SPERT® Suite Architecture Overview

Approval Kit

Components, data flows and trust boundaries for architecture, security and privacy reviewers · spertsuite.com/approval-kit

SPERT® Suite Approval Kit v1.0.0 (Release) · Content reviewed 2026-10-25 · Built 2026-10-26

For: [REVIEWING TEAM] at [ORGANIZATION] (architecture, information security, privacy)
Prepared by: [PROJECT MANAGER NAME], requesting approval to use [APPS REQUESTED]
Reviewed against: the app versions in section 1; shared backend in the spert-landing repository v2.7.2 (commit `52b7fbf`, 2026-10-25; Connect AI server version 1.14.1, confirmed live 2026-10-25); Terms of Service v1.4 and Privacy Policy v1.4 (both effective 2026-10-26); AI Privacy Notice v2.3 (effective 2026-10-26); AI Connectivity Consent Notice v2.2 (effective 2026-10-26). **Review date:** 2026-10-25.

This document supports an approval process. It is not legal, compliance or security advice, and it makes no commitments beyond those in the published legal documents.

THE SHORT VERSION

SPERT® Suite is six browser-based planning apps. Each one is a web page that runs entirely in the user's browser: the planning calculations happen on the user's device, and by default every project is saved in that browser's own storage. No account is needed, and no project data goes to the operator or to anyone else.

Two optional features add cloud services, and each can be approved or declined on its own. The numbers are labels, not a ladder: tiers 2 and 3 are separate decisions.

Tier	On by default?	What leaves the user's device	Suggested approval scope
1. Local-only use	Yes	Page requests to the web host (IP address, browser type). SPERT® Scheduler's optional holiday lookup sends a country code and year. No project data.	[APPROVE / DECLINE]
2. Cloud storage and sharing	No; the user signs in and accepts the Terms	Name, email and the projects the user saves, stored in Google Cloud in the United States	[APPROVE / DECLINE / NOT IN SCOPE]
3. Connect AI	No; started by the user each session, in three apps only	AI edits pass through a relay in Google Cloud. With Read Mode on, a copy of the open project goes to the relay and to the user's chosen AI provider.	[APPROVE / DECLINE / NOT IN SCOPE]

[ORGANIZATION] can approve tier 1 alone. Figure 1 (section 1) shows every component and which tier uses it; sections 2 to 4 show the data flows for each tier, with trust boundaries drawn as dashed regions.

1. THE SUITE IN BRIEF

App	What it is for	Address	Version reviewed	Connect AI
SPERT® Scheduler	Build and maintain a project schedule that accounts for uncertainty	`scheduler.spertsuite.com`	0.79.26 (`86b3e49`)	Write and Read
SPERT® Story Map	Map and size release scope before the first sprint	`storymap.spertsuite.com`	0.53.14 (`5e613c4`)	Write and Read

App	What it is for	Address	Version reviewed	Connect AI
SPERT® Forecaster	Forecast release dates and how confident to be in them	forecaster.spertsuite.com	0.47.4 (ff7ac3b)	Read only
SPERT® CFD	See where work is piling up (cumulative flow)	cfd.spertsuite.com	0.15.23 (c7f9c68)	Not offered
GanttApp™	Build and share a project timeline with uncertainty included	ganttapp.spertsuite.com	0.29.8 (64d1647)	Not offered
MyScrumBudget™	Plan and reforecast a project or team budget	myscrumbudget.spertsuite.com	0.44.4 (0825168)	Not offered

Descriptions are from the suite's home page, spertsuite.com. SPERT® Scheduler and SPERT® Forecaster use the Statistical PERT® estimation technique; the other apps do not depend on it. The legal documents also cover SPERT® AHP, which is not listed on the home page and is not part of this overview.

How it is built. Every app is a single-page React application (SPERT® Scheduler and SPERT® Story Map are built with Vite; the other four with Next.js), served by Vercel. None of the app repositories contains a server-side API route: the code that handles a user's projects runs in the browser. All six apps share one Google Firebase project for the optional features: Firebase Authentication for sign-in, Cloud Firestore for saved projects and the Connect AI relay, and Cloud Functions for sharing invitations, Connect AI pairing, and the Connect AI server itself. Invitation emails are sent through Resend. The suite has no analytics, telemetry or error-tracking libraries (none appears in any app's dependencies), and the Privacy Policy (§7.7) says it uses no tracking cookies.

Who runs it. famousdavis, LLC, a Florida limited liability company, has operated the apps since 2026-09-12 (Privacy Policy §1.1). It is a single-operator business. The source code is public on GitHub under GPL-3.0 with additional terms, so [REVIEWING TEAM] can read the code that runs.

