

PRIVATE FAMILY COLLABORATION AND ANCESTRY

CollabKin

One private space where a whole family builds one living record — the tree, the photographs, the stories and the timeline.


FAMILY TREE

MEMORY ARCHIVE

ORAL HISTORY

ROLES AND CONSENT

AI THAT ASKS PERMISSION



Pre-launch. Launching 2026, waitlist open. Capabilities marked planned or in progress are named as such throughout.

collabkin.com

01 The short version

CollabKin is a private home for your family's history — a family tree, a photo and document archive, the stories your elders carry, and a shared timeline, all built by the whole family in a space only the family can see.

THE PROMISE

Nothing is public, nothing is sold, and nothing is one person's private file. The record your family builds is meant to be inherited, not orphaned.

Why it exists

Most families keep their history in half a dozen places at once. A genealogy subscription that concentrates on the dead. A phone full of unsorted photographs. A group chat where yesterday scrolls away. A relative's hard drive of scans. And the part that exists nowhere at all — what one grandparent still remembers and nobody has written down. None of those places is somewhere a family can work together, and none of them is private.

CollabKin puts all of it in one place. You create a Family Tree — the product's word for a private family space — and add the relatives and ancestors in it, the ties between them, the photographs and documents that belong to each person, the stories that were only ever spoken aloud, and the dated moments that make up the family's history. Then you invite the family in. The cousin who has the albums, the aunt who knows the maiden names, the grandparent who can name everyone in the photograph: each of them adds what only they hold, to the same record.

Two decisions shape everything else. The first is that the living and the ancestral share one record, so the cousin you message this week and the great-grandparent you trace back to sit in the same tree with no re-keying between a social app and a genealogy tool. The second is that privacy is built into the data rather than promised in a policy: every relative holds a role, every photograph and story carries its own visibility, and anything filed about a living relative is tightened automatically before anyone has to think about it.

01 In four numbers

1

One family record

The relatives you talk to and the ancestors you trace share a single tree. There is no second system to keep in step.

4

Roles inside a tree

Owner, Steward, Contributor and Viewer decide what each relative may add, change and see.

3

Visibility tiers per item

Every memory and every story is scoped to the whole tree, to stewards only, or kept private.

Free

The plan you can start on

The Free plan is the one that can be provisioned today. Paid tiers are written and priced but not yet on sale.

The family organiser

The middle-generation relative who already collects the photographs and keeps everyone in touch. CollabKin gives that work somewhere to live, and turns it from one person's burden into something the whole family carries.

Keepers of memories

Anyone determined to rescue the box of fading prints, the brittle certificates and the stories an elder still carries, while there is still time to do it.

Family historians and genealogy enthusiasts

People tracing the line back who want a collaborative alternative to a private file on one laptop — with relationship certainty, approximate dates and standard import and export so the research is honest and portable.

Privacy-conscious families

Families who will not put photographs of their children, the details of living relatives or heirloom documents onto a public, advertising-funded network — and want the controls written into the product, not the terms.

Families spread across countries

Relatives in different cities and time zones who need somewhere meaningful to stay connected. The interface is multilingual and works on a phone, so distance and language are not the barrier.

Family associations and heritage organisations

Clan and lineage societies, family offices and archives that need governed, auditable collaboration over a large record, with dedicated tenancy and self-hosting available at the top tier.

Large and multi-branch families

Households where the father's line, the mother's line and a branch a marriage brought in are genuinely separate records. Each can be its own tree inside one private space.

None of these problems is exotic. They are the ordinary way family history is lost — slowly, through scattering and silence rather than through any single disaster.

The history lives in one person's head

One relative becomes the unofficial family archivist. When they stop — through illness, distance or simply time — the project stops with them, and whatever was never written down goes with it.

- The record is a shared space many relatives contribute to at once, each under their own role, with every change attributable to the person who made it.

The photographs are in a box, and the box is fragile

Prints curl, ink fades, a house move loses half of them, and nobody left alive can name the faces in the ones that survive.

- Photographs, scans, recordings and heirloom artifacts become memories held in durable storage, dated to when the moment actually happened and tagged with the people in the frame.

The stories exist only while the teller does

How the grandparents met, why the family left the village, the reason a dish appears at every celebration: none of it is written anywhere, and one funeral ends it.

- Stories are captured as written lore or recorded audio and video with a transcript and a named narrator, attached to the relatives and moments they concern.

Genealogy tools are built for solitary research

Trees are exchanged as files, corrections are emailed around, and the family ends up with several versions of itself that disagree.

