, and describe what happened. The dialog states the first reply target for the priority you chose.
2. Watch for the reply: it lands as a notification (and an email doorbell), and the ticket flips to **Waiting on you** when we need something from you.
3. When the ticket turns **Resolved**, confirm with **Confirm resolved & close**, or reply to reopen it if the issue is not gone.
4. A closed ticket can always be reopened; the thread continues where it left off.

Related: Chat · Notifications

NOTIFICATIONS

Where: the bell in the top bar opens the inbox panel · `/notifications` is the full wire
Who: everyone

Two surfaces share the bell. The bell's panel is the triage inbox for everything that happened while you were away: notifications (delivered reports, notice windows opening, replies, and mentions), watchtower changes, and money-critical signals, filtered to what needs you. The full page at `/notifications` is **The wire**: the complete record of every alert, printed as one timeline. Money items are flagged calmly on both, never alarmist.

The bell panel

- **The badge** — counts your unread signals plus unread team chat.
- **Team chat row** — pinned at the top of the panel, with your unread chat count; it opens Chat.
- **Needs decision filter** — narrows the panel to the items actually waiting on a decision, so signals you only need to be aware of drop away.
- **Per-signal mute** — mute any signal kind you do not want to keep hearing about, without turning off its whole category. A mute silences that kind on the bell, push, and email together; muted kinds still print on the full wire, and you can unmute from the panel's muted chips or under Settings › Notifications.
- **Mark all read** — settles the panel, and a link at its foot opens the full wire.

The wire (`/notifications`)

- **The ticker** — a navy band with your unread count, one verdict sentence on the week, a 14 day pulse of arrivals per day, and Mark all read.
- **The timeline** — every alert grouped by day on one threaded timeline, each entry carrying its desk's color and its kind. Urgent unread items carry a cream wash and a **Needs attention** mark.
- **Reading** — opening an entry marks it read and takes you to the record it is about; read entries settle to a single archived line, so the page stays a record without shouting.

How to use it

1. Triage from the bell: scan the unread items, switch to **Needs decision** to see only what is waiting on you, and mute the kinds you have handled.
2. Open the full wire when you want the complete history, day by day.
3. Click through any item to the record behind it to act on it.

Related: My work · Chat

Next: Chapter 06 · Settings & account

06 · Settings & account

Everything about you and your organization's setup lives here: your profile and how you sign in, what the AI has learned about you, your plan and billing, your teammates and their roles, the security and compliance controls that govern your data, and the connections to your other systems. Settings are grouped into four sections, You, Workspace, Security & compliance, and Connections. Some pages are visible only to organization owners; each section below says who can see it.

The settings workspace. Every settings page opens inside the same frame. An index rail runs down the left holding all twenty three pages under their four groups, with your organization, your role, and your plan at the top of it, and the Trust center on a navy plate at the foot. Pages that carry a live fact show it on the entry: the number of people, the plan name, how many systems are connected. A dot marks a setting that needs attention, such as two factor still being off.

- **Filter** — type in the box at the top of the rail to search every settings page by name and by description; press the forward slash key from anywhere in settings to jump into it, and Escape to clear it.
- **Nested sections** — the Security page's own sections (Posture, Single sign on, Baseline, Data controls, Audit log) appear indented under Security while you are inside it.
- **Narrow screens** — below the desktop breakpoint the rail collapses into an **All settings** button at the top of the page; press it to open the same grouped map.
- **Page masthead** — every page opens with the icon from the rail entry you pressed, the group it belongs to, the page name, and one sentence on what the page decides.
- **Owner-only pages** — if a page is managed by owners and you are not one, it says who can change the setting and what to ask for, rather than showing you an empty form.

The Security & compliance section is deliberately short: one **Security** page carries account security, posture, single sign on, the baseline, data controls, and the audit log as nested sections, and one **Trust center** page carries the overview and the figures ledger as underline tabs. Each keeps its own section below, and old bookmarks to the former standalone pages redirect to the right place.

SETTINGS HOME

Where: Settings · `/settings`
Who: everyone

The overview of your workspace: its current state first, then the map. Use it when you want to know where things stand; use the index rail when you already know the page you want.

On the screen

- **Masthead** — your organization name, and on the right the account you are acting as, with your role. Settings behave differently for owners and members, so the page states which you are.
- **Vitals** — four live numbers, each a link to the page behind it: **People** holding a seat, the **Plan** paying for the workspace (with days left when you are on a trial), whether **Two factor** is on for your account, and how many systems are **Connected**.
- **Worth doing now** — anything genuinely waiting on a decision, computed from the live state of your workspace rather than a fixed list, written in plain words with the page that resolves it. It ranks eight possible items and shows at most four: two factor still off; renewal alerts off on every channel while agreements carry renewal dates; a trial inside its last week; an access review past the quarterly cadence; an API key expiring inside 30 days; an empty security baseline; negotiation memory still empty after four weeks in the workspace; and no systems connected. Access reviews, API keys, and integrations are owner-only, so a member never sees those three. Nothing appears here when there is nothing to do. Every row is a link that states where you stand: the panel itself never changes a setting.
- **The four groups** — You, Workspace, Security & compliance, and Connections, each page with its one-line description and its live fact.
- **Trust center** — the navy panel closing the page, leading to how your data is held.

How to use it

1. Read the vitals. Anything reading Off or 0 is a gap worth closing.
2. Work through Worth doing now; each row goes straight to the page that fixes it.
3. Use the index rail on the left for everything else, or the filter if you are not sure which page owns the setting.

Related: Security · Billing · Trust center

PROFILE

Where: Settings › You › Profile · `/settings/account`
Who: everyone

Your personal details: your name, job title, timezone, photo, the vendors you follow, and your password. This is also where you connect your Microsoft account.

On the screen

- **Name, title, and timezone** — how you appear to teammates and the timezone used for dates and digests.
- **Photo** — drag an image on or click to upload; used for your avatar across the app.
- **Email** — your sign-in address, with a **Change** button. Enter the new address and your current password; we email a confirmation link to the new address, and nothing changes until it is clicked.
- **Followed vendors** — the vendors you want updates about; they drive your weekly report and vendor alerts. Email subscriptions themselves live under Notifications.

- **Password** — set a new password without leaving the page. You confirm your current password first, the new one is checked against known breaches, and every other session is signed out when it changes.
- **Your data** — download everything the platform holds about you as one JSON file (profile, memberships, preferences, assistant memory, sign-in devices, your own activity trail), or file a delete-my-account request our team processes within 30 days. You can cancel the request until it is processed.
- **Microsoft account** — connect Microsoft so contract files and calendar data can flow in (when your organization has it configured).
- **Invite a colleague** — a quick invite card; only owners can actually send invites.
- **Leave this workspace** — remove your own seat from the current workspace. Your account and any other workspaces stay; an owner can invite you back. If you are the only owner, promote someone else first.

How to use it

1. Fill in your name, title, and timezone, then save.
2. Add the vendors you care about so your alerts and weekly report are relevant.
3. Use **Change password** to rotate your password whenever you need to; use **Change** next to your email when you move addresses.

Related: Security · Members · Notifications

SECURITY

Where: Settings › Security & compliance › Security · `/settings/security`
Who: everyone

In plain words: your account's locks: two-factor, sessions, and the org-wide policy.

When to reach for this

- Day one, three minutes: turning on two-factor is the highest-value security step you personally control.
- Something feels off: the sign-in history and active sessions answer "was that me?"
- Owners setting the org-wide requirement: the handbook's security defaults sequence it.

The landing section (Account security) of the one Security page. Protect your account with two-factor authentication and see the devices where you are signed in; owners also set the organization-wide two-factor policy here. The other five sections, Posture, Single sign on, Baseline, Data controls, and Audit log, sit indented under Security in the index rail while you are inside it, and each is documented below.

On the screen

- **Two-factor authentication** — register an authenticator app and turn on the rotating six-digit code. Once on, that code becomes your sign-in: the sign-in screen asks for it in place of your password.
- **Backup codes** — one-time recovery codes to use if you lose your authenticator; the count remaining is shown.
- **Active sessions** — the sessions signed in to your account, with the option to sign out everywhere else.

- **Sign-in history** — every device your account has signed in from, with first seen, last active, and the IP address. A device you have never used before always triggers the new-device email; an entry you do not recognize means change your password (which signs out every other session) and review the active sessions above it.
- **Organization policy (owners)** — require two-factor for every member of the organization.

How to use it

1. Open your authenticator app (Google Authenticator, 1Password, Authy, and the like) and scan the code shown.
2. Enter the six-digit code to confirm, then save your backup codes somewhere safe.
3. Review active sessions and sign out any you do not recognize.

Tips

- Turning on two-factor is the single highest-value thing you can do for your account. If your organization requires it, you will be held here until you set it up.

Related: Security posture · Two-factor authentication

AI PERSONALIZATION

Where: Settings › You › AI personalization · `/settings/ai`
Who: everyone

What the platform has learned about how you like output, in plain language and fully under your control. The more you use Vera AI, the more its answers match your preferences.

On the screen

- **Learned preferences** — each thing the AI has picked up about how you like answers, newest first.
- **Remove** — delete any single preference; it stops shaping your answers immediately.

How to use it

1. Review the list to see what the AI has inferred.
2. Remove anything that is wrong or that you would rather it not use.

Tips

- This list is private to you. Your organization's shared negotiation memory is separate, on the Memory page.
- Learning feeds on what you do with AI output, not on how much you chat: rating an answer, exporting a deliverable, applying or rejecting a suggestion Vera staged, turning an answer into a watch, regenerating, or dismissing an insight. The learning pass runs weekly and writes a preference only after it sees the same pattern a few times, so a new account starts empty. A written feedback comment is the fastest way to teach it something.

Related: Memory · Vera AI

MEMORY

Where: Settings › You › Memory · `/settings/memory`
Who: everyone

What your organization learned in past negotiations and how it likes to negotiate. Briefs, mandates, and Ask all open with this record, so it is worth keeping current.

On the screen

- **Negotiation posture** — your organization's stance and preferences, editable by owners.
- **Facts** — durable conventions about your workspace, for example how a vendor's contracts are classified. Some are added by hand; some are learned automatically when your reviewers correct the same extracted field the same way twice.
- **Remembered facts** — the individual things learned across past negotiations, each of which you can edit or remove.

How to use it

1. Check the posture reflects how your organization actually negotiates.
2. Prune or correct any remembered fact that is out of date.
3. Corrections you make in the extraction review queue also teach the extractor directly: the next document from that vendor is read with your past corrections in hand.

Related: AI personalization

BILLING

Where: Settings › Workspace › Billing · `/settings/billing`
Who: everyone (only owners can change the plan or pay)

See what plan your workspace is on, how much of your allowance you have used, your benchmark credits, and, if you own the organization, upgrade.

In plain words: your plan, your meters, and the upgrade path.

When to reach for this

- A feature says an allowance ran out: the meter here says when the window frees up (Chapter 10 defines every allowance).
- The quarterly owner check: does the plan still fit how the team actually uses the platform?

On the screen

- **Current plan** — your plan, its status, and price, with contract, benchmark, AI, and Vera allowances at a glance.
- **Usage this period** — how many AI reports and agent requests you have used against your cap, plus your personal AI token usage over the last window.
- **AI allowance meter** — on seat-priced plans, the percent of your monthly allowance used and roughly how many runs are left; the allowance is yours alone, not pooled.
- **Plans** — the sellable plans side by side, with an upgrade button on each.

How to use it

1. Check where you sit against your allowance for the period.
2. If you own the organization and need more, pick a plan and message our team; upgrades are set up the same day and billed by invoice, net 30.
3. Not an owner? Ask an owner to upgrade, or message the team.

Related: Members

DESKTOP APP

Where: Settings › Workspace › Desktop app · `/settings/desktop-app`
Who: everyone

Install Vera AI as its own app on your Windows laptop: its own window, taskbar icon, and Start menu entry, on the same account and data as the browser.

On the screen

- **Download for Windows** — the installer, with version and size shown; it installs for your user only, no admin rights needed.
- **How to install** — the numbered steps, including what to do if Windows shows a "Windows protected your PC" screen.
- **What you get** — drop a contract onto the window or from Explorer, native Windows notifications, Vera on top of any app, Windows Hello lock, and automatic updates.
- **Removing it** — uninstall any time from Windows Settings; your account and data are untouched.

How to use it

1. Click **Download for Windows** and run the downloaded installer.
2. If Windows warns about a new app, choose **More info**, then **Run anyway**.
3. Sign in with the same email, password, and two-factor code you use in the browser.

Tips

- A macOS version is planned. Until then, Mac users get the same workspace in the browser.

MEMBERS

Where: Settings › Workspace › Members · `/settings/members`
Who: org owners only

Invite teammates, set their role, and manage everyone in your organization. Owners also see pending invitations, seat usage against the plan, and any join requests waiting for them. A non-owner who opens the page sees a locked plate naming the workspace owners to ask, not a bare access line.

In plain words: who is in, what they can do, and the clean way in and out.

When to reach for this

- Bringing people in (the onboarding recipe) and, just as importantly, out (when someone leaves).
- The monthly hygiene glance: stale invitations and unaccounted seats.

On the screen

- **Seats** — the verdict leads with seats in use against your plan's cap: 2 on Free, 5 on Trial, 2 on Growth, 5 on Professional, unlimited on Enterprise (the full table is in Plans & limits). The cap is a hard limit on the roster, checked when someone is added by invite, domain join, or SCIM: pending invites and blocked members hold seats too, and nobody already on the roster is ever removed by it. When the cap is full, remove a member, revoke a pending invite, or upgrade.
- **Member list** — everyone in the organization with their role, last sign-in, and recent activity trail. Export the whole roster to CSV for access reviews from the link above the grid.
- **Invite** — send an email invitation and set the new person's role. Invitations expire after 7 days; resending restarts the clock and refreshes the sign-in details.
- **Invite several at once** — paste up to 50 addresses (commas, spaces, or new lines) and give them all one role in a single batch.
- **Pending invitations** — invites not yet accepted, with the days until expiry and a delivery status when an email bounced or was delayed. Revoking a pending invite removes the reserved seat entirely.
- **Roles** — change a member's role inline, or block a member from the workspace. Blocking or removing someone signs them out everywhere immediately. The workspace always keeps at least one active owner: promote a second owner before demoting or removing the last one.
- **Sign out** — end every session a member holds without blocking or removing them, for a lost laptop or a departed contractor whose seat stays. They sign back in with their password (and MFA) next time.
- **Join requests and domain capture** — approve people who asked to join, pick their role at approval, and optionally let anyone on your email domain request to join.

How to use it

1. Click **Invite**, enter the person's email, choose their role, and send. For a whole team, use **Invite several people at once** instead.

2. Watch pending invitations for delivery problems and resend if needed; an expired invite just needs a resend.
3. Adjust roles as responsibilities change; block anyone who should no longer have access.
4. Before an owner leaves the company, promote their replacement to owner first; the platform refuses changes that would leave no active owner.

Tips

- If an invite shows as bounced, check the address; the person never received the email.
- Any member can leave a workspace themselves from the bottom of Settings › Profile, no owner action needed.

Related: Team activity · Deal approvals

TEAMS

Where: Settings › Workspace › Teams · `/settings/teams`
Who: org owners only

Named groups you can grant confidential contracts and agreements to as one unit. Membership in a granted team is the access itself: joining the team grants its documents at that moment, leaving revokes them, with nothing to clean up per document. A non-owner who opens the page sees a locked plate naming the workspace owners to ask.

In plain words: grant a group once, and joining or leaving the group is the access change.

When to reach for this

- Before the first sensitive document arrives, not after: create the teams the org chart implies.
- Restricted work (a Vault project, a diligence) where offboarding must be one removal, not a document sweep.

On the screen

- **Create a team** — name it after how your company actually organizes: Procurement, Finance, Legal.
- **Team cards** — each team with its people and how many confidential grants it currently holds. Add or remove people inline; removal takes effect immediately.
- **Delete** — deleting a team revokes every grant it holds for its members in the same stroke.

How to use it

1. Create a team and add the people who share an access need.
2. Open a restricted contract and grant the team from its **Access** panel, next to individual grants.
3. When someone changes roles, move them between teams; their document visibility follows automatically.

Tips

- Prefer team grants over long lists of individual names: the grant list stays readable and offboarding is one removal, not a document sweep.

Related: Members

TEAM ACTIVITY

Where: Settings › Workspace › Team activity · `/settings/activity`
Who: org owners only

What each member actually does on the platform, which features they use, and an on-demand report that turns the numbers into an adoption plan. A non-owner who opens the page sees a locked plate naming the workspace owners to ask.

On the screen

- **Usage verdict** — how alive the workspace is this week: active members, page views, most-adopted feature, and features nobody has touched.
- **Per-member table** — each member's last sign-in, page views over 7 and 30 days, recorded actions, features tried, and last screen.
- **Recent activity** — the last recorded actions across the workspace, named, newest first.
- **The usage report** — an analyst read of the numbers that you can generate on demand.

How to use it

1. Scan the verdict to gauge overall adoption.
2. Use the per-member table to spot who has not started and which features go unused.
3. Generate the usage report when you want a written adoption plan to act on.

Related: Members · Audit log

DEAL APPROVALS

Where: Settings › Workspace › Deal approvals · `/settings/approvals`
Who: org owners only

An optional sign-off chain: up to ten approvers who must clear a deal before it is signed. Off by default; turn it on when your organization needs a documented approval step. A non-owner who opens the page sees a locked plate naming the workspace owners to ask.

On the screen

- **Enable and mode** — turn the chain on and choose how sign-off works.
- **Approvers** — add up to ten people from your organization as the required approvers.

- **Recent sign-off requests** — every request negotiators have sent through the chain, with how many approvers have cleared it.

How to use it

1. Turn on deal approvals and pick your approvers.
2. Negotiators request sign-off from a deal's page under Negotiations; approvers vote from their queue.
3. Track progress in the recent requests list.

Related: Approvals

BRANDING

Where: Settings › Workspace › Branding · `/settings/branding`
Who: org owners only

Put your organization's logo on the documents you export. Once a logo is set it co-brands the exported CFO executive brief and the boardroom deck and brief alongside the Vera AI brand. With no logo set, exports are unchanged.

On the screen

- **Logo upload** — upload your logo; a preview shows how it will appear on an exported document.

How to use it

1. Upload a logo image.
2. Export a CFO brief or a boardroom deck to see it applied, alongside the Vera AI brand.

DISPLAY CURRENCY

Where: Settings › Workspace › Display currency · `/settings/currency`
Who: org owners only

The default currency for amounts your team enters. Benchmarks stay in USD so every market comparison holds.

On the screen

- **Workspace display currency** — the currency amounts render in when the row that holds them did not record one of its own. Contract values, spend, and invoice lines follow this setting.
- **Save** — applies to the whole workspace, not just to you.

How to use it

1. Pick your currency and press **Save**.

2. Leave rows that already carry their own currency alone; they keep it, whatever this is set to.

Tips

- This setting never converts a number and never relabels a benchmark. Market figures are stored in USD and stay labeled USD.

Related: Billing · Branding

NOTIFICATIONS

Where: Settings › Workspace › Notifications · `/settings/notifications`
Who: everyone (owners get the extra alert and webhook controls)

In plain words: what may interrupt you, on which channel; everything renewal-shaped starts silent until you ask.

When to reach for this

- Week two of the ramp: **Turn on money-critical alerts** is the one press that makes the calendar audible.
- An email you expected never came, or one you did not want keeps coming: this page is the truth for both.
- Wiring alerts into Slack or Teams DMs (the chat recipe).

The page opens with **What reaches you now**: a summary computed from your saved switches listing every email, briefing, and chat DM that will currently reach you, each naming its trigger, with renewal notices showing your organization's lead days. When nothing is on it says so plainly: everything is opt in until you turn it on.

Tune which alerts reach you, on which channel, and when. One preference model governs every channel: a setting made here (or a mute made from the inbox bell) applies to the in-app inbox, push, and email together, so nothing silences you on one surface while interrupting you on another. Renewal alerts are opt in: nothing about a renewal reaches you, on any channel, until you switch it on. Owners can also post alerts to Slack or Teams and set the AI copilot's alert rules.

On the screen

- **Weekly briefings** — three vendor-intelligence emails, each at most once a week, each strictly opt-in: price list updates per vendor (what moved, by how much, what to do), vendor licensing news (the Monday wire across the 50 most-negotiated vendors, your vendors leading), and the weekly knowledge briefing (new analyst research and what shipped). Every one of these emails carries a one-click unsubscribe that works without signing in, and this card is the standing in-platform switch. When a briefing is on, a link under it leads to its vendor picker (Watch for price lists, the Weekly report page for the wire).
- **Tips and lifecycle emails** — one switch over the onboarding guidance, trial check-ins, and come-back notes. It is off until you switch it on, so nobody receives these unasked. Nothing here touches invites, security notices, alert emails you enabled, or report deliveries.

- **Quiet hours** — pick a window (say 10:00 pm to 7:00 am, in your local timezone) and inside it phones stay silent while non critical alert emails wait until the window ends. Money critical alerts (renewal deadlines, deadline escalations, policy checks, above-market warnings, uplift breaches, anomalies) always come through immediately.
- **Every alert type is listed** — anything that can reach your bell has a row here, so there is always somewhere to switch it off. Six types that nobody addressed to anyone (commitment tracking, invoice findings, vendor health warnings, ecosystem shifts, sourcing events, and renewal engine events) start off until you turn them on, the same way the renewal family does.
- **Channel posts** — a Slack or Teams channel an owner wires up is the organization consenting for that channel, because the people in it cannot each be asked. What an owner chooses is which alerts belong there; a channel subscribes to the kinds you pick rather than to everything. Workflow agents posting to that channel have to clear the same switch.
- **Workflow agent output** — the memo a workflow agent mails when a run finishes. Off until you switch it on, because any colleague with edit rights can build a workflow that emails the whole organization, so the person receiving it decides rather than the person who wrote it. Approval requests addressed to you are the deliberate exception and always reach you, because the run waits on your answer.
- **Alert preferences** — for each kind of alert, up to four toggles: in-app, email, weekly digest, and Chat DM. Email is off until you switch it on, with four exceptions that start on because silence would be a failure: the report delivery email, benchmark status updates, the one-time outcome follow-up 30 days after a delivered benchmark, and @mentions. Every renewal alert that fires on its own goes further and starts off on every channel, in-app included: renewal notice alerts, deadline escalations, policy compliance checks, the AI copilot digest, above-market warnings, stale benchmark reminders, uplift cap breaches, auto-drafted briefs, renewal opening moves, and renewal prep packages say nothing until you turn them on here. The **Turn on money-critical alerts** button switches on the renewal set in one click. AI renewal agent briefs are the exception: switching the renewal agent on for a contract is itself the opt in, so those briefs send, and this toggle stops them across every contract at once. Weekly digest bundles that alert type into one Monday roll-up email instead of individual sends; money critical kinds cannot be bundled. Alert types you muted from the inbox bell appear here with a Muted chip and are one click to unmute; a mute silences that type on every channel unless you split channels with these toggles.
- **Chat DM** — send your own renewals and money alerts (renewal notices, deadline escalations, price drift findings, buyer window alerts, and above-market warnings) to your personal Slack or Teams direct messages. Link once by sending the Vera bot a direct message in your company Slack or Teams (an owner connects the bot under Settings › Integrations); the account link happens automatically and the note above the table confirms it. The DM lane is strictly opt in per alert type, only these five types can use it, a bell mute silences it like every other channel, and quiet hours hold the non critical pings. The shared org channel owners configure below is separate and unchanged.
- **The daily leverage note** — one email a morning covering everything that moved under your deals, in place of each watchdog mailing you separately. It is ranked on a single measure: how much the thing that happened changed your position. Money at stake and days left to act are what move that ranking, because between them they decide whether you can still do anything about it, so a price list change on an agreement that renews in two years sits below a notice deadline nine days out even when the price

change is the larger number. Findings that scored below the line are counted at the foot of the note rather than dropped silently, so you can see what was judged and set aside. On a day when nothing moved, nothing is sent. This rides on your renewal alert setting above: switch renewal alerts on and the note starts arriving, and every note carries the same one-click Turn it off link.

- **Monday brief sections** — your weekly emails arrive as one Monday brief with four sections: the vendor market wire, renewals coming up, watchtower changes, and savings and wins. Turn any section off here while keeping the rest.
- **AI copilot preferences (owners)** — the digest, anomaly watch, board report, community and outcome network sharing, and dispute alert thresholds.
- **Webhooks (owners)** — post alerts to a Slack or Teams channel.

How to use it

1. Pick your weekly briefings: price lists, licensing news, knowledge, any mix. Off is always one toggle (or one email click) away.
2. Set quiet hours if you carry the mobile app: nights stay silent and nothing important is lost, it waits for morning.
3. Turn on the renewal alerts you want. They start off, so a workspace stays silent about renewal dates until someone here asks for them, and the standard renewal checklist tasks are not filed either; **Turn on money-critical alerts** switches the whole set on at once. Renewal notice alerts also switch on the Monday renewal radar, one roll-up email covering everything entering or inside its notice window over the next 60 days, and the daily leverage note described above.
4. Turn off any alert kind you do not want, move it to a single channel, or send it to the weekly digest. Every renewal alert email also carries a one-click **Turn it off** link that works without signing in.
5. Want your renewals in Slack or Teams? DM the Vera bot once in your company workspace, then flip the Chat DM toggle on the alert types you care about.
6. Review the Muted chips: anything you muted from the bell in passing is listed here, so no alert type stays silenced invisibly.
7. Choose which Monday brief sections you receive; the others still land in the app.
8. Owners: set the copilot alert rules and connect a chat channel if the team wants alerts there.

SECURITY POSTURE

Where: Settings › Security & compliance › Security › Posture · `/settings/security/posture`
Who: everyone (owners see and manage every control; members see the policies that apply to them and their own devices)

The Posture section of the Security page (formerly the standalone Security Center, whose old address redirects here). Your organization's live security posture on one view: two-factor coverage, single sign-on, retention, API credentials, the audit stream, recent security events, and your sign-in devices. Live state, not marketing.

On the screen

- **Posture tiles** — two-factor coverage across members, sign-on method, active API keys, and audit events in the last 30 days.
- **Controls** — each security control with its current state and, for owners, a link to manage it.
- **Recent security events** — sign-ins, permission changes, and credential events.
- **Your sign-in devices** — the devices seen on your account, with a link to manage active sessions.

How to use it

1. Read the posture tiles for a one-glance health check.
2. Owners: work down the controls list and tighten anything that reads as merely "Available" rather than "In force."
3. Check your devices and sign out anything unfamiliar from the Account security section.

Related: Security · Trust center · Audit log

CLAUSE POSITIONS

Where: Settings › Security & compliance › Clause positions · `/settings/clause-library`
Who: everyone

The summary of your negotiation playbook: the positions, pre-approved language, and walk-away lines that every AI review holds your contracts to. There used to be two editors for the same library, one here and one on the desk. There is now one. **Editing lives in the Clause library desk, next to the comparison and redline tools that read the positions; this page is the summary.**

On the screen

- **Positions** — how many clause positions are on file. Each one is held against every contract the AI reviews.
- **Must have** — how many of them your organization will not sign without. These form the template the Coverage check sweeps across every contract.
- **Last edited** — when the library last changed. It reads **Never** until a position has been saved.
- **What reads it** — proposal reviews, redline drafting, the DPA check, and the counter-draft linter all hold contracts to these positions and insert your pre-approved language verbatim.
- **Starter library** — 22 proven positions to begin from. Restoring adds only the ones your library is missing: it never overwrites an edit and never duplicates.

How to use it

1. Read the three figures to see where the library stands.
2. With nothing on file yet, press **Restore starter library** to begin with the 22 starter positions. It takes effect immediately, with no confirmation step, and reports how many it added.
3. Press **Open the Clause library desk** to add, edit, or delete positions and to set the must-have flags.

Tips

- Restoring is safe on a library you have already shaped: it fills gaps only, so you can run it again after adding a position type you had skipped.

Related: Clause library · Security baseline

SECURITY BASELINE

Where: Settings › Security & compliance › Security › Baseline · `/settings/security/baseline`
Who: everyone

The Baseline section of the Security page. The data-protection and security requirements every vendor contract should meet. These are the rules the DPA check enforces when it reviews a contract.

On the screen

- **Baseline requirements** — your written data-protection and security requirements; a default is offered as a starting point.
- **Rulepacks** — pre-built requirement sets you can switch on.
- **Network and session policy** (owners only) — pin the workspace to known networks and bound session lifetimes: an IP allowlist that takes exact addresses and CIDR ranges, a maximum session age in hours, and an idle timeout in minutes.

How to use it

1. Edit the baseline to match what your organization requires of vendors.
2. Turn on any rulepacks that apply.
3. The DPA check then measures each contract against this baseline.
4. Owners setting a network policy: the card shows your current IP with one click to add it, and the allowlist refuses to switch on unless your own address is covered, so you cannot lock yourself out at enable time. Every denial lands in your audit log, and a policy change takes up to a minute to reach every server.

Note: if your office IP later changes, everyone including owners is blocked until the policy is updated from an allowed network; prefer ranges over single addresses if your egress IP rotates.

Related: Clause positions

SINGLE SIGN ON

Where: Settings › Security & compliance › Security › Single sign on · `/settings/security/sso`
Who: org owners only

The Single sign on section of the Security page. SAML single sign-on and SCIM provisioning for your organization, for teams that manage access through an identity provider.

On the screen

- **SSO configuration** — connect your SAML identity provider and choose whether sign-on is enforced for your domain.
- **Break-glass policy** — a plain statement of the one exception when Require SSO is on: org owners keep password sign-in so a broken identity provider cannot lock the whole tenant out. Every break-glass sign-in is written to your audit log as `auth.sso_breakglass`. Members and admins have no bypass.
- **SCIM tokens** — provisioning tokens so your identity provider can create and deactivate accounts automatically.

How to use it

1. Enter your identity provider's SAML details.
2. Decide whether to enforce single sign-on for everyone on your email domain.
3. Issue a SCIM token if you want automatic user provisioning.
4. If your identity provider supports group push, push the groups that should map to teams; members must be provisioned as users first.
5. Optionally map groups to roles: members of a mapped group get that role on every push (approver, analyst, member, or viewer; ownership is never assigned this way), and leaving the group takes the role with it.

Related: Security posture

DATA CONTROLS

Where: Settings › Security & compliance › Security › Data controls · `/settings/security/data`
Who: org owners only

The Data controls section of the Security page. Retention, export, and deletion for your organization's data, the document retention switch for the analysis tools, plus your consent choices for the community and outcome network. Every hard deletion writes a deletion certificate you can hand to an auditor.

In plain words: how long data lives, where it lives, how it leaves, and the certificates that prove deletion.

When to reach for this

- The two policy decisions of week one: retention and community consent (the handbook).
- A security questionnaire's data-location and deletion questions (the recipe).

On the screen

- **Retention policy** — automatically delete contracts a set number of months after they are uploaded, rows and files both. Leave it blank to keep contracts until you delete them. The sweep runs daily, and each automatic removal writes a deletion certificate.
- **Documents from analysis tools** — on by default: files handed to the decoders, the Verdict Bar, Second opinion, the Vera Deal-Room, Proposal comparison, the vendor playbook intakes, and now the Time

machine (both documents), Quote Face-Off (PDF quotes only) and the Rate Card Benchmark (PDF and DOCX only) are saved to your workspace as contracts, private to your organization, and each joins your benchmarking queue. Photo captures, XLSX rate cards, and anything handed to Snap a quote are never stored whatever this switch says. Benchmarking extracts only high level data points (the vendor, the price, the term, the key commercial terms); the document itself never leaves your organization, and every saved document can be deleted from its contract record. Turn the switch off and those tools read files in memory and keep nothing; documents already saved stay until you delete them.

- **Data location** — a read-only statement of where data physically lives: the United States (Vercel and Supabase US regions). EU residency is on the roadmap and tracked publicly on the security page's claim ladder.
- **Deletion certificates** — one row per completed hard deletion: date, scope (contract or organization), what was deleted, and how many files were removed. "View certificate" opens a print-ready document listing rows removed by table, the storage paths removed, the actor, and the UTC timestamp; use your browser's print dialog to save it as a PDF.
- **Community and outcome network consent** — opt in or out of contributing anonymized data to the community benchmarks and outcome network.
- **Reseller quotes** — recent quotes shared for benchmarking.
- **Export and deletion** — request an export or a deletion of your organization's data.

How to use it

1. Set a retention period if your policy requires one; blank keeps contracts until you delete them.
2. Choose whether to contribute to the community and outcome network.
3. Request an export or deletion when you need it; our team processes deletion requests and confirms by email.
4. After any deletion, open its certificate from the Deletion certificates list to verify what was removed.

Related: Trust center · Audit log

AUDIT LOG

Where: Settings › Security & compliance › Security › Audit log · `/settings/security/audit`
Who: org owners only

The Audit log section of the Security page. Who did what in your organization, including views and downloads of contract files. A transparency and compliance record you can filter, search, and export. Old links and saved views under the former standalone address redirect here with their filters intact.

On the screen

- **Filters** — narrow to file downloads, file views, uploads, report links, deletions, or member changes. A file view is a document opened in the in-app viewer; a file download is a copy saved to the reader's machine. The log records both separately.

- **Activity table** — each event with the actor, the source device and IP, and when it happened.
- **Saved views** — save a filter you return to often.
- **Export and stream** — download the full log as CSV, or stream events to your SIEM with an audit webhook.

How to use it

1. Pick a filter or search for an actor or activity.
2. Save a view if you audit the same slice regularly.
3. Export to CSV, or set up an audit webhook under Integrations for continuous streaming.

Related: Security posture · Integrations

ACCESS REVIEWS

Where: Settings › Security & compliance › Access reviews · `/settings/access-reviews`
Who: org owners only

In plain words: a quarterly, exportable proof that someone looked at who can reach what.

When to reach for this

- The quarterly governance clock (the cadence): thirty minutes, four times a year.
- An auditor or customer asks for your access-review process: the completed reviews are the answer.

A quarterly attestation of who can reach what. Starting a review snapshots your organization's entire access surface at that moment: every member with their role, every team membership, and every restricted document grant. You then confirm or revoke each line, and the completed review is stored as a durable record you can hand to an auditor.

On the screen

- **Start a review** — one action that snapshots the access surface. Only one review can be open at a time.
- **Review lines** — grouped into members and roles, team memberships, and restricted document grants, each with Confirm and Revoke buttons.
- **Complete review** — enabled only once every line is decided, so a finished review always covers everything.
- **Completed reviews** — past attestations with line counts and how many grants were revoked or flagged, each exportable as CSV.

How to use it

1. Open Settings › Access reviews and choose Start a review.
2. Work through each line. Confirm access that is still right; revoke a team membership or document grant that is not. Revocation takes effect immediately.
3. A member who should leave the organization is flagged rather than removed here. Finish the removal on the Members page, which protects the last owner and revokes their sessions.

4. When nothing is pending, choose Complete review. Every decision is written to the audit log.
5. Export any completed review as CSV for an auditor or a security questionnaire.

Related: Members · Teams · Audit log

TRUST CENTER

Where: Settings › Security & compliance › Trust center · `/trust`
Who: everyone

How your data is protected, told in three commitments: kept apart (tenant isolation), kept sealed (encryption and access), and kept yours (audit, deletion, and eyes-off AI). The masthead proves the isolation claim with live numbers. Two underline tabs make this one page: **Overview** (this view) and the **Figures ledger**.

On the screen

- **Kept apart** — tenant isolation at the database, least-privilege access, and roles matched to clearance.
- **Kept sealed** — private contract storage with short-lived signed links, encryption in transit and at rest, and multifactor authentication.
- **Kept yours** — the audit record, data deletion, and how the AI handles your data.
- **Data journey** — a diagram of the path a contract actually takes through the platform.
- **Print or save as PDF** — prints the page as an executive security overview: a cover carrying your organization name, the month, and the same live figures, then the controls. The animated masthead and the data journey diagram are left off the printed pack, so what a reviewer receives is the record rather than the illustration. The figures on the cover are read under your own access scope at the moment you print, and the footnote on the cover says so.

How to use it

1. Read the three chapters to understand the controls applied to your data.
2. Share this page with security reviewers who ask how the platform protects information.
3. When a reviewer wants a document rather than a link, press **Print or save as PDF** and send the brief.

Related: Security posture · Data controls · Figures ledger

FIGURES LEDGER

Where: Settings › Security & compliance › Trust center › Figures ledger · `/trust/figures`
Who: everyone

Every number the platform has cited for you, in one place. The ledger lists each figure across your reports and briefs, linked to the source fact it traces to and the day that fact was current, so any number you forward can be stood behind.