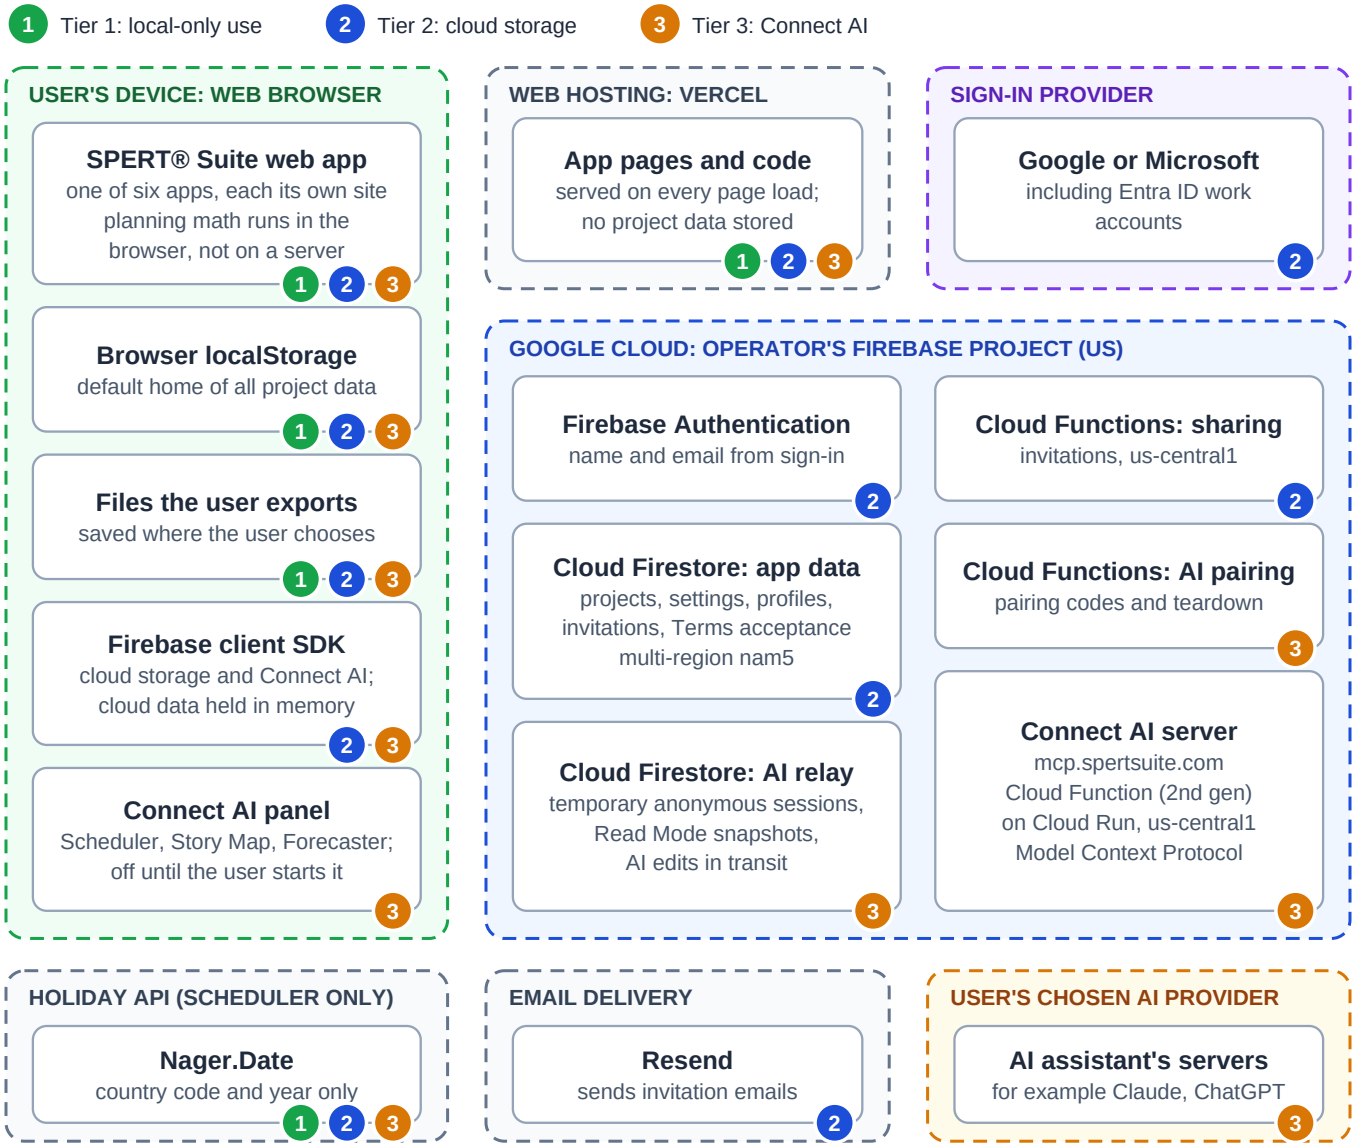

Figure 1. SPERT® Suite components. The numbered marks on each box show the tier or tiers that use it.