- One tree, edited by many hands, with relationship certainty from confirmed through disputed so the record shows what is known and what is merely likely.

The alternative is a public feed

Putting a family's most intimate material — living relatives' details, photographs of children, heirloom documents — onto an advertising-funded network is a decision most families regret only later.

- A members-only space with no public profiles, no advertising and no sale of family data. Access begins with an invitation and ends when the family removes it.

Nobody knows who is allowed to see what

Sensitive material gets shared with everyone or with nobody, because there is no middle setting — and the awkward compromise is usually silence.

- Each memory and story carries its own visibility, and anything filed about a living relative is narrowed automatically before it is ever shared too widely.

Old dates are treated as certainties

A guessed year gets copied into someone else's tree, hardens into family truth, and sends the next researcher down a dead end.

- Dates can be marked approximate so imprecise history reads as approximate, and every relationship carries a certainty flag that travels with the claim.

The record dies with the account

A subscription lapses, a service shuts down, a login is lost, and years of a family's work becomes unreachable — which is exactly the fate it was meant to escape.

- Standard-format export is available on every plan including Free, over-quota records go read-only rather than being deleted, and the family's own copy is always retrievable.

The product is built out of a small number of pieces that fit together. Learn these and you know what CollabKin is.

SPACE

Family Tree

The private family space itself, with a surname or house name, a cover image and its own tags. A family can keep several — one per branch — inside one account, and retire a tree without destroying it.

PEOPLE

Relatives and ancestors

Given, family and maiden names, nickname, gender, life status, birth and death dates and places, a biography and a portrait. A person can be linked to the relative's own account so the record and the member are one identity.

TIES

Relationships with certainty

Directed, typed links — parent, child, spouse, partner, sibling, guardian, grandparent, grandchild — each carrying a certainty flag of confirmed, probable, unverified or disputed, plus optional start and end dates.

ARCHIVE

Memories

Photographs, scans, audio, video and artifacts attached to the person and the moment they belong to. A memory records when the moment happened, not when it was uploaded, so a 1962 wedding photograph sits in 1962.

VOICE

Stories and oral history

Written family lore, or recorded audio and video with a transcript and a named narrator. Stories carry their own visibility and attach to the relatives and moments they are about.

TIME

The family timeline

Births, deaths, marriages, divorces, immigration, military service, education, residences and milestones, per person or family-wide, with approximate dates rendered honestly rather than rounded into false precision.

ROSTER

Membership and consent

The tree's guest list. Every relative is invited, then active, declined or removed, with a role attached and a record of who invited whom. A membership can be linked to the person it represents.

DISCOVERY

Relationship paths and lineage


Ask how two relatives are connected and get the computed path through the recorded relationships, plus ancestor and descendant views grouped by generation.

OUTPUT

Portability and keepsakes

Standard GEDCOM import and export, with a control that excludes living relatives' sensitive fields, and generated keepsake books that compile stories, memories and timeline events into a downloadable PDF.

Fourteen situations a family actually finds itself in, and what CollabKin does about each one. The roles named are the ones the product ships.




FAMILY ORGANISER

Start the tree, and let the family finish it

The bind. Somebody has to begin, and beginning alone is exactly why most family trees stall at forty names and a spreadsheet.

You create a family space, add yourself and the nearest relatives as people, and link the ties between them. Everything you enter is immediately something the rest of the family can build on rather than a private file you would have to share later.

- The first evening's work becomes the family's foundation instead of one person's hobby.



FAMILY ORGANISER

Invite relatives with a role that actually fits them

The bind. Handing a distant cousin full edit rights feels reckless; giving a devoted aunt read-only access feels insulting.

Each relative is invited with one of four roles. Owners run the space, Stewards curate and manage the roster, Contributors add without being able to delete anyone else's work, and Viewers browse. Every invitation moves through invited, then active, declined or removed, with a record of who invited whom.

- Everyone gets in at the level that suits them, and nobody has to be trusted with more than they asked for.




FAMILY HISTORIAN

Trace the line back with the doubt left in

The bind. Serious genealogy is full of maybes, and most tools force a maybe to be entered as a fact.

Ancestors go in generation by generation with whatever is known — a year, a village, a relationship. Each tie is marked confirmed, probable, unverified or disputed, and imprecise dates are flagged approximate so the interface shows them as approximate rather than exact.

- A lineage the next researcher can trust, because it says out loud which parts are still guesses.




FAMILY HISTORIAN

Ground each claim in the family's own documents

The bind. A date is a rumour until there is a certificate behind it, and the certificate is usually in a different drawer from the tree.