On the screen

- **Cited figures** — one row per number the platform has printed for you, with the document it appeared in, the source fact behind it, and the date that fact was current.
- **Search and sort** — narrow to a vendor, a document, or a value, and sort by any column.
- **Not-yet-traceable note** — the ledger is honest about the figures it cannot trace to a stored fact yet, rather than hiding them.

How to use it

1. Look up any number a report or brief shows to see the fact it came from and when it was current.
2. Search for the vendor or figure a reviewer is questioning and hand them the source in one step.

Related: Trust center · Audit log

INTEGRATIONS

Where: Settings › Connections › Integrations · `/settings/integrations`
Who: org owners only

Connect the systems your team already uses: contract intake, tickets and calendars, spend and identity data, your warehouse, and the SAM and ITAM tools that know what you own and use.

In plain words: the plumbing between this platform and everything else you run.

When to reach for this

- The highest-leverage owner hour (the recipe); each connection has a step-by-step walkthrough in Chapter 19.
- Checking a connector's health: Recent activity shows every sync's last result.

On the screen

- **Workflow hub** — contract intake (five email-in rails, DocuSign, Ironclad, Drive/SharePoint/Box folder watch), renewal workflow (Jira, ServiceNow, calendar invites, Slack, Teams), spend and usage sources (Coupa, Okta, Entra ID, the CSV importers, Smart import), warehouse export, and signed JSON event webhooks.
- **Agent Negotiation Protocol** — the opt-in card for structured agent-to-agent negotiation, with a link to the public spec.
- **Procurement and ERP (P2P)** — purchase orders, requisitions, and supplier records from your procure-to-pay systems, with maverick spend flagged.
- **License and deployment discovery (SAM / ITAM)** — connect the systems that know your entitlements and deployments; the license intelligence table shows entitled vs deployed vs active per vendor, with estimated shelfware.
- **Recent activity** — the latest sync results for each connected source.

How to use it

1. Pick a system, enter its credentials or copy the intake address it needs, and connect.
2. Let the first sync run, then review the license and procurement intelligence tables.
3. Add an audit webhook if you want to stream events to your SIEM.

Related: Audit log · API · The integration catalog

API

Where: Settings › Connections › API · `/settings/api`
Who: org owners only

API keys, webhooks, and the MCP connector, so you can connect Claude, ChatGPT, or any other tool to your benchmark data.

On the screen

- **API keys** — create, view, and revoke keys; keys expire within a year.
- **Connect** — the base URL and details for connecting an external tool or the MCP connector.

How to use it

1. Create an API key and store it securely; you see the full value only once.
2. Point your tool or the MCP connector at the connection details shown.
3. Revoke any key you no longer use.

Related: Integrations

That covers the settings. When you are not sure where a screen belongs, start from the manual index and follow the chapter that matches the rail group you are looking at.

Next: Chapter 07 · Common workflows

07 · Common workflows

The chapters before this one document the platform screen by screen. This chapter is the other way in: start from what you are trying to get done, and follow the steps across whatever screens it touches. Forty recipes cover the jobs buying teams actually run, and every recipe links back to the reference sections for detail.

Each recipe carries the same equipment: **For** (who runs it), **Time** (an honest estimate of your effort, not the platform's), **You need** (what to have before starting), the **Steps**, the **Forks** (the decision points where real runs diverge), and **Done when** (the observable state that means you can stop). Time estimates are for a prepared run; your first pass through any recipe takes longer, and that is normal.

The recipes at a glance

IF YOUR JOB IS…	START WITH
Price and value	Benchmark a proposal end to end · Get an instant price read, no upload · Choose between competing quotes · Get a second opinion before you sign · Turn a delivered report into the vendor conversation
Renewals and deals	Track agreements and never miss a notice window · Prepare a renewal 12 months out · Prepare a renewal 90 days out · Rescue a renewal 2 weeks out · Plan and run a renewal negotiation · Rehearse a negotiation before the real one · Hand your renewal calendar to Vera · Run your Monday renewal desk · Team up with other buyers on a renewal
Money and proof	Prove and share your savings · Report savings at quarter end · Dispute an invoice that does not match your contract · Right-size licenses before a renewal or true-up · Consolidate duplicate tools · Prepare the CFO one-pager · Build the board pack
Paper and protection	Set your clause standard and sweep the estate · Respond to a vendor audit letter · Respond to a mid-term price increase · Run diligence for an acquisition · Sunset a vendor cleanly · Answer a security questionnaire
Market and vendors	Watch a vendor's market moves · Prepare for a vendor meeting · Plan the year's vendor strategy · Run a proof of concept bake-off · Take a purchase request from intake to award
Team and workspace	Set up your workspace (owners) · Onboard a new team member · Take over a colleague's vendor portfolio · Open a workspace for a subsidiary · Connect your data and close the gaps · Bring the platform into email and chat
Working with the AI	Ask Vera well · Get unstuck

BENCHMARK A PROPOSAL END TO END

Get an analyst benchmark on a vendor quote or contract, from upload to a report you can forward.

For: anyone · **Time:** about 10 minutes of yours; the report follows per the analyst queue · **You need:** the quote or contract as PDF or DOCX, and one sentence on what you want answered

Steps

1. Open Proposals and choose **Upload a proposal**. Drop the file (up to 20 at a time, 40 MB each); vendor, value, term, and seat chips fill in as it is read.

2. Correct anything shaded gold or red, then add your benchmark focus: the question you actually want answered, for example "is this renewal uplift fair?" The one-click starter template matched to the category is a good skeleton.
3. Confirm. The proposal moves along the path Uploaded, In review, Benchmarked, Delivered, and every stage change reaches you by alert and email, including the expected delivery date when an analyst picks it up.
4. While you wait, get an immediate read on the same paper: an instant benchmark for the market position, or the AI redline scan for risky clauses.
5. When the report lands, read it the way it argues: the finding first (your discount, the peer range, the rank), then the evidence. Delivered reports explains every part, and the downloads (PDF, Word) live on the report itself.
6. About 30 days later, answer the one-touch **What happened?** follow-up. A won reduction, with its annual figure, files itself into your savings proof.

Forks

- **If the upload ends as Saved, not queued** — your plan's benchmark allowance is used up. The proposal page resolves it in place: an owner adds a benchmark credit ($5,000 per analyst benchmark, billed by invoice), or a teammate asks their owner in one click. See Proposals.
- **If the paper shows no list price** — the report will say the discount cannot be verified from the document, name exactly what to ask the vendor for, and price the sensitivity instead of substituting a different metric. That is the report working as designed, not a gap.
- **If you need the answer today** — do not wait for the queue: run Get an instant price read now and let the analyst report deepen it later.

Done when: the report is delivered, forwarded to whoever negotiates, and the 30-day outcome question is answered so the saving is on the record.

See also: Proposals · Delivered reports

GET AN INSTANT PRICE READ, NO UPLOAD

When you need a number in minutes rather than a full analyst report.

For: anyone · **Time:** 2 to 10 minutes · **You need:** the vendor name and roughly what you pay or were quoted; the vendor's paper makes it sharper but is optional

Steps

1. Open Instant benchmarks, type the vendor, and answer the questions. You get the three answers that matter: the discount you are being offered, what comparable customers get as a range with its midpoint, and where your deal ranks.
2. Have the vendor's own paper? Drop it anywhere on the page instead (Drop a quote, paid plans): the platform reads the deal off the document and returns a six page negotiation brief in a deal desk room,

usually in about three minutes.

3. For the fastest possible answer to "what is a good discount for this vendor," ask exactly that in the Ask the market bar: the quick ask pills carry the question, and the answer comes back grounded and cited.
4. To judge a specific document rather than a deal, drop it into the Verdict Bar (verdict and red flags in about a minute) or run Contract decode for the clause-by-clause read with paste-ready asks.

Forks

- **If the vendor has several benchmark tools** — the search bar lists them beneath your query; pick the product line that matches your deal. Matching vendor and product both is what makes the cohort honest.
- **If the range comes back directional** — too few comparable deals sit in your size band and region. Widen the region for the deeper cohort rather than forcing a thin one.
- **If you are on the trial** — instant benchmark runs are capped (see the plan table), so spend them on the vendor you are actually negotiating.

Done when: you can say, in one sentence, your discount, the peer range, and the rank, and you know which of the three the vendor would least like you to know.

See also: Instant benchmarks · Verdict Bar · Contract decode

TRACK AGREEMENTS AND NEVER MISS A NOTICE WINDOW

Set an agreement up once and let the platform watch the deadline for you.

For: anyone; the alert choices are per person · **Time:** a minute per agreement · **You need:** the vendor, the expiry date, the notice period, and ideally the annual value

Steps

1. Open Agreements and choose **Add agreement**: vendor, expiry date, and a 30, 60, or 90 day notice period. Add the annual value so the agreement also feeds your spend and portfolio analytics, and the uplift cap if the contract carries one.
2. Attach the signed paper, up to five files per vendor, for safe keeping.
3. Open Renewals to see every agreement sorted by expiry, next deadline on top, and the Renewal calendar for the year at a glance.
4. Decide what may interrupt you: renewal alerts are strictly opt-in and start silent. Press **Turn on money-critical alerts** under Notifications to switch the renewal set on in one click; teammates who turned them on are emailed when an agreement enters its notice window.

Forks

- **If an agreement has an expiry but no notice period** — it has no window to miss, so nothing can be watched. The notice period is the field that buys you time; hunt it down in the contract (the termination or renewal clause) rather than leaving it blank.

- **If you are past 40 tracked vendor agreements** — you have outgrown hand entry. Move to the folder or estate door in Migration & data-in.
- **If the vendor asks for more than your uplift cap at renewal** — the ledger masthead already counts that exposure, and Anomalies flags the breach.

Done when: the next four quarters are visible on the ledger, every agreement that matters carries both dates, and the notice-window email has a recipient.

See also: Agreements · Renewals · Vendor calendar

PLAN AND RUN A RENEWAL NEGOTIATION

Turn an upcoming expiry into a worked deal with a plan, a mandate, and a record of what was agreed.

For: whoever owns the deal · **Time:** an hour to stand the deal up; the negotiation itself runs weeks · **You need:** the renewal on your calendar and a market read (see the fork below if you lack one)

Steps

1. From Renewals, open the renewal you want to work.
2. Build the plan in the Renewal playbook: the step-by-step negotiation plan, the model, and the exec report. Turn on the agent if you want it to prep and remind.
3. Set your leverage before the first call: your positions in the Clause library, and the Fair Renewal Rider attached so every renewal carries your standard protections.
4. Open a tracked Negotiation and set the mandate: target, walk-away, and approvals if your organization runs a sign-off chain.
5. Work the rounds from inside the record: paste vendor emails for tactic analysis, log every give and get in the concession ledger, and capture each vendor promise under Commitments so it reaches the contract rather than dying in an inbox.
6. Close it properly: run the post-mortem, and record the outcome so the saving lands in Savings proof.

Forks

- **If you have no market read yet** — stop and get one first (Get an instant price read). A mandate set without the peer range is a guess with a signature line.
- **If the vendor's first response concedes little** — escalate the pressure honestly: price the credible alternative with Build competitive scenarios, check the timing against their quarter on the Vendor calendar, and consider handing the grind to a Vera-run case.
- **If deal approvals are enabled** — request sign-off from the deal's page before signature; approvers clear it from their Approvals queue.

Done when: the mandate is set, the ledger reflects every round, the commitments list matches the final paper, and the post-mortem is written.

See also: Renewal playbook · Negotiations · Opportunities

PROVE AND SHARE YOUR SAVINGS

Turn a claimed saving into one proven against real invoices, then a shareable badge.

For: whoever answers to finance · **Time:** minutes per saving; the proof accrues as invoices arrive · **You need:** a closed outcome (a negotiated reduction, a corrected invoice) and the invoices that followed

Steps

1. Open Savings proof: each negotiated saving is reconciled against the invoices that followed, so a claim becomes proven or disproven rather than asserted.
2. Feed it the outcomes it needs: answer the 30-day **What happened?** question after each delivered benchmark, and record negotiation outcomes when deals close.
3. Once a saving is invoice-verified, go to Verified wins, mint the public, revocable badge, and watch your Deal Score against the calibrated corridors.
4. Owners can list the organization on the opt-in board under whatever name you choose.

Forks

- **If the contract has no stored annual value** — the follow-up card asks for the pre-reduction baseline, because without it the saving cannot be reconciled against invoices. Take the minute; an unreconciled claim is the kind finance discounts.
- **If later invoices contradict the claim** — the row says disproven, and that is the feature: your ledger only carries numbers that survived contact with billing. Investigate on Invoices before anyone forwards the figure.

Done when: the saving reads proven, the badge exists if you want one, and the figure you would show a CFO traces to invoices.

See also: Savings proof · Verified wins · Invoices

PREPARE FOR A VENDOR MEETING

Walk into the call with the numbers already in hand.

For: anyone taking a vendor call · **Time:** 15 minutes, or zero once the calendar is connected · **You need:** the meeting on your calendar and the vendor's paper on file

Steps

1. For a single deal, drop the contract into Vera Deal-Room: one upload assembles the benchmark, the decoded asks, an analyst brief, the strategy playbook, live vendor intel, and the credible alternative priced.
2. To have it done for you automatically, connect your calendar so Meeting war room preps a pack before each vendor meeting.

3. Time the ask: the Vendor calendar shows the vendor's fiscal quarter end, when discounts run deepest.
4. Skim the vendor's last 90 days on Terms watch and the Weekly report archive, so nothing they announce in the call is news to you.

Forks

- **If the meeting is in an hour** — skip the full pack: run a quick instant benchmark for the range and the targets, and open the vendor's page under Vendors for your own position. Ten minutes buys you the two numbers that carry a first call.
- **If no calendar is connected** — meeting packs cannot appear on their own; Troubleshooting covers the connection.

Done when: you hold your position, their market moves, and the range, and the meeting's first surprise belongs to the vendor.

See also: Vera Deal-Room · Meeting war room

SET UP YOUR WORKSPACE (OWNERS)

The one-time setup an org owner does so the whole team gets value. These steps need owner access.

For: org owners · **Time:** about an hour for the core · **You need:** the teammate list with roles, your logo, and your plan facts

Steps

1. Work the owner's handbook top to bottom; it is this recipe at full depth. The condensed pass:
2. Invite your team under Members and set each person's role; create Teams for confidential access before the first sensitive document needs one.
3. Connect your systems under Integrations: contract intake, SAM and ITAM sources, Slack and Teams.
4. Set the security defaults: org-wide two-factor under Security, retention and consent under Data controls, and the Single sign on decision if you are on Enterprise.
5. Set Branding so exported reports carry your identity, pick the Display currency, and set your own Notifications; everyone else chooses their own.
6. Load the sample portfolio from the dashboard to explore, then remove it before entering real data.

Forks

- **If the team manages access through an identity provider** — read the SSO decision before inviting everyone by hand; SCIM may do the inviting for you.
- **If you inherited a half-set-up workspace** — start from Settings home: the vitals and Worth doing now list exactly what is missing.

Done when: the Settings home vitals read clean: two-factor on, people seated, systems connected, nothing urgent under Worth doing now.

See also: The owner's handbook · Settings & account

RESPOND TO A VENDOR AUDIT LETTER

An audit letter rewards a considered first move and punishes a fast one. Slow the clock down and answer from your contracts, not from memory.

For: IT asset management and whoever owns the vendor · **Time:** an hour on day one; the defense runs as long as the audit · **You need:** the letter itself, and your entitlement data as connected as it gets

Steps

1. The day the letter lands, drop it into the Audit letter decoder: who is actually auditing you, the deadline as a business-day countdown, what is demanded versus merely implied, and the three things not to do first.
2. Do not reply to the auditor yet, and do not run any vendor-supplied measurement tool until the plan says so.
3. Follow the decoder into the pre-filled defense playbook, and open the audit-defense analyses from the proposal record's right rail if the vendor's paper is on file.
4. Check your position on Data health: the audit exposure estimate tells you how defensible your entitlement picture is before you engage.
5. Answer on your timeline, from your records, with every claim the auditor makes checked against your entitlements rather than accepted.

Forks

- **If the exposure is large** — bring in a vetted specialist from the Expert Bench or book Advisory before the first substantive reply. The letter cost the vendor nothing; your reply can cost you leverage.
- **If the audit arrives near a renewal** — that is rarely a coincidence. Read the vendor's playbook in Vendor playbooks and treat audit resolution and renewal as one negotiation, not two.

Done when: the reply deadline is on your calendar with margin, the response plan exists in writing, and nobody has run a vendor script on your estate.

See also: Audit letter decoder · Expert Bench

DISPUTE AN INVOICE THAT DOES NOT MATCH YOUR CONTRACT

Turn a suspicion of overbilling into a documented, dollars-first dispute.

For: finance and vendor owners · **Time:** minutes per invoice once forwarding is set up · **You need:** the invoice, and the contract terms on file to check it against

Steps

1. Get the invoice in: upload it on Invoices, or set up the forward-in address so invoices process automatically from now on.

2. Read the line verdicts: price above contract, billed over entitlement, off-contract item, or uplift-cap breach, each with the dollars at stake.
3. Cross-check Billing watch: if the vendor moved its price list, the corrective draft is already attached to the finding.
4. Send the correction, then watch the following invoices: a correction only counts once later invoices confirm it, and the confirmed saving lands on Savings proof.

Forks

- **If the entitlement figures came from a connector** — they are held out of alerting until a human confirms the contracted numbers (the rows marked "confirm" on Billing watch). Confirm them once and the daily check has teeth.
- **If the vendor disputes your dispute** — the case is your contract language. Pull the exact clause with Document search and let the paper argue.

Done when: the correction is sent, the following invoice confirms it, and the saving reads proven rather than claimed.

See also: Invoices · Billing watch

HAND YOUR RENEWAL CALENDAR TO VERA

Let the AI run covered renewals under mandates you approve, with fees only on invoice-verified savings.

For: org owners activate; whoever owns renewals supervises · **Time:** an hour to activate; minutes a week to supervise · **You need:** renewals on the calendar with real dates, and an owner willing to sign the fee agreement

Steps

1. Read the offer on Managed renewals: Vera discloses itself as an AI, you approve the moments that matter, and claimed figures never bill.
2. An org owner activates the program, signs the fee agreement, and sets the standing mandate defaults.
3. Approve each deal's mandate before anything reaches a vendor; the standing defaults never skip this step.
4. Clear the command queue as Vera works: every interruption is typed (a mandate, a formal offer, a parked decision) and opens the deal it waits on.
5. Track the net position on the pairing meter, savings banked against fees paid, with every figure traced to a vendor invoice.

Forks

- **If a run pauses and will not continue** — that is the guardian doing its job: something in the mandate says a person decides here. The Control room shows exactly what is waiting on you and why.
- **If you want automation without delegation** — stop short of managed renewals: the renewal playbook's agent preps and reminds without ever contacting a vendor.

Done when: mandates are approved, the Waiting on you queue is routinely empty, and the net position line is one you would show finance.

See also: Managed renewals · Savings proof

RUN DILIGENCE FOR AN ACQUISITION

Deal-grade contract review with the sensitive documents locked down.

For: legal and the deal team · **Time:** an hour to set up; the register works through in days, mostly the platform's · **You need:** the target's contracts, and the short list of people allowed to see them

Steps

1. Create a restricted Vault project and grant only the deal team access; documents in a restricted project stay hidden from everyone else.
2. Add the target's contracts to the project, in bulk if needed.
3. Run Diligence scoped to that project: every contract is read for change of control, assignment, exclusivity, termination, liability, and data terms.
4. Work the risk register from the highest severity down; each finding cites the clause it rests on.
5. Read the analyst report for what needs consent or renegotiation before the deal, and what is clean.

Forks

- **If the deal team changes mid-process** — grant and revoke through a team rather than person by person: membership is the access, and offboarding is one removal.
- **If a finding is deal-critical** — do not settle it on a summary. Open the cited clause in the document itself, and put the specific question to the whole estate with Review tables if you need it answered across every contract.

Done when: the register is worked to the bottom, every consent-required contract has an owner and a plan, and the report is in the data room.

See also: Vault · Diligence

TEAM UP WITH OTHER BUYERS ON A RENEWAL

When several buyers renew the same vendor in the same window, one term sheet beats four separate negotiations.

For: whoever owns the vendor · **Time:** 30 minutes to evaluate and join; the bloc runs on the window's clock · **You need:** the renewal on your calendar, so the matchmaking has a date to line up

Steps

1. Check the matchmaking strip on Buying blocs: it shows where your renewal calendar lines up with other buyers of the same vendor.
2. Open an invitation to read the bloc's shape: buyers, combined ACV, window, and term sheet version. You stay anonymous until mutual opt-in.
3. Join the bloc to enter its room, and negotiate on the shared Fair Renewal Rider term sheet.
4. For a private one-to-one instead, use Buyer intros to talk to another buyer who runs the same vendor.

Forks

- **If no bloc matches your window** — do not force one. Run the renewal solo on the standard play (Prepare a renewal 90 days out) and let the matchmaking keep watching; windows align more often than calendars suggest.
- **If your terms diverge from the bloc's sheet** — the shared sheet is the point; a bloc where everyone rides their own paper is four negotiations wearing one name. Take your special terms to a private intro instead.

Done when: you are inside a bloc negotiating one sheet, or you have deliberately chosen the solo play with the market read in hand.

See also: Buying blocs · Fair Renewal Rider

CHOOSE BETWEEN COMPETING QUOTES

Two or three quotes for the same need, one defensible decision.

For: whoever recommends the award · **Time:** about 30 minutes · **You need:** the quotes as PDF or DOCX, ideally covering the same scope

Steps

1. Drop the quotes into Quote Face-Off: equivalent SKUs are aligned line by line, every line is benchmarked, and the best mix across all papers is priced as the Franken-quote.
2. Read the analyst memo: which quote to take, which terms to pull from the losers, and the counter to send.
3. For one quick quote on its own, Snap a quote gives the line-by-line check and the counter in about a minute.
4. Running a formal process? Work it as a Sourcing project through to the award decision.

Forks

- **If the quotes cover different scopes** — align the scope before trusting the comparison: the Proposal comparison read calls out what each paper includes that the others do not, which is often the real decision.
- **If the winning quote still benchmarks poorly** — winning the face-off is not the same as being fairly priced. Take the winner into Get an instant price read and counter before you award.

Done when: the award has a written why, the losing quotes' best terms are in your counter, and the winner has been benchmarked, not just compared.

See also: Quote Face-Off · Sourcing

SET YOUR CLAUSE STANDARD AND SWEEP THE ESTATE

Decide once what your organization signs, then measure every contract against it.

For: legal, with owner support for the settings side · **Time:** two hours to set the standard; the sweep is the platform's work · **You need:** agreement on your non-negotiables, even a first draft of it

Steps

1. Build your positions in the Clause library: the preferred language, the fallback, and the walk-away for each clause that matters, and mark the non-negotiables as must-have.
2. Run Coverage to sweep the must-have template across every contract: Present, Deviates, Missing, or Unclear per cell.
3. Chase the red Missing cells first: agreements without a protection you require.
4. Ask bespoke questions across the estate with Review tables: one question per column, a quote-backed answer per document.
5. From now on, every AI review argues for your standard instead of a generic one, inserting your pre-approved language wherever a contract deviates.

Forks

- **If legal cannot yet agree the standard** — start with the security side: the Security baseline ships a default and rulepacks you can adopt today, and the clause standard can follow it.
- **If a Missing cell belongs to a renewal inside 90 days** — that is not a sweep finding, it is a negotiation item: move it onto the deal's ask list now, while you still have the leverage of an open renewal.

Done when: the must-haves are marked, the sweep has run, and every red cell has either a fix in flight or a written acceptance.

See also: Clause library · Coverage

CONNECT YOUR DATA AND CLOSE THE GAPS

The platform gets sharper with every source you connect. This is the highest-leverage hour an owner spends.

For: org owners · **Time:** an hour for the first pass; each later gap is minutes · **You need:** credentials or admin access for the systems you are connecting

Steps

1. Open Data health: the completeness score, and every gap ranked by the feature filling it unlocks.
2. Connect the systems that know what you own under Integrations: SAM and ITAM sources, procurement, contract intake, and the invoice forward-in address.
3. Add annual values to your Agreements so Portfolio, Spend, and the Estate map fill in.
4. Index your documents on Document search, then harvest your paid-for rights into the Rights register.
5. Come back to Data health and work the next gap; the plan reranks as you close them.

Forks

- **If a connector's first sync looks thin** — check Recent activity on the Integrations page before debugging the source; the last sync result per connection is stated there.
- **If you cannot get credentials for a source** — the CSV importers cover most of the same ground manually; a monthly export beats an empty picture while access is negotiated.

Done when: the completeness score has moved, the top-ranked gap has an owner, and license data is flowing rather than uploaded once.

See also: Data health · Integrations

TAKE A PURCHASE REQUEST FROM INTAKE TO AWARD

One front door for new spend, so nothing is bought twice.

For: procurement; requesters just email or file a form · **Time:** minutes per request to triage; the sourcing project runs on its own clock · **You need:** the intake address set up by an owner, and the request itself

Steps

1. Raise the request on Requests, or let teammates email requests straight in through the intake address an owner set up.
2. Read the AI triage verdict: net-new, a renewal of something you own, a duplicate of existing capability, or over the approval threshold.
3. Route it onward: a Sourcing project for net-new spend, a renewal for something you own, or an approval.
4. In sourcing, run requirements, suppliers, and bids side by side through to the award decision.

Forks

- **If triage says duplicate** — that is the recipe paying for itself: point the requester at the existing capability, and if they insist the existing tool falls short, that is a consolidation conversation, not a purchase.
- **If the request is really a renewal** — do not source it; work it as a renewal with the leverage that brings (Prepare a renewal 90 days out).

Done when: the request is routed, the decision has a written trail, and nothing was bought that the estate already owned.

See also: Requests · Sourcing

BUILD THE BOARD PACK

The quarterly leadership story, told from your live data instead of a hand-built deck.

For: whoever presents to leadership · **Time:** about an hour, mostly reading · **You need:** annual values on your agreements, and savings recorded as they closed

Steps

1. Start from Portfolio for the spend and concentration picture.
2. Export the Estate map as a slide: every vendor sized by spend and colored by market position.
3. Generate the annual Strategy document: category strategies, vendor posture, and the year's negotiation calendar.
4. Add the proof: the savings ledger from Savings proof exports as CSV for finance.
5. For the monthly cadence in between, the Monthly briefing is the five minute spoken version.

Forks

- **If the map looks empty** — the pack is only as good as the values on the agreements; spend an hour on Connect your data and close the gaps before presenting thin charts as the estate.
- **If leadership wants one page, not a pack** — build the CFO one-pager instead and keep this recipe for the quarterly.

Done when: the pack tells one story from live data, every number in it traces to a record, and next quarter's version is a refresh, not a rebuild.

See also: Estate map · Strategy

WATCH A VENDOR'S MARKET MOVES

Know what a vendor did to its pricing before they tell you at renewal.

For: whoever owns the vendor · **Time:** 10 minutes to set up; the watching is the platform's job · **You need:** the vendor's name; holding a contract with them makes the alerts sharper

Steps

1. Follow the vendor on Terms watch and turn on the price-change digest; every published price-list change arrives diffed and priced at your seat counts.
2. Subscribe to the Monday Weekly report so licensing, audit, and M&A news on your vendors leads your copy.
3. Before a renewal, read the vendor's full history in the archives, its grade on Vendor scorecards, and the latest Research on it.
4. Check Billing watch to see whether your invoices already reflect a move.

Forks

- **If the vendor is not on the tracked roster** — request it from the card on the Price list check tab; the list joins the roster with its own weekly verification sweep.
- **If a price rise lands mid-term** — that is its own recipe: Respond to a mid-term price increase, starting from the case file the change itself opens.

Done when: the digest is on, the vendor's strip is quiet or understood, and no vendor announcement at renewal is news to you.

See also: Terms watch · Weekly report

RIGHT-SIZE LICENSES BEFORE A RENEWAL OR TRUE-UP

Walk into the true-up with the surplus already found, instead of paying for shelfware another year.

For: IT asset management · **Time:** 30 minutes per estate once the data source is connected · **You need:** usage or deployment data, connected or exported

Steps

1. Run the optimizer that matches the estate: M365 optimization for Microsoft seats, Cloud cost optimizer for hyperscaler spend, VCF licensing for VMware cores, or the Claude and OpenAI optimizers for AI commitments.
2. Review the inactive and oversized assignments the tool surfaces, and the per-user reassignment plan where the tool builds one.
3. Fold the findings into the renewal: the matching vendor playbook turns them into the negotiation case.
4. For the budget-wide view across every vendor, run Spend optimization.

Forks

- **If you have no usage data yet** — the tools are free to explore with sample data on any plan; running them on your own exports is what needs a Professional or Enterprise plan (see Chapter 10). Connect a SAM source and the picture stays current instead of exported.
- **If the surplus is real but the renewal is far away** — bank the finding: the reassignment plan saves money now, and the negotiation case waits for the window. Prepare a renewal 12 months out is the long game.

Done when: the surplus is quantified, the reassignment plan is running, and the renewal case opens with the shelfware already counted.

See also: Spend optimization · Vendor playbooks

ONBOARD A NEW TEAM MEMBER

From invite to their first useful result in one day, without you standing over their shoulder.

For: org owners send the invite; anyone can be the buddy · **Time:** 5 minutes of yours; their ramp is theirs · **You need:** their email and their role

Steps

1. Invite them under Members with the right role; the invitation expires after 7 days, so send it when they can act on it.
2. Add them to the Teams their access requires; team membership is the access itself, so this one step covers every confidential grant the team holds.
3. Point them at two links and nothing else: their first 30 days for the ramp, and their role's page for the five screens that matter to their seat.
4. Let the platform do the rest: their first sign-in walks them into a first verdict, the getting started course matches a track to their goal, and the desk survey arranges the rail around their answers.
5. A week later, glance at Team activity: if they have not started, the numbers say so before the habit hardens.

Forks

- **If the invite never arrived** — check the pending invitations list for a bounce and resend; Troubleshooting covers the delivery cases.
- **If they are drowning in alerts** — they are not: everything renewal-shaped starts silent. If anything, the common miss is the opposite; remind them alerts are opt-in under Notifications.

Done when: they have run a first verdict on real paper and can name the three keys (Ctrl/⌘ K, Ctrl/⌘ J, the Guide door) without looking.

See also: Members · Your first 30 days

REPORT SAVINGS AT QUARTER END

The recurring finance ritual: the quarter's savings, invoice-verified, in the format finance actually wants.

For: whoever answers to finance · **Time:** 30 minutes, four times a year · **You need:** outcomes recorded as deals closed (the quarter-end pain is proportional to what was skipped in-quarter)

Steps

1. Open Savings proof and read the quarter's rows: proven, pending confirmation, and disproven. The proven set is your number.
2. Chase the pending rows before you report: a correction or reduction still waiting on its confirming invoice is real money in limbo, not a smaller number.
3. Export the ledger as CSV for finance, and pull the badges from Verified wins for anything you want to celebrate publicly.

4. Put the number in context with one chart: Spend for the denominator, and the Monthly briefing if leadership prefers it spoken.
5. Book the next quarter's version now: 30 minutes on the calendar, because this recipe is only painless when outcomes were recorded as they happened.

Forks

- **If finance disputes a figure** — do not argue from memory: the Figures ledger traces every cited number to its source fact and date, and the saving's own row names the invoices that confirmed it.
- **If the quarter looks empty but the work happened** — outcomes were not recorded. Sweep the quarter's closed negotiations and delivered reports for unanswered What happened? questions before declaring a zero.

Done when: the CSV is with finance, every number in it reads proven, and nothing pending is silently missing from the story.

See also: Savings proof · Verified wins

PREPARE A RENEWAL 12 MONTHS OUT

The long game: at a year out, you have every lever, and the work is cheap.

For: whoever owns the vendor · **Time:** two hours now; it repays itself at the table · **You need:** the renewal date on the calendar, a year away

Steps

1. Benchmark now, not at renewal: an instant benchmark or an analyst report a year out tells you whether this is a repair job or a defense job, while there is still time for either.
2. Put the vendor on watch: Terms watch and the Weekly report. A year of diffs and news is a negotiation file that writes itself.
3. Start the usage clock: run the matching optimizer now and again at 6 months, so the true usage trend, not a snapshot, walks into the room with you.
4. Check the paper early: what does the contract say about uplift caps, notice, and renewal mechanics? Fix gaps at the midpoint amendment if one arises, and log your positions in the Clause library.
5. Mark the geometry on the Vendor calendar: their fiscal quarter ends, your notice deadline, and the date you want the first conversation, which is earlier than instinct says.
6. Budget honestly: the uplift the vendor will ask for is visible in their published moves; give finance the range now instead of a surprise later.

Forks

- **If the benchmark says you are at or above market** — the year is for defense: lock the rate, cap the uplift, and spend your effort on terms rather than price. A fine deal defended is a win.

- **If the vendor is consolidating or was acquired** — the calculus changes; read the Research on them and consider whether the credible alternative needs a real proof of concept (Run a proof of concept bake-off) this year.

Done when: the watch is on, the baseline benchmark exists, the calendar geometry is marked, and finance holds your range, not the vendor's.

See also: Prepare a renewal 90 days out · Vendor calendar

PREPARE A RENEWAL 90 DAYS OUT

The standard play: enough time for every move, none to waste.

For: whoever owns the deal · **Time:** the stand-up is an hour; then it is a running deal · **You need:** the renewal 90 days out with its notice period known

Steps

1. Confirm the geometry first: the expiry, the notice deadline, and the days between today and each, on Renewals. The notice deadline, not the expiry, is your real clock.
2. Get the market read this week: instant benchmark at minimum, the deal desk brief from the vendor's own renewal paper if you have it.
3. Stand the deal up per Plan and run a renewal negotiation: playbook, clause positions, rider, mandate.
4. Book the vendor conversation against their calendar, not yours: the Vendor calendar shows where their quarter pressure sits inside your window.
5. Open with your ask, grounded in the range, before they open with theirs.

Forks

- **If the vendor already sent the renewal quote** — even better: drop their paper into Drop a quote and negotiate from the six page brief it returns rather than from their framing.
- **If 90 days shrinks under you** — deadlines slip into the two-week regime faster than teams expect; the moment the notice deadline is inside a fortnight, switch to Rescue a renewal 2 weeks out.

Done when: the mandate is set, the first conversation is booked on favorable geometry, and the notice deadline has margin around it.

See also: Plan and run a renewal negotiation · Renewal playbook

RESCUE A RENEWAL 2 WEEKS OUT

Triage: with a fortnight left, sequence beats thoroughness.

For: whoever just found the deadline · **Time:** today, one focused hour; then daily attention · **You need:** the contract, and honesty about which deadline you are actually against

Steps

1. Establish the real clock in the first ten minutes: is the **notice deadline** still ahead of you, or already behind you? Check the agreement on Renewals. Everything else depends on this answer.
2. If notice is still open and you might want out or leverage: file the protective notice conversation now, before any price talk. An open exit is your only real lever at this range.
3. Get the number today, not this week: an instant benchmark gives you the range in minutes; drop the vendor's renewal paper into the Verdict Bar for the red flags.
4. Make one ask, not five: the range's strong ask on price, plus the Fair Renewal Rider attached so the structural protections ride along without separate negotiation.
5. If the timeline is genuinely impossible, negotiate time itself: a short extension at current terms is a legitimate ask, and better than signing a bad year under the gun.

Forks

- **If the notice deadline has already passed** — stop mourning the exit; it is gone. Your play is price and terms inside the renewal: the uplift cap if the contract carries one, the market range, and next year's geometry fixed so this never repeats.
- **If the paper auto-renews** — confirm what it auto-renews INTO (term and price) before deciding whether the rescue is even needed; sometimes the status quo year is the best available outcome, bought with zero leverage spent.

Done when: the deal is signed or extended on terms you chose knowingly, and next year's renewal is on the calendar with its notice period, 90 days of runway, and an owner.

See also: Track agreements and never miss a notice window · Instant benchmarks

CONSOLIDATE DUPLICATE TOOLS

Two tools doing one job is a negotiation you have not run yet.

For: procurement and IT asset management together · **Time:** a half day of analysis; the consolidation runs a quarter · **You need:** the estate on file, values on agreements

Steps

1. Find the overlaps: Portfolio and the Estate map show category crowding, and Document clusters groups similar paper you may not have known was similar.
2. Quantify each overlap: what does each tool cost annually, who actually uses it, and what would the survivor cost at combined volume? The Workload comparator prices workloads across platforms where that is the question.
3. Choose the keeper on evidence, then benchmark the keeper at the combined size: consolidation is a volume story, and the peer range at the bigger deal size is your new target.
4. Run the two moves in parallel: sunset the loser cleanly, and renegotiate the keeper with the volume as leverage (Plan and run a renewal negotiation).