2. TIER 1: LOCAL-ONLY USE (THE DEFAULT)

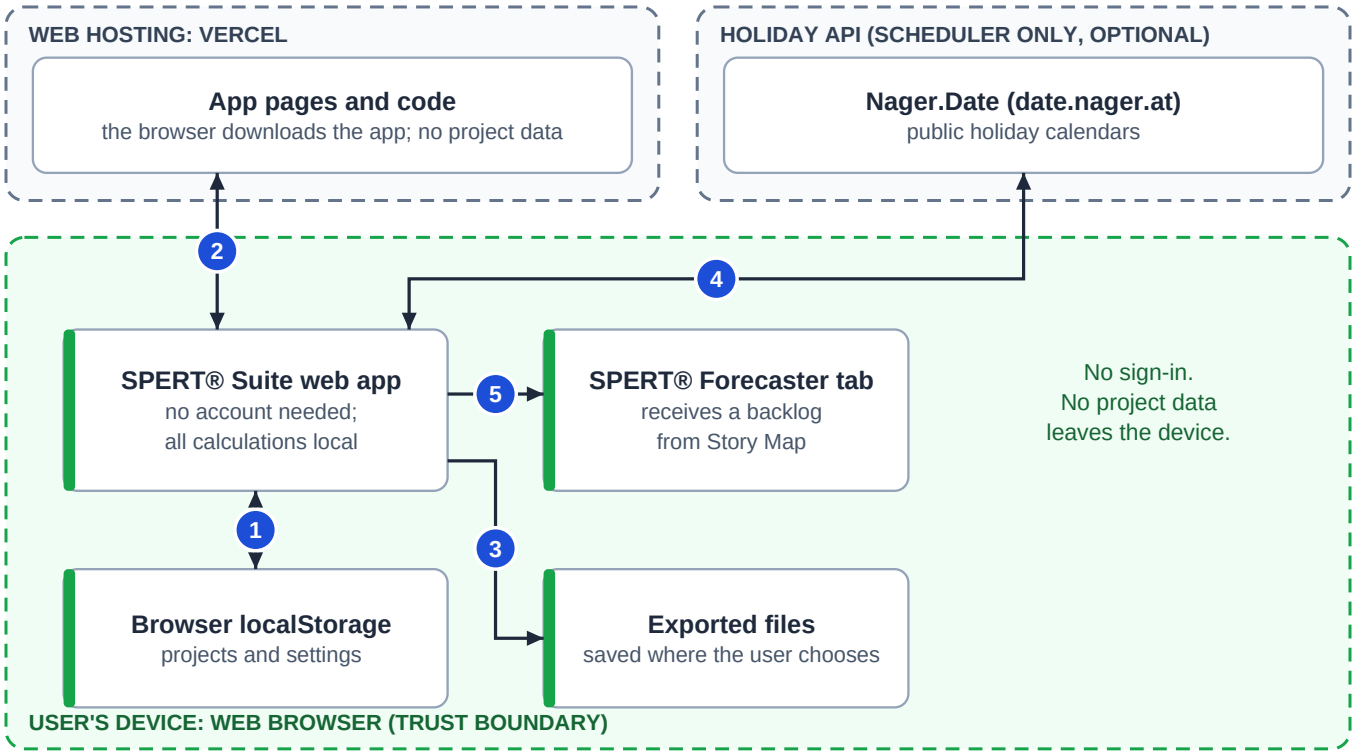

Figure 2. Tier 1 data flows. Nothing in this tier sends project data off the device.

#	From and to	What travels	When
1	App and browser localStorage	Projects and settings, read and written on the device	Always
2	Vercel to browser	The app's pages and code. Vercel receives standard request data, such as IP address and browser type (Privacy Policy §1)	Every page load
3	App to the user's files	Export files, saved where the user chooses	When the user exports
4	SPERT® Scheduler to Nager.Date (`date.nager.at`)	The list of supported countries when the Calendar page opens; a two-letter country code and a year when the user loads holidays. The service also sees the request's IP address	Calendar page only
5	SPERT® Story Map to a SPERT® Forecaster tab	A backlog, handed from one open browser tab to the other (`window.postMessage`). It does not pass through a server	Only when the user sends a backlog

What this means for a reviewer

— **The Privacy Policy (§1) says so directly:** core functionality "keeps your project data entirely within your local browser; the Apps do not send project data to the Operator or any third party."

— **Local storage belongs to the browser profile, not the person.** Anyone using the same browser profile can see its local projects, and clearing the browser's site data deletes them. Users should export regularly if a project matters.

— **Firebase is not involved.** No sign-in happens, and the Privacy Policy (§7.7) says no authentication cookies are set for users who do not enable cloud storage.

— **The holiday lookup is the one outbound call besides page loads,** and only in SPERT® Scheduler. Nager.Date is a public holiday API, described in Privacy Policy §1; it receives no project or account data.