Birth and marriage certificates, scanned records and research notes are filed as memories against the person they concern, so the evidence sits with the claim. Anyone reviewing the tree can see why a date is what it is, and the certainty flag records how far the evidence goes.

- The difference between we think and we know stays visible instead of dissolving over time.



FAMILY HISTORIAN

See the whole family laid out in time

The bind. A list of dates hides the shape of a family; you cannot see the decade everyone scattered or the year three cousins were born.

Births, marriages, migrations, military service, education, residences and milestones take their places on one chronology, either against a person or family-wide. The timeline assembles itself from what the family has already entered rather than needing separate curation.

- The family's story becomes something you can scroll through and show a child, not a table you have to interpret.




KEEPER OF MEMORIES

Record an elder before the stories go quiet

The bind. Everyone means to sit down with a grandparent and record them. Almost nobody does it in time, and the reason is usually not knowing what to ask.

A library of interview prompts, organised by life stage and relationship, gives you the questions. You record the answer as audio or video, or write it down, name the narrator, add a transcript, and attach the story to the relatives and moments it is about.

- What lived in one memory becomes part of the permanent record, in the teller's own words.




KEEPER OF MEMORIES

**Keep the tender records
for the right eyes only**

The bind. A final letter, a fragile certificate, a photograph of a child: too important to lose, too private to put in front of everybody.

Every memory and story carries its own visibility — the whole tree, stewards only, or private — and the setting is applied inside the query that fetches it, so a restricted item is simply absent from another relative's results rather than hidden behind an interface control.

- The precious things are preserved against loss without being exposed to the whole family.




FAMILY ORGANISER

**Turn a gathering into
part of the record**

The bind. After a reunion the photographs live on a dozen phones, and the complete set of the day exists nowhere.

Everyone who was there adds their photographs to the same family space, tags the relatives who appear, and the gathering takes its place on the timeline as a dated, located event. Relatives who could not attend see the day as it happened.

- The event becomes a chapter in the family's history rather than a fading remember-when.




FAMILY MEMBER

Keep a directory that maintains itself

The bind. Somebody's number changes, a cousin moves, a baby arrives, and the family address list is wrong within a month of being written.

Each relative keeps their own profile current — contact details, important dates, a recent photograph — and because their membership can be linked to their person record in the tree, the living directory and the family tree are the same data rather than two lists to reconcile.

- Birthdays stop being missed and addresses stop going stale, without anyone maintaining a spreadsheet.




FAMILY MEMBER

Stay close in the gaps between milestones

The bind. In a group chat the photograph a cousin just scanned is gone by morning; on a public network it competes with strangers and advertising.

Comment threads sit on the memory or story itself. You name a face, correct a date, or add the part only you remember, and the conversation stays attached to the thing it is about — inside the family space, with no audience and no ranking algorithm.

- A static archive turns into a living one, and the family talks more rather than less.



FAMILY ADMINISTRATOR

Decide who sees what, deliberately

The bind. Family records are not uniformly sensitive, and treating them as if they were means either over-sharing or over-locking everything.

Roles govern what a relative may do; per-item visibility governs what they may see; and anything filed about a living relative is narrowed to stewards-only by default, or to private when the recorded birth date makes that person a child. Defaults only ever tighten.

- Privacy becomes a property of the record rather than a setting somebody discovers too late.




FAMILY ADMINISTRATOR

Keep several branches without forcing them into one

The bind. A large family is not one neat tree, and flattening the father's line, the mother's line and an in-married branch into a single chart produces a tangle nobody wants to maintain.

Each branch is its own tree inside the same private space, with its own people, memories, stories and timeline, cared for by the relatives who know it best. You move between them from one home, and none of it spills into a public space.

- The extended family stays organised the way it actually is — many branches, one private home.




FAMILY HISTORIAN

Answer how am I related to her, properly

The bind. Every reunion produces the question and nobody can answer it; working it out by hand means counting generations and arguing about what second cousin means.

Pick two people and the connecting path is computed by walking the relationships the family recorded, alongside ancestor and descendant views grouped by generation. The traversal is ordinary computation, so it is repeatable and cannot be invented; AI is only ever asked to put the result into plain words.

- A real answer, with the shaky links in the chain visible rather than smoothed over.



TECHNICAL RELATIVE

Point your own AI agent at the family record

The bind. Researchers already have an assistant open, and the only ways to connect it are pasting screenshots or handing a script an API key nobody will maintain.


An agent connects over the Model Context Protocol and gets the family domain as named tools — trees, relatives, relationships, invitations, memories, stories, timeline events and comments, alongside files, tags and roles. It signs in as you, and every action is re-checked against your role on that tree.