5. Record both outcomes in Savings proof: the retired spend and the improved rate are separate, provable numbers.

Forks

- **If users defend the duplicate** — the workloads may genuinely differ; that is what the comparator is for. Consolidating two real needs into one wrong tool costs more than the duplicate did.
- **If both contracts renew far apart** — sequence the sunset to the natural expiry rather than paying termination costs; the consolidation can take a year and still be the right move.

Done when: one tool remains per job, the survivor is priced at its true volume, and both savings are on the ledger.

See also: Estate map · Sunset a vendor cleanly

RESPOND TO A MID-TERM PRICE INCREASE

The vendor moved the price list mid-term. What you do next depends on paper you already hold.

For: whoever owns the vendor · **Time:** an hour from alert to sent response · **You need:** the change alert, and your contract on file

Steps

1. Open the change on Terms watch and open its case file: the change with both published entries side by side, the impact brief priced at your own seat counts, and a response sheet weighing your three options: grandfathering, renewal timing, or acceptance.
2. Check what your contract already says: an uplift cap or price-protection clause may make the increase simply inapplicable to you until renewal. The agreement's terms on Agreements and a Contract decode settle it.
3. Verify your invoices have not moved early: Billing watch checks billed against entitled daily, and an increase applied before it contractually can be is a dispute, not a negotiation.
4. Send the response from the case: the grandfathering draft opens in your own inbox (nothing sends from the platform), asking the vendor to hold your terms.
5. Whatever the outcome, reprice your renewal math: the new list changes next year's baseline, and your renewal preparation should absorb it now.

Forks

- **If your contract is silent on mid-term increases** — the response sheet's honest read applies: your leverage is the renewal, so bank the grievance and spend it there, in writing.
- **If the increase already hit an invoice** — jump recipes: Dispute an invoice that does not match your contract, with the price-list diff as your exhibit.

Done when: you know whether the increase lawfully touches you, the response is sent, and the renewal model reflects the new list.

See also: Terms watch · Billing watch

RUN A PROOF OF CONCEPT BAKE-OFF

Two vendors, one structured trial, and a decision that survives scrutiny.

For: procurement running it; the evaluating team scoring it · **Time:** the setup is an afternoon; the PoC runs weeks by design · **You need:** written success criteria before any vendor is invited

Steps

1. Open a Sourcing project and write the success criteria first: the requirements, weighted, agreed by the evaluating team before a vendor demo can charm anyone.
2. Invite the vendors and run the trial period with scores kept against those criteria in the project, not in inboxes.
3. Price the endgame before the end: build the competitive scenarios so each vendor's likely pricing behavior under competition is modeled while both still believe they can lose.
4. When the bids land, run them through Quote Face-Off: lines aligned, every line benchmarked, the best mix priced.
5. Award in the project with the written why, and take the loser's best terms into the winner's final round.

Forks

- **If one vendor knows they are the incumbent favorite** — the bake-off's value collapses; competition only prices honestly when it is believed. Keep the scoring blind to preference and the timeline symmetric.
- **If the PoC proves neither fits** — that is a successful PoC. The requirements were wrong or the market is not there; revisit the request rather than awarding the least bad option.

Done when: the award is scored against pre-agreed criteria, the price reflects real competition, and the file would survive an auditor or a losing vendor's protest.

See also: Sourcing · Quote Face-Off

TAKE OVER A COLLEAGUE'S VENDOR PORTFOLIO

Inherit a book of vendors without inheriting its surprises.

For: the new owner of the portfolio · **Time:** half a day for a real handover · **You need:** the list of vendors now yours; the departing colleague for an hour if you can get them

Steps

1. Walk the book vendor by vendor on Vendors: each page holds the contracts, benchmarks, market intel, and notes in one place, which is most of the handover document people never write.

2. Audit the clocks first: on Renewals, check every inherited agreement for expiry and notice dates, and rank by what shuts soonest. Anything inside 90 days jumps the queue (Prepare a renewal 90 days out).
3. Read the live deals: open Negotiations and Cases for anything in flight, the mandate, the ledger, and what was promised under Commitments.
4. Point the alerts at yourself: alert choices are per person, so the predecessor's subscriptions did not transfer. Turn on money-critical alerts under Notifications and follow your new vendors on Terms watch.
5. Ask Vera for the briefing: "catch me up on our Microsoft position" against each major vendor is the fastest first pass, cited, from your own records.

Forks

- **If tasks and approvals were assigned to the predecessor** — reassign them: Tasks for the open items, and the owner should update any approver seat they held under Deal approvals.
- **If the predecessor already left** — the records are the handover: the concession ledgers, the audit trail of what was agreed, and the vendor pages' notes were built for exactly this reading.

Done when: every inherited deadline has you as its owner, the in-flight deals have continuity, and the alerts reach you, not a departed inbox.

See also: Vendors · When someone leaves

PREPARE THE CFO ONE-PAGER

One page, five numbers, every one of them defensible.

For: whoever briefs the CFO · **Time:** 30 minutes · **You need:** values on agreements and savings recorded; without those, do Connect your data first

Steps

1. Pick the five numbers that carry the story: annual estate spend (Spend), renewals closing inside 90 days with their value (Renewals), proven savings this year (Savings proof), the estate's market position (Portfolio), and the one risk worth a sentence (Anomalies or the audit exposure on Data health).
2. Generate the document through the Reporting studio rather than hand-building a slide: the output lands in the house executive-brief format, and Branding co-brands it with your logo.
3. Check every figure's provenance before it ships: the Figures ledger traces each cited number to its source fact and date, which is the difference between a brief and a claim.
4. Send it yourself, from your own inbox, with the one-sentence ask you want the CFO to grant.

Forks

- **If the CFO asks for the detail behind a number** — that is the design working: the ledger row, the savings row's confirming invoices, or the benchmark's evidence section answers without a scramble.
- **If this becomes a monthly ask** — stop hand-running it: the Monthly briefing covers the cadence, and a workflow agent can deliver the recurring version.

Done when: the page is sent, every number on it traces, and the ask was explicit.

See also: Build the board pack · Figures ledger

OPEN A WORKSPACE FOR A SUBSIDIARY

A second legal entity gets its own sealed workspace, set up like the first one, faster.

For: org owners · **Time:** the request is minutes; the setup is the first-hour recipe again · **You need:** the entity's name and who will own its workspace

Steps

1. Ask for the workspace rather than improvising one: message our team from Chat with the entity and the intended owner. Workspaces are sealed tenants, and if your domain has capture on, a colleague self-signing up routes to your existing workspace admin instead of creating a new tenant.
2. When it exists, switch between entities with the org chip in the top bar; your account can belong to more than one workspace, and each stays sealed from the others.
3. Run the owner's first hour in the new workspace: people, teams, notification stance, security defaults. It goes faster the second time, and nothing carries over automatically, by design.
4. Set the Display currency to the entity's own currency if it differs; benchmarks stay in USD everywhere so market comparisons hold across your workspaces.
5. Bring in that entity's estate through the intake doors, scoped to its own paper.

Forks

- **If you are consolidating entities rather than adding one** — the question inverts: one workspace with Teams separating access may serve better than two tenants. Talk it through with our team before migrating paper twice.
- **If the same person owns both workspaces** — their alerts and personal settings are per account, not per workspace; they choose notification stances in each.

Done when: the new workspace passes the same first-hour finish line as your first one did, with its own estate inside.

See also: The owner's handbook · Migration & data-in

ANSWER A SECURITY QUESTIONNAIRE

Security reviews reward the team that answers from live state instead of last year's document.

For: whoever fields the questionnaire, with an owner for the exports · **Time:** an hour for a standard questionnaire · **You need:** the questionnaire, and owner access for the audit and review exports

Steps

1. Start from the Trust center: tenant isolation, encryption and access, audit and deletion, and the data journey diagram, written to be shared with reviewers directly.
2. Answer the "prove it" questions from live state, not prose: Security posture shows two-factor coverage, sign-on method, API credentials, and the audit stream as they stand today.
3. Export the evidence the reviewer wants: the audit log as CSV (or offer the SIEM webhook for continuous streaming), completed access reviews as CSV, and deletion certificates from Data controls.
4. Answer the data-location question precisely: Data controls states where data physically lives, as a read-only fact, currently the United States.
5. For the AI questions, the Trust center's "kept yours" chapter covers how the AI handles your data; answer from it rather than paraphrasing from memory.

Forks

- **If the questionnaire demands a control the platform states as roadmap** — say so plainly with the public claim ladder as the reference; a precise "not yet, tracked publicly" outlasts an optimistic yes.
- **If the reviewer wants a live walkthrough** — share the Trust center page itself and walk Security posture on a call; live state is the demonstration.

Done when: every answer either quotes live state or attaches an export, and nothing in the questionnaire was answered from hope.

See also: Trust center · Audit log

GET A SECOND OPINION BEFORE YOU SIGN

The deal is agreed, the paper is drafted, and one hour of checking is still cheap insurance.

For: whoever holds the pen · **Time:** about an hour · **You need:** the final draft as PDF or DOCX

Steps

1. Run Second opinion on the draft: the market read on the finished deal, while there is still a signature between you and it.
2. Run Contract decode for the clause-by-clause read: risks in plain English and asks with paste-ready language for anything worth one last round.
3. Check the draft against your own standard: your Clause library positions and must-haves, so the deviations are knowing ones.
4. Verify the commitments made it in: everything the vendor promised during the negotiation (your Negotiation record's Commitments list) should appear in the paper, because a promise not in the contract is a memory.
5. If the stakes warrant a human, book Advisory or an Expert Bench specialist for the final read.

Forks

- **If the checks surface something real** — reopen narrowly: one precise ask with the clause language attached, not a renegotiation. Late-round precision is respected; late-round breadth burns goodwill.
- **If everything comes back clean** — sign it and say so in the record: a deal verified as fine is a result, and the post-mortem is short.

Done when: the draft has survived the market read, the clause read, your standard, and the commitments check, and the signature is an informed one.

See also: Second opinion · Contract decode

TURN A DELIVERED REPORT INTO THE VENDOR CONVERSATION

A benchmark that stays a PDF saved nobody money. This is the handoff from finding to negotiation.

For: whoever takes the report into the deal · **Time:** 30 minutes from report to sent email · **You need:** a delivered report (analyst or deal desk brief)

Steps

1. Read the report the way it argues: the three answers first (your discount, the peer range, the rank), then the levers and the order to play them, then the talking points, two of which carry lines you can say out loud.
2. Use **Send it on** beneath the report: the follow-up email is drafted for you, opens in your own inbox, and nothing sends from the platform.
3. Open the deal as tracked work: from a deal desk room, **Open a negotiation case** carries the position into a timeline with owners; from an analyst report, stand it up per Plan and run a renewal negotiation.
4. Commission the follow-on documents you need from the commission rail under the verdict, and manage exports and sharing from the deliverable toolbar.
5. After the conversation, log what the vendor said in the negotiation record, and let the 30-day outcome question close the loop.

Forks

- **If the report says the deal is already fine** — use it as a defense brief: lock the rate, play the structural levers, and spend the goodwill on terms. Fine is a finding.
- **If the vendor disputes the numbers** — every figure is grounded and citation-tagged; offer the basis rather than debating impressions, and let the Figures ledger back any number they push on.

Done when: the email is sent, the deal is tracked, and the report's levers have owners and dates instead of readers.

See also: The deal desk room · Send it on

RUN YOUR MONDAY RENEWAL DESK

The weekly operating rhythm, in under 30 minutes, that makes everything else in this chapter rarer.

For: whoever owns renewals · **Time:** 20 to 30 minutes, Mondays · **You need:** the calendar live and your Monday brief sections chosen

Steps

1. Read your Monday brief email first if you switched it on: the vendor market wire, renewals coming up, watchtower changes, and savings and wins, in one send (Notifications chooses the sections).
2. Open the Control room: the renewal pressure timeline puts every notice window shutting inside 120 days on two axes, and the first circle inside the red 14-day bracket is your morning.
3. Clear what waits on people: your own My work queue, and the board's **Needs you** lane if agents or workflows run in your workspace.
4. Scan the week's market: the Wire with the estate filter on, and any Terms watch strip that struck upward in brick.
5. Move one deal forward deliberately: the nearest renewal gets its next action booked before the meeting-driven week eats the calendar.

Forks

- **If the timeline shows nothing** — verify it is true quiet rather than missing data: an agreement without an expiry and notice period cannot appear. An empty control room means the rows are missing, not that nothing is happening.
- **If the desk keeps surfacing the same stuck deal** — stop touching it weekly and change its shape: escalate per its recipe (90 days, 2 weeks, or hand it to Vera).

Done when: nothing sits unowned inside its notice window, the Needs you lanes are clear, and one deal moved.

See also: Control room · The Desk

BRING THE PLATFORM INTO EMAIL AND CHAT

Meet the team where they already work: contracts in by email, alerts out to Slack or Teams, Vera in the channel.

For: org owners connect; each person links their own DMs · **Time:** 30 minutes for the owner; two minutes per person · **You need:** admin rights in your Slack or Teams to install the bot

Steps

1. An owner connects the channels under Integrations: the Slack or Teams bot for the workspace, and the org alert channel if the team wants shared alerts (Notifications, webhooks).
2. Turn on the intake rails while you are there: the email-in addresses for contracts and requests, the invoice forward-in address, and folder watch on Drive, SharePoint, or Box, so documents file themselves from the places they already land.

3. Each person who wants renewals in their DMs links once: send the Vera bot a direct message in your company Slack or Teams, then flip the **Chat DM** toggle on the alert types you want under Notifications. The DM lane is strictly opt-in per alert type.
4. Use Vera where the conversation is: the chat clients carry the same assistant into the channel, and email works as an interface for teams that live in their inbox.

Forks

- **If IT balks at the bot install** — start with the intake rails alone: email-in requires nothing installed in your tenant and delivers most of the daily value.
- **If the shared channel gets noisy** — that is a configuration smell: move individual alert types back to personal DMs or the in-app inbox, and keep the shared channel for money-critical only.

Done when: paper flows in without anyone uploading, the right alerts reach the right surface, and nobody misses a deadline because it lived in the wrong tool.

See also: Integrations · Clients: what runs on your machines

REHEARSE A NEGOTIATION BEFORE THE REAL ONE

The first time you make your ask should not be in front of the vendor.

For: whoever takes the call · **Time:** an hour, ideally the day before · **You need:** the position (targets, walk-away) already set

Steps

1. Run Self play on the deal: the negotiation played out against itself, so you see the likely rounds and counters before any of them is real.
2. Spar with the Vendor Twin: the vendor's side, simulated, for practicing the conversation you are about to have.
3. Rehearse the two lines that matter out loud: your opening ask (grounded in the range) and your walk-away response. The deal desk brief's talking points carry lines written to be spoken.
4. Pressure-test the mandate against the rehearsal: if the practice rounds keep breaking through your walk-away, the walk-away is not real yet; fix it with the deal's owner before the call, not during it.
5. Book the real conversation on your geometry (Vendor calendar) and go make the ask you practiced.

Forks

- **If rehearsal exposes a hole in the case** — better now: the missing number is usually one instant benchmark or one decode away, and the call can wait a day.
- **If you rehearse well but freeze live** — bring the paper: the brief's talking points are designed to be read from, and there is no penalty for a buyer who consults their notes.

Done when: you have heard yourself make the ask, the counters have answers, and the walk-away survived practice.

See also: Self play · Vendor Twin

PLAN THE YEAR'S VENDOR STRATEGY

The annual sit-down that turns next year's renewals from a series of surprises into a calendar of planned negotiations.

For: the function's lead · **Time:** a half day, once a year · **You need:** the estate on file with values and dates; a year of watching helps

Steps

1. Generate the annual Strategy document: category strategies, vendor posture, and the year's negotiation calendar, built from your live records.
2. Read the estate's shape before accepting the plan: Portfolio for concentration and market position, and the Portfolio benchmark to rank where the money and the gaps actually are.
3. Mark the year's geometry on the Vendor calendar: every renewal, every notice deadline, and every vendor fiscal quarter end that falls inside your windows.
4. Assign the year: each significant renewal gets an owner and its preparation recipe now (12 months out for the far ones, 90 days queued for the near ones).
5. Put the watching in place for the year: Terms watch follows and the Monday Weekly report across your top vendors, so the file builds itself between sit-downs.

Forks

- **If the strategy document reads thin** — it reflects the data: values and dates missing from agreements starve it. An hour of Connect your data before the sit-down pays for the afternoon.
- **If leadership wants the strategy presented** — that is the board pack; this recipe is the working session that makes the pack honest.

Done when: every significant renewal in the next 12 months has an owner, a recipe, and a date, and the calendar knows things the vendors will not tell you.

See also: Strategy · Vendor calendar

SUNSET A VENDOR CLEANLY

Leaving a vendor is a project with deadlines, not an email. Done well, it ends with money on the ledger and your data in hand.

For: whoever owns the vendor, with legal on the exit terms · **Time:** an hour to plan; the exit runs to the contract's clock · **You need:** the decision to leave, made; the contract on file

Steps

1. Find the real exit door first: the notice deadline on Renewals is the date that matters, and missing it usually buys an unwanted year. If it is close, file the non-renewal notice before perfecting the plan.
2. Read the exit terms from the paper, not from memory: run Contract decode on the agreement for the termination mechanics, data return and deletion obligations, and any surviving clauses, and check the Rights register for paid-for rights (transition assistance, export help) you already own.
3. Plan the overlap: the replacement's timeline against the incumbent's end date, with the gap or overlap priced deliberately (Sourcing if the replacement is not chosen yet).
4. Execute the data exit as its own workstream: what comes back, in what format, by when, and the deletion confirmation you want in writing from the vendor.
5. Close the money: the retired spend is a saving with a start date; record it so Savings proof can verify it against the invoices that stop arriving.
6. Keep the paper: the expired agreement stays in Your Archive as history, correctly judged as such, where audits and comparisons can still reach it.

Forks

- **If the vendor makes leaving hard** — the decoded exit terms are your script; obligations they carry (data return, transition help) are asks you make in writing, on the clause.
- **If the sunset is part of a consolidation** — run it inside Consolidate duplicate tools so the keeper's renegotiation and the loser's exit share one timeline.

Done when: notice was on time, the data is back, the last invoice is the last invoice, and the saving is on the ledger.

See also: Contract decode · Rights register

ASK VERA WELL

The assistant is an analyst, not a search box. Asked well, she answers with numbers you can forward.

For: anyone on a plan that includes AI · **Time:** this recipe saves time; that is its point · **You need:** a real question

Steps

1. Ask from where the context is: **Ctrl/⌘ J** opens Vera over whatever you are looking at, and the screen you are on is context she can use.
2. Name the vendor and the decision, not just the topic: "am I overpaying?" is weaker than "is 18 percent off list good for Salesforce at our size, and what should I ask for?" The question that names a number gets a number back.
3. Attach the paper when the question is about your paper: the Ask bar takes a contract, and grounded means her answer cites your document and the benchmark evidence rather than gesturing at them.
4. Follow the citations before you forward: every figure carries its basis, and thirty seconds of checking is what makes the answer forwardable.

5. Teach her when she misses: a written feedback comment is the fastest way to shape future answers, the learned preferences are inspectable and deletable under AI personalization, and your organization's negotiation Memory is worth keeping current, because briefs and mandates open with it.

Forks

- **If the answer seems generic** — the question probably was: add the vendor, the size, and what you will do with the answer, or ask from the screen that holds the record instead of from a blank page.
- **If you hit the daily fair-use meter** — it warns before it bites, and the meter near the input says exactly what remains today (Chapter 10).

Done when: the answer is cited, checked, and specific enough to act on, and the correction, if one was needed, was taught rather than swallowed.

See also: Vera AI · AI personalization

GET UNSTUCK

Every help surface, in the order to try them.

For: anyone · **Time:** the whole ladder takes minutes · **You need:** nothing

Steps

1. The round help button at the bottom right explains the current screen and answers questions about it in place, and the **Guide** door in the top bar opens this manual at that screen's own section.
2. Press **Ctrl/⌘ J** to open Vera over whatever you are looking at; ask about a price, a clause, or how the platform works.
3. Check Troubleshooting for the symptom-first fixes, and Plans, allowances & limits if a feature looks locked.
4. Message our team from Chat: the "Vera AI team" thread reaches our analysts directly, and Support tracks the tickets.
5. For deal help rather than product help, book Advisory or engage the Expert Bench.

Forks

- **If a page you expect is missing from your rail** — nothing is broken: the short rail hides the long tail. Troubleshooting explains the three places it went, and Ctrl/⌘ K always finds it.
- **If it is genuinely broken** — report it from the account menu (**Report a bug**): the page, browser, and viewport attach automatically, you get a reference code, and bugs are typically fixed within 24 hours.

Done when: you have your answer, or a reference code and a human on the other end.

See also: Troubleshooting · Chat

Next: Chapter 08 · Glossary & FAQ

08 · Glossary & FAQ

The words the platform uses, and the questions people ask most. If a term on screen is unfamiliar, it is probably defined here.

GLOSSARY

Proposal — a vendor quote or contract you upload for our analysts to benchmark. It moves along a path: Uploaded, In review, Benchmarked, Delivered. Lives under Proposals.

Agreement — a signed contract you track for spend and renewals, not necessarily uploaded as a document. Lives under Agreements.

Instant benchmark — the self-serve price check you run yourself against market data. No upload, no analyst, a result in about two minutes. See Instant benchmarks.

Analyst benchmark — the deliverable our team produces from an uploaded proposal: metrics, market ranges, and a written narrative. Delivered to Delivered reports.

Report — any delivered narrative document you read or forward: the analyst benchmark, board reports, strategy reports.

Deliverable — the umbrella for every generated artifact: reports, briefs, playbooks, memos, and scans. They collect in Deliverables.

Brief — a short AI-generated executive summary of one thing: a contract, a vendor, or a renewal.

Playbook — a structured negotiation guide, either generated by an agent or curated for a specific vendor. See Playbooks.

Vendor — the company you are buying from. Called a supplier only inside the Sourcing module, where RFP language applies.

Verdict — the fast, classified read the Verdict Bar returns on any dropped quote, contract, or invoice.

Notice window — the period before an agreement expires during which you must tell the vendor if you are cancelling or renegotiating. Set as 30, 60, or 90 days. Teammates who opted into renewal alerts are emailed when it opens.

Percentile — where your deal sits against a cohort of comparable buyers. A lower percentile on price usually means a better deal for you.

Corridor — the calibrated range of renewal outcomes behind our benchmarks. Vendor grades and your Deal Score are read from it.

Deal Score — where your invoice-verified deals land inside the corridors, with 100 being the strong end. Tracked on Verified wins.

Buying bloc — an opt-in group of buyers of the same vendor who renew in a coordinated window under one term sheet, anonymous until mutual opt-in. See Buying blocs.

Fair Renewal Rider — the versioned standard addendum you attach to every renewal: uplift cap, benchmark repricing, audit conduct, and change-of-control protection. See Fair Renewal Rider.

Vera — the AI analyst, grounded in your contracts and our market data. The feature is branded Vera AI and opens anywhere with Ctrl or Cmd J. See Vera AI, and the Vera handbook for the working relationship.

The three answers — the order every benchmark answers in: the discount you are being offered, what comparable customers get as a range with its midpoint, and where the deal ranks. Everything else in a benchmark is supporting material.

Rank — the one-word verdict on a deal: Poor, Below market, At market, Good, or Top of the market, always stated with the gap in points and in money over the term.

Grounded — backed by a checkable source. A grounded figure carries a citation to the record or evidence it came from; a figure the data cannot support is refused ("not on file") rather than estimated.

Mandate — the written authorization a negotiation runs under: the target, the walk-away, what may be revealed, and how much rope. Agent-run deals never contact a vendor without one you approved.

Walk-away — the boundary at which you genuinely stop; the mandate's honest edge. A walk-away that rehearsal keeps breaking through is not real yet.

Concession ledger — the running record of every give and get on both sides of a negotiation, so the trade balance is a fact rather than a feeling.

Commitments — vendor promises captured on the deal record so they reach the final paper. A promise not in the contract is a memory.

Case — an agent-run deal in Case Mode: Vera negotiates under your mandate, you decide every offer, and one rule holds throughout: at most one pending human decision at a time.

Case file — a record opened as a full-screen set of numbered executive-brief documents, paged like a dossier, always ending in a What now sheet whose moves run on the record's own figures.

Workroom — the working surface a finished run opens into: the document on stage, the dock of moves beside it, and Ask Vera grounded on the whole case.

Deal desk room — where a dropped quote's finished brief lands, and the surface you take into the pricing round. See the deal desk room.

The Wire — the always-on market feed: pricing moves, licensing changes, research, and policy diffs, each item read for what it means for your estate. See The Vera Wire.

Watchtower — the standing watch surfaces: Terms watch on the vendor's published prices, Billing watch on your own invoices against entitlements.

Decode — the clause-by-clause plain-English read of a document, with paste-ready asks. The family: Contract decode, the SOW Decoder, the MSA Decoder, the Audit letter decoder.

Staging — where a Migration Studio run holds everything it loads: invisible to alerts, digests, and agents until an owner signs the estate off, in one transaction.

Binding — assembling a master agreement and its amendments into one contract whose terms in force are computed across the stack, newest document winning on each term separately. Binding is what makes a renewal date true.

Superseded — paper replaced by a newer agreement, judged as history: kept, searchable, and never polluting the renewal calendar.

Sourcing runway — the working time between today and the notice window: the stretch in which you are negotiating, before the contract starts deciding for you.

Benchmark credit — one full analyst benchmark added to a plan without included ones: $5,000, billed by invoice, requested by an owner from the proposal page itself.

Licensing and negotiation jargon

The vendor-side vocabulary the chapters use. You do not need to be a licensing veteran to use the platform, but these twelve terms come up constantly.

List price — the vendor's published price before any discount. **Net price** is what you actually pay; the gap between them is your discount, and benchmarks often measure it.

ACV — annual contract value: what an agreement costs per year, the number most benchmarks and blocs coordinate on.

Uplift — the percentage increase a vendor applies at renewal. An **uplift cap** is a contract clause limiting it; the Fair Renewal Rider indexes it.

True-up — the periodic reconciliation where you pay for usage above what you licensed. A **true-down**, reducing counts to actual usage, is rarer and usually only possible in the renewal window.

Shelfware — licenses you pay for but nobody uses. The optimizers in Chapter 03 exist largely to find it.

EA / ELA — an Enterprise Agreement (Microsoft's term) or Enterprise License Agreement (IBM's and others'): a large multi-year bundle covering many products, renegotiated as one deal.

ULA — Oracle's Unlimited License Agreement: deploy as much as you like during the term, then **certify** your deployment counts at exit to bank the licenses. The certification date drives the whole Oracle ULA playbook.

MACC — the Microsoft Azure Consumption Commitment: a contracted amount of Azure spend, negotiated inside the EA.

FUE — full user equivalent, the license metric RISE with SAP prices on; rebuilding the FUE count from real usage is the first move of the RISE playbook.

SOW — statement of work: the services document that carries day rates and deliverables. Priced against the market by the Rate Card Benchmark.

Third-party support — maintenance for a vendor's product bought from an independent provider instead of the vendor, typically far below the OEM rate. See the third-party support playbook.

Co-termination — aligning several agreements to end on the same date so they renew as one negotiation instead of dribbling out through the year.

TCV — total contract value: what the agreement costs over its whole term, the figure a multi-year commitment is really negotiated on. A three year deal quoted per year hides two thirds of the number.

True forward — the committed discount applying to consumption above the commitment, instead of overage billing at list. One of the reference brief's standard structural asks.

Rollover — unused committed spend carrying into the next period instead of expiring. "Use it or lose it" is the absence of this clause, and it is negotiable.

Quarter close — the end of the vendor's fiscal quarter, when discount authority runs deepest. The Vendor calendar maps every tracked vendor's.

Perpetual license — a license owned outright rather than subscribed. Old perpetual paper is a current asset, and the archive judges it as one rather than filing it as expired.

Reseller / CSP — an intermediary selling the vendor's product, sometimes on its own paper and price. Competing a direct quote against reseller and CSP quotes is what Quote Face-Off aligns.

Franken-quote — the best mix across several competing quotes, priced as if one paper. What you negotiate toward after a face-off.

Contract status path

Every proposal moves through four fixed statuses:

1. **Uploaded** — received, waiting for an analyst to pick it up.
2. **In review** — an analyst is benchmarking it.
3. **Benchmarked** — the final quality check before delivery. Still internal.
4. **Delivered** — the report is ready and you have been emailed.

FAQ

Who can see the documents I upload?
Only your own organization and our analysts. Files are stored privately, never on a public URL, and downloads use short-lived signed links. Data is isolated per organization at the database level, so another customer can never see your records.

What file types and sizes can I upload?
Proposals and contracts accept PDF and DOCX. The Verdict Bar also accepts photos and screenshots. Limits are shown on each upload screen (proposal and contract files up to 40 MB, Verdict Bar files up to 30 MB, images up to 3.5 MB). Files are checked on type, size, and signature, not just the extension.

How is my account protected?
With your password, or, once you turn on two-factor authentication under Security, with the rotating six-digit code from your authenticator app, which then replaces the password at sign-in. Accounts with elevated

access are required to use two-factor authentication. The sign-in screen also offers **Remember this device for 30 days**; untick it on shared machines and the session ends when the browser closes.

Why does a benchmark say "insufficient data"?
We never estimate a market position without data to defend it. When a vendor lacks reference coverage, the report says so plainly and names the one step that closes the gap, usually adding a list price or seat count in that vendor's benchmark tool.

What is the difference between an instant benchmark and an analyst benchmark?
An instant benchmark is self-serve: you answer a few questions and get a market comparison immediately. An analyst benchmark is a deliverable our team produces from a document you upload, with metrics, ranges, and a written narrative.

Can I try the platform without entering real contracts?
Yes. Load the sample portfolio from the dashboard to explore populated screens, and remove it in one click. Sample rows never mix with your real data.

How do I move around quickly?
Press **Ctrl or Cmd K** anywhere to open the command palette and jump to any page, tool, or vendor. Press **Ctrl or Cmd J** to open Vera over whatever you are looking at. The round help button at the bottom right explains the current screen.

Which actions are limited to org owners?
Inviting and managing members, connecting integrations, configuring single sign on, setting branding and notification rules, and data controls. Regular members can use every workspace and tooling feature but not these org-level settings.

Do the AI features cost extra?
AI-powered features (marked "requires an AI plan" in this manual) depend on your plan including AI, and each run counts against a periodic allowance. Everything else works on any plan. Plans, allowances & limits defines every gate and cap, and your own usage is always visible under Billing.

Which plan do I need for a feature?
The plan matrix in Plans, allowances & limits lists what each plan includes, from the Free single-contract decoder to Enterprise's included analyst benchmarks and advisory sessions. Every tool is free to explore with sample data on any plan.

How do referrals work?
Every account has a referral code under Referrals. Refer a company and earn 10% of their spend for their first 12 months.

Something is broken, or my question is not here.
Start with Troubleshooting, which walks the most common symptoms. Otherwise message our team from Chat (the "Vera AI team" thread) or email info@vendorbenchmark.com.

Next: Chapter 09 · Record & detail pages

09 · Record & detail pages

Lists get you to a record; the record page is where the work happens. Proposals, vendors, agreements, and reports all open into a detail page built from the same parts, so once you know the pattern you can read any of them. This chapter covers that shared anatomy and then the three detail pages you open most.

THE RECORD PAGE PATTERN

In plain words: learn one page shape and you have learned them all.

Every detail page in the platform is laid out the same way, top to bottom:

- **Breadcrumb** — where you are, and one click back to the list. Beside it sit two live controls: quiet avatar pips showing who else is viewing the record right now, and a **Watch** button. Watching subscribes you to the record: comments, vendor messages, sign-off activity, and Vera's work land on your notification wire; commenting or acting on the deal subscribes you automatically.
- **One banner at most** — when the record has something to tell you (a report correction notice, a just-submitted note, a suggested lineage link to related paper), only the most urgent one shows above the header, never a stack.
- **Highlights header** — a vendor logo, the record title and status, and a row of the key facts (the numbers you would read out on a call). Actions for the record sit here on the right.
- **Tabs** — the main content, as underline tabs (the active one carries a blue border). The same tab pattern is used everywhere, never boxed or segmented.
- **Related records strip** — below the tabs: one hop to everything connected, the vendor, the matching renewal, the delivered report, the contract this one renews, plus any connected workspaces.
- **Right rail** — supporting tools and context beside the tabs: renewal tracking, AI insights, and comments, with the less frequent tools in a **More for this record** fold.
- **Research for this vendor** — where the library covers the vendor, the record carries pointers to it, led by **Read before the meeting**: the negotiation prep brief, three pages on what the account team is compensated to sell you, the facts about your own estate to hold before they quote, the talking points to use, and what they will say with the counter to each. Beside it sit the renewal outlook (what the next twelve months will cost and what to lock in now) and the vendor assessment (where the vendor is taking pricing and packaging). Nothing appears for a vendor the library does not cover yet, so what you see is either right or absent. These pointers show on the contract, negotiation and renewal records, the renewal playbook, and the vendor page.
- **Timeline** — one chronological stream of everything that happened on the record: comments, vendor messages, sign-off requests and decisions, and Vera's work, each entry with who did it and when, refreshing live while teammates act. Under it, **Discussed in** lists every chat message and Board post that referenced the record, with the snippet and a jump link.

Because the pattern is fixed, the tab names change but the way you move around does not. Negotiation deal rooms open with the same highlights header (target, latest quote round, concession balance, open commitments, last activity), so a live deal reads like every other record.

The workroom

Most records carry a navy band that opens the record as a **workroom**: one room, the same on every screen that has it, so once you have worked in one you can work in all of them.

- **The stage** — the record's papers as numbered executive brief documents on a dark stage, one per screen, paged with the arrow keys, with a read meter across the top and Esc to close.
- **The left rail** — the document list, then two shelves: a keyword and semantic search over your own uploaded documents, and your shelf of recent and starred workrooms, so what you read last is one click away.
- ⚡ **What now** — the icon dock on the right edge splits the stage and puts the record's moves beside whichever paper is open: the drafts Vera composes from this record's own figures, the copies, the exports, and the handoffs into other screens. The same moves are also the last document in the file, so you can read them in the flow of the paper, and they appear in the command palette. **Every move is real.** Nothing in a dock is a button that does nothing.
- **✦ Ask Vera** — the second panel in the same split, grounded on this record alone rather than on the platform in general, with one chip per document that jumps the stage. Questions asked here land in the same saved conversation every time you reopen the record.
- **Drafts** open in a window for you to copy or edit, and nothing sends from a workroom. A vendor-facing draft also asks whether it is the start of a negotiation; saying yes files it as the opening move of a tracked thread (see Negotiations).
- **The room saves itself.** Opening a room records it, and the paper you are reading, so **Back to the room** on your dashboard takes you back to exactly where you left off, on any machine. Only you can see your own trail. **Keep this room on my desk**, at the foot of the What now panel, turns it off and clears what was already saved.
- **The sibling desks** — pills beside the record's title: **Case file** first, then the desks this record can hand off to. They swap what is on the stage without closing the record. The rail, the What now dock, Ask Vera, your shelf and every figure stay exactly where they are. Each desk's room says what that desk does, lists the runs you already have on it with the vendor you are working on first, so you can reopen one instead of starting cold, and carries one control that hands this work over. That last control does leave the record, because those desks read a document and run for a few minutes as their own job, but it is a press you choose after seeing what you are choosing. The way back is the Case file pill, any document in the rail, the button at the foot of the room, or Esc. On a narrow window the same pills ride the end of the document strip. The same moves in the What now dock open these rooms too, so reaching for another desk never costs you this one.
- **Esc unwinds one layer at a time** — the open dock panel first, then the open room, then the record itself.

Rooms have addresses: `?open=case` on the record's own route, so a workroom can be bookmarked and forwarded, and the browser Back button closes it.

The screens that open a workroom: the vendor record, the renewal record, a delivered benchmark report, an invoice review, the savings ledger, a sourcing project, The Desk desk file, a diligence review, the board record, your monthly briefing, the org estate, the Fair Renewal Rider, a Terms watch price change, the weekly report, and every optimizer and decoder in Chapter 03.

THE PROPOSAL RECORD

Where: open any row under Proposals · `/contracts/<id>`

Everything about one uploaded proposal, from the moment you submit it to the delivered report.