3. TIER 2: CLOUD STORAGE AND SHARING

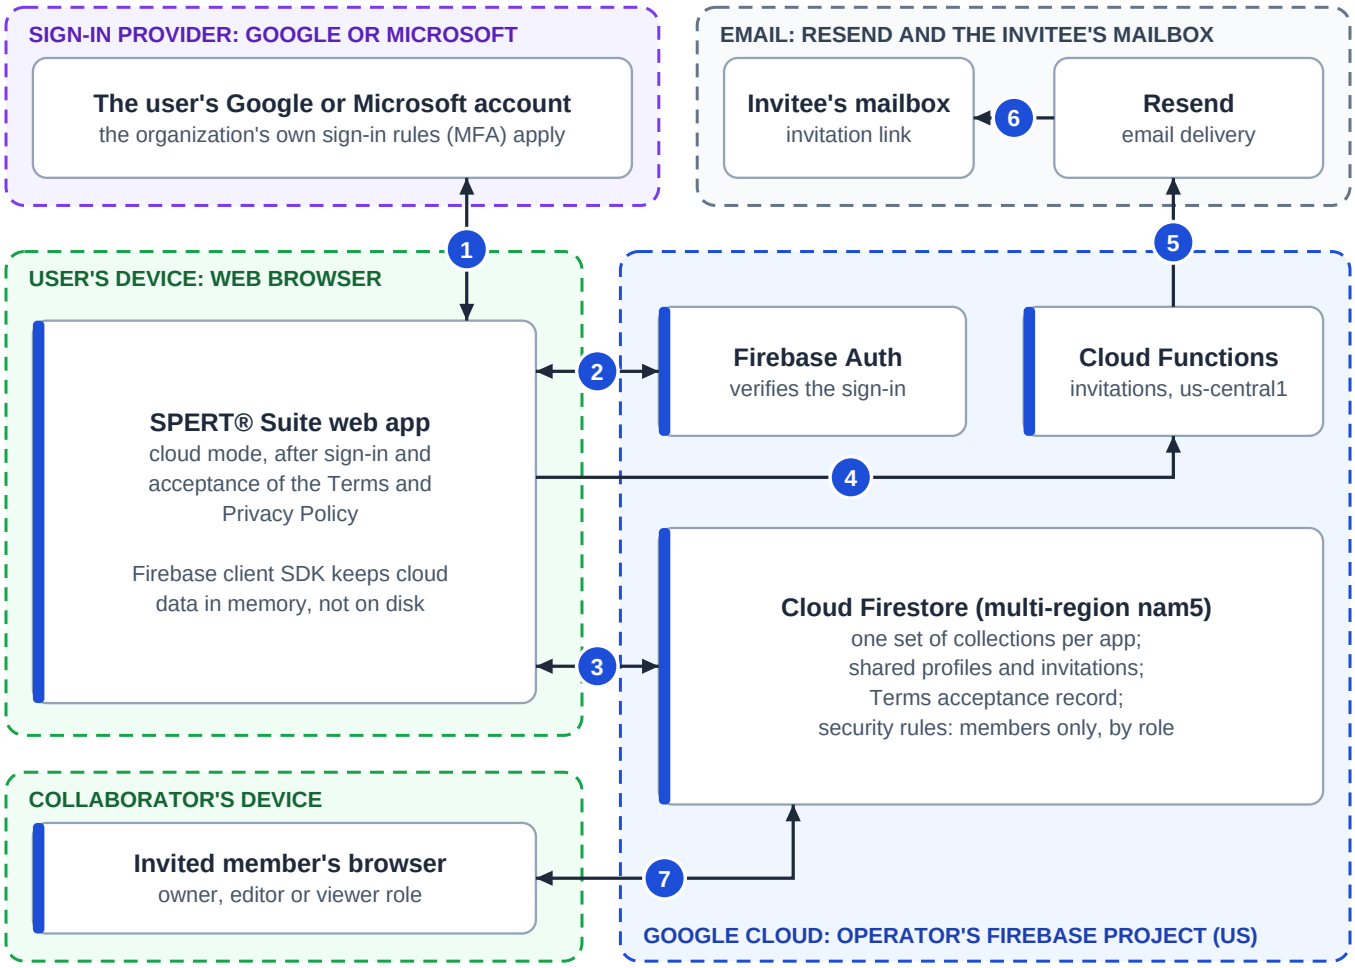

Figure 3. Tier 2 data flows. Cloud storage needs sign-in and acceptance of the Terms and Privacy Policy.