- Agent work lands in the archive as ordinary, attributable, metered changes with no side channel.

06 Built AI-native

AI that can only ever see the part of the family record you are allowed to see

Most claims about AI and privacy are a paragraph in a policy. That is not adequate for a family archive, because the sensitive material here is not a payment token — it is a living relative's birth date, a photograph of a child, and a story an aunt asked to be kept between stewards. CollabKin built the permission boundary before the intelligence, so AI arrives as another caller subject to rules that already existed rather than as a system that has to be persuaded to behave.




The two rules

The mechanism is deliberately unglamorous. Every memory and story carries its own visibility, and that tier is combined with your role on the tree inside the database query that fetches the record. What you may not see is absent from the result rather than hidden on screen, which means a cleverly worded prompt has nothing to work around. The assistant reaches your record through the same door your own screens use, signed in as you: a Viewer never receives a Steward's answer, and there is no separate index of your family that quietly ignores the settings.


The second rule is that AI proposes and a relative decides. Naming a face in a photograph or asserting a relationship is a claim about a real person, very often a living one, and an archive that makes such claims on a model's word becomes untrustworthy in a way that is hard to repair. So suggestions arrive as things to accept, adjust or reject, and an accepted suggestion is saved as an ordinary change attributed to the human who approved it. Where the machinery genuinely is not built yet, this document says so rather than describing the plan in the present tense.

06 Four pillars



Permission comes first, always

Your role on the tree and each item's own visibility decide what AI can read. The assistant, your own agent and your own screens obey the same single check — there is no second, looser path.




A real record, not a pile of files

Relatives, relationships, memories, stories and dates are properly linked, so an answer can be computed from the record rather than guessed from prose. Relationship paths are walked, not generated.



Suggestion-only, by design

Captions, interview questions and person links arrive as proposals. A relative accepts, and the change is saved under their name with the ordinary audit trail behind it.



Built on shared foundations

The assistant, the agent tools, the roles and the record of changes are platform capabilities the whole product line shares. CollabKin adds the family meaning on top rather than maintaining a private AI stack.

An assistant on every screen

A floating panel you can open anywhere in the app, so a question never costs you your place in the tree. Answers stream as they are produced.

Shipped. A question needing several lookups shows progress instead of a spinner, and a hiccup part-way through is retried once.

Visibility-scoped reading

Whole tree, stewards only or private, weighed against your role, applied inside the query rather than after the fact.

Shipped. This is the actual AI boundary — a restricted item is absent from the result, not filtered out of the display.

Every request runs as you

The assistant reaches the family record through the same gateway your own screens use, signed in as you. What you can see, it can see; what you cannot, it cannot.

Shipped. Nothing the model is told can widen your permissions, because the check happens before any record is returned.

Stricter defaults for the living

File a memory or story about a relative recorded as living without saying who may see it and it narrows to stewards only; to private when the recorded birth date makes them a child.

Shipped. Defaults only ever tighten, which matters most for agents and integrations — exactly the callers that forget to set a visibility.

Credits checked before the question runs

The AI credit balance is checked before a question starts and reported afterwards, so spending is visible rather than discovered on an invoice.

Shipped. Running out is explained as an out-of-credits message rather than an opaque failure.

Family tools for the agent you already use

Thirty-three named tools cover trees, relatives, relationships, invitations, memories, stories, timeline events and comments, alongside files, tags and roles.

Shipped. Your agent connects over the Model Context Protocol and every action is re-checked against your role on that tree.

Audited writes with details stripped

Every change — yours, the assistant's or an agent's — records who made it and what they did, while birth dates, birthplaces, maiden names and biographies are kept out of that trail.

Shipped. You get a trail you can hold someone to without it becoming a second, less-protected copy of the family's details.

Photo caption suggestions

Ask the assistant for a caption on a memory and it proposes text from the details already recorded against it.

Shipped, and narrow on purpose — it drafts words, it does not identify anybody.

Follow-up questions for an interview

From the prompt library and what a relative has already said, the product suggests the next questions worth asking.

Shipped. The blank page and not knowing what to ask are the two reasons family stories never get recorded.

Relationship paths computed, not invented

The connection between two people is found by walking the recorded relationships; ancestor and descendant views are grouped by generation.

Shipped. Deterministic where it must be — AI's only job is turning the path into a sentence and naming the weakest link.

A suggestion queue you decide on

Proposals are listed with a status, so a relative can work through them and accept or reject each one rather than finding the record changed underneath them.