In plain words: one proposal's whole life on one page: status, files, and eventually the report.

A worked example. The sample Northwind CRM proposal's record shows the pattern at its best: the status path across the top, the files beneath, and once delivered, the Report tab carrying the finding ($28,650 a year in reach) with its case file one press away. While the proposal is anywhere before Delivered, the page refreshes itself within half a minute of a status change; delivery announces itself with a toast.

On the screen

- **Status path** — the four-step path Uploaded, In review, Benchmarked, Delivered, with the date each step was reached and, while it is with our team, who is working on it and how long it has been in the current step.
- **Highlights** — vendor, category, status, who uploaded it, and the delivered date. If the AI redline found the quote below the vendor's usual discount band, a "below band" badge appears here.
- **Deal brief** — the one-glance read of the deal, directly below the header.
- **First read** — when you have an instant AI read on the document, it leads the page so you get a number before the analyst report lands.
- **Workspaces strip** — below the tabs, with the related records: the deeper analyses that already have a result, the redline review, the SOW rate review, the walk-away alternative, and the negotiation brief. On the brief, the call prep pack now fronts itself with a cover: the asks counted, the pushbacks answered, a contents list that jumps to each section, and your first three moves lifted from the call plan; a freshly generated pack lands with a short delivery reveal. Editing the brief is truly co-edited: teammates in the editor at the same time share one live draft, a banner names who is editing with you, and each field shows who is typing in it, so nobody overwrites anyone. Saving still appends a new version.
- **SOW rate review** — for services documents, the review extracts every rate line and benchmarks each day rate against the market band for the role. When you have a saved MSA rate card for the vendor (see Rate Card Benchmark), a **Your MSA** column appears beside the market band: each line billed above your contracted rate is flagged as a contract-compliance gap, with the per-day delta and the overage across the stated volume computed in code. The market band is context; the MSA rate is the ceiling you hold. A **Save as rate card** button on the rate table records the extracted lines as your contracted rates in one click, and any blended or offshore-blended line carries a **Deconstruct this rate** panel showing the plausible senior/standard/junior mixes at that price and the margin range each read implies.

The tabs — always in the same order, Report, Plan, Ask, Files, Details, at every stage; only the tab that opens first changes.

- **Report** — once delivered, this opens first, and a small line points here from any other tab: the white-paper narrative, the negotiation playbook, and the metrics table with percentile bars (see The benchmark report below). Before delivery it shows the request-a-benchmark card with the estimated delivery date.
- **Plan** — the step-by-step engagement plan for this vendor, with AI guidance on each activity.
- **Ask** — ask questions about this specific document and get answers grounded in its text.
- **Files** — the PDF or DOCX files you uploaded, opened with a short-lived signed link.
- **Details** — the extracted fields (seats, annual value, uplift cap, term dates, source language), your notes, and what you asked us. You can re-extract or export a redline from here.

The right rail — three cards, then a fold

- **Track renewal** — when the record has a renewal to track, the tracking card leads the rail.
- **AI insights** — the AI analyses, grouped as Review (proposal review, SOW review, second opinion, renewal vs current), Model (what-if scenarios, entitlements, credible alternative), and Compliance (audit defense, audit workbench, security & DPA check). Rows with a stored result open first; the rest sit behind one entry. In the what-if modeler you can save a scenario by name, reload it any time, compare two side by side with the levers and the difference between them, and pin the chosen case to the negotiation plan.
- **Comments** — a thread on the record, and you can comment on a specific metric or clause.
- **More for this record** — a fold below the three cards holding the negotiation toolkit, the renewal outlook, and access management (whether the proposal is visible to the whole organization or restricted to named people, managed by the owner, the uploader, or an admin).

See also: Proposals · Delivered reports · Contract decode

THE VENDOR RECORD

Where: open any row under Vendors · `/vendors/<vendor>`

Your whole relationship with one vendor on a single page.

On the screen

- **Highlights** — the record opens as a dossier: a gold kicker over the vendor name, then a gold rule over one facts strip holding the annual value, the renewal countdown and date, proposals (and how many delivered), and tracked research, with an "in notice window" badge when a deadline is close. For the 176 tracked publishers the masthead also states the company identity (headquarters, listing or ownership, headcount) and the vendor's commercial model chips (cloud, subscription, perpetual). The deeper tabs head their sections with navy bands, the same grammar as the delivered briefs.
- **Company tab** — the vendor intelligence card leads the record: who the vendor is, a Financials section with the latest reported revenue, margins, and segment mix (refreshed each time the vendor reports earnings) closed by a Buyer's angle read on what those numbers mean at your table, the commercial model

(revenue mix, contract vehicles, true up mechanics, and an audit posture rating), the product families we track with a one click benchmark on each, and the vendor's latest signals from the Vera Wire.

- **War room mode** — when a renewal crosses into its notice window, the record leads with a live countdown, this week's playbook step, the negotiation brief, and the owner and approvers pinned. The same deal also jumps to the top of My work, so a live renewal never runs off screen.
- **Take it to market** — the renewal record's challenge control opens a competitive sourcing event with the incumbent's current annual price as the baseline. Alternative vendors bid against it, every priced response reads as a delta against the incumbent, and the best credible bid becomes the leverage line the renewal negotiation cites.
- **Open the workroom** — a navy band under the related records opens the vendor as a workroom. Up to six documents build from what this vendor's record holds: your agreements ledger, the pricing band at your deal size, the vendor's terms and conduct intel, a negotiation angles brief fitted to your footprint, and a What now sheet. ⚡ **What now** carries the moves that run on this vendor's live records: open the renewal record, open the latest proposal, read the delivered benchmark report, start tracking the agreement, open a Vera Deal Room, and copy the vendor one pager. ✦ **Ask Vera** answers from this vendor's own records, so a question about pricing is measured against what you actually pay rather than against the market in general.
- **Proposal pipeline** — this vendor's proposals across the status path.
- **Vendor profile** — your position and percentile, the typical discount off list once a benchmark exists, and the plays to expect from this vendor (quarter-end pressure, high anchor, bundling misdirection, and more).
- **Workroom** — the renewal, any live negotiation, watchdog findings, and the deliverables on this vendor, each one click away.
- **QBR pack** — the quarterly business review pack for this vendor opens with a delta cover built from your published benchmarks: the three moves that matter, your position change, savings realized, and the percentile trajectory since last quarter, on one page. It reads calmly in a vendor's first quarter, before there is a history to compare. A freshly generated pack lands with a short delivery reveal before the paper opens, and it downloads as a PDF in the executive brief theme alongside the Word export.
- **Deeper tabs** — the tabs are the whole map: position and risk reads live under Analysis, intelligence briefs under Research, and the activity log under Log. The relationship analysis under Analysis opens as an executive brief: a vitals row computed from your own records (spend on file, estate size, days to expiry, notice deadline, with anything not captured flagged), the full estate table, an annual value chart when more than one item carries a value, and the analyst narrative in short sections.
- **Present the intel briefing** — the market intel briefing's reading dialog carries a Present button for the fifteen minutes before a vendor meeting: it walks the scanned developments one at a time with the negotiation angle on each, and closes on your stance, with your annual spend, benchmark position, and the market risk read pinned to the top of the stage.

How to use it

1. Read the highlights and the vendor profile to see whether this vendor looks above or below market and what plays to expect.

2. Check Intelligence before a renewal for the latest pricing and licensing moves.
3. Use the related records strip to jump to the vendor's agreements, proposals, and delivered reports, or open the war room straight from the renewal alert.

See also: Vendors · Vendor scorecards · Terms watch

THE RENEWAL RECORD

Where: open any renewal from Renewals or Agreements · `/renewals/<id>`

One renewal cycle as a command room. The record opens on a navy masthead that states the position in figures before the first row, then splits into a tabbed working canvas and a context rail that stays beside you while you work.

On the screen

- **The command masthead** — the vendor, one sentence of figures (what you pay, where the cycle is expected to land, how many days remain to give notice), and the record's vitals: current value, expected at renewal under the uplift cap, the annual savings on the table at the should-cost target, the notice window countdown colored by urgency, and your benchmark standing percentile once one exists. **Watch** and **Request benchmark** sit beside them.
- **The stage path** — Upcoming, In negotiation, Benchmarked, Closed drawn as one path with the current stage lit and finished stages checked. The room's controls ride the same band: **Take it to market** (a competitive event with the incumbent's price as the baseline), snooze, the renewal agent, the calendar file, close cycle, and the managed service. The line beneath states the owner, who is watching, the notice date, whether the agreement auto renews, and the uplift cap.
- **The workroom band** — opens this renewal as six exec-brief documents in the full-screen workroom: the stored terms, the position, the dated path to the notice deadline, and the asks they ground.
- **Deal position** — the default working tab: the projected renewal cost with its range and reasoning, the should-cost target with its derivation, the step-down waterfall from list price to your ask, and the true-down demand when seat utilization data supports one. With no value or benchmark on file the tab says exactly what to add to build the math.
- **The other tabs** — Playbook (the work-back plan and its progress), Benchmark standing (your market position for this vendor), Proposals (every proposal on file), and Files. The renewal checklist sits below the tabs; each step is a task with its own owner and due date.
- **The context rail** — how the last cycle with this vendor ended, the renewal run (the autonomous arc of this cycle, with the approval surface for anything waiting on you), the research pointers for the vendor, the renewal discussion, and the full timeline and activity streams. On a wide screen the rail pins beside the work and scrolls on its own.

How to use it

1. Read the masthead: the vitals state what you pay, what the cycle will cost if nothing changes, and how much runway the notice clock leaves.
2. Work the **Deal position** tab to set your target, then open the workroom when you need the position as documents.
3. Move the stage forward as the cycle progresses; the path, the run rail, and My work stay in step.

See also: Renewals · Renewal playbook · Negotiations

THE BENCHMARK REPORT

Where: the Report tab on a delivered proposal, or open it from Delivered reports

The deliverable itself: a senior-analyst read of your deal against the market, not a bare table of numbers.

In plain words: the report argues a finding; the metrics beneath it are the evidence, never the report.

When to reach for this

- Reading a delivered report properly: finding first, evidence second, levers third (the handoff recipe turns it into the conversation).
- Explaining to a stakeholder why the deliverable is prose with a table, not a table alone.

On the screen

- **The reveal** — the very first time you open a delivered report, a short full screen reveal announces it: your market standing stamps in, the money on the table counts up, and a band draws where you sit between market low and high before the stage lifts onto the report. One click skips it, it respects your reduced-motion setting, and it never appears in print or PDF output. It plays once per report; replay it any time from the "Replay the reveal" control on the Delivered banner. Shared report links get the same first-open moment, once per link.
- **The headline** — page one answers three questions in one screen: the negotiable value found across your benchmarked terms (or "At or better than market" when there is nothing to chase), the signature position bar for your single biggest lever showing where that term sits in the cohort and the dollars on the table, and a report-level confidence chip stating how solid the data is (the deals referenced, the freshest data date, and the sources).
- **The narrative** — the white-paper report leads: an executive summary, your market position, section-by-section analysis, prioritized recommendations, and a bottom line. Every figure is grounded in stored data, with its sources listed at the foot of the document.
- **Trust strip** — an AI-generated report opens with one verification story for the whole page: the grounding state ("Figures verified against your data", or a review notice when a figure needs one), the source count, the generation date, and feedback thumbs.
- **Section chips** — a long report keeps a sticky row of section chips across the top of the document on every screen size, shared links included: the chip for the section you are reading stays highlighted, and tapping one jumps straight to that section. On wide screens with free margin, a slim side rail beside the document lists the sections instead and tracks where you are as you read.

- **Your standing, metric by metric** — inside the "Where you stand versus the market" section of the AI narrative report, a bar chart with one bar per metric: the bar's length is the share of comparable deals your terms beat, and your value is printed at the bar.
- **Note from our team** — on a delivered report, the analyst note renders as a signed note on a cream panel with a gold rule, "The analyst desk" byline, and the delivery date.
- **Negotiation levers** — your top opportunities ranked by annual dollar value, at most three so it reads as an agenda: each card states the finding in one sentence, the money, the move, and the recommended ask. Every figure repeats on the metric rows below.
- **Negotiation playbook** — the metrics turned into specific asks you can take into the room.
- **Metrics table** — each metric with your value, the market low, median, and high, a percentile, and notes. The signature percentile bar shows where each one sits: the cohort's middle half shaded, the median ticked, your deal flagged, and the dollar gap to median on the bar. You can pick the peer cohort the comparison uses. When the narrative leads the page, this section sits below it under the title "Appendix · Metric detail".
- **Overall score** — one number summarizing how the deal stacks up.
- **Delta and trend** — how the report changed against earlier snapshots, when there is a history.
- **Section notes** — pin a comment thread to any section of the report, reply, and resolve it before you forward. Notes stay internal to your team and never appear on a shared link, a PDF, a Word file, or the slides.
- **Open the workroom** — the navy band on the Report tab opens the delivered benchmark as a six document workroom: the cover, the source proposal record, the analyst brief, a metrics atlas with the market ranges and the verdict table, an ask sheet priced from the stored figures, and a What now sheet. ⚡ **What now** drafts the CFO brief and the vendor safe meeting request from the report's own findings, copies a chat ready summary or the one line verdict, downloads the PDF, and carries the report into the negotiation dossier, a second opinion, or an analyst booking. ✦ **Ask Vera** is grounded on this report, so a question about where a figure came from is answered from the report's own evidence.
- **Downloads and sharing** — export the report as PDF, Word, a PowerPoint deck, or an Excel workbook, hand it to a teammate, or create a share link. The PDF and the deck open with the same headline and position bar as page one on screen, so the forwarded document matches what you read.
- **Present** — the Present control above the report runs it as a full screen show for a live meeting: the headline and market position first, then the evidence, then the play. Arrow keys move between screens, G shows every screen at a glance, N shows a one line presenter note, and the print control produces a leave-behind of exactly the screens you walked. The show reopens where you left off on this report.
- **Reading panel** — once you share a report link, this panel shows how it is being read: total opens, unique readers, the sections read most, and time spent on the recommendations. It carries no IP or personal data and honors Do Not Track.
- **What happened?** — thirty days after your report is delivered, the Report tab asks one question: did you win a reduction, did the vendor hold, or are you still negotiating? You answer once per contract. A won reduction with an annual figure opens a reconcilable savings record against your invoices automatically, and your answer sharpens the benchmark over time.

How to use it

1. Read the narrative first; it interprets the numbers rather than restating them.
2. Drop to the metrics table to see the exact gaps, and switch the cohort if a tighter peer set applies to you.
3. Take the negotiation playbook and the recommendations into your negotiation, and export or share the report with your budget owner.
4. Pin notes to the sections your team needs to weigh in on, resolve them, then forward the clean report.
5. When the prompt appears a month on, tell us what happened; a won reduction flows straight to Savings proof.

See also: Delivered reports · Benchmark a proposal end to end

---

Next: Chapter 10 · Plans, allowances & limits

10 · Plans, allowances & limits

Every plan gate and platform limit mentioned anywhere in this manual, collected on one page. When a chapter says a feature "requires an AI plan" or is "metered by plan", this is the page that defines what that means. Your own workspace's plan, usage, and remaining allowances are always visible under Billing.

HOW THIS MANUAL DESCRIBES PLAN GATING

The manual uses a small, fixed vocabulary so you always know what a gate means:

- **requires an AI plan** — the feature needs a plan that includes AI: Growth, Professional, or Enterprise. The Free plan has no AI features; the 30 day Trial includes a small starter allowance instead (see the table below).
- **metered by plan** — the feature works on your plan, but each run counts against a periodic allowance (for example AI analyses per month). The Billing page shows the meter.
- **your own data requires a Professional or Enterprise plan** — the tool is free to explore with sample data on any plan, but uploading your own exports or documents and running the AI over them needs Professional or Enterprise.
- **org owners only** — not a plan gate at all: the feature is on every plan but only organization owners can use it.

Allowance windows are rolling, not calendar months: "per month" means the last 30 days, "per week" the last 7. An allowance frees up on its own as older runs age out of the window.

THE PLANS

Five plans, priced annually. Owners can compare them under Billing; upgrades are set up by our team and billed by invoice.

	FREE	TRIAL (30 DAYS)	GROWTH	PROFESSIONAL	ENTERPRISE
Contracts tracked	1	2	50	250	1,000
Seats (hard cap)	2	5	2	5	Unlimited
Self-serve instant benchmarks	—	Salesforce Sales + Service Cloud and Adobe Experience Cloud only, 2 runs	5 / month	25 / month	No cap
Research downloads (reading is always included)	—	—	✓	✓	✓
Conversational AI (Vera, Ask, briefs)	—	—	✓	✓	✓
Benchmarks by analysts	—	—	—	—	Custom
Own-data uploads to tools	—	—	—	✓	✓
AI tool reports on your own data	—	—	—	✓	✓
Live voice (call copilot, sparring room)	—	—	—	—	✓
SSO, SCIM, and the API	—	—	—	—	✓

Seats are a hard cap, not a billing meter. The number is how many people can be on the roster at once, and everyone counts toward it: active members, pending invites, and blocked members alike. The cap is checked whenever someone is added, by invite, domain join, or SCIM provisioning, and it never removes anyone already on the roster. When it is full, remove a member, revoke a pending invite, or upgrade; Members shows seats in use against the cap.

Plans unlock apps. An app your plan does not carry shows an in-place panel naming the plan it joins on, with a link to Billing; it never opens partially. Workspaces created before July 29, 2026 keep every app they could already reach, and trial workspaces see every app, with the trial's own allowances as the only limit.

Plans without included analyst benchmarks (or with the allowance used up) can still send proposals for expert review with benchmark credits: an organization owner requests one directly on the proposal page ($5,000 per analyst benchmark, billed by invoice), our team confirms and adds it, and the Billing page shows the balance. See Proposals.

The trial is built around the portfolio scan: uploads are unlimited for the 30 days and the scan runs across everything you upload, so the number you see is your own. Two trial limits worth knowing before they surprise you: trial workspaces read research online but do not download it, and instant benchmark runs are capped at 10 in total. Partner-program firms also read research online without downloads, per program terms, whatever their plan.

AI ALLOWANCES BY PLAN

Each run of an AI feature counts against the matching allowance. All counts are per rolling window as described above.

ALLOWANCE	TRIAL	GROWTH	PROFESSIONAL	ENTERPRISE
AI analyses (reports)	10 in total	No monthly cap	No monthly cap	20 / month
AI agent requests (playbook agents)	1 in total	—	20 / month	100 / month
Proposal redline scans	1 in total	2 / month	5 / month	10 / month
Side-by-side compare scans	1 in total	2 / month	5 / month	10 / month
Cloud cost analyses (own uploads)	—	—	5 / month	10 / month

When an allowance runs out, the feature says so in place and the Billing meter shows when the window frees up. Owners can upgrade at any time; the higher allowance applies immediately.

Fair-use daily caps also protect the conversational surfaces (asking Vera) and the instant tools (Contract decode, Snap a quote). You will rarely meet them, and you will never meet one unwarned: once less than half of a day's cap remains, a quiet meter appears near the input ("7 of 10 questions left today") with the upgrade path attached. Trial workspaces, where asking Vera is metered at ten questions a day, always see the meter. Caps reset at midnight UTC.

PLATFORM LIMITS

The fixed limits that apply regardless of plan. Each links to the section where the limit matters.

LIMIT	VALUE	WHERE IT APPLIES
Proposal and contract files	PDF and DOCX only, 40 MB each	Upload a proposal
Verdict Bar documents	up to 30 MB	Verdict Bar
Photos and screenshots	up to 3.5 MB	Snap a quote, Verdict Bar images
Tracked vendor agreements	up to 40	Agreements
Files per vendor agreement	up to 5	Agreements
Research file downloads	10 per day, 30 per week per person	Research
Review table columns	up to 12 questions	Review tables
Deal approval chain	up to 10 approvers	Deal approvals
Benchmark share links	valid 30 days	The benchmark library
API keys	expire within one year	API
Spend optimization vendor set	top 10, 25, 50, or 100 by spend	Spend optimization

File uploads are validated on type, size, and file signature, not just the extension, and every upload screen states its own limits in place.

Related: Billing · Troubleshooting · Glossary & FAQ

Next: Chapter 11 · Troubleshooting

11 · Troubleshooting

When something does not behave the way the manual says it should, start here. Each section opens with the symptom and gives you the checks in the order most likely to fix it, and where it matters, the one habit that prevents the repeat. If none of them do, the last section tells you how to reach a human.

Find your symptom by where you are standing

WHERE YOU ARE	THE SYMPTOMS
Getting in	Invite never arrived · Invite link expired · Held at two-factor setup · Signup email declined · Wrong workspace · Lost authenticator · SSO sign-in fails · Network policy lockout · Session keeps ending
Uploading	Upload rejected · Contract reads with no terms · Re-dropped folder skips files · Interrupted upload · Nothing to benchmark · Emailed-in document missing
Benchmarks and AI	Not enough data · Feature greyed out · No list price · Not derivable · Run seems stuck · Review notice · Not on file · Currency looks wrong · Daily question meter
Renewals, alerts, email	Emails not arriving · No renewal alert ever · One agreement never alerts · Control room empty · Alert arrived late · Leverage note went quiet · Meeting packs missing
Sharing and exports	Share link stopped working · Confidentiality scan blocked it · Word or PDF link missing · Download limit
Data and views	Page missing from rail · Sample data showing · Charts thin · Vendor appears twice · Rows marked confirm · Shared view locked · Folds reset on a new device
Plan and billing	Seats full · Saved, not queued · Trial limits

---

A PAGE FROM THIS MANUAL IS NOT IN YOUR RAIL

By default the rail shows the core set of destinations most people use weekly; everything else still exists, it is just not pinned. Three ways to reach any page in this manual:

1. Press **Ctrl/⌘ K** and type the page name. This always works.
2. Open **More** at the foot of the Workspace group for the full searchable catalog.
3. Switch to the full rail with the control at the foot of the rail if you prefer to see everything.

Sections whose destination sits outside the core rail carry a "via More" marker on their **Where** line throughout this manual.

YOUR UPLOAD WAS REJECTED

1. Check the type: proposals and contracts accept **PDF and DOCX only**. Rename tricks do not help, because the file signature is checked, not the extension.
2. Check the size: 40 MB per proposal or contract file, 30 MB in the Verdict Bar, 3.5 MB for photos and screenshots. All limits are listed in Plans, allowances & limits.

3. Very large files upload straight from your browser to storage, which can take a moment on a slow connection. Leave the tab open until the file shows as attached.
4. A scanned contract inside an image-only PDF still uploads fine; the AI tools read it with vision.

A BENCHMARK SAYS THERE IS NOT ENOUGH DATA

That message is deliberate: the platform never estimates a market position it cannot defend. To get a ranked verdict:

1. Widen the region filter to all regions, or check that the deal size is realistic.
2. Add a list price or seat count in that vendor's benchmark tool; the data coverage page names the one step that closes each gap.
3. For a vendor outside the library's coverage, upload the proposal for analyst benchmarking instead.

AN AI FEATURE IS GREYED OUT OR ASKS YOU TO UPGRADE

1. Check whether your plan includes AI: the Free plan has none and the Trial has a small starter allowance. See the plan table.
2. If the plan is right, the allowance may be spent: the Billing meter shows what you have used and when the rolling window frees up.
3. Only org owners can upgrade. If you are not one, the feature names who is, or check the workspace summary at the foot of the rail.

YOU HIT THE RESEARCH DOWNLOAD LIMIT

Downloads are capped at 10 per day and 30 per week per person so access stays fair across organizations. Reading papers in the viewer is unlimited. The allowance frees up on its own as recent downloads age out; if you need more right away, message the team from Chat and an admin can reset it.

AN INVITE NEVER ARRIVED

1. Owners: check **Pending invitations** under Members. A bounced or delayed delivery status means the email never landed; fix the address and resend.
2. Ask the invitee to check their spam folder for a Vera AI sender.
3. If your organization has domain capture on, the person can also sign up directly with their work email and be approved as a join request.

ALERT OR DIGEST EMAILS ARE NOT ARRIVING

1. The Terms watch and Weekly report emails are **off until you turn them on**, with the switch at the top left of each page. Nothing is sent by default.
2. Check Notifications settings: each alert kind can be in-app, email, both, or off.
3. A one-click unsubscribe in any email turns off only that email, but check you did not use one by accident.

4. Still nothing? Check spam, then message the team.

MEETING PACKS ARE NOT APPEARING

1. You do not have to wait for the agent: on the Meeting war room, name the vendor and time (or drop the .ics invite) and the pack assembles on the spot.
2. For automatic packs, the calendar connection must be a read-only ICS feed (Outlook: Publish calendar; Google: Secret address in iCal format), connected under Settings › Integrations by an owner.
3. The agent checks the feed about every 20 minutes, so a meeting created moments ago may not have a pack yet.
4. The meeting title must mention a tracked vendor by name. "Sync with Dave" produces nothing; "Salesforce renewal call" produces a pack.

YOU LOST YOUR AUTHENTICATOR DEVICE

On the sign-in page choose **Email me a link** and sign in with the one-time link; at the six-digit code prompt choose **Use a recovery code** and enter one of the one-time **backup codes** you saved when you set up two-factor; each works once. Then register the new device under Security. If the backup codes are gone too, ask an org owner to contact our team from Chat and we will verify you another way.

YOU CANNOT SIGN IN THROUGH SINGLE SIGN ON

If your organization enforces SSO for your email domain, the password form will not work and you must sign in through your identity provider. Access problems there (a deactivated IdP account, a group you were removed from) are fixed by your own IT team, not in Vera AI. Org owners can review the configuration under Single sign on.

A SHARE LINK STOPPED WORKING

Benchmark share links are valid for 30 days unless they were created as kept-live links. Open the scenario or report again and create a fresh link; the old one cannot be revived. Research share links can also be revoked by the person who created them.

SAMPLE DATA IS SHOWING IN YOUR VIEWS

Sample rows never mix with your real data and never feed your alerts; they exist so screens are not empty while you explore. Remove the whole sample portfolio in one click from the dashboard. If you see a row you believe is sample data after removal, it is not: check who uploaded it on the record's activity tab.

CHARTS AND ANALYTICS LOOK EMPTY OR THIN

Portfolio, Spend, and the Estate map build from proposals and agreements that carry an **annual value**. Add the missing values on Agreements and the charts fill in immediately. Data health ranks every gap by what filling it unlocks.

YOUR INVITE LINK EXPIRED

Invitations expire after 7 days. Ask an owner to resend from Members: resending restarts the clock and refreshes the sign-in details. Prevention: send invites when people can act on them, not before a holiday.

YOU ARE HELD AT TWO-FACTOR SETUP

Your organization requires two-factor for every member, so sign-in holds you at registration until it is on. Scan the code with any authenticator app (Google Authenticator, 1Password, Authy), enter the six digits, and save the backup codes somewhere real; the hold is the policy working, not a fault.

YOUR SIGNUP EMAIL WAS DECLINED

Two address families are declined by design: personal domains (Gmail, Outlook, Yahoo and the like), because the platform licenses to organizations, and the corporate domains of software vendors and resellers, because it is licensed to buyers only (Terms of Service, Section 2). Use your work address; if your company already runs the platform with domain capture on, signup routes you to your team's admin instead of creating a duplicate workspace.

YOU LANDED IN THE WRONG WORKSPACE

Your account can belong to more than one organization; the org chip in the top bar switches between them, and the workspace summary at the foot of the rail always names the tenant you are in. If a workspace you expect is missing from the chip, its owner has not invited that email address; there is no cross-workspace visibility to fall back on, by design.

THE WHOLE TEAM IS BLOCKED BY A NETWORK POLICY

An IP allowlist under Security baseline pins the workspace to known networks, and an office IP change blocks everyone, owners included, until the policy is updated from an allowed network (home office on the list, or the range widened). Every denial lands in the audit log, and changes take up to a minute to reach every server. Prevention: prefer ranges over single addresses when egress IPs rotate.

YOUR SESSION KEEPS ENDING

Three legitimate causes, in likelihood order: **Remember this device for 30 days** was unticked at sign-in, so the session ends with the browser; an owner has set a maximum session age or idle timeout under Security baseline; or your password changed, which signs out every other session deliberately.

A CONTRACT READS WITH NO TERMS

Almost always a scan with no text layer. The platform reports it as needing OCR rather than silently filing an empty contract: the Archive queues it under **Needs a look**, and Migration Studio raises it in the exceptions queue. Fill in the vendor, dates, and value by hand and the record refiles instantly; your correction always outranks the automatic read.

A RE-DROPPED FOLDER SEEMS TO SKIP FILES

It found duplicates, and that is the feature: documents are deduplicated by content, not file name, so byte-identical files are counted and skipped, never imported twice. Two different documents sharing a name are both kept. The run's results name everything skipped and why; nothing is dropped silently.

AN INTERRUPTED UPLOAD DID NOT RESUME

A reloaded tab cannot resume files that never finished uploading, because the browser lost its handle on them. In the Archive, everything already uploaded still scans and files; re-drop the folder for the remainder (dedupe makes that safe). In Migration Studio, the work list survives on our side: **Re-pick the folder** so the new tab can find your files, and **Re-queue stalled files** if a run was cut mid-upload.

THE DROPPED PAPER HAS NOTHING TO BENCHMARK

A master agreement with no figures gives the benchmark nothing to price, and the run says so rather than guessing. Drop the document that carries the money, usually the order form, the quote, or the renewal proposal. The clause-reading tools (Contract decode, the MSA Decoder) are the right engines for the figure-less paper itself.

AN EMAILED-IN DOCUMENT NEVER APPEARED

1. Confirm the rail is provisioned and you used the right address: each of the five email-in rails has its own, listed under Integrations.
2. The attachment must be PDF or DOCX; a reply whose attachment cannot be read is flagged for a human look rather than filed wrong.
3. Each organization's inbound addresses carry unguessable tokens, so mail forwarded from outside your team to the wrong tenant's address is dropped, not misfiled.

YOU HAVE NO LIST PRICE FOR THE DISCOUNT READING

Leave it blank and proceed: the range still computes from the cohort, and the list value only adds your own placement to the scale. Where the vendor's paper prints net but no list, a dropped-quote brief states a working baseline, explains why it is credible, tells you exactly what to request from the vendor to verify it, and prices the sensitivity if it is wrong. That page is behaving correctly; the ask for the list-price breakdown belongs in your next vendor email.

THE RECOMMENDED OUTCOME READS NOT DERIVABLE

On Renewals, the market-median outcome carries the **recommended** flag only when a delivered benchmark supports it; without one, the desk refuses to guess and says *Not derivable*. Run an instant benchmark or upload the paper for analyst benchmarking and the outcome fills in.

A LONG AI RUN SEEMS STUCK

Most long runs are working, not stuck: briefs take about three minutes and show a progress meter with elapsed seconds against the usual duration, archive scans continue after you close the tab and notify you when done, and migration reading runs over hours in passes. The honest checks: a real progress meter that has genuinely frozen for several minutes deserves a reload (saved runs and cases keep your place), and a run that ended in failure states its reason rather than spinning.

THE REPORT OPENS WITH A REVIEW NOTICE

The trust strip on AI-generated reports opens with the grounding state. "Figures verified against your data" is the normal case; a review notice means one or more figures could not be fully verified and deserves your eye before forwarding. It is a flag, not a failure: the sources and generation date are beside it, and regenerating after fixing the underlying record usually clears it.

A FIGURE READS "NOT ON FILE"

The report or answer needed a number your records do not hold, and said so instead of estimating. The fix is data, not retry: add the missing annual value, date, or list price on the owning record. The Data gaps report in the Reporting studio catalog names exactly which missing values to fix first.

YOUR CURRENCY SHOWS DIFFERENTLY THAN YOU EXPECTED

Three rules explain nearly every case. Rows keep the currency they were entered in; the workspace display currency covers only rows that did not record one, and never converts anything. Benchmarks are held in USD and stay labeled USD so market comparisons hold. In benchmark tools that support it, picking a currency carries through the whole page and its exports at a fixed conversion rate printed under the selector with its date, so a saved position replays identically instead of drifting.

YOU HIT THE DAILY QUESTION METER

Fair-use caps protect the conversational surfaces, and you never hit one unwarned: once less than half a day's cap remains, a quiet meter appears near the input with the count left and the upgrade path. Caps reset at midnight UTC. Trial workspaces meter Vera at ten questions a day and always show the meter; Chapter 10 has every number.

NO RENEWAL ALERT HAS EVER FIRED

Working as designed until someone opts in: every renewal alert starts off, on every channel, so a workspace stays silent about deadlines until asked. Press **Turn on money-critical alerts** under Notifications to switch the set on in one click. This is per person; your colleague's alerts say nothing about yours.

ONE AGREEMENT NEVER ALERTS

Check the agreement carries **both** an expiry date and a notice period on Agreements. Without both there is no window to open, so nothing can fire, and the same gap keeps the agreement off the renewals queue and the control room timeline. The notice period hides in the contract's termination or renewal clause; a decode finds it fast.

THE CONTROL ROOM IS EMPTY

An empty Control room means the rows are missing, not that nothing is happening: only agreements with an expiry date and a notice period can appear on the pressure timeline. Fix the dates and the windows appear. The board's lanes fill only when agents, workflows, or mandates actually run in your workspace.

AN ALERT ARRIVED LATER THAN THE EVENT

Usually quiet hours: inside your chosen window, phones stay silent and non-critical alert emails wait until it ends, by design. Money-critical alerts (renewal deadlines, deadline escalations, uplift breaches, above-market warnings, anomalies) always come through immediately. Also check whether the alert kind is set to the weekly digest, which bundles it into one Monday email instead of individual sends.

THE DAILY LEVERAGE NOTE WENT QUIET

Two silent-by-design cases before any fault: on a day when nothing moved under your deals, nothing is sent; and the note rides your renewal alert setting, so switching renewal alerts off stops it. Every note also carries a one-click **Turn it off** link, worth checking you did not use in passing. Findings that scored below the line are counted at the note's foot rather than dropped, so a thin note is still a complete one.

THE CONFIDENTIALITY SCAN BLOCKED YOUR SHARE

Vendor-facing drafts and public share links pass a confidentiality scan, and a block lists exactly what to remove: an internal figure, a name, or a detail that should not leave the workspace. Edit the flagged items and share again. The scan is the reason a shared document can be trusted not to leak your percentile, your cohort, or your internals.

A WORD OR PDF LINK IS MISSING ON A ROW

Not every document type exports to every format: rows in the library show a Word link in the Formats column only where the type supports it, and reports carry their own PDF and Word downloads on the report page. If a type you need regularly lacks the format, tell the team; the export catalog in the reference card lists what exports to what today.

A VENDOR APPEARS TWICE

Different legal-entity spellings of one company (Salesforce, Salesforce UK Limited, SFDC Ireland Limited) can arrive as separate rows from different documents. The archive files one folder per vendor across entities, migration reconciliation proposes same-company merges with evidence, vendor aliases live under the renewal calendar's governance settings, and Vera OS's vendor view folds spellings into one negotiating party. If a duplicate survives all that, merge is a support ask away in Chat.

ENTITLEMENT ROWS ARE MARKED "CONFIRM"

Connector-reported entitlement figures are held out of alerting until a human confirms the contracted numbers, so a sync error can never fire a money alert. Confirm the figures against the contract on Billing watch and the daily check starts enforcing them.

A SHARED VIEW CANNOT BE EDITED

A view an owner shared with the whole organization stays owner-managed: it is pinned in everyone's rail with a **Shared** marker, and only its owner can rename, unshare, or delete it. Save your own copy under a new name if you need a variant.

YOUR FOLDS AND LAYOUTS RESET ON A NEW DEVICE

Some choices are deliberately per device rather than per account: rail group folds, the covers-or-table library toggle, reading-mode lighting, and your recent-apps trail. Your theme, preset, and notification choices follow your account. Two minutes of refolding on a new machine is the whole cost.

INVITING FAILS BECAUSE SEATS ARE FULL

The Members verdict states seats in use against your plan's allowance, and inviting past it says exactly what is full. Free a seat by revoking a pending invitation (the reserved seat returns entirely) or removing a departed member, or an owner upgrades the plan under Billing.

YOUR UPLOAD ENDED SAVED, NOT QUEUED

The upload worked; the analyst queue is what it could not enter, because your plan's benchmark allowance is used up. The proposal page resolves it in place: an owner adds a benchmark credit ($5,000 per analyst benchmark, billed by invoice), or a teammate asks their owner in one click, and the request tracks under Support while it is in flight.