#	From and to	What travels
1	Browser and the sign-in provider	The user signs in with Google or Microsoft, including Microsoft Entra ID work accounts. SPERT® Suite never sees or stores a password
2	Browser and Firebase Authentication	The sign-in result. The operator receives only name and email (Privacy Policy §2.1)
3	Browser and Cloud Firestore	Projects, settings, a shared profile (name and email) and the record of which Terms the user accepted. Firestore security rules check every read and write
4	Browser to Cloud Functions	A sharing invitation: invitee email, role, project id and name
5	Cloud Functions to Resend	The invitation email to send
6	Resend to the invitee's mailbox	An email with a link to the shared project
7	Collaborator's browser and Cloud Firestore	The shared project, limited to the role the owner chose (owner, editor or viewer)

Where cloud data lives. The Privacy Policy (§9) says cloud storage data and Connect AI relay data "are currently stored at rest in the United States (Google Cloud Firestore multi-region nam5)", and that the operator "will give notice by updating this Policy before changing that location". Cloud Functions run in Google's `us-central1` region. On backups, §6 says: "Copies may remain in the Operator's database backups until they expire: up to 98 days for daily backups and up to 7 days for point-in-time recovery data. Backups are stored in the same United States location as the live database." On logs, §2.3 says operational logs are "retained for 30 days in Google Cloud Logging, whose default storage location is not restricted to a specific country or region."

Who can see a project. Firestore security rules let a user list only the projects they are a member of, read a project only if they are a member, and delete it only if they own it; viewers cannot edit, and each user's settings are visible to that user only. The rules are kept in one file in source control, deployed automatically when a change is merged, and compared with the deployed copy every six hours to catch drift.

What stays on the device. The Firebase client in every app keeps cloud data in memory only, not in the browser's on-disk database; Firebase keeps sign-in state in the browser (IndexedDB). Switching back to local-only use is always possible (Privacy Policy §7.4), but it does not delete data already saved to the cloud.

Deletion. The Privacy Policy (§6) says deletion is done by the operator on request, "within a commercially reasonable period, not to exceed thirty (30) days under normal circumstances." There is no self-service delete-account control. Invitations expire 30 days after they are sent (§2.6).

4. TIER 3: CONNECT AI

Connect AI lets a user connect an AI assistant they already use, such as [AI PLATFORM], to one open project, using the Model Context Protocol (MCP). None of the apps has AI built in, and Connect AI is off unless the user starts it. Only three apps offer it:

App	Write Mode (AI edits the project)	Read Mode (AI reads the project)
SPERT® Scheduler	Offered; required to connect	Offered; optional
SPERT® Story Map	Offered; required to connect	Offered; optional
SPERT® Forecaster	Not offered: read-only by design	Offered; the only permission asked
SPERT® CFD, GanttApp™, MyScrumBudget™	Not offered	Not offered

Source: AI Privacy Notice §2.3, which matches the code reviewed.