Shipped. It is the mechanism that makes suggestion-only real instead of aspirational.

Portability an agent can drive

Standard-format family-tree import and export are ordinary operations, with a control that excludes living relatives' sensitive fields on the way out.

Shipped. Bulk work does not need a private back door, because the public operations already do it.

Knowing which tree you are looking at

The assistant knows which screen you are on but not yet which tree or relative is in front of you, so add a memory for him does not resolve today.

In progress. It is the smallest change with the largest effect on how useful the assistant feels.

Fluency in the product's own vocabulary

Teaching the assistant trees, relatives, certainty flags, memories, stories and timeline events, plus the privacy rules it must never route around.

Roadmap. Until it lands, treat it as a capable general assistant that happens to live inside CollabKin.

Hardened isolation between families

One family's records are sealed off from every other family's and the seal fails shut — with no family context established, a request sees nothing rather than everything.

In progress. A final hardening step, a dedicated low-privilege runtime identity, is still outstanding.

Confirmation before an agent deletes

Removing a relative, a memory or a tree currently takes no confirmation when an agent asks for it.

Roadmap, and the reason we would not yet recommend giving an unattended agent write access to a tree you care about.

Photo intelligence over the archive

Planned help with the tedious part of a scanned shoebox: this face recurs across fourteen photographs, this group shot looks late-1970s, these three scans are the same print.

Roadmap. Grouping would stay inside one family's space and never become an index that spans families.

Life-story drafts from your own records

Planned: a first draft of a relative's life drawn from the timeline, photographs and what other relatives already wrote, saying where the record is silent instead of filling the gap.

Roadmap. A draft would stay a draft until a relative edits and saves it under their own name.

Permission before reasoning

Your role on the tree and each item's visibility decide what exists for a question before a single word is generated. Filtering happens in the retrieval, not in the answer.

One gate, no second path

The same check applies whether you are browsing, asking the assistant, or driving your own agent. There is no privileged AI route and no separate copy of the family to query.

Defaults that only tighten

A memory or story about a living relative defaults to stewards only, and to private for a child. Nothing in the system widens a record's reach as a side effect of being organised or linked.

Guardians can demand a correction or removal

A verified guardian of a child, or a steward acting for one, can formally file a correction or removal request that an owner or steward must resolve, and the decision is recorded.

Proposals, never verdicts

Identifying a person or asserting a relationship is a claim about a real human being. AI proposes and a relative decides; the accepted change is saved under the approver's name.

AI changes are ordinary changes

An agent adding a relative passes the same checks as the dialog in the app. Deletion stays with owners and stewards, removals are recoverable, and everything is metered against the family's quotas.

Identifying fields kept out of the trail

The change record captures who did what, deliberately excluding birth dates, birthplaces, maiden names and biographies so the audit trail does not become a second copy of the family.

Spending is checked before it happens

AI credits are verified before a question runs and reported after it completes, so a family cannot quietly run up a bill by leaving an agent working.

No training on your family record

Family records are not used to train models and family data is not sold. The product is funded by the families that pay for it.

No cross-family index

Any future face grouping stays inside a single family space. There is no index of faces or people spanning families, and building one is not on the roadmap.

No pinned model, no accuracy claims

CollabKin pins no model of its own; model choice sits with the shared platform assistant. Bring-your-own-key is the committed direction rather than a shipped feature, and no benchmarks or accuracy figures are published.

The gap we will not paper over

There is no AI-specific read ledger and no evaluation set yet, so which answer used my mother's record is not answerable today. AI activity is covered by the general change trail, not by AI-specific logs.

Reachable from, and in numbers

REACHABLE FROM

- The floating assistant panel, on every screen of the web application
- Memory pages, where caption suggestions are offered against an uploaded photograph
- Story and interview flows, where follow-up questions are suggested from the prompt library
- The suggestion queue, where proposals wait for a relative to accept or reject them
- Your own AI client, connected over the Model Context Protocol
- The command-line client, for scripted and bulk work against the family record
- The Node.js SDK, for integrations and automation
- The Python SDK, for archival and research workflows
- The in-app AI section — a placeholder today, and described as one rather than dressed up

IN NUMBERS

- Every query and mutation in the product carries an access-control directive that is enforced at the gateway before the request reaches the family record.
- Three visibility tiers — whole tree, stewards only, private — are pushed into the database query itself rather than applied to the results afterwards.
- A memory or story filed about a living adult defaults to stewards only; about a relative whose recorded birth date makes them a child, to private.
- Thirty-three family tools are exposed to external AI agents, alongside the shared file, tag and role surfaces.
- Deleting a relative, a memory or a story is reserved for a tree's Owner and Stewards, whether a person or an agent is asking.
- Removals are soft: the record is recoverable and the removal is recorded against the actor who made it.
- Birth dates, birthplaces, maiden names and biographies are deliberately excluded from the stored change metadata.
- Relationship paths are computed by traversing recorded relationships; the model is asked only to describe the result and name the weakest link.
- The product pins no AI model and publishes no accuracy benchmarks, having no results it is willing to stand behind.