TRIAL LIMITS SURPRISED YOU

The two that catch people: trial workspaces read research online but do not download it, and instant benchmark runs are capped at 2 in total (Salesforce Sales and Service Cloud and Adobe Experience Cloud only), so spend them on a vendor you are actually negotiating. Vera is metered at ten questions a day with the meter always visible. Everything else on the trial is the full platform with starter allowances; the plan table has it all.

SOMETHING IS BROKEN, OR NOTHING HERE HELPS

- Message our team from Chat: the "Vera AI team" thread is private between you and our analysts.
- Or email **info@vendorbenchmark.com**.
- For screen-by-screen help in place, the round help button at the bottom right explains the current page, and **Ctrl/⌘ J** opens Vera over whatever you are looking at.

Next: Chapter 12 · Integrations, clients & deployment

12 · Integrations, clients & deployment

The technical chapter, written for the IT, security, and procurement-ops people who wire Vera AI into an estate. It answers three questions: how data gets in and out (the integration catalog), what runs on your machines (thin and thick clients), and where the platform itself can run (the four deployment models). Nothing here is required to use the platform; everything here makes it stronger.

HOW THE PLATFORM IS DELIVERED

Where: the standard service · `app.vendorbenchmark.com`

Vera AI is delivered as a managed cloud platform (SaaS). The application, database, storage, and backups are operated by us; your team signs in from a browser and maintains nothing. The technical posture, in the terms a security review asks about:

- **Tenant isolation** — every organization's rows are isolated at the database layer with row-level security, tested with real cross-organization attempts, not filtered in application code.
- **File handling** — contract files live in private encrypted storage. There are no public URLs; every download uses a signed link that expires in about a minute, and every view and download lands on the audit trail.
- **Encryption** — TLS in transit, encryption at rest for the database and file storage.
- **Identity** — SAML and OIDC single sign-on, SCIM provisioning, and TOTP multi-factor are available on the standard service (see Single sign on).
- **Auditability** — an exportable audit log records who did what, including file views and downloads, and can stream to your SIEM.
- **Updates** — continuous. Features, security fixes, and fresh benchmark data ship without any action from you.

The Trust center tells this story with live numbers, and the Security posture shows your own workspace's posture against it.

CLIENTS: WHAT RUNS ON YOUR MACHINES

The platform is thin-client first: the browser is the primary interface and nothing needs to be installed. Around it sit a thick-client desktop app and several headless surfaces.

The browser (thin client)

Who: everyone

Any modern browser is a full client. All processing happens server side, with one deliberate exception: the browser-local calculators. Several tools (Cloud cost optimizer, VCF licensing, M365 optimization) parse your usage export as a file **in the page itself**, so raw usage data never leaves your machine unless a step explicitly sends approved figures for analysis; each tool states which side of that line it works on. Files sent to the AI tools are read in memory and never stored.

The Windows desktop app (thick client)

Who: everyone · install is per user, no admin rights

A native shell around the same account and data, for teams that live in the platform daily: its own window and taskbar identity, drag a contract from Explorer straight onto it, native Windows notifications for deliveries and notice windows, Vera summonable over any app, Windows Hello lock, and automatic updates. Install it from Settings › Desktop app. It is a client, not a data store: uninstalling it touches nothing in your workspace. A macOS build is planned; until then Mac users get the identical workspace in the browser.

Chat clients: Vera in Slack and Microsoft Teams

Who: connected by an org owner

Vera answers inside your own Slack workspace or Teams tenant: ask about a price, a clause, or a renewal without leaving chat, with every figure tagged to its source exactly as in the app. Connected under Integrations.

Email as an interface

Who: addresses are provisioned by an org owner

Five inbound rails give every teammate an integration that needs zero training, because it is just email:

- **Contract email-in** — forward a vendor proposal to your intake address and it lands in Proposals.
- **Deal email-in** — forward a live vendor thread onto the deal record, where the negotiation ghost writer can draft the reply.
- **Invoice email-in** — forward invoices and they reconcile billed against contracted on Invoices.
- **Purchase request email-in** — teammates raise Requests by mail, straight into triage.
- **Ask an agent by email** — send a question plus a document and the answer comes back to your inbox.

The API, MCP, and agent-to-agent

Who: org owners only (keys)

- **REST API** — create keys under Settings › API; keys expire within a year and are shown once.
- **MCP connector** — point Claude, ChatGPT, or any MCP-capable tool at your workspace so your own AI assistants can query your benchmark data with your permissions.
- **Agent Negotiation Protocol** — the open, published protocol for structured buyer-agent-to-vendor-agent negotiation, with a hash-chained transcript; the spec and artifacts are public at `app.vendorbenchmark.com/agent-protocol`.

Vera OS, the autonomous procurement companion

The lens now starts work as well as reading it: on the grounding screen, **Queue the top gaps** creates up to five open tasks titled Benchmark that vendor (never duplicated) and links your Tasks runway, and every deal row in the Control Tower and the deal rooms carries **Hand to a case**, opening the Vera AI-run negotiation intake with that agreement preselected.

> **Where:** the Vera OS page · `/vera-os` · reachable from the Renewals & deals rail group and the launcher home
> **Who:** everyone; Vera OS is in beta and included with your plan, nothing to request

Vera OS is a separate product: an autonomous procurement operating system that runs as its own workspace with its own sign-in. Agents there work renewals and deals end to end under mandates you sign once, evaluated in deterministic code on every step that touches money: a renewal watchtower that opens a deal room for every term ending inside 180 days with the benchmark preloaded, cited decision briefs you approve with one tap against a 24 hour SLA, co-pilot drafting where a human always sends, and an autopilot lane that closes tail spend inside two deterministic gates with a human veto window before anything settles. Every deal is readable as a case file with a replayable trace.

How it relates to this platform. Vera OS consumes your VendorBenchmark benchmarks, playbooks, and Terms Watch through the platform API, and contributes closed-deal datapoints back anonymized, released only once at least five exist for a vendor. It replaces nothing here: every view in Vera AI stays exactly where it is.

Using the page. Vera OS is open to every workspace with no access request. The navy stage leads with **Launch the Control Tower**, which opens immediately.

The Control Tower preview. · `/vera-os/tower`

The Control Tower reads your own estate the way Vera OS reads it, live and read-only: a metric strip (deal rooms to open inside 180 days, the exposure in that window, auto-renews, the estate run rate, vendors with a delivered benchmark) over a deal-rooms table of every agreement with an expiry, soonest first. The suggested lane column is a stated deal-size rule, command at $250,000 and above, co-pilot at $25,000 and above, autopilot below, never a model's judgment, and a row opens that agreement's renewal war room. Nothing is written back. Once your full workspace link is online, an Open the full workspace button appears on both pages and opens Vera OS in its own tab.

The Vera OS menu. The dark menu down the left of the Control Tower is the frame around the whole workspace, not part of one page: it stays with you as you move between screens, stays put under the top bar while a long table scrolls, and marks the screen you are on. On a narrow screen it becomes one scrolling row of the same map above the page. It holds three groups.

- **Workspace** — Control Tower, and the three screens below.
- **Grounding** — what the lanes are allowed to stand on: Benchmarks and The Vera Wire. The delivered reports used to sit here as a third screen; they now live inside Benchmarks, on the vendor each one grounds.
- **Full workspace** — Decision queue, Mandates, Studio · policy, Outcome, Network. These five run in the Vera OS workspace itself rather than here, so they read **soon** until your workspace link is connected, at

which point each one opens it in a new tab.

Deal rooms. · `/vera-os/deals`

Every term Vera OS would open a room for, soonest first, with the lane its size puts it in, the notice date that actually binds, the exposure, and whether a delivered benchmark grounds it. The lane chips and the **Show every term** link filter the table and are written into the address, so a filtered view can be sent to a colleague. Rows open the renewal war room in Vera AI. When nothing ends inside 180 days the table falls back to the whole term book and says so.

Agreements. · `/vera-os/agreements`

The whole estate, largest first, with the annual value, the term end, the notice date, whether it auto-renews, and the benchmark behind it. Agreements carrying no term date are marked in red and counted under the table: no room can ever open for them, so an auto-renewal on one would pass unseen.

Vendors. · `/vera-os/vendors`

The estate rolled up to the party you negotiate with. Legal-entity and country spellings of one company (for example Salesforce, Salesforce UK Limited, and SFDC Ireland Limited) fold into a single row with the spellings listed beneath it, and the lane follows the total held with that party rather than the smallest line.

Benchmarks. · `/vera-os/benchmarks`

The grounding, and the gap. A lane may only quote a number a delivered benchmark supports, so this screen leads with the vendors that have none, ranked by what they are worth and flagged when their term ends inside the window.

Beneath the gap sits **Grounded vendors, the delivered reports**: one row per grounded vendor with its annual value, when the term ends, when the report was delivered, a link to the report itself, and the headline figure a lane would quote. That figure is the overall score the analyst gave the agreement, out of 100, or **unscored** where the report carries no score; it is never a price or a saving. A metric tile above counts the delivered reports this workspace holds, and drafts are never counted and never quotable. **Open the document library in Vera AI** goes to the full reports shelf.

This absorbed the separate delivered-reports screen. `/vera-os/reports` now forwards here, so an old link or bookmark still lands somewhere useful.

The Vera Wire. · `/vera-os/wire`

The market feed filtered to vendors you actually hold a term with, each item with its verdict, Vera's read, and the date, and the rest of the market kept short beneath it. The full feed with its lenses stays at The Vera Wire in Vera AI.

Items on your own estate carry two extras the market items do not. **The deal behind it** states the clock on that vendor's nearest term: how many days until notice closes or how long ago it closed, and how long the term still has to run, reading **no term date on the paper** where the agreement holds no dates. **Open the war room** goes straight from the market move to the renewal record it lands on. Both appear only where a term end is on file and not more than 30 days past, so a vendor with undated paper gets neither.

THE INTEGRATION CATALOG

Where: Settings › Connections › Integrations · `/settings/integrations`
Who: org owners only

The page is one calm scroll, arranged the way this section is: contract intake first, then renewal workflow, spend and usage data, data out, and automation, with the two deep estates, procurement (P2P) and license discovery (SAM / ITAM), at the foot. Every connector ships with the same operating rules: credentials are stored encrypted and used only server side, read-only scopes are requested wherever the source system offers them, every sync is logged under **Recent activity**, and connector-reported entitlement figures are held for human confirmation before they can drive an alert (you will see them marked "confirm" on Billing watch). Disconnecting a source stops its syncs immediately.

Contract intake

Get the paper in without anyone uploading by hand. Half of these connectors are just private email addresses, provisioned in one click and shared with your whole team.

CONNECTOR	WHAT IT DOES
Email-in	A private intake address; forwarded proposals become tracked contracts
Deal email-in	Forward a live vendor thread onto the deal record; the ghost writer can draft the reply
Invoice email-in	Forwarded invoices reconcile billed against contracted on Invoices
Purchase request email-in	Teammates raise Requests by mail, straight into triage
Ask an agent by email	Send a question plus a document; the answer returns to your inbox
DocuSign	Completed envelopes flow in as signed agreements
Ironclad	Executed contracts sync from your CLM
Google Drive folder watch	Drop files in a folder, they arrive here
SharePoint folder watch	Same, for Microsoft estates
Box folder import	Bulk import from a Box folder

Renewal and deal workflow

Push the work to where your team already lives.

CONNECTOR	WHAT IT DOES
Jira	Renewal tasks and findings as issues in your project
ServiceNow (tickets)	The same, as ServiceNow tickets
Calendar invites	Notice windows and renewal deadlines as calendar events
Calendar feed (meeting packs)	A read-only ICS feed; vendor meetings get a pre-call pack
Slack	Alerts to a channel, plus Vera in your workspace
Microsoft Teams	Alerts to a channel, plus Vera in your tenant

Spend and usage data

The money and the activity behind it: what you pay, and who actually uses what you pay for.

CONNECTOR	WHAT IT DOES
Coupa	Spend actuals by API
Okta usage	Who actually signs in to which app
Microsoft Entra ID usage	The same, for Microsoft estates
Spend actuals import	A CSV from Coupa, SAP Ariba, Zip, or a NetSuite GL export, or any generic spend CSV
Seat usage import	A CSV of per-app seat activity
Smart import	One-off bulk loads: the AI reads whatever columns your export has and maps them for you

Data out and automation

Your data is yours to take with you, continuously.

CHANNEL	WHAT IT CARRIES
Warehouse export	Contracts, renewals, savings, and spend tables to Snowflake or BigQuery
Event webhooks	Signed JSON events (deliveries, status changes, findings) to any endpoint, including Zapier or Make; verify the signature before trusting a payload
Audit webhook	The audit log streamed to your SIEM
File exports	Every report as PDF, Word, or PowerPoint; ledgers and tables as CSV or Excel

The integrations page also carries the **Agent Negotiation Protocol** card: an org owner opts your workspace into structured agent-to-agent negotiation there, and the public spec lives at `app.vendorbenchmark.com/agent-protocol` (see The API, MCP, and agent-to-agent).

Procurement and ERP (P2P)

Purchase orders, requisitions, and supplier records, so spend that never touched a contract still shows up, with maverick spend flagged. Beneath the connector cards, the **Procurement intelligence** table turns the synced records into a supplier-by-supplier read.

CONNECTOR	SOURCE
Coupa	API connection
SAP Ariba	API connection (US, EU, APAC endpoints)
NetSuite	Token-based API connection
Oracle Fusion Cloud Procurement	API connection
SAP S/4HANA	API connection
File exports	Coupa, NetSuite GL, Ariba, Zip, or CSV when an API connection is not approved

License and deployment discovery (SAM / ITAM)

Connect the systems that know what you own and what is deployed; the **License intelligence** table beneath the cards then shows entitled against deployed against active per vendor, with estimated shelfware. This is the data that powers Data health, Billing watch, and audit exposure. Each connector card opens on a click to show its credential form and carries a mode badge: API connection, file import, or both.

CONNECTOR	SOURCE
ServiceNow SAM Pro	Entitlements and license positions
Flexera One	Entitlements and deployments (US, EU, APAC endpoints)
Snow License Manager / Snow Atlas	Entitlements and usage
USU / Aspera SmartTrack	Entitlements
Microsoft 365	Seat assignments and last-activity via Microsoft Graph
Microsoft Intune	Device and app inventory
Jamf Pro	Apple estate inventory
ServiceNow CMDB	Deployed software inventory
Oracle LMS / ULA position	Your measured Oracle deployment position
SaaS management platforms	Zylo, Zluri, Torii, or Productiv, via their CSV exports
Cloud billing	AWS Cost & Usage Report, Azure cost export, or Google Cloud billing export, feeding Spend and the Cloud cost optimizer

How to use it

Every connection named in this catalog has a step-by-step setup recipe, with a verification line and a teardown, in Chapter 19 · Integration recipes. The order that pays best:

1. Start with identity: SSO and SCIM first, so access is governed from day one.
2. Open the cheapest intake rails: email-in addresses and a folder watch cost minutes and remove the manual-upload habit.
3. Connect one SAM or usage source and one spend source, then follow Data health, which ranks every remaining gap by the feature it unlocks.
4. Turn on data out last: the warehouse export and the audit webhook make the platform a citizen of your data estate rather than an island.

Related: Integrations (settings reference) · API · Connect your data and close the gaps

VERA AI WITH YOUR SAM TOOL

Who: owners connect the sources · everyone benefits from the joined data

Your SAM or ITAM platform and Vera AI know different halves of the same estate. Flexera One, ServiceNow SAM, Snow License Manager, USU, and Xensam know what you own and what is deployed: entitlements, installations, active users, license positions. Vera AI knows what it all costs and what it should cost: the contracts themselves, the prices against the market, the renewal mechanics, and the negotiation. Connect the two and each side covers the other's blind spot.

Data flows in. Under Settings › Integrations, the discovery connectors pull license positions from your SAM tool, live over its API where one is exposed (ServiceNow, Flexera One, Xensam, Microsoft 365, Intune, Jamf) or from its standard export file (Snow, USU, the SaaS management platforms). Every source lands in the same License intelligence table: entitled versus deployed versus active per vendor, priced where a unit rate is known. That is where shelfware stops being a suspicion and becomes a number.

Data flows out. Report extract writes your contract register back the other way, in the agreement layout your tool imports: one workbook shaped for Xensam, Flexera One, ServiceNow, Snow, or USU, or your own template filled under its own headers. The ServiceNow connector goes further on its own: it can create contract records here from entitlements it discovers, and write delivered benchmark verdicts back onto the source records as work notes.

TOOL	WHAT FLOWS IN	WHAT FLOWS OUT
Flexera One	License names, entitlements, and consumption over the API, or an ELP export	The contract register in the Flexera One contract layout
ServiceNow SAM	Entitlements from alm_license, installs from the CMDB, optional auto created contracts	The register in the ast_contract layout, plus benchmark verdicts as work notes
Snow License Manager	Compliance report exports	The register in the Snow agreement layout
USU (Aspera)	License position exports	The register in the USU contract layout
Xensam	The software list over the Xupervisor REST API, or a CSV export	The register in the Xensam agreement layout

The goal when both are connected. One closed loop instead of two half informed systems.

1. **Shelfware gets priced, then negotiated.** The SAM side proves 1,500 seats entitled and 900 active; the Vera side knows the unit price and the renewal date. Together that is a true down argument with a number on it, delivered before the notice deadline rather than after the auto renewal.
2. **Renewals argue from usage evidence.** A benchmark says where the price should land; the deployment data says how much of the product you actually consume. A vendor hears both at once, and the second one is the harder to argue with.
3. **Neither register goes stale.** The SAM tool stays the operational system of record for the estate; Vera stays the commercial one for what it costs. The connectors keep the first flowing in, Report extract keeps the second flowing out, and nobody re keys agreements by hand in either direction.

Start with whichever direction closes your bigger gap: pull the inventory in if you negotiate blind on usage, push the register out if your SAM tool's agreement list is empty. Most workspaces end up running both.

Related: Report extract · The integration catalog · Integrations (settings reference)

DEPLOYMENT MODELS

Who: Cloud is the standard service; the other three are Enterprise conversations

Most teams run the standard Cloud service and are live the same day. For organizations whose policies require more, three further models exist; they are scoped with your security team rather than sold as checkboxes. Bring your requirements to the Advisory desk or the team Chat, and the full comparison lives at vendorbenchmark.com/deployment.

01 · Cloud (available today)

The managed multi-tenant platform described at the top of this chapter. Live the same day, always current, full benchmark network access, Enterprise controls (SSO, SCIM, MFA, audit) on by default, and the lowest cost of the four. The trade-offs are the ones inherent to shared infrastructure: isolation is enforced in the database rather than by separate hardware, the hosting region is ours, and contract files are processed outside your network perimeter, encrypted and access-logged.

02 · Dedicated cloud (Enterprise agreements)

A single-tenant deployment provisioned and operated exclusively for you: dedicated database, dedicated file storage, dedicated encryption keys, your own subdomain or custom domain, hosted in the region your policy requires. It answers the two requirements shared multi-tenancy cannot: data residency mandates, and security policies that rule out a shared database regardless of how it is isolated. Still zero operations work for your team; updates arrive on an agreed maintenance cadence rather than continuously, and it is priced as a dedicated environment above the standard plans. Time to first benchmark: days.

03 · Hybrid connector (Enterprise early access)

For a hard "documents cannot leave the network" policy. A lightweight connector (one container) runs inside your environment and reads contracts locally: it extracts the commercial terms, shows your team exactly what it found, and sends **only the figures you approve** to the cloud for benchmarking. The documents themselves never cross your boundary, and you still get the full market comparison, negotiation targets, and renewal calendar. Document-centric features (clause-level review, full-text document search) run reduced or stay local, and your team operates and updates the one connector. Time to first benchmark: weeks, including your security review.

04 · Self-hosted (by arrangement)

The full platform inside your own boundary, for defense, government, and the strictest financial mandates. We deliver versioned releases your team deploys into your own cloud account or data center, with our runbooks and support; your team owns upgrades, backups, scaling, and restores, audit events stream straight into your SIEM, and market data arrives as periodic snapshots rather than a live network connection. It is the most control and the most work of the four, which is why it is scoped as a joint engagement. Time to first benchmark: months.

The four at a glance

	CLOUD	DEDICATED CLOUD	HYBRID CONNECTOR	SELF-HOSTED
Best for	Most teams	Residency mandates, no shared database	Documents cannot leave the network	Air-gapped and sovereign environments
Who operates it	Us	Us, on your dedicated infrastructure	You run one connector, we run the rest	Your team, with our support
Contract files live	Our encrypted cloud storage	Your dedicated storage, your keys, your region	Inside your network, never uploaded	Your infrastructure, entirely
Updates	Continuous	Agreed cadence	Cloud continuous, connector by you	Versioned releases you apply
Benchmark network	Full	Full	Full, for approved figures	Periodic snapshots
Time to first benchmark	Same day	Days	Weeks	Months

Related: Trust center · Plans, allowances & limits

Next: Chapter 19 · Integration recipes for the step-by-step setup of everything cataloged here, or back to the manual index.

13 · Your AI agents

The platform runs a workforce on your behalf: analysts that write documents, watchdogs that monitor your portfolio and the market, autopilots that act under a mandate you set, and the workflows your own team composes. This chapter explains the four classes, lists every agent, and walks you through building one. The one rule behind all of them: an agent must show its work. Every run is logged, every figure is grounded in your stored data, and nothing leaves the platform without a gate you chose.

THE AGENT CENTER

Where: Team & AI › Agents · `/agents`
Who: everyone

The operations desk over every agent working for your organization: who is on duty, what they did while you were away, and where every piece of work lands. The Center is an application with five views on one masthead, and it presents agents as a staff: a roster you manage, a bench you hire from, and a work ledger that proves the output. Press 1 to 5 to switch views and / to jump to Ask Vera; each view is addressed (`/agents?view=roster`), so a forwarded link lands on the right one.

In plain words: the staff page for your AI workforce: who works for you, on what clock, with what proof.

When to reach for this

- The weekly "what did the machines do?" pass: the work ledger answers it with output, not activity.
- Deciding what to delegate next: the bench and the Ask Vera box on this page suggest from your own workload.
- Something automated went quiet: the roster shows every clock and its last run.

The masthead

- **Your workforce** — the headline is **Returned a year**: real money, in currency, from verified outcomes only, meaning closed case certificates and confirmed outcomes your agents drafted. It is counted per currency and never summed across currencies, and the line beneath states how many verified outcomes are behind it. Where nothing has been verified yet the band shows **Completed runs** in its place and says money lands here once an outcome is verified, rather than putting an estimate where a measured figure belongs. Beside it: **On duty**, split between the standing agents and the ones your team built; **Runs, 30 days**; and **Waiting on you**, which is runs paused until somebody decides. Three controls sit beside them: **Which agents do I need?** (a four question interview that recommends a starting roster), **Meet your staff** (the introduction, replayable any time), and **New agent**.
- **The time estimate** — one labeled line on the work ledger, never a headline figure: your completed runs valued at 20 minutes of manual work each. It says on its face that it is an estimate rather than a

measurement, and points at the verified money on the ledger rows themselves.

- **What "on duty" counts** — the standing watchdogs and autopilots sweep your estate on the platform's clock whether or not you have rostered them. What the roster decides is whether they may tell you anything: an agent you have not put on your roster is silent on every channel, in-app, email, and chat alike. Nothing alerts you until you ask for it.

The five views, on the masthead's band

- **Overview** — the floor at a glance. **Ask Vera** leads: describe your situation in a sentence and Vera answers with a numbered shortlist of the agents that cover it (see Ask Vera which agents you need). Under it, the four departments, one tile per class with its verb: Analysts write, Watchdogs monitor, Autopilots act, Workflows compose. Then the latest entries of the work ledger and **Today on the floor**, the agents scheduled to run on today's date.
- **Your roster** — the agents you chose to lead with, plus every workflow your team built. Workflow cards carry a 14 day heartbeat sparkline and their run count. The tab's badge counts them.
- **Hire agents** — **Start from your situation** first: six named situations ("A renewal lands inside 90 days", "I want every invoice checked") that put the right bundle of agents on your roster in one click. Then the bench: every other agent, grouped by class, each one add away. A benched watchdog still sweeps your estate, but it reports nothing to you until you roster it, so the bench is silence rather than absence.
- **The week** — a week with your agents, read from the schedule the platform actually deploys, not from a label, so a day with nothing scheduled says so. Agents that run monthly, on an event, or only when you ask are counted in the line beneath rather than pinned to a weekday they do not keep.
- **The ledger** — everything your agents did, newest first: watchdog findings and workflow runs, each row wearing its agent's monogram, with a link into the full run history.

The rail, in reach from every view

- **Waiting on you** — appears only when workflow runs are paused on an approval, naming each run and how long it has waited, with a link to review and approve.
- **What is being watched for you** — the coverage card: price lists, billing, standings, anomalies, vendor health, and paper hygiene, each with its cadence and a link to the detail.
- **How staffing works** — the three rules in short: the sweeps run regardless, the roster is your alert list, and the autonomy dial sets how far an agent may go alone.

Every card shows the agent's monogram, class, and cadence, an arrow sentence (what it reads, what it does, where the work lands), an example find, a run receipt, and links to its dossier and its output page. Agents that can act (the autopilots) carry the autonomy dial: **Watch**, **Draft**, or **Act**, one vocabulary everywhere.

The run receipt is the card's proof of work: when the agent actually last ran, read from the platform's own scheduler heartbeats, and when it runs next, computed from its real schedule ("Last ran 2 hrs ago · next today at 06:00"). Event driven agents state their trigger instead ("Runs when you forward a vendor email"), and the on demand analysts say they run when you ask. When a last run cannot be derived yet, the card shows only the next run; a timestamp is never invented, and an agent nothing actually runs is not advertised in the Center.

How to use it

1. Scan the masthead; anything waiting on an approval is your queue, not the machine's, and it stays pinned in the rail on every view.
2. Read the ledger, on the Overview or in full on its own view, to see what ran while you were away, and open any row.
3. Not sure where to start? Tell **Ask Vera** what is going on, use **Which agents do I need?**, or pick a situation card on the Hire agents view; all three put a sensible bundle on your roster.
4. Use New agent to add to the workforce.

Related: Workflow agents · Vera AI

Ask Vera which agents you need

Where: Team & AI › Agents, the box under the workforce band · `/agents`
Who: everyone; putting agents to work needs edit access

Both staffing doors now ground their recommendations in your workspace counts (renewals inside 90 days, agreements with no renewal owner, open anomaly findings, vendors on file), and each recommended agent states its reason from those counts. With no data on file the recommendations stay generic.

Most people arrive holding a problem rather than shopping for agents. Say the problem in your own words and Vera staffs it: "our Microsoft renewal lands in 60 days", "I think we are being billed for seats we dropped", "leadership asks about software spend every month".

On the screen

- **The question box** — one line, with four starter situations underneath it until you have asked something.
- **Vera's answer** — two to four sentences on why these agents fit your situation and what will land in your inbox.
- **The numbered shortlist** — up to five agents, best fit first, each row showing its number, monogram, what it does, and how often it runs. An agent already on your roster is marked **On roster** and cannot be added twice. Every other row carries a checkbox, ticked by default.
- **Put N on my roster** — activates the ticked rows. The count follows what you tick.
- **Built for you** — appears only when nothing on the shelf covers what you asked for: Vera drafts a custom workflow agent, saves it, and shows you its chain of steps. It arrives switched off, so nothing runs until you press **Turn it on**. **Edit every step** opens it on the workflow desk.

How to use it

1. Type your situation and press **Ask Vera**, or tap one of the four starters.
2. Read the shortlist. Open any agent's dossier first if you want the detail before deciding.
3. Answer with the numbers you want, for example `1, 2, 3 and 4`, or `all of them`. Ticking the checkboxes and pressing **Put N on my roster** does exactly the same thing.
4. The agents join your roster immediately and start on their next scheduled run. Tune the leash or bench any of them from the roster below.

Good to know

- Vera recommends only agents that exist and are genuinely running; she cannot invent one or offer one the Center hides.
- She never recommends an agent you already have. If you pick one anyway, she says it is already covered rather than adding it twice.
- Answering with numbers is resolved against the list on your screen, so the agents you name are always the ones you get.
- Ask Vera needs the AI plan. Without it, the four question front desk interview still recommends a starting roster.
- Viewers see the same recommendations but cannot switch agents on; ask an owner or a member with edit access.

AGENT DOSSIERS

Where: Team & AI › Agents › any agent · `/agents/{agent}`
Who: everyone

Every agent has a home page: what it reads, when it runs, what it delivers, and your settings for it, in the same record page grammar as the rest of the app.

On the screen

- **The masthead** — monogram, class, cadence, the agent's one line, and the See its work button.
- **Highlights** — what it reads (as chips), what it does, where its work lands, and your setting: roster membership and, for agents that can act, the autonomy dial.
- **Example output** — a curated miniature of the agent's deliverable, always labeled as an example, never your data.
- **Recent findings** — for watchdogs, what the agent flagged for you lately, straight from your notifications.
- **Part of these situations** — the situation bundles this agent serves.

How to use it

1. Open a dossier from any card, ledger row, or week strip pill.
2. Read the example output to understand the deliverable in ten seconds.
3. Add the agent to your roster, or set its autonomy, from the highlights panel.

BUILD AN AGENT

Where: Team & AI › Agents › New agent · `/agents/new`
Who: everyone · building workflows needs edit access · drafting needs an AI plan

One wizard for all three buildable classes. Pick the job, and the wizard routes you: document writing goes to the custom analyst builder, monitoring shows you the watchdogs already on duty, and routine automation drafts a workflow from one sentence.

In plain words: describe the job; the wizard picks the right kind of agent and builds it.

When to reach for this

- The recurring chore worth a sentence: "every Monday, summarize what changed on our vendors and post it to Slack."
- A document you keep commissioning by hand deserves its own analyst.
- Check the roster first: the watchdog you are about to build may already be on duty.

On the screen

- **Three doors** — Write documents for me (Analyst), Watch something and tell me (Watchdog), Do the steps I always do (Workflow).
- **Describe it in a sentence** — for workflows: state the goal in plain language and the AI drafts the chain: trigger, steps, and the approval gate where anything would leave the platform. Three template chips above the box (review every proposal upload, the weekly renewal memo, routing vendor emails) prefill a proven goal you can edit, so you never start from a blank box.
- **Review the draft** — the drafted chain as a step map and a numbered list. Rename it, change the trigger, or rewrite the goal and draft again. Full per-step settings live on the workflow desk after saving.
- **Test with my data** — runs the chain end to end without sending anything outside the platform, so you read the output before the agent earns a trigger.
- **Turn it on** — enables the agent and shows you exactly where its work will land.

How to use it

1. Pick the door that matches the job.
2. For a workflow: describe the goal in one sentence. A goal that names the vendor, the cadence, and the audience drafts best.

1. Press **Draft the agent** and read the drafted chain: the name, the trigger, the step map, and where the approval gate sits. Rename it, change the trigger, or rewrite the goal and draft again.

1. Save the agent, test with your data, read the run on the workflow desk, then turn it on.

Tips

- Drafts are conservative on purpose: anything that emails a person outside a pure internal note gets a Wait for approval step. Remove it on the workflow desk only when you trust the chain.
- A goal that names the vendor, the cadence, and the audience drafts best: "Every Monday, review our Microsoft position and email the owners a short memo when anything moved."

THE FOUR CLASSES

Analysts write documents. On demand (the six playbook specialists plus up to six custom analysts your org defines) or on a schedule (the Monday vendor wire, the meeting pack before each vendor call, the monthly board report, the spoken monthly briefing). Their output lands in your documents library or on the page named on their card.

Watchdogs monitor and flag. They run on the platform's clock, read only the data named on their card, and surface findings on their home pages, in My work, and by the email digests you turn on. You do not build watchdogs; you tune what each one tells you from its own page.

Autopilots act under a mandate. The Renewal Autopilot works your renewal calendar at the autonomy you choose: watch and remind, draft for your approval, or negotiate under a mandate with fees only on invoice-verified savings. The negotiation ghost writer drafts vendor replies from forwarded email threads. Every action an autopilot takes is logged on the record it touched.

Workflows are yours. A trigger plus up to eight steps: AI analysis, classification, extraction, an if then gate, a wait for approval, then email, Slack, a reply to the sender, or a task. Build them in the wizard or by hand on the workflow desk.

THE ROSTER

Every standing agent the platform runs for you, the clock it runs on, and where its work lands. Workflows are not listed because they are yours: the Center shows each one your team builds, with its own trigger and run history.

Watchdogs

WATCHDOG	CADENCE	FINDINGS LAND ON
Portfolio watchdog	Weekly, Tuesday	Anomalies
Terms watch	Daily	Watch
Billing watch	Daily	Billing watch
Benchmark alerts	Daily	Notifications
Earnings radar	Weekly, Wednesday	Vendors
Vendor viability	Weekly, Wednesday	Vendors
Ecosystem watch	Weekly, Thursday	Vendors
Contract hygiene	Weekly, Saturday	Data health
Commitment chase	Weekly, Monday	Negotiations
Standing answers	Daily	Vera AI

Autopilots

AUTOPILOT	RUNS	WHAT IT DOES
Renewal Autopilot	Daily	Works your renewal calendar at the autonomy you set (Managed renewals)
Negotiation ghost writer	On every forwarded email	Drafts vendor replies from forwarded threads (Workflow agents)

Analysts

ANALYST	RUNS	OUTPUT LANDS ON
The six playbook specialists: Benchmarking, Contract Flexibility, Legal Terms, License Specialist, Executive Brief, Negotiation Talking Points	On demand	Your documents library
Your custom analysts, up to six	On demand	Your documents library
Weekly vendor report	Every Monday morning	Weekly report
Meeting pack analyst	Before each vendor meeting	Meeting war room
Board reporter	Monthly, on the 1st	Delivered reports
Monthly briefing, spoken	Monthly, on the 2nd	Monthly briefing

TRUST AND GUARDRAILS

- **Grounded output.** Analyst and workflow memos read your live portfolio data; figures are checked against stored records, and missing facts come back empty, never invented.
- **Approval gates.** A Wait for approval step pauses any workflow run until a named human decides; rejections end the run. Autopilot mandates cap what can be conceded before a human is pulled in.
- **Run records.** Every workflow run keeps a per-step record you can open. Watchdog findings link to the evidence behind them.
- **Plan limits.** AI-marked agents need the org's plan to include AI, and runs count against your plan allowance (Plans, allowances & limits).

TROUBLESHOOTING

- **It did not run.** Check the agent's card in the Center: a paused status dot means it is switched off; a workflow with a weekly trigger runs about once a week, not at a fixed hour.
- **It ran but the output is wrong.** Open the run record and read the step where it went wrong; tighten the instructions on the workflow desk, or for analysts, edit the custom analyst's brief.
- **Too noisy.** Watchdog digests have per-page switches; workflows can gain an if then gate so only significant runs notify anyone.
- **Nothing sends.** A run showing Waiting for approval is paused on a human; the approvers were emailed and can also decide from the run page.

Which tool do I use?

The platform has more than a hundred features, and nobody should have to learn that map. This chapter is the finder: every feature listed under the job it does, with the situations it answers phrased the way you would say them. Scan the group that matches what you are trying to do, find your situation in the left column, and the right columns name the tool and where it lives.

Five faster ways to the same answer:

- **The Tool finder screen.** Open Tool finder from the Tooling group in the rail: the same index as this chapter, with a job dropdown, a situation search, and a Start button on every result.
- **Ask Vera.** Press Ctrl/⌘ J anywhere, or the Ask Vera button at the bottom right, and say what you are trying to do in your own words. She names the one tool that does it and links you straight there.
- **The command palette.** Press Ctrl/⌘ K and type what you want. The palette matches tools by what they do, not only by their names.
- **The Verdict Bar.** If your starting point is a document and you do not know which analysis it needs, drop it at Verdict Bar. It is classified and routed to the matching engine automatically.
- **The More catalog.** The **More** entry in the left rail lists every destination in the workspace, grouped and searchable.

Routes below are workspace paths: open them from the rail, the palette, or by appending them to `app.vendorbenchmark.com`.

The index currently covers 113 features and 222 situations.

CHECK A PRICE OR A QUOTE

You have a number, a quote, or a proposal, and you want to know how it compares to the market.