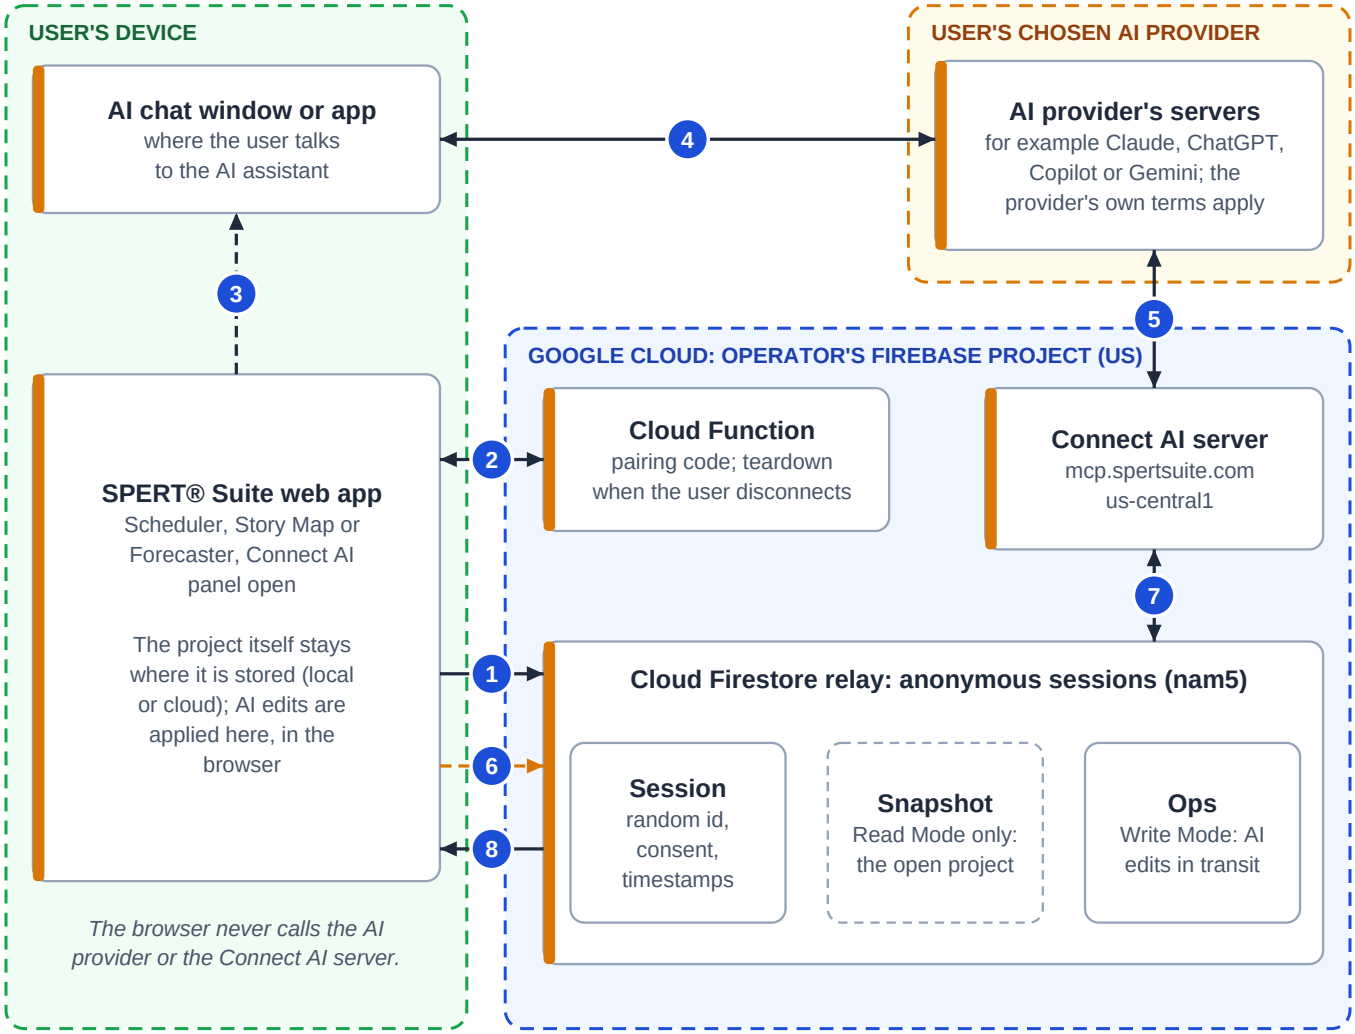

Figure 4. Tier 3 data flows. The orange dashed flow happens only when the user turns Read Mode on.

#	From and to	What travels
1	Browser to the relay	A new session record: a random identifier "with no connection to your name, email, or account" (AI Privacy Notice §4), plus consent and activity timestamps
2	Browser and a Cloud Function	A request for a pairing code (one word plus 10 digits, valid 15 minutes, closed about a minute after the AI first uses it); at the end, a request to tear the session down
3	The user, from the app to the AI chat	The user copies the pairing code into their AI assistant
4	AI chat and the AI provider	Everything the user types, plus the AI's replies. This goes to the AI provider whatever SPERT® Suite's settings are
5	AI provider's servers and the Connect AI server	MCP tool calls over HTTPS to `https://mcp.spertsuite.com`, a Google Cloud Function (second generation, running on Cloud Run) in `us-central1`; and the tool results, which include the project snapshot when Read Mode is on
6	Browser to the relay (Read Mode only)	A snapshot of the one open project, refreshed after edits. The AI Privacy Notice (§2.2) is explicit: "Read Mode uploads a copy of your open project regardless of your storage setting."

#	From and to	What travels
7	Connect AI server and the relay	The server reads the snapshot (Read Mode) and writes the AI's edits as operations (Write Mode)
8	Relay to browser	The AI's edits, which the browser applies to the open project. SPERT® Scheduler first checks each one against bounds

What the AI provider receives is outside SPERT® Suite's control. The Privacy Policy (§4.5) says the provider's "own terms of service and privacy policy govern their handling of content transmitted to them." The AI provider is chosen by the user and is not one of the operator's service providers.

What the AI can change. Each session is bound to one app; the server refuses tool calls meant for a different app. The counts below are from the Connect AI server's source (spert-landing, `functions/src/mcp/tools/`).

App	Tools	Can overwrite existing content	Can delete
SPERT® Scheduler	25 (1 read, 24 write)	Yes: names, estimates, descriptions (an empty value clears one), milestones, dependency settings, checklist state, activity order	Dependency links only. No tool deletes an activity, milestone, scenario or project
SPERT® Story Map	21 (1 read, 20 write)	Yes: names, descriptions, categories, notes (an empty value clears one), sizes, release allocations, location in the map	No entity deletes. Tools can remove release allocations
SPERT® Forecaster	3 (all read)	No	No

Built-in limits. Browsers cannot write AI operations into the relay, and cannot read a snapshot back out of it (Firestore rules), so only the Connect AI server can send edits and only that server can read a snapshot. The server rejects batched requests and applies a request-rate limit. SPERT® Scheduler's browser code re-checks every edit against bounds before applying it, and the server refuses out-of-range estimates.