06 What is not true yet

What is not true yet, stated plainly. The assistant does not know CollabKin's own vocabulary, so a question like which relatives have no photograph is answered by a capable general assistant rather than a family historian, and it does not yet know which tree or relative is on screen. Photo intelligence — face grouping, era hints, duplicate spotting — does not exist; memories are labelled by hand today. Story drafting does not exist; there is no generation endpoint. An agent's destructive actions take no confirmation step, which is why unattended agent write access to a tree you care about is not something we would recommend. There is no AI-specific read ledger, no citations on answers and no privacy evaluation set. Bring-your-own-key is a committed direction, not a shipped feature.

The product ships six family roles plus a plan manager, and each one is a real permission boundary rather than a label. These are the roles seeded into every family space.

Family Owner

That the record survives, that the right people are in it, and that nothing is deleted by accident.

Owns the family space end to end, including who is in it and whether a branch is retired. Two acts are irreversible and reserved here: retiring a tree, and removing a relative from the roster.

- Create and retire family trees
- Set roles at invite time
- Resolve guardian correction requests
- Invite and remove relatives
- Delete people, ties, memories and stories
- Manage tags across the whole space

Family Steward

Keeping the record accurate and the roster current without needing the owner for routine work.

The trusted relative who administers the space day to day, curates the record and manages invitations on the owner's behalf. Matches the owner everywhere except retiring a tree.

- Curate people, ties, memories and stories
- Create and rename branches
- See stewards-only memories and stories
- Invite and remove relatives
- Resolve guardian requests
- Approve or reject AI suggestions

Family Historian

That the tree is right, that certainty is recorded honestly, and that a wrong date can be taken down.

Builds and corrects the lineage — relatives, ancestors, the ties between them and the dated events. Reads the archive but does not curate it, which is what separates this role from the memory keeper.

- Add ancestors generation by generation
- Remove wrong or disputed ties
- Import and export the tree
- Set and correct relationship certainty
- Add and remove timeline events
- Trace relationship paths and lineage

Memory Keeper

That the box gets scanned, the elders get recorded, and the archive stays organised.

Curates the family archive — photographs, scans, recordings and stories — and dates the moments they hang on. The mirror image of the historian, with read-only access to the lineage itself.

- Upload photographs, scans and recordings
- Set per-item visibility
- Tag memories and stories
- Record and transcribe stories
- Add and date timeline moments
- Prune duplicate uploads

Contributing Relative

Being able to add the part only they hold, without fear of breaking anything.

The ordinary cousin, aunt or parent in the shared record. Adds what they know without being able to remove anyone else's contribution — which is exactly what many hands, one record requires.

- Add relatives and ties
- Write and record stories
- Comment on memories and stories
- Upload photographs and documents
- Add milestones to the timeline
- Tag people and memories

Viewing Relative

Being included and being able to browse without any risk of altering something.

The distant relative or elder invited to look rather than edit. Deliberately holds nothing whose grant would let them change the record or invite anyone.

- Browse the tree and its people
- Read stories and transcripts
- Read comment threads
- View the archive and the timeline
- See who is in the space
- Receive family notifications

Family Plan Manager

Predictable cost, enough storage and members for the family's size, and no surprise over-quota behaviour.

Holds the billing relationship for the family account — the plan, the quotas and the add-on capacity — separately from the family record itself, so paying for the space is not the same as running it.

- Choose and change the plan
- Review usage against quotas
- Set the audit retention window
- Buy storage, people and member packs
- Manage invoices and billing contacts
- Request enterprise terms

Under the hood

CollabKin is API-first. Product state is served over a federated GraphQL API organised into two domains: the kin side holds family trees, people, relationships and memberships; the memories side holds memories, stories, timeline events and comments. Every type in the schema is namespaced so it cannot collide with any other product sharing the same composed graph, and every client — the web application, the command-line client, the SDKs and the agent tool server — speaks that same API.

Authentication and authorisation happen at the API gateway, never inside the product service. The gateway validates identity, reads the access-control directive attached to each query and mutation from the composed schema, checks it against the caller's actor and scope, and only then forwards the request. The service trusts the signed identity headers it receives and never re-implements authentication, which is why the same rules apply identically to a browser, a script and an AI agent.