IF THIS IS YOU	USE	ROUTE
I have a document and no idea which tool fits	**Verdict Bar**	`/verdict`
I want one fast verdict on whatever the vendor just sent	**Verdict Bar**	`/verdict`
I want the platform to pick the right analysis for me	**Verdict Bar**	`/verdict`
I want to know if a vendor's price is fair before I reply	**Instant benchmarks**	`/benchmarking`
I want the market discount range for a specific vendor	**Instant benchmarks**	`/benchmarking`
I am budgeting and need a sanity price for a product we do not own yet	**Instant benchmarks**	`/benchmarking`
A deal is on my desk and I want the fast market read before I engage	**Second opinion**	`/tooling/second-opinion`
I think the price is wrong but I need evidence to say so	**Second opinion**	`/tooling/second-opinion`
I want a brief I can forward to my boss about a live quote	**Second opinion**	`/tooling/second-opinion`
I only have a screenshot or photo of the quote	**Snap a quote**	`/tooling/snap-a-quote`
I am in a meeting and need a verdict in under a minute	**Snap a quote**	`/tooling/snap-a-quote`
I want the counter to send back without retyping the quote	**Snap a quote**	`/tooling/snap-a-quote`
I have quotes from the vendor and a reseller and want to know who wins where	**Quote Face-Off**	`/tooling/quote-faceoff`
I want to build the best combined deal from several quotes	**Quote Face-Off**	`/tooling/quote-faceoff`
I need a memo saying which quote to take and what to pull from the losers	**Quote Face-Off**	`/tooling/quote-faceoff`
The vendor sent a revised proposal and I want to know if it actually improved	**Proposal comparison**	`/tooling/proposal-comparison`
I am choosing between two competing vendors' proposals	**Proposal comparison**	`/tooling/proposal-comparison`
I want proof of what quietly changed between draft one and draft two	**Proposal comparison**	`/tooling/proposal-comparison`
I want to know if these consulting day rates are above market	**Rate Card Benchmark**	`/tooling/rate-benchmark`
I want blended rates deconstructed into the implied staffing pyramid	**Rate Card Benchmark**	`/tooling/rate-benchmark`
I recorded our MSA rates and want new SOWs checked against them	**Rate Card Benchmark**	`/tooling/rate-benchmark`
I want every vendor's market position in one place, not one check at a time	**Portfolio benchmark**	`/benchmarking/portfolio`
I want to know which contracts are furthest above market	**Portfolio benchmark**	`/benchmarking/portfolio`
I want one report covering the whole estate's market position	**Benchmark all my vendors**	`/reports/vendor-benchmark`
I need the annual state-of-the-stack read for leadership	**Benchmark all my vendors**	`/reports/vendor-benchmark`
The deal is big enough that I want a human analyst on it	**Analyst benchmark (upload a proposal)**	`/contracts/new`
I need a delivered report I can put in front of the board	**Analyst benchmark (upload a proposal)**	`/contracts/new`

The tools in this group

- **Verdict Bar** (`/verdict`). Drop any document, quote, contract, order form, or invoice; it is classified, routed to the matching engine, and judged against the market in about a minute.
- **Instant benchmarks** (`/benchmarking`). Self-serve price benchmarks for hundreds of vendors and product lines: type your deal, see where it lands against the market, no upload needed.
- **Second opinion** (`/tooling/second-opinion`). Drop a contract and get a benchmarked second opinion in minutes: your placement against the peer set, the money on the table priced, five arguments that work, and three priced paths.
- **Snap a quote** (`/tooling/snap-a-quote`). Photograph or screenshot any vendor quote: every line item is read off the paper, checked against published list prices, placed against comparable deals, and countered.
- **Quote Face-Off** (`/tooling/quote-faceoff`). Drop two or three competing quotes for the same need: equivalent SKUs aligned into one board, every line benchmarked, and the best mix across all papers priced.
- **Proposal comparison** (`/tooling/proposal-comparison`). Drop two vendor proposals and see exactly what changed: every pricing change priced, every term change read, each marked improved, worse, or unchanged.
- **Rate Card Benchmark** (`/tooling/rate-benchmark`). Upload a services rate card or SOW and see every day rate against the market band per role, region, and vendor tier, with the annualized overpay per line.
- **Portfolio benchmark** (`/benchmarking/portfolio`). Your whole vendor portfolio placed against the market in one standing view, with scenarios and the savings each would unlock.
- **Benchmark all my vendors** (`/reports/vendor-benchmark`). One AI-written report that benchmarks every vendor you track and ranks the gaps.
- **Analyst benchmark (upload a proposal)** (`/contracts/new`). Upload a contract or quote and a human analyst benchmarks it and delivers a written report.

REVIEW A CONTRACT BEFORE YOU SIGN

You have paper to read: a contract, an MSA, a statement of work, or two versions to compare.

IF THIS IS YOU	USE	ROUTE
I have one contract or quote to dissect clause by clause	**Contract decode**	`/tooling/contract-decode`
I want the redlines drafted for me before I reply	**Contract decode**	`/tooling/contract-decode`
I want to know which terms are already in my favor	**Contract decode**	`/tooling/contract-decode`
I have an MSA to sign and no lawyer hours to spare	**MSA Decoder**	`/tooling/msa-decoder`
I want the legal clause checklist judged before legal sees it	**MSA Decoder**	`/tooling/msa-decoder`
I have a consulting or SI statement of work to sign	**SOW Decoder**	`/tooling/sow-decoder`
I want to know where this SOW will leak into change orders	**SOW Decoder**	`/tooling/sow-decoder`
I want to know what changed between last year's contract and this renewal	**Contract time machine**	`/tooling/time-machine`
I suspect the renewal quietly dropped protections we had	**Contract time machine**	`/tooling/time-machine`
I want two uploaded documents diffed side by side	**Compare**	`/contracts/compare`
I want the comparison theatre for documents we already track	**Compare**	`/contracts/compare`
I want to know which contracts are missing our must-have clauses	**Coverage**	`/contracts/coverage`
I need a gap list for the next round of renewals	**Coverage**	`/contracts/coverage`
I need our standard language for a clause the vendor is pushing back on	**Clause library**	`/clause-library`
I want a starting position for a term I have never negotiated	**Clause library**	`/clause-library`
I want one question, for example the termination notice period, answered across all contracts	**Review tables**	`/review-tables`
Legal asked for a report over the whole estate and I do not want to open forty PDFs	**Review tables**	`/review-tables`

The tools in this group

- **Contract decode** (`/tooling/contract-decode`). Upload any vendor contract and watch it decoded live: every consequential clause in plain English, up to ten changes worth asking for with paste-ready language, and a term-by-term market benchmark.
- **MSA Decoder** (`/tooling/msa-decoder`). A buyer-side legal review of a master services agreement: liability cap, indemnities, IP, data protection, termination, and payment terms judged against standard negotiated positions, with replacement language.
- **SOW Decoder** (`/tooling/sow-decoder`). A buyer-side review of a consulting statement of work: scope leaks, deliverables without acceptance criteria, junior-heavy staffing, open-ended fees, and rate terms benchmarked.
- **Contract time machine** (`/tooling/time-machine`). Drop last year's paper and this year's for the same vendor and get the real story: terms won, protections that quietly disappeared, and the price walk line by line.

- **Compare** (`/contracts/compare`). A side-by-side AI redline scan of two documents already in your workspace.
- **Coverage** (`/contracts/coverage`). Must-have clause coverage across the whole estate: which contracts carry the protections you require and which are missing them.
- **Clause library** (`/clause-library`). Preferred clause language, fallback positions, and walkaway lines, ready to paste into any negotiation.
- **Review tables** (`/review-tables`). Ask one question and have it answered across every contract you store, as a table.

RUN A RENEWAL

A renewal is coming, and you want the calendar, the plan, and the vendor-specific playbook.

IF THIS IS YOU	USE	ROUTE
I want to see every renewal coming in the next twelve months	**Renewals**	`/renewals`
I never want to miss a notice window again	**Renewals**	`/renewals`
I want the team emailed automatically when a renewal goes live	**Renewals**	`/renewals`
A renewal is twelve months out and I want the plan now	**Renewal playbook**	`/renewals/playbook`
I want talking points and a timeline for a specific contract's renewal	**Renewal playbook**	`/renewals/playbook`
I want someone else to run these renewals for me	**Managed renewals**	`/renewals/managed`
We have no procurement team and renewals keep auto-renewing	**Managed renewals**	`/renewals/managed`
I want to time the negotiation to the vendor's quarter end	**Vendor calendar**	`/vendor-calendar`
I want to know which of my vendors close their year this quarter	**Vendor calendar**	`/vendor-calendar`
I want an uplift cap in writing before I sign	**Fair Renewal Rider**	`/rider`
I want one standard set of renewal protections across all vendors	**Fair Renewal Rider**	`/rider`
My Microsoft EA renewal is inside twelve months	**Microsoft EA renewal playbook**	`/tooling/ms-ea-renewal`
I need the CFO business case for the Microsoft renewal	**Microsoft EA renewal playbook**	`/tooling/ms-ea-renewal`
My Adobe ETLA renewal window is opening	**Adobe ETLA renewal playbook**	`/tooling/adobe-etla`
I want to know if VIP would beat renewing the ETLA	**Adobe ETLA renewal playbook**	`/tooling/adobe-etla`
My IBM ELA renewal quote just landed	**IBM ELA playbook**	`/tooling/ibm-ela`
I want to drop shelfware support without losing the licenses	**IBM ELA playbook**	`/tooling/ibm-ela`
Anthropic put an enterprise agreement on the table	**Anthropic EA playbook**	`/tooling/anthropic-ea`
I want the commit sized from our real Console usage, not the sales forecast	**Anthropic EA playbook**	`/tooling/anthropic-ea`
My Oracle ULA is approaching its certification date	**Oracle ULA certification playbook**	`/tooling/oracle-ula-certification`
I want to know whether to renew the ULA or certify out	**Oracle ULA certification playbook**	`/tooling/oracle-ula-certification`
SAP is pushing us from ECC to RISE and I want our own numbers first	**RISE with SAP playbook**	`/tooling/rise-with-sap`
I want the FUE count verified before we sign	**RISE with SAP playbook**	`/tooling/rise-with-sap`
My Salesforce SELA renewal is on the horizon	**Salesforce SELA playbook**	`/tooling/salesforce-sela`
I want the add-on shelfware out of the wrap before we renew	**Salesforce SELA playbook**	`/tooling/salesforce-sela`

IF THIS IS YOU	USE	ROUTE
My ServiceNow renewal is coming and the uplift looks steep	**ServiceNow renewal playbook**	`/tooling/servicenow-renewal`
I want to true down idle seats before the renewal locks them in	**ServiceNow renewal playbook**	`/tooling/servicenow-renewal`
My Workday renewal quote bundles more than we asked for	**Workday renewal playbook**	`/tooling/workday-renewal`
I want our worker bands corrected in the window they re-open	**Workday renewal playbook**	`/tooling/workday-renewal`
We pay OEM support on stable products we barely call	**Third-party support playbook**	`/tooling/third-party-support`
I want third-party support quotes without disclosing what we pay the OEM	**Third-party support playbook**	`/tooling/third-party-support`

The tools in this group

- **Renewals** (`/renewals`). Every agreement sorted by expiry with its notice window, status, and automatic team alerts when a window opens.
- **Renewal playbook** (`/renewals/playbook`). A guided, dated plan for one renewal: timeline, talking points, and the negotiation moves in order.
- **Managed renewals** (`/renewals/managed`). Hand the renewal calendar to the VendorBenchmark team; fees are charged on verified savings.
- **Vendor calendar** (`/vendor-calendar`). Every vendor's fiscal year and quarter ends on one timeline, so you time asks for when reps need deals.
- **Fair Renewal Rider** (`/rider`). A standard addendum that caps uplifts and locks renewal protections, ready to attach to any deal.
- **Microsoft EA renewal playbook** (`/tooling/ms-ea-renewal`). The Microsoft EA renewal run end to end: scope the cycle, load the estate, work E5, Copilot and Azure MACC savings, and finish with the CFO case.
- **Adobe ETLA renewal playbook** (`/tooling/adobe-etla`). The Adobe ETLA renewal end to end: Admin Console evidence, inactive seats cut in the only window allowed, GenAI padding decomposed, the VIP alternative priced.
- **IBM ELA playbook** (`/tooling/ibm-ela`). Renew the bundle or right-size out: IBM's quote decomposed line by line, shelfware S&S dropped while keeping perpetual licenses, ILMT remediated before it becomes audit leverage.
- **Anthropic EA playbook** (`/tooling/anthropic-ea`). The Anthropic enterprise agreement run on your own usage: model mix right-sized, caching and Batch wired, the commitment sized off your forecast with rollover and price protection.
- **Oracle ULA certification playbook** (`/tooling/oracle-ula-certification`). Renew or certify, decided on numbers: assess the position, maximize deployments before the count, work the certification rules, and exit with the licenses banked.
- **RISE with SAP playbook** (`/tooling/rise-with-sap`). The move from ECC or on-premise S/4HANA to RISE, end to end: the FUE count rebuilt from usage evidence, SAP's offer decomposed, alternatives priced, a five-year TCO case.

- **Salesforce SELA playbook** (`/tooling/salesforce-sela`). The SELA renewed on your numbers: the Unlimited base segmented against real usage, add-on shelfware cleared, the Agentforce commitment sized to measured burn.
- **ServiceNow renewal playbook** (`/tooling/servicenow-renewal`). Renew on your numbers, not theirs: the discount benchmarked against the corridor, idle seats trued down in the one window allowed, the uplift capped.
- **Workday renewal playbook** (`/tooling/workday-renewal`). The Workday renewal negotiated on the market: the discount benchmarked against the peer corridor, the bundled expansion decomposed, worker bands trued to reality.
- **Third-party support playbook** (`/tooling/third-party-support`). Move stable OEM support to third-party support on your terms: eligibility decided product by product and a blind tender run across providers.

NEGOTIATE THE DEAL

The negotiation is live, or about to be, and you want preparation, leverage, and support in the room.

IF THIS IS YOU	USE	ROUTE
I want the full prep pack for one vendor from one upload	**Vera Deal-Room**	`/tooling/dossier`
The negotiation starts next week and I have nothing prepared	**Vera Deal-Room**	`/tooling/dossier`
I want every prep document in one room I can forward	**Vera Deal-Room**	`/tooling/dossier`
I want one place to run a live negotiation from open to close	**Negotiations**	`/negotiation`
I want a record of every concession we and the vendor made	**Negotiations**	`/negotiation`
I need reply drafts for the vendor's latest email	**Negotiations**	`/negotiation`
I want the AI to run this negotiation for me end to end	**Vera AI-run negotiation (Case Mode)**	`/cases/new`
I have too many small renewals to negotiate personally	**Vera AI-run negotiation (Case Mode)**	`/cases/new`
I want the long tail of small vendors negotiated without my time	**Competitive Quotes (the Quote Fleet)**	`/negotiation/fleet`
I want competing bids collected for a set of vendors automatically	**Competitive Quotes (the Quote Fleet)**	`/negotiation/fleet`
I need a credible alternative on the table before I renew	**Build competitive scenarios**	`/tooling/competitive-scenarios`
I want to know what switching would really cost, migration included	**Build competitive scenarios**	`/tooling/competitive-scenarios`
I want the vendor to know we are evaluating options, in writing	**Build competitive scenarios**	`/tooling/competitive-scenarios`
I want to see the negotiation played out before I have it	**Negotiation self play**	`/tooling/self-play`
I want to know which strategy and sequencing wins for this vendor	**Negotiation self play**	`/tooling/self-play`
I want to rehearse the hard conversation before the call	**Vendor Twin**	`/vendor-twin`
I want to test how the vendor will respond to our ask	**Vendor Twin**	`/vendor-twin`
I have a vendor meeting tomorrow and want the prep pack	**Meeting war room**	`/meetings`
I want live support during the vendor call	**Meeting war room**	`/meetings`
I want to team up with other buyers renewing the same vendor	**Buying blocs**	`/blocs`
Our own volume is too small to move this vendor alone	**Buying blocs**	`/blocs`
I want one queue of everything in flight and what needs me today	**The Desk**	`/desk`
I want to see what the AI progressed while I was out	**The Desk**	`/desk`

The tools in this group

- **Vera Deal-Room** (`/tooling/dossier`). Drop one contract and the whole negotiation deal room assembles itself: the deal placed against the market, the decoded ask list, an analyst brief, the tactics playbook, live vendor intel, and the priced walk-away alternative.

- **Negotiations** (`/negotiation`). Each live deal tracked with a mandate, counters, tactic analysis, a concession ledger, and reply drafts.
- **Vera AI-run negotiation (Case Mode)** (`/cases/new`). Pick the vendors, hand over the contacts, approve the mandate, and Vera runs the negotiation herself, checking in for decisions.
- **Competitive Quotes (the Quote Fleet)** (`/negotiation/fleet`). An agent negotiates a set of tail vendors at once, collecting bids and pushing each toward the benchmark.
- **Build competitive scenarios** (`/tooling/competitive-scenarios`). Name the product you are renewing and get the credible alternatives researched and three to five cost scenarios priced against the renewal, with ready-to-send emails.
- **Negotiation self play** (`/tooling/self-play`). The machine plays your renewal to completion eight times, four strategies against two rep temperaments, and shows the outcome distribution and the sequencing move that mattered.
- **Vendor Twin** (`/vendor-twin`). A simulation of your vendor's rep to spar against: practice the conversation before the real one.
- **Meeting war room** (`/meetings`). Prep, live support, and debrief for vendor meetings, grounded in your contracts and benchmarks.
- **Buying blocs** (`/blocs`). Join other buyers renewing the same vendor in the same window to negotiate from combined volume.
- **The Desk** (`/desk`). The working queue for deals in flight: what needs you now, what the AI moved overnight, and the next move on each.

CUT SOFTWARE SPEND

You own a budget number and need to find, plan, and capture the savings.

IF THIS IS YOU	USE	ROUTE
I own the software budget and need the savings plan	**Spend optimization**	`/tooling/spend-optimization`
I was told to cut ten percent and need to know where it comes from	**Spend optimization**	`/tooling/spend-optimization`
I want the whole board on one screen: spend, renewals, market position	**Spend optimization**	`/tooling/spend-optimization`
I have the raw AP export and want it read end to end	**Spend Intelligence**	`/tooling/spend-intel`
I want duplicate payments and quiet price increases found	**Spend Intelligence**	`/tooling/spend-intel`
I want a spend forecast with the buying windows marked	**Spend Intelligence**	`/tooling/spend-intel`
I want to see what we spend, by vendor and category	**Spend**	`/spend`
I need the spend number for a vendor before the meeting	**Spend**	`/spend`
I want the savings pipeline ranked by size	**Opportunities**	`/opportunities`
I want to know which renegotiation to start first	**Opportunities**	`/opportunities`
I want to be told when a vendor's billing quietly changes	**Anomalies**	`/anomalies`
I want spend spikes explained before finance asks	**Anomalies**	`/anomalies`
I want invoices checked against what the contract says we owe	**Invoices**	`/invoices`
I think we are being overbilled and want the evidence	**Invoices**	`/invoices`
I want to use rights we already paid for and forgot about	**Rights register**	`/obligations`
I want every contract's obligations and deadlines in one register	**Rights register**	`/obligations`
An AI vendor wants a committed-spend number from us	**Spend Studio**	`/tooling/spend-studio`
I want the AI commit sized from real usage, not optimism	**Spend Studio**	`/tooling/spend-studio`

The tools in this group

- **Spend optimization** (`/tooling/spend-optimization`). Pick your top vendors by spend and get a written optimization plan: about twenty prioritized moves across right-sizing, renegotiation, consolidation, terms, and timing, anchored to your renewal calendar.
- **Spend Intelligence** (`/tooling/spend-intel`). Upload an AP or ERP spend export and get the ledger read back: every supplier classified, duplicate payments and price steps flagged, a six-month forecast with buying windows marked.
- **Spend** (`/spend`). Spend by vendor and category, built automatically from the proposals and agreements you track.
- **Opportunities** (`/opportunities`). Every benchmarked vendor above market, ranked by the annual prize, ready to turn into a tracked deal.
- **Anomalies** (`/anomalies`). Spend and billing anomalies surfaced across the estate: spikes, drift, and things that changed without a signature.
- **Invoices** (`/invoices`). Invoice line-item intelligence: invoices reconciled against contracts, overbilling flagged, recovery tracked.
- **Rights register** (`/obligations`). The rights and entitlements you already paid for, harvested from your contracts, with deadlines to use them.

- **Spend Studio** (`/tooling/spend-studio`). Size Claude, OpenAI, Copilot Cowork, and Agentforce commitments from your own usage exports, and price the committed-spend deal before you sign.

LICENSES AND CLOUD

Usage-driven right-sizing: seats, cores, and cloud workloads measured against what you pay for.

IF THIS IS YOU	USE	ROUTE
I suspect M365 seats are idle or over-licensed	**M365 License Optimization**	`/tooling/m365-optimization`
I want the E5 case tested against real usage before the renewal	**M365 License Optimization**	`/tooling/m365-optimization`
Broadcom is repricing our VMware estate to VCF cores	**VMware VCF licensing calculator**	`/tooling/vcf-licensing`
I want to know our billable core count before Broadcom tells us	**VMware VCF licensing calculator**	`/tooling/vcf-licensing`
Oracle quoted Fusion ERP modules and I need the net price a year	**Oracle ERP Cloud pricing calculator**	`/tooling/oracle-erp-pricing`
I want to know whether our Fusion discount is good for a deal our size	**Oracle ERP Cloud pricing calculator**	`/tooling/oracle-erp-pricing`
I need to show the CFO what the renewal uplift costs over the term	**Oracle ERP Cloud pricing calculator**	`/tooling/oracle-erp-pricing`
Oracle quoted Fusion HCM modules and I need the net price a year	**Oracle HCM Cloud pricing calculator**	`/tooling/oracle-hcm-pricing`
I want to price Global HR, recruiting, learning, and payroll as one order	**Oracle HCM Cloud pricing calculator**	`/tooling/oracle-hcm-pricing`
The cloud bill keeps growing and I want to know what is recoverable	**Cloud cost optimizer**	`/tooling/cloud-cost`
I need a FinOps assessment without hiring a FinOps team	**Cloud cost optimizer**	`/tooling/cloud-cost`
I wonder what our workloads would cost on another cloud	**Cloud workload comparator**	`/tooling/workload-compare`
I want a placement report before committing to one provider	**Cloud workload comparator**	`/tooling/workload-compare`

The tools in this group

- **M365 License Optimization** (`/tooling/m365-optimization`). Upload your M365 usage export to find inactive accounts, right-size E5/E3/E1, and size Copilot, with a per-user reassignment plan.
- **VMware VCF licensing calculator** (`/tooling/vcf-licensing`). Upload your RVTools export to size the VCF cores Broadcom will bill, build three pricing scenarios against real discount bands, and get a negotiation playbook.
- **Oracle ERP Cloud pricing calculator** (`/tooling/oracle-erp-pricing`). Pick the Fusion ERP cloud services on your order from Oracle's published price list, set the quantities, apply a discount, and see the net fees to Oracle a year, the term cost under an uplift, and the deal's tier on the Fusion discount ladder.
- **Oracle HCM Cloud pricing calculator** (`/tooling/oracle-hcm-pricing`). Pick the Fusion HCM cloud services on your order from Oracle's published price list, set the quantities, apply a discount, and see the net fees to Oracle a year, the term cost under an uplift, and the deal's tier on the Fusion discount ladder.

- **Cloud cost optimizer** (`/tooling/cloud-cost`). Upload a cost export from AWS, Azure, GCP, or OCI and get a FinOps assessment: identified savings by workstream and a phased action plan.
- **Cloud workload comparator** (`/tooling/workload-compare`). Describe your workloads or import an inventory and see what each costs on AWS, Azure, GCP, and OCI, and the optimal placement across the portfolio.

HANDLE A VENDOR AUDIT

An audit letter arrived, or you want to be ready for one.

IF THIS IS YOU	USE	ROUTE
We received a software audit letter and I do not know what to do first	**Audit letter decoder**	`/tooling/audit-decoder`
I want to know what the letter actually obliges us to do	**Audit letter decoder**	`/tooling/audit-decoder`
I want the response plan before anyone replies to the auditor	**Audit letter decoder**	`/tooling/audit-decoder`

The tools in this group

- **Audit letter decoder** (`/tooling/audit-decoder`). Drop the audit letter the moment it lands: who is actually auditing you, the deadline as a business-day countdown, what is demanded versus implied, the three things not to do first, and the response plan.

KNOW YOUR VENDORS AND THE MARKET

Standing intelligence: how your vendors behave, what they are changing, and what other buyers see.

IF THIS IS YOU	USE	ROUTE
I want everything we know about one vendor on one page	**Vendors**	`/vendors`
I am new to the account and need the vendor's history fast	**Vendors**	`/vendors`
I want to know how this vendor typically behaves at renewal	**Vendor scorecards**	`/scorecards`
I want a grade and a one-page read on a vendor I rarely touch	**Vendor scorecards**	`/scorecards`
I want the written playbook for a vendor before I engage	**Research**	`/research`
I want market context I can cite in the business case	**Research**	`/research`
I want vendor news that actually affects my contracts	**The Vera Wire**	`/wire`
I want to know about a price change before the rep tells me	**The Vera Wire**	`/wire`
I want to be told when a vendor changes its licensing fine print	**Policy watch**	`/wire/policy`
I need the before-and-after of a product terms change	**Policy watch**	`/wire/policy`
I want a standing watch on my vendors' terms and price lists	**Your watchtower (Terms Watch)**	`/wire/watch`
I want billing drift on my own accounts flagged	**Your watchtower (Terms Watch)**	`/wire/watch`
I want a weekly digest on the vendors I follow	**Weekly vendor report**	`/wire/weekly`
I want something scannable to forward to the team each week	**Weekly vendor report**	`/wire/weekly`
I want one monthly read on the whole estate for leadership	**Monthly briefing**	`/briefing`
I need talking points for the monthly ops review	**Monthly briefing**	`/briefing`
I want to know what uplift other buyers are being quoted	**Community polls**	`/polls`
I want a quick pulse from peers before I commit to a position	**Community polls**	`/polls`

The tools in this group

- **Vendors** (`/vendors`). Every vendor you track as a briefing: position, spend, paper on file, and how they behave in negotiations.
- **Vendor scorecards** (`/scorecards`). A renewal scorecard per vendor: how they price, how they negotiate, and what a strong outcome looks like.
- **Research** (`/research`). The research library: negotiation toolkits, market outlooks, terms benchmarks, and vendor assessments, downloadable.
- **The Vera Wire** (`/wire`). The vendor news feed: pricing moves, licensing changes, and market events filtered to vendors like yours.
- **Policy watch** (`/wire/policy`). Vendor licensing and product-terms documents monitored for changes, diffed daily.
- **Your watchtower (Terms Watch)** (`/wire/watch`). Your vendors' terms, billing, and price lists watched for drift, with alerts tied to your own estate.
- **Weekly vendor report** (`/wire/weekly`). A weekly written report per vendor you follow: what moved, what it means, what to do.
- **Monthly briefing** (`/briefing`). A monthly executive briefing over your whole estate: what changed, what is coming, and where to act.
- **Community polls** (`/polls`). What other buyers are seeing, asked and answered in structured polls.

ORGANIZE CONTRACTS AND DOCUMENTS

Getting the estate in: filing, tracking, searching, and keeping the data healthy.

IF THIS IS YOU	USE	ROUTE
I have a folder of old contracts and want them filed and read	**Your Archive**	`/archive`
I inherited a contract mess and need it organized this week	**Your Archive**	`/archive`
I want an intake brief on what we actually signed over the years	**Your Archive**	`/archive`
I want every signed agreement tracked with its expiry and notice period	**Agreements**	`/agreements`
I want renewals and spend views built from what we actually signed	**Agreements**	`/agreements`
I want to see the status of everything we submitted	**Proposals**	`/contracts`
I want one list of every document under analysis	**Proposals**	`/contracts`
I need a room for deal documents only some people may see	**Vault**	`/vault`
I want diligence files kept away from the general workspace	**Vault**	`/vault`
I need to find which contract contains a specific clause or phrase	**Document search**	`/doc-search`
I want to search inside the PDFs, not just their titles	**Document search**	`/doc-search`
I want duplicate and near-duplicate documents found	**Document clusters**	`/doc-clusters`
I want to see which documents belong to the same deal family	**Document clusters**	`/doc-clusters`
I want to know which records are missing values before I report on them	**Data health**	`/data-health`
I want the extraction review queue worked down	**Data health**	`/data-health`
I want the portfolio overview for a leadership meeting	**Portfolio**	`/portfolio`
I want to see vendor concentration at a glance	**Portfolio**	`/portfolio`
I want one picture of the whole estate for a slide	**Estate map**	`/portfolio/map`
We are moving off spreadsheets or another tool and want everything carried over	**Migration Studio**	`/migration`
I want a dry run and reconciliation before the data goes live	**Migration Studio**	`/migration`
We are acquiring a company and need its software contracts read	**Diligence**	`/diligence`
I need change-of-control exposure across a target's estate	**Diligence**	`/diligence`
I want our category strategy written down where the team works	**Strategy**	`/strategy`
I want the org's entities and their agreements mapped	**Your organization**	`/org`

The tools in this group

- **Your Archive** (`/archive`). Drop a whole contract folder and Vera files it by vendor, extracts the terms, and briefs you on what the estate contains.
- **Agreements** (`/agreements`). The ledger of signed agreements: vendor, value, expiry, notice period, and the files, feeding renewals and spend automatically.
- **Proposals** (`/contracts`). Every uploaded contract and quote with its status from upload to delivered report.
- **Vault** (`/vault`). Secure project rooms for sensitive documents, with per-project access.

- **Document search** (`/doc-search`). Full-text search inside every contract and file you store.
- **Document clusters** (`/doc-clusters`). Your documents grouped automatically by similarity, so duplicates and families surface.
- **Data health** (`/data-health`). Extraction quality across the estate: what is complete, what is missing, and what needs review.
- **Portfolio** (`/portfolio`). The whole estate in one view: vendors, values, concentrations, and risk.
- **Estate map** (`/portfolio/map`). The estate drawn as a map: every vendor, sized and colored by what matters.
- **Migration Studio** (`/migration`). A guided migration of your existing contract records and files into the workspace, with a dry run, reconciliation, and go-live.
- **Diligence** (`/diligence`). An M&A diligence run over a target's software contracts: exposure, change-of-control, and synergy, in a clean room.
- **Strategy** (`/strategy`). The standing vendor strategy: category plans, positions, and the moves agreed for the year.
- **Your organization** (`/org`). The organization's estate page: entities, structure, and what belongs where.

PROVE SAVINGS AND REPORT UP

Turning outcomes into evidence: verified savings, deliverables, and executive reporting.

IF THIS IS YOU	USE	ROUTE
I need to prove the savings we claimed actually landed	**Savings proof**	`/savings`
Finance wants savings tied to invoices, not estimates	**Savings proof**	`/savings`
I want a shareable record of a verified win	**Verified wins**	`/savings/wins`
I want the team's wins visible and credited	**Verified wins**	`/savings/wins`
I want the guarantee terms and status of a guaranteed engagement	**Guarantees**	`/guarantees`
I want every report we ever received in one place	**Documents (the studio shelf)**	`/studio`
I need to re-download or share a past deliverable	**Documents (the studio shelf)**	`/studio`
I need a board-ready report built from our own data	**Reporting studio**	`/studio/generate`
I want the quarterly procurement report written for me	**Reporting studio**	`/studio/generate`

The tools in this group

- **Savings proof** (`/savings`). Savings reconciled against what vendors actually billed, not what the deal memo claimed.
- **Verified wins** (`/savings/wins`). Badges, deal scores, and an opt-in wins board for outcomes verified against billing.
- **Guarantees** (`/guarantees`). Engagements with the fee at risk: the guarantee terms, progress, and the certificate at the end.
- **Documents (the studio shelf)** (`/studio`). Every delivered report, brief, and export on one shelf, with sharing and history.

- **Reporting studio** (`/studio/generate`). Generate an executive report from your data: spend, renewals, savings, or a custom brief, grounded and citation-checked.

BUY SOMETHING NEW

Net-new purchases: requests, sourcing, bids, and approvals.

IF THIS IS YOU	USE	ROUTE
We are buying something new and I want the process run properly	**Sourcing**	`/sourcing`
I want supplier bids collected and scored in one place	**Sourcing**	`/sourcing`
I want one desk where every new deal request lands	**Deal desk**	`/sourcing/desk`
I want the business to request software through one door	**Requests**	`/requests`
I want to stop purchases we already own licenses for	**Requests**	`/requests`
I approve deals and want the queue with recommendations	**Approvals**	`/approvals`
I want approval rules enforced before anything is signed	**Approvals**	`/approvals`

The tools in this group

- **Sourcing** (`/sourcing`). The sourcing pipeline from request to award: suppliers fielded, bids collected and scored, the decision documented.
- **Deal desk** (`/sourcing/desk`). The intake and triage desk for deals: what came in, who owns it, what happens next.
- **Requests** (`/requests`). Purchase requests from the business, captured and routed.
- **Approvals** (`/approvals`). The approvals queue: deals waiting on a decision, with the negotiator's recommendation attached.

WORK WITH YOUR TEAM

Coordination: who does what, where discussions live, and what is waiting on you.

IF THIS IS YOU	USE	ROUTE
I want the team's work visible and assigned	**Tasks**	`/tasks`
I want one inbox of everything waiting on me	**My work**	`/my-work`
I want to know who owns each vendor relationship	**Team**	`/team`
I want deal discussions out of email and next to the data	**Board**	`/board`
I want to message a colleague about a deal without leaving the page	**Chat**	`/chat`
I want to catch up on what changed since I was last in	**Notifications**	`/notifications`
I want to find something I bookmarked last month	**Saved**	`/saved`
Something is wrong and I want to raise it with the team	**Support**	`/support`
I want to refer another company and track what it earns	**Referrals**	`/referrals`
I want the status of our advisory engagements in one place	**Your engagements**	`/delivery`
I want an analyst on a call about a specific deal	**Advisory**	`/advisory`

The tools in this group

- **Tasks** (`/tasks`). The team's task list across deals, renewals, and projects.
- **My work** (`/my-work`). Everything waiting on you personally: tasks, approvals, mentions, and renewals.
- **Team** (`/team`). Who covers what: vendor ownership and coverage across the team.
- **Board** (`/board`). The internal discussion board: questions, decisions, and wins, kept where the work is.
- **Chat** (`/chat`). Direct messages with your team inside the workspace.
- **Notifications** (`/notifications`). The notification feed: everything that changed that you asked to hear about.
- **Saved** (`/saved`). Your reading room: every report and research paper you bookmarked.
- **Support** (`/support`). Support tickets with the VendorBenchmark team.
- **Referrals** (`/referrals`). Your referral code and earnings: refer a company and earn a share of their first-year spend.
- **Your engagements** (`/delivery`). Service engagements you run with the VendorBenchmark team, in one workspace shared with your members.
- **Advisory** (`/advisory`). The concierge desk: book analyst time and advisory sessions.

PUT AI TO WORK

The assistant and the agent workforce: standing watches, scheduled reports, and automated jobs.

IF THIS IS YOU	USE	ROUTE
I have a question about our contracts and want an answer with sources	**Vera AI**	`/ask`
I do not know which tool to use and want to be pointed to it	**Vera AI**	`/ask`
I want a standing question re-answered as the data changes	**Vera AI**	`/ask`
I want a watchdog on something so I never check it manually again	**Agent Center**	`/agents`
I want a report generated on a schedule without asking each time	**Agent Center**	`/agents`
I want to describe a job in a sentence and get an agent that does it	**Agent Center**	`/agents`
I want a multi-step job automated end to end	**Workflow agents**	`/workflows`
I want the ghost writer drafting vendor correspondence for me	**Workflow agents**	`/workflows`

The tools in this group

- **Vera AI** (`/ask`). The assistant over your whole workspace: ask anything about your contracts, vendors, and the market, and be routed to the tool that does the job.
- **Agent Center** (`/agents`). Your AI workforce: standing watchdogs, scheduled analysts, autopilots, and the wizard to build a new agent from a sentence.
- **Workflow agents** (`/workflows`). Multi-step automations chained from the platform's tools, run on triggers or on demand.

COMMUNITY AND OUTSIDE HELP

Other buyers, peers who have done your deal, and vetted vendor-side experts.

IF THIS IS YOU	USE	ROUTE
I want to ask other buyers what they are seeing from a vendor	**Buyer Network**	`/buyer-network`
I want to talk to someone who has negotiated this vendor recently	**Buyer intros**	`/network`
I want an hour with someone who used to sell for this vendor	**Expert Bench**	`/experts`
The deal is big and I want insider context on how they discount	**Expert Bench**	`/experts`

The tools in this group

- **Buyer Network** (`/buyer-network`). The trading floor: threads with other buyers about vendors, tactics, and what deals look like right now.
- **Buyer intros** (`/network`). Introductions to peers who have done the deal you are about to do.
- **Expert Bench** (`/experts`). Vetted insiders from the vendor side, bookable for the deals where inside knowledge changes the outcome.