How long relay data is kept. Per the AI Privacy Notice (§2.2, §4): relay data is deleted when the user disconnects, and the snapshot when Read Mode is turned off; otherwise it expires seven days after the session's last activity, including requests from the AI assistant. Backups may keep copies for up to 98 days. Signing out ends the session on the server in all three apps (since SPERT® Scheduler 0.79.25, SPERT® Story Map 0.53.13 and SPERT® Forecaster 0.47.3); if the server cannot be reached at that moment, for example when offline, the app tells the user and keeps retrying. Disconnecting also ends the session on the server, but if the server cannot be reached, SPERT® Scheduler and SPERT® Story Map do not tell the user. Until a later attempt succeeds, the relay data stays until the session expires seven days after the session's last activity, including requests from the AI assistant; because each AI request extends the session, one that an AI keeps using does not expire. The AI provider's own retention applies to anything it received.

How access works, plainly stated. The Connect AI server has no API key or user login. Access rests on the single-use pairing code, then on the session's random identifier, plus a rate limit. This "capability token" design means anyone holding a live session identifier could act on that session. The pairing code was strengthened on 2026-10-25 (from 4 to 10 digits, single use). Stronger protection of session identifiers and abuse limits on unauthenticated relay writes are planned in the code but not yet built.

For SPERT® Scheduler, the kit's *Connect AI Data Explainer* lists exactly which fields a Read Mode snapshot contains.

5. TRUST BOUNDARIES

Each dashed region in the figures is a trust boundary: a place where data passes to a party with its own controls.

Boundary	What crosses it	How it is protected	Who controls the far side
User's device and Vercel	Page requests; app code	HTTPS	Vercel, under its own terms
User's device and the sign-in provider	Sign-in	The provider's own sign-in, including [ORGANIZATION]'s MFA or conditional access for work accounts	Google or Microsoft (for work accounts, [ORGANIZATION]'s tenant settings)
User's device and the operator's Google Cloud project	Cloud projects, profile, invitations (tier 2); relay records and snapshots (tier 3)	HTTPS; Firebase Authentication; Firestore security rules	The operator, on Google Cloud
Operator's Google Cloud and Resend	Invitation details	HTTPS; called only by server code	Resend, under its own terms
AI provider and the Connect AI server	MCP tool calls and results, including snapshots in Read Mode	HTTPS; pairing code, session identifier, rate limit	The AI provider the user chose
User's device and the AI provider	Everything typed into the AI chat	The AI provider's account controls	The AI provider, under the user's account terms
User's device and Nager.Date (SPERT® Scheduler only)	Country code and year	HTTPS; responses checked against a schema	The holiday API's operator

6. DATA INVENTORY

Data	Tier	Where it is held	How long	Source
Projects and settings (local)	1	The user's browser	Until the user deletes them or clears site data	Privacy Policy §1
Request data (IP address, browser type)	All	Vercel	Under Vercel's policies	Privacy Policy §1
Name, email, profile-photo link	2	Firebase Authentication; shared profile in Firestore	While the account is active; deleted on request within 30 days (Firebase Authentication may keep sign-in data up to 180 days)	Privacy Policy §2.1, §6
Projects and settings (cloud)	2	Firestore, United States (`nam5`)	As above; backups up to 98 days	Privacy Policy §2.2, §6, §9
Terms acceptance record	2	Firestore	As above	Privacy Policy §2.1
Invitation records	2	Firestore; the email goes through Resend	Expire after 30 days; records kept until the operator purges them	Privacy Policy §2.6
Technical data (IP address, browser, timestamps)	2, 3	Google Firebase; Google Cloud Logging	Logs 30 days, location not restricted	Privacy Policy §2.3

Data	Tier	Where it is held	How long	Source
Connect AI session record and AI edits in transit	3	Firestore relay	Deleted on disconnect, or 7 days after the session's last activity (AI requests count as activity); backups up to 98 days	AI Privacy Notice §2.1, §4
Read Mode snapshot of the open project	3	Firestore relay, then the AI provider	Relay: deleted when Read Mode is turned off or as above. Provider: its own terms	AI Privacy Notice §2.2
Contact and training forms on `spertsuite.com`	Website only	Formspree, then the operator's email	As long as needed to respond	Privacy Policy §2.7

Not collected (Privacy Policy §2.4): payment data, government identifiers, phone numbers, addresses, job titles or organizational affiliation, and any data from local-only use other than standard web-hosting request data.

Not allowed. The Terms of Service (§5.3) and Privacy Policy (§2.2) prohibit storing third-party personal information, protected health information, payment card data, government identification numbers or other regulated data. [ORGANIZATION]'s data classification for planning data should be checked against this: [DATA CLASSIFICATION APPROVED FOR USE].

7. SERVICE PROVIDERS

Provider	Role	Receives	Tier	Source
Google LLC (Firebase, Google Cloud)	Sign-in, database, server functions, Connect AI server, logs	Account, cloud project, relay and technical data	2, 3	Privacy Policy §2.3, §4.1
Vercel Inc.	Hosts the apps' web pages	Standard request data	All	Privacy Policy §1, §4.6
Resend	Sends invitation emails	Invitation details	2	Privacy Policy §2.6, §4.6
Microsoft (Exchange Online)	Hosts the operator's spertsuite.com mailbox	Email sent to spertsuite.com addresses; form messages	Website and email only	Privacy Policy §2.7, §4.6
Formspree	Delivers website form messages	Name, email, message	Website only	Privacy Policy §2.7, §4.6
Stripe	Payments, "if subscription payments are implemented"	Nothing today: there is no paid tier	n/a	Privacy Policy §4.2
The user's AI provider	Not a service provider of the operator; chosen by the user	Prompts; Read Mode snapshots	3	Privacy Policy §4.5

Google Cloud and Vercel publish their own security and compliance information. Those reports cover their infrastructure, not SPERT® Suite's code, and this document does not restate them: [REVIEWING TEAM] should read the current versions directly. The operator holds no certification of its own.