Domain state lives in a relational database with a graph extension used for lineage traversal, reached through pooled connections and covered by versioned schema migrations. Media bytes are never stored in the domain database: photographs, scans and recordings live in the shared file service and are served through signed, expiring links, while CollabKin holds the family meaning around them. Events flow over a durable stream, so notifications, activity and usage metering happen without coupling services together.

THE SPLIT THAT MATTERS

The domain owns context, the file service owns bytes, and the gateway owns access. Because those three responsibilities are genuinely separated, a photograph can be re-scoped, an actor's permissions can change, and the family's record can be exported — without any of the three having to reach into the others.

Family graph

Directed, typed relationships across eight kinds with a four-value certainty flag and optional start and end dates. Living, deceased and unknown life status is a first-class field rather than a convention, because privacy behaviour depends on it.

Media handling

Five memory kinds — photo, video, audio, document and artifact — with the moment's true date and place recorded separately from upload time, and bytes held in shared storage behind signed links.

Lineage traversal

Shortest-path computation between any two people in a tree, plus ancestor and descendant views grouped by generation, executed over the recorded relationships rather than inferred from text.

Access control

An access-control directive on every query and mutation enforced at the gateway, layered with in-tree role checks and per-item visibility pushed into the database query for both single reads and lists.

Consent and guardianship

A full membership lifecycle, automatic tightening of defaults for living relatives and children, and formal guardian correction and removal requests with a recorded resolution.

Publishing

Keepsake books compiled from stories, memories and timeline events into a PDF rendered on demand, generated under the same visibility scope as any other read, so a book cannot contain what a screen would not show.

Portability

Standard GEDCOM 5.5.1 import and export, with an explicit control to exclude living relatives' sensitive fields on the way out, plus full data and media export on every plan.

Metering

Trees, people, members, memories, stories, storage and AI operations are metered per family and enforced at the action layer, which is what makes one price per family possible instead of per seat.

08 Standards and integrations

OPEN STANDARDS IT SPEAKS

- GraphQL, with a federated schema composed across services
- Apollo Federation v2 for cross-service entity resolution
- GEDCOM 5.5.1 for family-tree import and export
- OpenID Connect for single sign-on
- OAuth 2.0, including the device-code flow for terminal clients
- Model Context Protocol for AI agent tooling
- Outbound webhooks for event delivery into your own systems
- PDF for generated keepsake books and printable output
- Accessibility semantics, keyboard navigation and a command palette throughout
- Locale-aware dates, numbers and right-to-left layout

PLUGS INTO

- Single sign-on and identity from the shared platform identity layer
- Shared file storage with signed, expiring links for every photograph, scan and recording
- Workspace-wide tagging shared across trees, people, memories and stories
- Comment threads on memories and stories via the shared conversations service
- Notifications by in-app message, email and webhook, with per-relative preferences
- Billing, plans, regional pricing bands, add-on capacity and usage quotas
- Activity and audit streams that are searchable and exportable
- Localisation across the interface so relatives read it in their own language
- Guided tours and contextual onboarding for less technical relatives
- Feature flags for gradual rollout of new capabilities per family space

How it deploys

Shared cloud

The default. Multi-tenant with strict per-family isolation enforced in the data layer, and no setup for the family beyond signing in.

Dedicated tenancy

Available on the top tier for family associations, societies and archives that require their own isolated environment rather than a logical boundary.

Self-hosting

Offered at the top tier for institutions with a hard requirement to run the service inside their own estate, under a negotiated agreement.

Data residency

Production data is hosted in India, in regions operated by providers holding their own physical-security certifications. Residency elsewhere is handled through enterprise terms.

Recovery objectives

Automated backups with a rolling hot tier of thirty days and a cold tier up to ninety. Enterprise agreements target a one-hour recovery point and an eight-hour recovery time; the standard service targets twenty-four hours for both.

Mobile

iOS and Android builds exist and run today but are not yet published to the app stores, so mobile access is currently through the web application.

09 Security & trust

1 / 3

CollabKin holds a family's most intimate record, including material about children. The controls below are the ones the product and the platform beneath it actually implement.

Members-only by construction

There are no public profiles and no public trees. Every operation in the product is scoped to a family space and gated by an access-control directive; the only publicly reachable field is a health check.

Per-item visibility applied in the query

Whole tree, stewards only and private are pushed into the database query for both single reads and lists, and the same scope governs generated keepsake books. A restricted item is absent from the result rather than hidden in the interface.