SET UP AND GOVERN THE WORKSPACE

Owner and security jobs: members, integrations, approvals, and the trust posture.

IF THIS IS YOU	USE	ROUTE
I want to invite my team and control who can do what	**Members**	`/settings/members`
I want usage and license data flowing in automatically	**Integrations**	`/settings/integrations`
I want to connect our SAM or ITAM tool	**Integrations**	`/settings/integrations`
I want to pull our data into our own reporting stack	**API**	`/settings/api`
I want MFA and SSO enforced for the workspace	**Security**	`/settings/security`
Our security team wants the audit log and data controls	**Security**	`/settings/security`
I want deals above a threshold to require sign-off	**Deal approvals**	`/settings/approvals`
I want to choose what the platform emails me about	**Notification settings**	`/settings/notifications`
Our security team is reviewing the platform and needs the posture	**Trust center**	`/trust`

The tools in this group

- **Members** (`/settings/members`). Invite teammates, set roles, and manage access (org owners).
- **Integrations** (`/settings/integrations`). Connect SAM, ITAM, SSO directories, spend systems, and email rails so the workspace fills itself.
- **API** (`/settings/api`). API keys and the MCP surface for wiring the platform into your own systems.
- **Security** (`/settings/security`). Two-factor, sessions, SSO, audit log, data controls, and the security posture view.
- **Deal approvals** (`/settings/approvals`). Approval rules: which deals need whose sign-off before they close (org owners).
- **Notification settings** (`/settings/notifications`). Exactly which alerts you receive, all strictly opt-in.
- **Trust center** (`/trust`). How your data is isolated, encrypted, and handled, written for your security team.

KEEPING THIS CHAPTER CURRENT

This chapter is generated from the platform's use-case index, the same source Vera uses to route "which tool should I use?" questions, so the manual and the assistant always give the same answer. It regenerates whenever the index changes.

15 · Your first 30 days

The reference chapters document every screen. This chapter is the ramp: what to actually do on day one, in week one, and by day 30, in the order that pays off fastest. Follow it as written and by the end of the month your contracts are on file, your renewal calendar is live and watching deadlines, your team is in, and you have taken at least one real position into a vendor conversation.

Two promises before you start. Nothing in this plan is busywork: every step leaves something standing that the platform keeps using. And nothing blocks you: every offer the platform makes (the introduction, the survey, the course) waits until you want it, so you can run this ramp at your own pace. Most days below cost 30 to 60 minutes; several cost nothing at all.

Who this is for. Anyone new to the platform. If you run the workspace (you invited the team, you hold the owner role), read this alongside the owner's handbook, which covers the setup only owners can do. If you joined a workspace someone else runs, this chapter alone is your ramp. For a reading path matched to your job, see the manual by role.

BEFORE DAY ONE

Have three things within reach; none is mandatory, all three make the first week faster.

1. **One live document.** A vendor quote, a renewal proposal, or a signed contract, as PDF or DOCX. Your first minutes on the platform are strongest on your own paper.
2. **Your contract folder.** Wherever your organization keeps its signed agreements today: a shared drive, a legal folder, a laptop. You will drop it in during week one. Subfolders and clutter are fine; the platform filters.
3. **Your renewal dates**, if you keep them anywhere (a spreadsheet, a calendar, memory). Week two turns them into watched deadlines.

Access itself is covered in Chapter 00: creating your account is reviewed by a person, usually within 1 business day, and while you wait, the waiting screen offers a live demo workspace you can explore without saving anything.

DAY 1: PROVE IT ON YOUR OWN PAPER

The goal of day one is not to learn the platform. It is to get one defensible result on a document you actually care about, so you know what the platform is for before you learn how it is laid out.

1. Sign in. Set up two-factor authentication now, not later: open Settings › Security, scan the code with your authenticator app, and store the backup codes. It takes three minutes and it is the single highest-value security step you control.

2. A brand-new account is walked into your first verdict: drop your document and watch the four steps land, Upload, Decode, Benchmark, Plan. Under five minutes, and the verdict is a real saved run, not a demo. This flow requires an AI plan; on the Free plan, read the example report instead to see what a delivered report contains.
3. No document handy? Load the sample portfolio instead, from the dashboard, so every screen has something to show while you explore. Remove it before you enter real data, and if it ever lingers, Troubleshooting has the fix.
4. Take the two offers when they appear, in either order: the platform introduction (a six page picture book, three minutes) and the desk survey (six questions that arrange the navigation around your work and put your vendors on watch).
5. Before you leave, learn the three keys you will use every day: **Ctrl/⌘ K** jumps anywhere, **Ctrl/⌘ J** opens Vera over whatever you are looking at, and the **Guide** door in the top bar opens this manual at the section for the screen you are on.

What is standing at the end of day one: a secured account, one real verdict on your own paper, and a workspace arranged around your answers.

WEEK 1: YOUR ESTATE ON THE TABLE

The platform is only as sharp as the paper it can read. Week one is about getting your contracts in, and it is shorter than it sounds: the archive does the reading.

1. **Drop your contract folder into Your Archive.** The whole tree at once, up to 300 files per run, 40 MB each, PDF and DOCX; everything else is reported, never silently dropped. Vera reads every document, files it by vendor, and judges each agreement current, due, expired, superseded, or perpetual. You can leave while it runs; a chip follows you and a notification brings you back.
2. **Work the Next moves rail when the scan finishes.** It asks three questions with the money named on each: which due renewals go on the calendar, whether to review your top spend, and which contracts deserve a benchmark. Answer the renewals one at minimum; that is your calendar seeded in one sitting.
3. **Estates in the hundreds or thousands** are a different door: ask your owner about Migration Studio, the guided run with staging, a register import, reconciliation, and a sign-off scorecard. It is owner-run and switched on per organization. A handful of documents needs neither door: upload a proposal for anything you want benchmarked, or add Agreements by hand. Chapter 18 compares all four intake doors, covers registers and spreadsheets, and answers the migration FAQs.
4. **Run your first instant benchmark.** Open Instant benchmarks, type a vendor you hold, and answer the questions. You get the three answers every buyer needs: the discount you are being offered, what comparable customers get as a range, and where your deal ranks. Faster still, drag the vendor's own quote onto the page (Drop a quote, on paid plans) and the paper is read for you. Allowances vary by plan; Chapter 10 has the table.
5. **Meet Vera properly.** Press Ctrl/⌘ J and ask something hard about your own estate: "which of my vendors renews next and what should I do first?" Answers are grounded in your data and cited; Vera AI explains

what she can see and do.

What is standing at the end of week one: your estate filed by vendor, the due renewals flagged, one market read on a real deal, and a working relationship with the assistant.

WEEK 2: THE CALENDAR AND THE TEAM

Week two turns a filed estate into a watched one, and brings in the people who share the work. If you are not an owner, steps 3 and 4 are an ask to your owner rather than something you do yourself.

1. **Finish the calendar.** Open Renewals and check every agreement you care about carries an expiry date and a notice period (30, 60, or 90 days). An agreement without both has no window to miss, so it cannot be watched. Add what is missing under Agreements; the Renewal calendar shows the year at a glance.
2. **Decide what is allowed to interrupt you.** Everything about renewals is opt-in: the workspace stays silent about deadlines until someone asks for them. Open Settings › Notifications and press **Turn on money-critical alerts** to switch on the renewal set in one click, then pick your weekly briefings (price list changes, the Monday wire, the knowledge briefing) and set quiet hours. Every email carries a one-click way off.
3. **Invite the team** (owners): Members, one invite or up to 50 addresses at once, each with a role. Invitations expire after 7 days; a resend restarts the clock.
4. **Set up confidentiality before you need it** (owners): create Teams named the way your company organizes (Procurement, Finance, Legal), so a restricted contract can be granted to a group once instead of person by person.
5. **Split the work.** Assign owners to the renewals closing soonest, and start the day on My work, the queue of everything waiting on you personally.

What is standing at the end of week two: a live calendar that emails the team when a notice window opens, alerts you chose rather than inherited, and the right people inside.

WEEK 3: YOUR FIRST LIVE DEAL

Everything so far was setup. Week three is the point of it: take the nearest real renewal and work it like the platform intends.

1. **Pick the deal.** The renewal closing soonest with real money on it. The dashboard or Control room will point at it; a gold renewals count means something is already on the clock.
2. **Get the market read first.** Run the vendor through Instant benchmarks, or send the paper for analyst benchmarking if your plan includes it (plans without included analyst benchmarks can add one with a benchmark credit; the proposal page handles it). You want the discount, the peer range, and the rank in hand before any vendor call.

3. **Build the plan.** Open the Renewal playbook for the step-by-step negotiation plan, then set your standard: your clause positions in the Clause library and the Fair Renewal Rider as your standing protections.
4. **Track the negotiation.** Open a tracked Negotiation: the mandate (target, walk-away), vendor emails analyzed for tactics, every give and get in the ledger, and vendor promises captured as commitments so they reach the contract.
5. **Watch the vendor while you negotiate.** Follow them on Terms watch and subscribe to the Monday Weekly report; a price-list move or an earnings wobble mid-negotiation is leverage you want the same week, not at signature.
6. **Before the first call**, walk the full recipe once: Plan and run a renewal negotiation in the cookbook stitches these screens into one sequence.

What is standing at the end of week three: one renewal being worked with a market position, a plan, a mandate, and a paper trail.

WEEK 4 AND DAY 30: THE OPERATING RHYTHM

The last week converts a good start into a routine the workspace sustains on its own.

1. **Close the loop on money.** When a deal lands, record what it saved: Savings proof reconciles each claim against the invoices that follow, which is what makes the number defensible in front of finance. About 30 days after any delivered benchmark, answer the one-touch What happened? follow-up; a won reduction files itself.
2. **Check your data health.** Open Data health: the completeness score and every gap ranked by what closing it unlocks. Work the top one or two; Connect your data and close the gaps is the recipe.
3. **Skim the wider desk.** You have not needed Portfolio, the Estate map, or the Monthly briefing yet; know they exist before leadership asks. The tool wall on All tools rewards ten minutes of browsing, and Which tool do I use? answers the question directly, by job.
4. **Set your weekly rhythm.** Monday: the Weekly report and the renewal radar, if you switched them on. Daily: My work. Before any vendor meeting: the pack, per Prepare for a vendor meeting.
5. **Day 30, run the health check below.** Anything unticked is your week five.

THE DAY 30 HEALTH CHECK

Tick these against your own workspace. Each line links to where it is done.

STANDING?	THE CHECK	WHERE
☐	Two-factor is on for your account	Security
☐	Your contracts are filed, by folder drop or migration	Your Archive
☐	Every agreement that matters carries an expiry date and a notice period	Agreements
☐	Renewal alerts are on, or deliberately off	Notifications
☐	The team is in, with the right roles	Members
☐	At least one deal has a market read: discount, peer range, rank	Instant benchmarks
☐	The nearest renewal has a plan and an owner	Renewal playbook
☐	Sample data is gone from your views	Troubleshooting
☐	Data health is green, or its top gap has an owner	Data health

IF YOU ONLY HAVE ONE HOUR

The whole ramp, compressed to its spine: sign in and set up two-factor (5 minutes), run your first verdict on a real document (5 minutes), drop your contract folder into Your Archive and answer the renewals question when it finishes (20 minutes, mostly the platform's), press **Turn on money-critical alerts** under Notifications (1 minute), and run one instant benchmark on the vendor renewing soonest (10 minutes). That hour buys you a watched calendar and a market position; everything else in this chapter deepens it.

Related: The manual by role · The owner's handbook · Common workflows

Next: Chapter 16 · The manual by role

16 · The manual by role

Nobody reads a manual cover to cover, and nobody should. This chapter gives each seat at the table its own way in: the five screens that role lives in, the three workflows to learn first, and the chapters worth reading in order. Find your row, follow it, and treat everything else in the manual as reference for the day you need it.

WHO THIS CHAPTER IS FOR

Six seats cover almost everyone who works in the platform. If you hold two of them (a procurement lead who also owns the workspace is common), read both sections; they are short by design. Whatever your seat, Chapter 15 is the shared ramp for your first month, and the **Guide** door in the top bar always opens this manual at the screen you are standing on.

THE PROCUREMENT AND SOURCING LEAD

You own vendor relationships and the deals themselves: renewals, negotiations, new purchases. The platform's center of gravity, the contract-to-close loop, is built around your week.

The five screens you live in

- Renewals — every deadline sorted by expiry, the next one on top.
- Instant benchmarks — the market read before any vendor call: your discount, the peer range, the rank.
- Negotiations — the tracked deal: mandate, tactics, the concession ledger, commitments.
- Vendors — the whole relationship per vendor on one page.
- Terms watch — what the vendor did to its price list, priced at your seat counts.

Learn these three workflows first: Plan and run a renewal negotiation, Choose between competing quotes, and Watch a vendor's market moves.

Read in this order: Chapter 01 (contracts and renewals, your home chapter), Chapter 02 (benchmarks), then Chapter 03 for the tool wall, one tool at a time as deals demand them.

THE IT ASSET AND LICENSING MANAGER

You know what the organization owns, deploys, and actually uses, and you carry the audit risk. Your leverage is the gap between entitlement and deployment, and the platform's optimizers exist to find it before the vendor does.

The five screens you live in

- All tools — the optimizer wall: M365, cloud cost, VCF, workload comparison, and the rest.
- Integrations — the SAM and ITAM connections that feed entitled against deployed against active.
- Data health — your completeness score and the audit exposure estimate.

- Billing watch — billed against entitled against market, checked daily.
- Rights register — the paid-for rights your contracts already grant you.

Learn these three workflows first: Right-size licenses before a renewal or true-up, Respond to a vendor audit letter, and Connect your data and close the gaps.

Read in this order: Chapter 03 (tooling, your home chapter), Chapter 12 (integrations and clients), then Chapter 01 for the watchtowers.

THE CFO AND FINANCE

You care whether the money is right: what the estate costs, whether claimed savings are real, and what the exposure is. The platform's answer to finance is that every figure traces to a source, and savings only count when invoices confirm them.

The five screens you live in

- Spend — what the estate costs, by vendor and category.
- Savings proof — claims reconciled against the invoices that followed: proven or disproven, never asserted.
- Invoices — line-level verdicts against contract terms, dollars first.
- Portfolio — concentration, market position, and where the leverage is.
- Figures ledger — every number the platform cited for you, traced to its source fact.

Learn these three workflows first: Prove and share your savings, Dispute an invoice that does not match your contract, and Build the board pack.

Read in this order: Chapter 04 (insight, your home chapter), then Chapter 02 for what a delivered benchmark contains, then Chapter 10 so plan and allowance questions never surprise you.

LEGAL AND CONTRACT OWNERS

You own what the organization signs: the clause standard, the deviations, the obligations the paper creates. The platform holds your positions as data and measures every contract against them, which turns review from rereading into exception handling.

The five screens you live in

- Clause library — your positions: ideal, fallback, and walk-away language per clause.
- Coverage — the must-have template swept across the estate: Present, Deviates, Missing, or Unclear.
- Contract decode — a contract read clause by clause, risks in plain English, asks with paste-ready language.
- Vault — restricted projects for the documents most people must not see.
- Rights register — obligations and entitlements harvested from the signed paper.

Learn these three workflows first: Set your clause standard and sweep the estate, Run diligence for an acquisition, and Track agreements and never miss a notice window.

Read in this order: Chapter 01 (your home chapter), Chapter 09 (record pages, where the evidence lives), then the decoders in Chapter 03 (Contract decode, the MSA and SOW decoders).

THE CIO AND LEADERSHIP

You do not work the platform daily; you consume what it proves. Your interest is the estate picture, the risk posture, and whether the function is delivering, each of which the platform states as a document you can forward rather than a dashboard you must visit.

The five screens you live in

- Portfolio and the Estate map — the whole estate, sized by spend, colored by market position.
- Monthly briefing — the five minute version of the month, spoken or read.
- Strategy — the annual document: category strategies, vendor posture, the negotiation calendar.
- Verified wins — what the function has proven, invoice-verified.
- Trust center — the answer you forward when security asks how the data is held.

Learn these three workflows first: Build the board pack, Prove and share your savings, and Prepare for a vendor meeting for the calls you personally take.

Read in this order: Chapter 04 (your home chapter), then the first pages of Chapter 02 to understand what a benchmark verdict means when one crosses your desk.

THE ORG OWNER

You run the workspace itself: people, roles, security, connections, and the governance record. Your seat has its own full chapter, the owner's handbook, which walks the whole setup in one sitting; this section is the short form.

The five screens you live in

- Members — invitations, roles, seats, and offboarding.
- Security — two-factor, the org-wide policy, sessions, and the nested security sections.
- Integrations — every connected system, from contract intake to the SAM sources.
- Data controls — retention, deletion certificates, export, and consent.
- Team activity — who is actually using what, and the adoption report.

Learn these three workflows first: Set up your workspace, Connect your data and close the gaps, and the quarterly access review.

Read in this order: The owner's handbook end to end, Chapter 06 as its reference, then Chapter 10 so you can answer every plan question your team brings you.

Related: Your first 30 days · The owner's handbook · Which tool do I use?

Next: Chapter 17 · The owner's handbook

17 · The owner's handbook

Chapter 06 documents every settings page. This chapter is the same territory walked as work: what an organization owner actually does, in order, to take a workspace from empty to properly run. The core setup is one sitting of about an hour; the rest of the chapter is the decisions that deserve more thought and the cadence that keeps the workspace governed after the first week.

Everything here needs the **owner** role. If a settings page is managed by owners and you are not one, the page itself says who to ask, and any member can check who owns the workspace at the foot of the navigation rail.

BEFORE YOU BEGIN

Five minutes of gathering saves the sitting from stalling. Have ready:

1. **The teammate list**: names, emails, and who gets which role. Roles are assigned per person at invite time and changed inline later.
2. **Your logo** as an image file, if you want exported briefs and decks co-branded.
3. **Your plan facts**: what the workspace is on and what that includes. Billing states the plan, allowances, and seats; Chapter 10 is the full matrix. Two gates worth knowing before you promise features to the team: single sign on, SCIM, and the API ship on the Enterprise plan, and analyst benchmarks outside a plan that includes them are added per proposal as benchmark credits.
4. **Your organization's answer to two policy questions**, even provisionally: how long contracts should be retained after upload, and whether you contribute anonymized data to the community benchmarks and outcome network. Both live under Data controls and both default to the conservative choice until you decide.
5. If SSO is on your path: whoever administers your identity provider.

THE FIRST HOUR

Work top to bottom. Every step is reversible except where marked, and the Settings home vitals will confirm your progress: anything reading Off or 0 is a gap this sitting closes.

1. **Secure your own account first.** Security: register your authenticator, turn on two-factor, store the backup codes. You are about to hold the keys to everyone's access; your account is the one that must not fall.
2. **Set the workspace identity.** Branding: upload the logo so exported CFO briefs and boardroom decks carry it. Display currency: pick the currency your team enters amounts in. It relabels nothing retroactively and never converts a number; benchmarks stay in USD so market comparisons hold.
3. **Bring the people in.** Members: invite one by one with a role each, or paste up to 50 addresses for a batch with one role. Invitations expire after 7 days and a resend restarts the clock. Decide **domain capture** while

you are here: letting anyone on your email domain request to join replaces ad hoc invite requests with a queue you approve.

4. **Create the teams you will need for confidentiality.** Teams: named the way the company organizes (Procurement, Finance, Legal). A team is an access unit: grant a restricted contract to the team once, and joining or leaving the team is the access change itself. Set these up before the first sensitive document arrives, not after.
5. **Decide whether deals need sign-off.** Deal approvals is off by default. If your organization requires a documented approval step before signature, turn it on and name up to ten approvers. If it does not, leave it off; an empty ceremony helps nobody.
6. **Set the notification stance.** Notifications: everything about renewals is strictly opt-in and starts silent, so the workspace never spams a team into filtering it. Turn on what you personally want (**Turn on money-critical alerts** switches on the renewal set in one click), then tell the team the same control exists for them; each person chooses their own. Owners additionally set the AI copilot rules and can post alerts into a Slack or Teams channel once one is connected.
7. **Load the sample portfolio, then plan its removal.** From the dashboard, sample data gives every screen something to show while the team explores. Remove it before real data lands; Troubleshooting covers the cleanup if it lingers.

The hour is done when: the vitals on Settings home show two-factor On, people holding seats, and nothing urgent under Worth doing now.

SECURITY DEFAULTS WORTH SETTING IN WEEK ONE

None of these blocks the first hour, and all of them are cheaper to set before habits form.

1. **Require two-factor for everyone.** The organization policy lives on the same Security page as your personal setup. Turn it on once the team has had a day to register authenticators; from then on new members are held at setup until they comply.
2. **Set retention deliberately.** Data controls: a retention period auto-deletes contracts a set number of months after upload, rows and files both, with a deletion certificate written for each sweep. Blank means keep until deleted, which is the right answer for most estates; set a number only if your records policy requires one.
3. **Answer the consent question.** The community and outcome network toggle is also under Data controls. Contributing anonymized data sharpens the benchmarks everyone draws on; opting out is one switch and entirely respected.
4. **Know your audit surface before you need it.** Audit log: who did what, including file views and downloads as separate events, exportable as CSV, streamable to a SIEM via webhook. Nothing to configure on day one beyond knowing it is there; security teams that want the stream set it up under Integrations.
5. **Network and session policy is available, and sharp.** Under Security baseline, owners can pin the workspace to known networks (IP allowlist) and bound session lifetimes. The allowlist refuses to switch on unless your own address is covered, but an office IP change later still locks everyone out until the policy is

fixed from an allowed network. Prefer ranges over single addresses, and skip this control entirely unless policy demands it.

THE SINGLE SIGN ON DECISION

Single sign on is worth a considered yes or no rather than a default. The short version:

- **Say yes** when your organization already manages access through an identity provider and the workspace is on the Enterprise plan (SSO, SCIM, and the API are Enterprise features). SAML connects the provider; SCIM tokens let it create and deactivate accounts automatically, which makes your joiner and leaver process the single source of truth.
- **The enforcement choice is the real decision.** Connecting SSO lets people use it; requiring it for your domain makes it the only door. Require it once you trust the provider setup, and read the break-glass policy first: org owners keep password sign-in even under Require SSO, so a broken identity provider cannot lock the whole tenant out, and every break-glass sign-in is written to the audit log.
- **Group push, used carefully, replaces manual team upkeep.** Pushed groups can map to Teams, and optionally to roles; ownership is never assigned through a mapping, by design. Members must be provisioned as users before group push places them.
- **Say not yet** when the team is small and stable. Two-factor for everyone plus disciplined offboarding covers a great deal, and SSO added later loses nothing.
- **When you say yes**, the walkthroughs are written: SSO with Okta, with Microsoft Entra ID, with Google Workspace, and SCIM provisioning.

BRINGING THE ESTATE IN

The owner's part of data-in is choosing the door and clearing the way; Chapter 15 covers the team's part, and Chapter 18 is the full guide: the four doors compared, spreadsheets and registers, coming from another tool, and the migration FAQs.

1. **A handful of documents**: no ceremony. Upload a proposal for anything that needs benchmarking; Agreements for signed contracts with dates to watch.
2. **A folder, up to a few hundred files**: Your Archive. Anyone can run it; nothing about it is owner-gated.
3. **A whole estate, hundreds to thousands**: Migration Studio, which is owner-run precisely because it writes thousands of records. Everything stays in staging, invisible to alerts and agents, until you sign it off against a quality bar; the sign-off writes to your audit log with the scorecard attached. It is switched on per organization, so ask your VendorBenchmark contact if the account menu does not show it.
4. **Then connect the systems that keep it current.** Integrations: contract intake rails, the SAM and ITAM sources that know entitlements and deployments, procurement, and the invoice forward-in address. Connect your data and close the gaps is the recipe, and Data health is the scoreboard that says when you are done.

THE GOVERNANCE CADENCE

A well-run workspace needs about an hour a month after setup, on three clocks.

- **Weekly, five minutes.** Your own My work queue, and a glance at Control room if agents or workflows run in your workspace: what is waiting on a person is listed by name.
- **Monthly, twenty minutes.** Team activity: who has not started, which features go untouched, and the on-demand usage report when you want an adoption plan written for you. Seats and pending invitations on Members while you are there: revoke what expired unaccepted.
- **Quarterly, thirty minutes.** An access review: snapshot the whole access surface, confirm or revoke every line, and store the attestation. Completed reviews export as CSV for auditors. Pair it with a skim of the audit log filtered to member changes and file downloads, and a check under API that no key is lingering unused (keys expire within a year regardless).

WHEN SOMEONE LEAVES

Offboarding is three moves, in this order, all on Members:

1. **If they are the only other owner**, promote a replacement first. The platform refuses any change that would leave the workspace without an active owner.
2. **End access**: remove them (or block, if the seat should stay reserved). Either signs them out everywhere immediately. For a lost laptop rather than a departure, **Sign out** alone ends every session without touching the seat.
3. **Let structure do the cleanup.** Team memberships die with the member, and every confidential grant held through a team dies with the membership; nothing to sweep per document. The next quarterly access review confirms the surface is clean.

A member leaving of their own accord needs nothing from you: anyone can leave a workspace from the bottom of their own Profile.

THE QUARTERLY OWNER CHECKLIST

The standing review, one screen per line. Thirty minutes, four times a year.

STANDING?	THE CHECK	WHERE
☐	Access review completed and stored	Access reviews
☐	Two-factor coverage is full, or the policy explains why not	Security posture
☐	No pending invitation is stale, no seat is unaccounted for	Members
☐	Team rosters match how the company actually organizes today	Teams
☐	Audit log skimmed: member changes and file downloads look right	Audit log
☐	API keys are all known and in use	API
☐	Connected integrations synced recently, none silently failing	Integrations
☐	Retention and consent settings still match policy	Data controls
☐	Plan and allowances still fit how the team actually uses the platform	Billing

---

Related: Set up your workspace · Settings & account · Your first 30 days

Next: Chapter 01 · Contracts & renewals

18 · Migration & data-in

Everything on the platform gets sharper with data: benchmarks read your real deals, renewals watch real dates, and Vera answers from your actual estate. This chapter is the complete guide to getting that data in, whether you are carrying one quote, a folder, a whole contract estate, or an export from the tool you are leaving. It ends where every migration should: at a green Data health score.

The honest headline first: **you do not migrate into this platform by retyping**. Documents are read, registers are mapped, and your corrections always outrank the machine's guesses. The work that remains for a person is choosing, checking, and deciding, not data entry.

CHOOSE YOUR DOOR

Four doors, sized to what you are carrying. They combine freely: most organizations use the folder door in week one and the connector door forever after.

THE DOOR	BEST FOR	YOUR EFFORT	WHAT IS STANDING AFTERWARDS
One by one	A handful of documents, or anything you want benchmarked now	Minutes per document	Proposals in the analyst queue; agreements with watched dates
The folder	Your contract folder as it exists today, up to 300 files a run	One drag, then answer three questions	The estate filed by vendor, renewals flagged, an executive brief
The estate	Hundreds to thousands of documents, plus your contract register	A guided run with a sign-off	A reconciled, quality-scored portfolio that goes live in one stroke
Connectors	Keeping everything current from your other systems	An owner connects each source once	Contracts, invoices, and license data flowing in on their own

Two properties of intake worth knowing before you pick:

- **Nothing ambushes your team.** Renewal alerts are strictly opt-in, so loading your estate never triggers a wave of email; deadlines start reaching people only when someone switches alerts on under Notifications. Migration Studio goes further: everything it loads is held in staging, invisible to alerts, digests, and agents, until an owner signs it off.
- **Duplicates are detected by content, not by file name.** Re-dropping the same archive is safe on every door: byte-identical documents are counted and skipped, never imported twice, while two same-named order forms a week apart are correctly kept as two real order forms.

THE ONE-BY-ONE DOOR: PROPOSALS AND AGREEMENTS

The right door when the point is the document itself, not the volume.

For anything you want benchmarked, use Upload a proposal: drop up to 20 files, watch the vendor, value, term, and seat chips fill in per file (shaded by how confident the read is), correct anything gold or red, and confirm the batch once. PDF and DOCX, 40 MB each, verified by contents rather than name.

For signed contracts with dates to watch, add Agreements: vendor, expiry date, notice period (30, 60, or 90 days), annual value, and up to five files per vendor for safe keeping. You can also type an agreement in beside the ledger and let the AI fill the details, or drop the signed paper straight into the upload panel. Up to 40 vendor agreements are tracked this way; past that scale you want the folder or estate door.

Time: minutes. A proposal batch reads as it lands; an agreement is a one-minute entry.

THE FOLDER DOOR: YOUR ARCHIVE

The door most organizations should start with. Drag your contract folder, subfolders and clutter included, onto Your Archive; Vera reads every document, files it into one folder per vendor, and judges each agreement current, due for renewal, expired, superseded, or perpetual. The judgments respect how vendors actually work: an old perpetual license is a current asset, and only your newest enterprise agreement counts as live.

What to expect during the run. Up to 300 files per run, 40 MB each, PDF and DOCX; every unreadable file type is named in the results rather than silently dropped. A live ticker narrates the filing, and you can leave at any point: a floating chip follows you with the percentage, and a notification brings you back when the archive is ready, even if you closed the tab mid-scan.

The three questions at the end. The Next moves rail asks, with the money named on each: which genuinely due renewals go on the calendar (never ten-year-old paper, and vendors you already track are advanced, not duplicated), whether to review your top 20 vendors by spend, and which contracts deserve an analyst benchmark. Nothing happens without your confirmation.

The Estate Intake Brief. Every scan also produces an executive document in the house brief style: the vendor table, the renewals inside the window, the lifecycle findings, and a bottom line. It prints clean, which makes it the fastest way to show a CFO what the estate actually holds, usually on day one.

What lands where. Archive uploads land as contracts, so everything you drop also feeds Document search, the spend portfolio, and the vendor pages automatically. Documents the scan could not read queue behind the **Needs a look** chip; fill in a vendor, end date, or value and the contract refiles instantly.

Time: one drag, then the scan works while you do something else. Plan for the three follow-up questions to take ten minutes of actual decisions.

THE ESTATE DOOR: MIGRATION STUDIO

When the estate is hundreds to thousands of agreements, plus a register spreadsheet someone maintains, use Migration Studio. Chapter 00 walks the full run step by step with screenshots; this section is the decision view.

Why this door is different. Three properties the folder door does not have:

1. **Staging.** Nothing a migration loads is visible anywhere on the platform until sign-off. No alert fires, no digest counts it, no agent reads it. Your estate becomes real in one transaction when an owner presses **Go live**, and the acceptance is written to the audit log with the scorecard attached.
2. **The register comes along.** Your master spreadsheet (CSV or XLSX) is mapped column by column with a dry run that shows exactly what it would create, update, and reject before anything is written. Columns the platform does not recognize are kept on each contract as custom fields, exactly as you wrote them; nothing in your register is thrown away.
3. **Reconciliation with evidence.** Vendors that look like the same company, byte-identical duplicates, and agreements that supersede one another arrive as proposals with evidence and confidence, each yours to accept or reject. Binding is where renewal dates become true: a master agreement and its amendments become one contract whose terms in force are computed across the whole stack, newest document winning on each term separately.

Who runs it. An owner, because a run writes thousands of records; it is switched on per organization, so ask your VendorBenchmark contact if the account menu does not show it. The owner can start the run and hand the working of it to someone else.

Time: the run generates its own plan from your six scoping answers, with each phase carrying a date range counted in working days, so a 300 document estate is not scheduled like a 5,000 document one. Reading a large estate takes several passes over hours and continues while the tab is closed. Budget real person-time for two moments: the register mapping decisions, and the reconciliation and exceptions queues, which are sorted by annual value so the top of the list is what matters.

COMING FROM A SPREADSHEET

Most estates arrive with a register: the spreadsheet that lists vendors, values, and dates. Three preparations make it map cleanly, all drawn from how the mapping step actually reads your file:

1. **Know what your value column means.** A column headed something like Contract Value is read as the total over the term and divided by the term length to get the annual figure. If your column is already annual, map it to **Annual value** instead. Getting this one mapping right is the difference between a correct portfolio and one inflated by every multi-year deal.
2. **Expect the date question.** If nothing in a date column proves whether 01/04/2024 means 1 April or 4 January, the mapping deliberately stops and asks, because guessing moves every date in your register by up to eleven months. Answering once unlocks the dry run.
3. **Keep your odd columns.** Internal owner, cost center, procurement reference: leave them in. Unrecognized columns are preserved on each contract as custom fields rather than discarded.

Run the **dry run** before applying: it states how many contracts would be created, how many updated, and how many rows cannot be used with the reason for each, and rejected rows download as a CSV you can fix and upload again. A register alone, with the documents to follow later, is a legitimate way in: rows create and update contracts, and paper attached later adds the evidence behind the terms.

COMING FROM ANOTHER TOOL

Leaving a contract management or SAM tool for this platform is an export plus one of the doors above, not a special project. What to expect:

- **What carries over cleanly:** the documents themselves (export them to a folder tree, ideally one folder per vendor, and use the folder or estate door), and your register (most tools export their contract list as CSV or XLSX; map it in Migration Studio). Terms are re-read from the paper itself rather than trusted from the old tool's fields, which is a feature: you arrive with evidence-backed data, and folder names become the first guess at each contract's vendor.
- **What is recomputed rather than imported:** renewal dates and lifecycle states. The old tool's flags are ignored in favor of what the bound document stack actually says, which is usually the moment teams discover renewal dates their previous system had wrong.
- **What you rebuild by choice, not necessity:** your clause standard in the Clause library, notification choices per person (strictly opt-in here), Teams for confidential access, and any saved views and workflows. The owner's handbook sequences all of this.
- **A SAM or ITAM tool you are keeping** should not be exported at all: connect it instead (next section), and its entitlement and deployment data flows in continuously.

THE CONNECTOR DOOR: KEEPING IT CURRENT

Migration is an event; currency is a pipeline. Once the estate is in, an owner connects the systems that keep it true under Integrations:

- **Contract intake** — five email-in rails, DocuSign and Ironclad, and folder watch on Drive, SharePoint, or Box, so newly signed paper files itself from now on.
- **Spend and usage** — Coupa, Okta, Entra ID, the CSV importers, and Smart import, feeding the spend and license pictures.
- **License and deployment discovery (SAM and ITAM)** — the systems that know what you own and run; the license intelligence table then shows entitled against deployed against active per vendor, with estimated shelfware.
- **Procurement and ERP (P2P)** — purchase orders, requisitions, and supplier records, with maverick spend flagged.
- **Invoices** — the forward-in address, so invoice lines get checked against contract terms as they arrive (see Invoices).

The full catalog, with what each connector feeds, is Chapter 12; the per-connection setup walkthroughs are Chapter 19; the working recipe is Connect your data and close the gaps.

DATA HEALTH: THE FINISH LINE

"Migrated" is not a feeling; it is a score. Data health scans on every load and states the share of data checks you have covered, with each open gap naming the feature it would unlock and a plan ranked by value. It also carries the audit exposure estimate, built from whatever entitlement data is connected so far, which is your early read on how defensible the license position is.

Treat the finish line as three checks, in order: every agreement that matters carries an expiry date and a notice period (without both, there is no window to watch), annual values are on the agreements so Portfolio and Spend mean something, and the top-ranked gap on Data health has either been closed or deliberately parked with an owner. From there, the day 30 health check covers the rest of the workspace.

MIGRATION FAQS

Will re-dropping the same folder create duplicates? No. Duplicates are detected by content, not file name: identical documents are counted and skipped on every door. Two different documents that share a name are both kept, because they are both real.

Will loading the estate spam my team with alerts? No. Renewal alerts are opt-in and start silent on every channel, so intake never generates email on its own. Migration Studio additionally holds everything in staging, where no alert, digest, or agent can see it until sign-off.

My register dates are ambiguous, like 01/04/2024. The mapping stops and asks rather than guessing, and stays locked until you choose, because the wrong answer moves every date by up to eleven months. Answer once; the dry run then shows you your own rows as the platform understood them.

My spreadsheet has columns the platform will not recognize. They are kept on each contract as custom fields, exactly as written. Nothing in a register is thrown away.

My value column is the total contract value, not annual. That is the expected case: a Contract Value column is read as the total over the term and divided by the term. If yours is already annual, map it to Annual value instead.

Should I exclude old and expired paper? Include it. History is judged as history (superseded and expired paper visibly sits in the past), it never pollutes your renewal calendar, and it powers the reads that need a past: year-over-year comparison, the Time machine, and binding, which needs the full stack to compute the terms in force.

Some documents are scans with no text layer. They are reported as needing OCR rather than silently filed as contracts with no terms: the Archive queues them under **Needs a look**, and Migration Studio raises them in the exceptions queue, sorted by annual value.