8. SECURITY CONTROLS

Control	How it works	Tier
Local-first design	Project data stays in the browser unless the user turns on cloud storage or Connect AI Read Mode	1

Control	How it works	Tier
No passwords	Sign-in only through Google or Microsoft; [ORGANIZATION]'s own MFA and conditional access apply to work accounts	2
Encryption in transit	HTTPS for all traffic (Privacy Policy §5)	All
Encryption at rest	The Privacy Policy does not make a separate statement; data at rest relies on Google Cloud's infrastructure controls (§5). No customer-managed keys	2, 3
Access control	Firestore security rules: members only, by role; owner-only deletion. One rules file, deployed from source control and drift-checked every six hours	2, 3
Connect AI consent	Opt-in each time; separate consent per app and per browser; Read Mode is a separate choice	3
Connect AI isolation	Session bound to one app; only the server writes AI edits; browsers cannot read snapshots back; SPERT® Scheduler re-checks edits in the browser	3
Code transparency	Public source code (GPL-3.0)	All
Supply chain and code scanning	Dependabot alerts, CodeQL code scanning, secret scanning with push protection, and protected main branches on all public repositories (GitHub settings, since 2026-10-25)	All
Vulnerability reporting	Email security@spertsuite.com, or use GitHub private vulnerability reporting on each public SPERT® repository. The Security page (https://spertsuite.com/security) says the operator aims to acknowledge reports within 3 business days, and publishes a short advisory after fixing an issue that could have affected users' data. Machine-readable contacts are at `https://spertsuite.com/.well-known/security.txt` and `https://scheduler.spertsuite.com/.well-known/security.txt`	All

9. WHAT THE ARCHITECTURE DOES NOT PROVIDE

— **No enterprise administration.** There is no admin console, single-sign-on enforcement, user provisioning (SCIM), organization-wide export or tenant separation. All users' cloud data shares one database, separated by rules.

— **No switch to turn Connect AI off for an organization.** Consent is per user, per app and per browser. Three apps never offer Connect AI, and SPERT® Forecaster's connection is read-only.

— **No audit log** of user or administrator actions.

— **United States only** for cloud data at rest; operational logs are not region-restricted.

— **No SLA.** The service is provided as-is by a single operator, and the Terms of Service (§9) cap liability.

— **No independent audit, certification or penetration test.** Security reviews to date are the developer's own.

Tier 1 is the compensating control: in local-only use, none of these applies to project data.

REVIEWER NOTES

— **"It's open source, so nobody has reviewed it."** Open source lets [REVIEWING TEAM] review it, but no independent assessment has been done. The developer runs self-reviews: a suite-wide self-review was completed on 2026-10-25, and its fixes are being released in stages. SPERT® Scheduler has had the most recent fixes, which is why the kit's *Security Questionnaire Answer Sheet* covers that app only.

— **"Data leaves our environment."** In tier 1 it does not, apart from page requests and Scheduler's optional holiday lookup. In tier 2, data goes to Google Cloud in the United States. In tier 3 with Read Mode, a copy of one

project goes to the user's AI provider; [ORGANIZATION] can limit this by allowing only an approved [AI PLATFORM] account, or by not approving tier 3.

- **"Can we block Connect AI?"** Not inside the apps. If [AI PLATFORM] lets administrators restrict which connectors users can add, that is the effective control. Blocking the Connect AI server (`mcp.spertsuite.com`) on [ORGANIZATION]'s network does not work, because the AI provider's servers call it, not the user's device. Starting Connect AI does need the user's browser to call SPERT® Suite's pairing function at `us-central1-spert-suite.cloudfunctions.net/generatePairingCode`, so blocking that path, or that whole host (which also stops sharing invitations), may stop Connect AI from starting. This is untested, so [ORGANIZATION]'s IT team should test any block it relies on.
- **"Where are the audit logs?"** There are none for users or administrators. Google Cloud keeps operational request logs for 30 days for the operator's use only.
- **"We've never heard of this vendor."** Start small: a tier 1 pilot limits continuity risk, because data stays on the device and can be exported at any time. The GPL also permits self-hosting, although no guide exists and the trademark terms require a renamed version.
- **Scope and limits of this review.** Checked against each app's source at the versions in section 1 and against the shared backend's source. Not verified: live security headers, export formats, Google Cloud's and Vercel's current certifications, and deployed logging settings.