We run several agreements with one vendor at once. So do most enterprises, and reconciliation knows it: overlapping agreements with one vendor are never proposed as superseded. Superseding is only proposed where the evidence supports it, and you decide each proposal.

Our contracts are in several currencies. Enter values in the currency the row carries; rows keep their own currency, and the workspace display currency covers rows that did not record one. Benchmarks stay in USD so market comparisons hold, and nothing is silently converted.

We only have the spreadsheet today; the documents are scattered. Start anyway: map the register, let rows create the contract ledger, and switch on the watching that dates enable. Attach paper as you find it, through Agreements or the Archive; binding recomputes terms as evidence arrives, and your corrections always outrank the automatic reading.

Related: Your Archive · Migration Studio · Connect your data and close the gaps · The owner's handbook

Next: Chapter 01 · Contracts & renewals

19 · Integration recipes

Chapter 12 is the map: what exists, what each connector feeds, and how the platform is delivered. This chapter is the wrench: one recipe per connection, walked step by step from both sides, with **You know it worked when** as the finish line and **To undo it** so nothing you try is a commitment. Most recipes need the **owner** role; each says so.

One honest note on the identity-provider recipes: the platform side is exact, and the identity-provider side follows each vendor's standard flow, whose menu labels shift slightly between versions. Where a label differs from what you see, trust the shape of the step: create a SAML app, exchange metadata, assign people.

BEFORE YOU BEGIN

1. **Know your gates.** Everything in this chapter is configured by an org owner. Single sign on, SCIM, and the API are Enterprise-plan features (the plan table); the email rails, chat connections, folder watch, and desktop app are not plan-gated.
2. **Know the operating rules**, which apply to every connector: credentials are stored encrypted and used only server side, read-only scopes are requested wherever the source offers them, every sync lands under **Recent activity** on the Integrations page, connector-reported entitlement figures wait for human confirmation before they can drive an alert, and disconnecting a source stops its syncs immediately.
3. **Sequence by leverage.** If you are wiring a workspace from scratch: identity first (SSO and SCIM), then the cheap intake rails (email-in, folder watch), then one usage source and one spend source, then data-out. That order is argued in Chapter 12 and it holds.

SINGLE SIGN ON WITH OKTA

SAML sign-on for the workspace, governed from your Okta org.

For: org owners, with your Okta administrator · **Plan:** Enterprise · **Time:** about 30 minutes including the test sign-in

Steps

1. In Okta, create the app: **Applications › Create App Integration**, choose **SAML 2.0**, and name it so your team recognizes it in their Okta dashboard (Vera AI works).
2. Open Settings › Security › Single sign on in the platform in a second tab, and copy the service-provider details the page shows into Okta's SAML settings, so the assertion comes back to the right place.
3. Back in Okta, finish creating the app and copy its **metadata URL** from the app's Sign On tab.
4. On the platform's SSO page, paste the metadata URL (or the metadata XML itself if your policy prefers a static copy), set your **email domain**, and save. Leave **Require SSO** off for now.

5. Assign yourself to the app in Okta and test: sign out, sign back in through Okta, and land in your workspace.
6. Assign the rest of the team in Okta, and only then decide enforcement. Turning **Require SSO** on makes Okta the only door for your domain, with one deliberate exception: org owners keep password sign-in as break-glass, and every break-glass sign-in is written to your audit log.

You know it worked when a teammate signs in from their Okta dashboard tile without a password and lands in the right workspace, and the sign-in appears in your audit log.

To undo it: switch **Require SSO** off first (never remove a required IdP's config while it is the only door), confirm password sign-in works, then remove the configuration.

SINGLE SIGN ON WITH MICROSOFT ENTRA ID

The same recipe with Entra (formerly Azure AD) as the identity provider.

For: org owners, with your Entra administrator · **Plan:** Enterprise · **Time:** about 30 minutes

Steps

1. In the Entra admin center, create the app: **Enterprise applications › New application › Create your own application**, choose the non-gallery option, then open its **Single sign-on** blade and pick **SAML**.
2. Copy the service-provider details from Settings › Security › Single sign on into Entra's Basic SAML Configuration (Entra calls the return address the Reply URL).
3. Copy Entra's **App Federation Metadata Url** from the SAML certificates card.
4. On the platform's SSO page, paste the metadata URL, set your email domain, save, and leave **Require SSO** off until testing is done.
5. Assign users or groups to the app in Entra, test a sign-in end to end, then decide enforcement. The break-glass rule is the same: owners keep password sign-in, audited.

You know it worked when the app appears in a teammate's My Apps and clicking it lands them signed in, no password prompt from us.

To undo it: Require SSO off, confirm password sign-in, then remove the configuration.

SINGLE SIGN ON WITH GOOGLE WORKSPACE

The same recipe with Google as the identity provider.

For: org owners, with your Google Workspace admin · **Plan:** Enterprise · **Time:** about 30 minutes

Steps

1. In the Google Admin console, create the app: **Apps › Web and mobile apps › Add app › Add custom SAML app**.
2. Google shows its IdP details first: download the **IdP metadata** file.
3. Copy the service-provider details from Settings › Security › Single sign on into Google's Service provider details step.

4. On the platform's SSO page, paste the downloaded metadata as XML (Google's flow hands you the file rather than a URL), set your email domain, and save with **Require SSO** off.
5. In Google Admin, turn the app **ON for everyone** (or the org units that need it), test a sign-in, then decide enforcement, with the same owner break-glass rule.

You know it worked when the app shows in a teammate's Google app launcher and the sign-in round-trips cleanly.

To undo it: Require SSO off, confirm password sign-in, then remove the configuration.

SCIM PROVISIONING AND GROUP MAPPING

Accounts created, updated, and deactivated by your identity provider, so your joiner and leaver process is the single source of truth.

For: org owners, after SSO is working · **Plan:** Enterprise · **Time:** 20 minutes, plus your IdP's first sync

Steps

1. On Settings › Security › Single sign on, issue a **SCIM token**. Treat it like a password; it is the credential your IdP will authenticate with.
2. In your identity provider, enable provisioning for the app you created above and paste the SCIM endpoint details and token.
3. Provision users first: membership must exist before group operations can place anyone.
4. If your IdP supports group push, push the groups that should map to Teams; joining and leaving a pushed group then grants and revokes that team's confidential access automatically.
5. Optionally map groups to roles: members of a mapped group get that role on every push (approver, analyst, member, or viewer), and leaving the group takes the role with it. Ownership is never assigned this way, by design; owners are promoted by owners, on Members.

You know it worked when a user created in your IdP appears on Members without an invite, and deactivating them there removes their access here.

To undo it: turn provisioning off in the IdP, then revoke the SCIM token on the SSO page. Existing accounts stay and fall back to manual management.

VERA IN SLACK

Alerts in the channel your team already watches, and the assistant on call in your own Slack workspace.

For: an org owner with Slack app-install rights; each person links their own DMs · **Time:** 15 minutes for the workspace; 2 minutes per person

Steps

1. An owner connects Slack under Settings › Connections › Integrations (it lives in the renewal and deal workflow group) and approves the install in Slack.

2. Pick the shared channel for org alerts under Notifications (the owners' webhook controls), and keep it money-critical only; a noisy alert channel trains people to ignore it.
3. Each person who wants renewals in their own DMs links once: send the Vera bot a direct message in your Slack, then flip the **Chat DM** toggle on the alert types you want under Notifications. The DM lane is strictly opt-in, per person, per alert type.
4. Use the assistant where the conversation is: ask Vera in Slack about a price, a clause, or a renewal, and the answer comes back with every figure tagged to its source, exactly as in the app.

You know it worked when a test question to the bot answers in Slack, and the note above the Chat DM table confirms your account link.

To undo it: disconnect Slack on the Integrations page; DM links and channel posts stop with it.

VERA IN MICROSOFT TEAMS

The same capability for Teams estates.

For: an org owner with Teams app permissions; each person links their own DMs · **Time:** 15 minutes for the tenant; 2 minutes per person

Steps

1. An owner connects Teams under Settings › Connections › Integrations and approves it in your tenant.
2. Choose the shared alert channel under Notifications, same discipline as Slack: money-critical only.
3. Each person links their DM lane by messaging the Vera bot once in Teams, then opting in per alert type under Notifications.
4. Ask Vera in the channel or the DM; figures arrive cited, as everywhere else.

You know it worked when the bot answers in Teams and your linked account is confirmed on the Notifications page.

To undo it: disconnect Teams on the Integrations page.

THE EMAIL-IN RAILS

Five private addresses that turn forwarding into filing. The cheapest integration on the platform, and usually the first one worth doing.

For: an org owner provisions; the whole team uses · **Time:** 10 minutes to provision and share all five

Steps

1. On Settings › Connections › Integrations, provision the rails you want; each is one click, and each address is private to your organization and unguessable by design, so mail sent to yours can only land in your workspace:

- **Contract email-in** — forwarded proposals become tracked Proposals.
- **Deal email-in** — forward a live vendor thread onto the deal record.
- **Invoice email-in** — forwarded invoices reconcile on Invoices.
- **Purchase request email-in** — teammates raise Requests by mail.
- **Ask an agent by email** — a question plus a document; the answer returns to the sender's inbox.

1. Put the addresses where people will actually find them: the team wiki, and as contacts in shared mailboxes. The invoice address belongs in your AP team's hands.
2. Send one real test down each rail you provisioned.

You know it worked when a forwarded proposal appears on Proposals reading as its document, and a forwarded invoice shows line verdicts.

To undo it: remove the rail on the Integrations page; the address stops accepting.

FOLDER WATCH ON DRIVE, SHAREPOINT, OR BOX

Documents file themselves from the folder your team already saves into.

For: org owners, with access to the folder being watched · **Time:** 15 minutes

Steps

1. On Settings › Connections › Integrations, pick your platform: Google Drive folder watch, SharePoint folder watch, or Box folder import for bulk loads.
2. Authenticate with an account that can read the folder, and choose the folder. Prefer a purpose-named folder (Signed contracts) over a busy shared drive root, so the watch ingests intent rather than noise.
3. Drop one test document in and watch it arrive; from then on, saving to the folder is filing.

You know it worked when the test file appears in your workspace and the sync shows under Recent activity.

To undo it: disconnect the watch; already-imported documents stay, and the folder is no longer read.

THE DESKTOP APP, SELF-SERVE AND IT-MANAGED

The Windows thick client: same account, same data, native window, drag-in from Explorer, Windows Hello lock.

For: anyone self-serve; IT for a managed rollout · **Time:** 5 minutes per machine

Steps

1. Self-serve: open Settings › Workspace › Desktop app, press **Download for Windows**, and run the installer. It installs per user, no admin rights needed.
2. If Windows shows a "Windows protected your PC" screen, choose **More info**, then **Run anyway**; the page documents this exact case.
3. Sign in with the same email, password, and two-factor code as the browser.

4. IT-managed rollout: because the install is per-user and needs no elevation, a managed rollout is distribution rather than packaging: push the installer through your software portal or endpoint tool and let it run in the user context. Updates are automatic after that; there is nothing to patch centrally.

You know it worked when the app has its own taskbar identity, a contract dragged from Explorer lands as an upload, and a notice-window notification arrives natively.

To undo it: uninstall from Windows Settings; it is a client, not a data store, so the workspace is untouched. Mac users get the identical workspace in the browser until the planned macOS build ships.

THE API QUICKSTART

From no key to your own data coming back over HTTPS in ten minutes.

For: org owners create keys; your developers consume them · **Plan:** Enterprise · **Time:** 10 minutes

Steps

1. On Settings › Connections › API, create a key. You see the full value once, so store it in your secrets manager on the spot. Keys expire within a year.
2. Make the first call, the one that proves the loop, against the REST API at `/api/v1` with the key as a bearer token:

```
curl -H "Authorization: Bearer <your-api-key>" \
  https://app.vendorbenchmark.com/api/v1/contracts
```

1. Read the JSON that comes back: your own organization's contracts, under your key's permissions, which is the whole authorization model demonstrated in one request. The connect panel on the API page carries the same snippet with your workspace's base URL filled in, plus endpoints for vendors, benchmarks, and Ask.
2. Build with the key server-side only; it is a credential, not a config value, and never belongs in a browser or a repository.

You know it worked when the response is your data, and the key's use appears in your audit trail.

To undo it: revoke the key on the API page; revocation is immediate.

CONNECT CLAUDE AND OTHER MCP CLIENTS

Your own AI assistants querying your benchmark data, with your permissions, over the Model Context Protocol.

For: org owners create the key; anyone with an MCP-capable client connects · **Plan:** Enterprise (rides the API) · **Time:** 10 minutes

Steps

1. Create an API key as in the API quickstart; the MCP connector authenticates with the same key.

2. Open the **Connect** panel on Settings › Connections › API and pick your client. The MCP server lives at `/api/mcp` over HTTP transport, and the panel shows a ready-to-paste block per client.
3. For Claude Code, the panel's command is the whole setup:

```
claude mcp add vendorbenchmark --transport http https://app.vendorbenchmark.com/api/mcp \
  --header "Authorization: Bearer <your-api-key>"
```

1. For ChatGPT, other MCP tools, or your own agents, use the matching block from the same panel; the Claude API's MCP connector and the Agents SDK take the server URL and key in their standard remote-server shape.
2. Test with a question only your data can answer: "which of our agreements renew in the next 90 days?" A grounded answer with your vendors in it is the proof.

You know it worked when your assistant answers from your workspace's records, not from general knowledge.

To undo it: remove the server from the client and revoke the key; either alone severs access.

WEBHOOKS AND DATA-OUT

Your data as a citizen of your estate: events pushed to your systems, tables in your warehouse, the audit stream in your SIEM.

For: org owners, with whoever runs the receiving system · **Time:** 20 minutes per channel

Steps

1. **Event webhooks:** on Settings › Connections › Integrations, add your endpoint (your own service, Zapier, or Make). Events arrive as JSON, HMAC-signed with a timestamped signature. Verify the signature before trusting any payload; the signing secret exists so your systems can authenticate us the same way we authenticate inbound mail and CLM webhooks.
2. **The audit webhook:** point it at your SIEM and the audit log streams continuously instead of being exported quarterly. Security teams should prefer this to CSV pulls.
3. **Warehouse export:** connect Snowflake or BigQuery and the contracts, renewals, savings, and spend tables land where your analysts already work.
4. Send a test event end to end before declaring any channel done, and alarm on signature failures rather than dropping them silently; a failing signature is either misconfiguration or someone knocking.

You know it worked when a real event arrives, its signature verifies, and your receiving system acted on it.

To undo it: remove the endpoint or disconnect the export; delivery stops immediately.

Related: The integration catalog · Integrations (settings reference) · Connect your data and close the gaps · The owner's handbook

Next: Chapter 12 · Integrations, clients & deployment for the architecture behind these recipes.

20 · Worked examples, start to finish

The reference chapters explain screens and the cookbook gives recipes. This chapter does the third thing a manual owes you: it works complete cases end to end, every screen named, every number real. Four of the five run on the sample portfolio, the same one the screenshots show and the same one **Load sample data** puts in your workspace, so you can follow along live. The fifth walks the platform's own reference deliverable, whose figures are illustrative by design.

Each case ends with **With your own vendor**: the two or three pointers that turn the example into your Monday.

CASE 1: THE NORTHWIND RENEWAL, FROM UPLOAD TO PROVEN SAVING

The spine of the platform, walked once in full: a renewal quote becomes a delivered finding, the finding becomes a negotiation, and the outcome becomes a number finance accepts.

The situation. The sample workspace holds a CRM subscription with Northwind Software: $142,800 a year for 250 seats, a 7 percent uplift cap, and an auto-renewal with a 60 day notice window. The vendor's renewal proposal asks $162,792, a 14 percent increase.

The work.

1. **The breach is visible before anyone reads anything.** The Agreements masthead counts the ask above the cap, and Anomalies flags it: 14 percent asked, 7 percent permitted. That alone reframes the first conversation.
2. **The paper goes in for expert review.** Upload a proposal with the focus question "is this renewal uplift fair?". The status path runs Uploaded, In review, Benchmarked, Delivered, with email at each stage.
3. **The finding lands.** The delivered report states it in the standard's order: the deal is above market with $28,650 a year in reach, evidenced by a seat price of $571 against a $504 market median (worth $16,750 a year) and a 12 percent discount off list against a 22 percent median (worth $11,900). The structural read compounds it: the 7 percent cap against a 5 percent median means the premium widens every year it is left alone.
4. **The finding becomes a mandate.** A tracked Negotiation opens with the target at the median seat price of $504, the discount rebuild to 22 percent as the second ask, and the cap at 5 percent as the closing concession, exactly the levers the report priced.
5. **The conversation runs on the record.** Vendor emails go through tactics analysis, every give and get lands in the concession ledger, and the vendor's promises are captured as commitments so they reach the final paper.
6. **The outcome is recorded, then proven.** The closed deal's reduction is answered into the 30 day What happened? follow-up, and Savings proof reconciles the claim against the invoices that follow. Only then is

it a number for the board.

What it comes to. A 14 percent ask, met with a benchmarked counter, closed on evidence, with the saving invoice-verified rather than asserted. Every figure in the story traces: the report's numbers to the benchmark panel, the saving to the invoices.

With your own vendor. Start where the sample starts: the agreement filed with both dates and the cap (Track agreements), the quote uploaded the day it arrives, and the mandate set from the report's levers, not from instinct (Plan and run a renewal negotiation).

CASE 2: THE RENEWAL BRIEF THE PLATFORM WRITES FROM ONE DROPPED QUOTE

What the deal desk brief contains and how to read it, walked against the platform's own reference deliverable: the MongoDB Atlas renewal benchmark. Its figures are illustrative, and they are exactly the shape your own dropped quote comes back in.

The situation. A vendor proposes a 3 year Atlas commitment of $2.5M total, roughly $833K a year. The buyer's realized discount is not stated on the paper.

The work.

1. **The quote goes in as paper, not typing.** Drop a quote reads the deal off the document; about three minutes later the brief opens in a deal desk room.
2. **The masthead answers the three questions before any prose.** Four figures lead: a 17 percent assumed baseline discount, a 20 percent peer median at this size, a 21 percent supported target floor, and a 3 year opportunity above $120K. Not one of them is a percentile.
3. **The load-bearing assumption is stated and priced, never smuggled.** The realized discount is unknown, so the brief says so, explains why 17 percent is credible for a first proposal, tells the reader how to verify it against public list rates, and carries a sensitivity table: if the true baseline is 14 percent the move to the floor is worth roughly $203K; if it is already 22 percent, the rate case is closed and the negotiation shifts to rollover, true forward overage, and a renewal cap.
4. **The peer table places the buyer inside it, by rank.** Twelve comparable commitments on a ladder, the buyer ninth on discount depth at a commitment slightly above the cohort average, with the honest reading attached: below the midpoint, not off the map, so the case is midpoint alignment plus term structure, not a claim of mispricing.
5. **The levers close as a package.** Rate to the 21 percent floor, pooled rollover, committed-rate overage, uplift waivers in writing, and a renewal cap at signature, with the sequencing stated and four talking points written to be said out loud.

What it comes to. A brief a senior analyst would sign: the discount, the range, and the rank in the first two seconds, the assumption priced, the ask anchored to the market's own trend line.

With your own vendor. Drop the paper that carries the pricing, and read your brief in this order: masthead, assumption, ladder, levers. If your paper shows no list price, the brief will say the discount cannot be verified and name what to request; that page is working correctly, and the handoff recipe turns it into the reply.

CASE 3: THE MICROSOFT TRUE-UP, DEFENDED

Walking into a true-up with the surplus already found, on the sample estate's own Microsoft position.

The situation. The sample estate runs 1,800 E3, 350 E5, and 900 F3 seats, and the flagged exposure on the E3 pool alone is $64,800 a year. The anniversary true-up is approaching, which is when unmanaged estates pay for their drift.

The work.

1. **The usage export goes into M365 optimization.** The estate scan colors every licensed user's verdict and ends in the money bridge from today's spend to the optimized mix; the E3 surplus is the headline finding.
2. **The reassignment plan runs now, not at renewal.** Inactive seats are reclaimed and oversized tiers stepped down per user, which is money this quarter and a smaller true-up base by definition.
3. **The finding becomes the negotiation case.** The Microsoft EA playbook folds the right-sized estate into the renewal engagement, and the pricing-era control matters: the tool models the November 2025 removal of EA volume discount levels directly, so the benchmark matches the rules your agreement actually renews under.
4. **The timing is checked against Microsoft's calendar,** not yours, on the Vendor calendar.

What it comes to. The true-up conversation starts from an estate you measured, with the surplus already off the table and the growth story priced under the correct discount regime.

With your own vendor. The same three moves in the same order, whatever the publisher: measure (Right-size licenses), fix what needs no negotiation, then negotiate the remainder from the playbook.

CASE 4: THE AUDIT LETTER, ANSWERED

The day-one hour that decides an audit, run against the decoder's own worked sample.

The situation. An audit letter arrives naming Oracle. The estate's flagged Java position is real money: 2,600 employees under the employee metric, with $93,600 a year of flagged exposure, which is exactly the kind of backdrop that makes a fast, unconsidered reply expensive.

The work.

1. **The letter goes into the Audit letter decoder before anyone replies.** Not sure what you would get? The page's **See a decoded sample** shows a fully decoded Oracle GLAS notice. The decode names who is actually auditing, turns the stated deadline into a business-day countdown, and separates what is demanded from what is merely implied.

2. **The three things not to do first are honored:** no reply to the auditor yet, no vendor measurement scripts run, no data volunteered beyond the demand.
3. **The defensibility is checked on Data health:** the audit exposure estimate says how strong the entitlement picture is before engagement, and connecting a SAM source strengthens it more than any letter-drafting can.
4. **The response runs on your timeline** through the pre-filled defense playbook, with the Expert Bench engaged if the exposure warrants a specialist, and the audit worked as one negotiation with any nearby renewal, never two.

What it comes to. A considered first move, a countdown with margin, and a response built from records rather than recollection. The full recipe is the standing version.

With your own vendor. The decoder is vendor-agnostic and takes a phone photo. The discipline transfers unchanged: decode first, slow the clock, answer from paper.

CASE 5: THE CAPLESS QUOTE, REPAIRED, AND THE YEAR PROVEN

The smallest case in money and the most common in practice: a quote with a structural hole, fixed before signature, with the outcome landing on the year's proof.

The situation. The sample Contoso Analytics renewal quote: 120 Creator seats at $450, $54,000 a year, net 45, an auto-renewal with only a 30 day notice window, and no uplift cap anywhere in the paper. Nothing about the price is scandalous; the structure is the finding.

The work.

1. **The gap is caught by standard, not by vigilance.** The Clause library comparison flags the quote's silence on uplift against your recorded position and drafts your pre-approved language; the Coverage sweep would catch the same hole across the whole estate.
2. **The fix travels as one attachment.** The Fair Renewal Rider generator produces the ready-to-send edition with the parties filled in, its indexed uplift cap doing the work the quote left undone, and the 30 day notice window argues for answering now rather than at the deadline.
3. **The negotiation is small and specific:** the cap, the notice window widened, and the price checked once against the market (a quick instant benchmark is enough at this size).
4. **The outcome joins the year's story.** The recorded result reconciles on Savings proof, and at quarter end the proven set, this case included, goes to finance as one CSV, with Verified wins minting the badge where one is deserved (the reporting recipe).

What it comes to. A deal that was "fine" on price and dangerous on structure, repaired for the cost of one attachment and one ask, and a savings ledger that ends the year answering for itself.

With your own vendor. Record your positions once, sweep the estate once, and let the standard catch what tired eyes miss. The rider goes on every renewal; the proof discipline goes on every outcome.

Related: Common workflows · Your first 30 days · Delivered reports

Next: Chapter 21 · The Vera handbook

21 · The Vera handbook

Vera is the platform's AI analyst, and like any analyst she does her best work for people who brief her well. This chapter is the working relationship in one place: where she lives, what she can and cannot see, how grounding works in plain language, the questions worth asking, how to teach her, where the humans take over, what the allowances mean in practice, and privacy stated plainly. The screen reference for the console itself is Vera AI; asking well is the compact recipe this chapter expands.

Everything conversational here requires an AI plan; Chapter 10 defines the gates.

WHERE VERA LIVES

One analyst, many doors, the same grounding behind all of them:

- **Ctrl/⌘ J**, anywhere: the overlay opens on whatever you are looking at, and the screen you are on is context she can use.
- **The console** at Vera AI, with the full chat, memory, and saved threads behind the Full chat button.
- **The Ask bar** on the dashboard and home screen: a question, an app search, and a contract attachment share one field.
- **The screen assistant**, the round button on every page: its question box is Vera herself, grounded in the screen you are on.
- **Inside the work**: the Ask panels in benchmark cases and workrooms, and the commission rail's question chip.
- **Where your team lives**: Slack and Teams, and email for a question with a document attached.

WHAT VERA CAN AND CANNOT SEE

The honest boundary, because trust in an assistant starts with knowing its limits:

- **She sees your organization's records**: contracts, proposals, agreements, benchmarks delivered to you, your savings ledger, and the documents you uploaded. She never sees another tenant's workspace; isolation is enforced at the database, and the Trust center states it with live numbers.
- **She sees the market evidence**: the benchmark cohorts, the published research library, the wire. This is what lets her answer market questions without your data.
- **She sees the screen you asked from**, when you ask through the overlay or the screen assistant, which is why asking from the record beats asking from a blank page.
- **She sees what you attach**, for that conversation: the Ask bar takes a contract explicitly, and nothing is attached unless you attach it.
- **She does not see the open internet on your behalf** except where a tool says so (the Deal-Room's vendor intel briefing is the named exception), and **she does not send anything to a vendor** outside Case Mode, where every send runs under a mandate you approved and you ride CC.

HOW GROUNDING WORKS, IN PLAIN LANGUAGE

Grounded means every figure she gives you has a source you can check, and a number without one is refused rather than invented.

- **Citations ride the answers.** A price, a clause, a date arrives tagged to the record or evidence it came from; thirty seconds of following the tag is what makes an answer forwardable.
- **AI-written reports open with their grounding state.** The trust strip says "Figures verified against your data" or raises a review notice, with the source count and date, before you read a word.
- **"Not on file" is an answer.** Where your records cannot support a number, the honest response is that sentence, plus what to add to change it. A generated report says not on file instead of estimating; hold conversations to the same bar.
- **The Figures ledger** is the audit trail: every number the platform has cited for you, traced to its source fact and the day it was current.

TWENTY QUESTIONS WORTH ASKING

Grouped by job. Each names the question and what comes back; the specific figures depend on your records, which is the point.

On price (the market questions)

1. *"Am I overpaying for Salesforce at 320 seats?"* — your position against the cohort: the discount implied, the peer range, the rank, cited.
2. *"What discount should I ask for on this renewal?"* — the range with the median, the strong ask, and the stretch, the same cohorts the instant benchmarks draw on.
3. *"What is a strong outcome for a deal this size with this vendor?"* — the target floor and stretch, with what each is worth a year.
4. *"What did this vendor do to its prices in the last year?"* — the terms-watch history, diffed, with the changes that touch your seats priced.
5. *"When should I time the ask?"* — the vendor's fiscal geometry against your window.

On paper (the clause questions)

1. *"What does our Northwind contract say about auto-renewal?"* — the clause itself, quoted and cited to the document.
2. *"Which of our contracts have no uplift cap?"* — the estate-wide answer, with Review tables as the exportable version.
3. *"What are the risks in this attached quote?"* — the redline-style read on the paper you attached.
4. *"What rights do we already hold with this vendor?"* — the harvested register: true-downs, price holds, exit assistance, with exercise windows.
5. *"What changed between last year's paper and this draft?"* — the differences called out, with the Time machine as the full instrument.

On your portfolio (the estate questions)

1. *"Which of my vendors renews next and what should I do first?"* — the calendar read with the recommended first move.
2. *"Where is the money across our stack?"* — spend by vendor with the above-market positions ranked by prize.
3. *"Catch me up on our Microsoft position."* — the takeover briefing: agreements, benchmarks, open deals, recent moves.
4. *"What did we prove in savings this quarter?"* — the invoice-verified set, distinguished from the claimed.
5. *"What in my estate is exposed if this vendor audits us?"* — the entitlement picture's defensibility, honestly bounded by what is connected.

On working (the platform questions)

1. *"How do I set up renewal alerts?"* — the steps, with the manual section cited; the Guide door goes to the same place.
2. *"Draft the follow-up email to my CFO about this report."* — the draft, built only from figures the report carries.
3. *"Which tool should I use to compare these three quotes?"* — the routing answer; Which tool do I use? is the standing version.
4. *"Turn this answer into a standing watch."* — the handoff into an agent or workflow, so the question keeps answering itself.
5. *"What would you need to answer that properly?"* — the meta-question that pays best: she names the missing record, and closing it upgrades every future answer.

TEACHING HER

She learns from what you do with her output, and you hold the eraser:

- **The learning is behavioral and slow by design**: rating an answer, exporting a deliverable, applying or rejecting a staged suggestion. The pass runs weekly and writes a preference only after a repeated pattern, so a new account starts empty. **A written feedback comment is the fastest teacher.**
- **Everything learned is inspectable and deletable** under AI personalization; removing a preference stops it shaping answers immediately.
- **Your organization's negotiation Memory** is separate and shared: posture, durable facts, and lessons from past deals. Briefs, mandates, and Ask open with it, so pruning it is high-leverage maintenance.
- **Corrections compound**: fixing the same extracted field the same way twice teaches the extractor, and the next document from that vendor is read with your corrections in hand.

WHEN A HUMAN TAKES OVER

Vera is the analyst on call, not the whole firm. The handoffs, in the order they usually arrive: uploaded proposals go to **human analyst benchmarking** and come back as delivered reports; Advisory books our analysts for the moments a report is not enough; the Expert Bench brings vetted deal specialists into a specific fight; Managed renewals puts the program under human-approved mandates with fees only on invoice-verified savings; and Chat reaches the team directly when the question is really for us.

THE ALLOWANCES, IN PRACTICE

The practical version of Chapter 10: conversation is protected by fair-use daily caps you will rarely meet and never meet unwarned, since a quiet meter appears near the input once less than half a day's cap remains. Trial workspaces meter Vera at ten questions a day and always see the meter. Tool runs count against their own windows, rolling rather than calendar, so allowance frees up on its own as older runs age out; the Billing page is the live meter for all of it. When something is greyed out, Troubleshooting names the fix.

PRIVACY, IN PLAIN WORDS

- **What is stored**: your workspace's records, the conversations and deliverables you create, and the learned preferences and memory described above, all inside your tenant's isolation.
- **What is never stored**: documents the instant tools read in the moment. A file dropped on the Verdict Bar or the quote lane is read in memory and not retained, and the tool pages say so where it applies.
- **What leaves the workspace**: nothing, unless you send it. Share links are minted by you and guarded by a confidentiality scan; vendor-facing drafts never reveal your percentile, your cohort, or where the numbers come from; contributing anonymized data to the community benchmarks is a consent switch under Data controls, yours to leave off.
- **The standing answer** for security reviewers is the Trust center, including how the AI handles your data, written to be forwarded.

Related: Vera AI · Ask Vera well · Your AI agents

Next: Chapter 13 · Your AI agents, for the standing workforce built on the same analyst.

22 · The reference card

The lookups, on one page: every keyboard shortcut, every email the platform sends with the setting that controls it, what exports to what, the fixed vocabularies, and the limits you actually hit. Nothing here is new; everything here is faster to find than in the chapter that teaches it.

KEYBOARD SHORTCUTS

Everywhere

KEYS	WHAT HAPPENS
Ctrl/⌘ K	The command palette: jump to any page, tool, or vendor, or run an action
Ctrl/⌘ J	Vera opens over whatever you are looking at
g then d / r / c / p / s / t / w / b	Jump without the palette: dashboard, renewals, proposals, portfolio, spend, tooling, watch, benchmarking
?	The shortcuts overlay, every jump grouped
Esc	Steps out one layer at a time: a dock panel, then an open room, then the case, then the stage

On specific screens

KEYS	WHERE	WHAT HAPPENS
/	Instant benchmarks, Research, the settings rail	Focus the search or filter field
j and k	The renewals queue	Step through the cases without the mouse
Enter and Tab	The benchmark search bar	Run the highlighted vendor; complete the name
Ctrl+Enter	Upload a proposal	Confirm the batch
Ctrl+V	Audit letter decoder	Paste a screenshot of the letter
Arrow keys	Readers, pagers, case files, the platform introduction, calendar Present mode	Turn the pages; walk the shelf
F	Reading mode	True fullscreen on the stage
N and U	The sourcing workspace	Leave a note on the wall; log what you did
N	Calendar Present mode	Presenter notes

EVERY EMAIL THE PLATFORM SENDS

Each with the switch that controls it. The standing rules: renewal-shaped email is strictly opt-in and starts silent; every recurring email carries a one-click unsubscribe that turns off only itself; quiet hours hold non-critical email, never the money-critical kind.

On by default, because silence would be a failure (each still has its toggle under Notifications):

EMAIL	WHEN IT ARRIVES
Report delivered	Your benchmark report is ready
Benchmark status updates	An analyst picks the proposal up (with the expected date); the report enters final review
The outcome follow-up	Once, about 30 days after a delivered benchmark: What happened?
@mentions	A teammate mentions you
New device sign-in	Your account signs in from a device it has never used
Invitations and account email	Invites, the change-of-email confirmation, password resets

Opt-in, per person (all under Notifications unless named otherwise):

EMAIL	THE SWITCH
Renewal notice alerts, deadline escalations, policy checks, above-market warnings, uplift breaches, stale benchmark reminders, drafted briefs, opening moves, prep packages	The renewal alert set; **Turn on money-critical alerts** flips it in one click
The Monday renewal radar	Rides the renewal notice alert
The daily leverage note	Rides the renewal alert set; sends nothing on a quiet day
The Monday brief (wire, renewals, watchtower, savings sections)	Per-section toggles
Price list changes, priced at your seats	The switch on Terms watch, plus your vendor follows
The Monday vendor wire	The subscribe card on the Weekly report
The weekly knowledge briefing	The weekly briefings card
Tips and lifecycle notes	One switch, which never touches invites, security notices, or deliveries
Chat DMs in Slack or Teams	DM the Vera bot once, then per-alert-type toggles
Your cases weekly digest, and the first-reply note per case	Arrive with open cases; the digest sends only while cases are open

EXPORTS BY FEATURE

YOU ARE LOOKING AT	IT LEAVES AS
A delivered benchmark report	PDF and Word, matching the on-screen brief
A Reporting studio run	Word, print-to-PDF, and a 30 day read-only forward link
An instant benchmark placement	PDF, and a share link (30 days, optionally kept live)
The Fair Renewal Rider	A ready-to-send Word document, or tracked-change insertions in the vendor's own DOCX
The Boardroom Brief	Word and PowerPoint, plus Present mode's printable leave-behind
Case Mode's portfolio	A three-slide PowerPoint board deck, and per-case executive briefs
The savings ledger, members roster, audit log, access reviews, review tables	CSV (review tables also Excel)
The estate map	A slide
The renewal calendar	An ICS feed your calendar app subscribes to; Present mode's printed packet
Research papers	Download (within the daily and weekly caps), a 90 day public link, or a forward to five colleagues
Contracts, renewals, savings, and spend tables	The warehouse export to Snowflake or BigQuery
The whole manual	One PDF, linked in the footer of every manual page

THE FIXED VOCABULARIES

VOCABULARY	THE WORDS, IN ORDER
Proposal status path	Uploaded · In review · Benchmarked · Delivered
Benchmark rank	Poor · Below market · At market · Good · Top of the market
Coverage verdicts	Present · Deviates · Missing · Unclear
Contract lifecycle	Current · Renewal due · Expired · Superseded · Perpetual · Needs a look
Vendor Twin verdicts	Lands · Traded away · Stalls · Reveals
Twin confidence	Evidenced · Informed · Directional
Report verdict tones	Favorable · Caution · Risk
Clause positions	Preferred · Fallback · Walk-away

LIMITS AT A GLANCE

The ones people actually hit, with the full table in Chapter 10:

LIMIT	VALUE
Proposal and contract files	PDF and DOCX, 40 MB each
Verdict Bar documents · images	30 MB · 3.5 MB
Archive runs	300 files per run
Tracked vendor agreements · files per vendor	40 · 5
Research downloads	10 per day, 30 per week per person
Benchmark share links	30 days
Invitations	Expire after 7 days
API keys	Expire within one year, shown once

Related: Plans, allowances & limits · Notifications · Troubleshooting

Next: back to the manual index.