Authorisation enforced at the edge

Identity is validated at the API gateway and forwarded as signed actor headers. The product service never validates tokens and never re-implements authentication, so a bug in the product cannot become an authentication bypass.

Stricter defaults for the living and for children

A memory or story attached to a relative recorded as living defaults to stewards only; where the recorded birth date indicates a minor, it defaults to private. Defaults only ever tighten.

Four in-tree roles, enforced on writes

Owner, Steward, Contributor and Viewer are checked against the specific tree an operation touches — not merely against the account — for creating, updating, deleting and reading the timeline.

Guardian correction and removal requests

A verified guardian, or a steward acting for one, can formally file a correction or removal request against a person, memory or story. Only an owner or steward may resolve it, and an approved person-removal soft-deletes the child's record.

Consent lifecycle on every membership

Invited, active, declined and removed, with an auditable record of who invited whom. Only the invited relative can accept their own invitation, and only from the invited state.

Tenant isolation that fails closed

Family data is logically segregated at the data layer. With no family context established a request returns nothing rather than everything, so a misconfigured integration gets zero records instead of somebody else's history.

Encryption in transit and at rest

Modern TLS with strict transport security on all public endpoints, AES-256 at rest for databases, cache and object storage, and application-layer secrets and signing keys encrypted separately with managed rotation.

Media held apart from the record

Photographs, scans and recordings live in shared file storage behind signed, expiring links. The domain database holds family context, never the bytes, so re-scoping an item does not require moving the media.

Least privilege for operators

Single sign-on with multi-factor authentication required for all administrative access, just-in-time elevation for production, break-glass accounts with full audit, quarterly access reviews and no shared accounts.

Audit with identifying fields redacted

Sensitive actions — record changes, media access and visibility changes — are auditable per plan retention, while birth dates, birthplaces, maiden names and biographies are deliberately excluded from the stored metadata.

09 Security & trust

3 / 3

Secure development lifecycle

Code review on every change, static analysis, dependency and secret scanning in the build, container image scanning for known vulnerabilities and misconfigurations, parameterised database access, and query depth and cost limits on the API.

Vulnerability management and disclosure

Continuous scanning of images and dependencies with a published patch cadence — critical within seven days, high within thirty, medium within ninety — plus a coordinated disclosure programme with a published reward scale.

Stated plainly

CollabKin holds no security certifications today. Practices are modelled on the ISO 27001 control families and designed against SOC 2 criteria, and the product is built to align with GDPR and India's data-protection regime, but no SOC 2 report and no ISO 27001 certificate exists — those are part of the path to general availability, not a claim about today. Annual third-party penetration testing is a target rather than a completed programme. Client-side or zero-knowledge encryption is planned and not built. Affirmative guardian-consent capture for living children remains an open product and legal decision; what exists today is the automatic tightening of defaults plus the guardian correction and removal mechanism. There is no AI-specific read ledger. CollabKin has no DNA product and no health-records product, and does not intend to build either. Data is hosted in India; residency elsewhere is handled through enterprise terms.

One product, many front doors

The web application is the complete product today. The other surfaces exist at different stages, and each is described at the stage it is genuinely at.

Web application

The full product in any browser — the tree, people, relationships, the roster, memories, stories, the timeline, keepsake books and the conversation around them. Available today.

Workspace experience

The same CollabKin experience composed into the shared application shell alongside the rest of the product line, so a family already using another product does not manage a second login. Available today.

Marketing site

Positioning, feature detail, use-case stories, honest comparisons against the incumbents and the launch waitlist, at collabkin.com. Available today.

Command-line client

A terminal client for bulk work — importing a tree from a legacy genealogy tool, exporting trees and memories, and running scripted audits against the family record. Published.

Agent tool server

A Model Context Protocol server exposing thirty-three family tools plus files, tags and roles, so an AI client signs in as you and works under your roles. Published.

Node.js SDK

A typed, promise-based SDK for integrations and automation in any Node.js or Bun project, with access control and audit applying exactly as they do in the app. Published.

Python SDK

A fully asynchronous Python SDK for archivists, genealogists and data engineers working with trees, sources and large memory archives. Published.

Documentation site

Product guides, API reference and integration tutorials. In progress and coming online ahead of launch.

Mobile applications

iOS and Android builds are complete and running, for camera-first capture, recording an elder and checking the timeline from anywhere. Not yet submitted to the app stores.

Browser extension

Planned. Clip ancestry records, obituaries, articles and scanned documents straight into the right person or timeline event, with the source citation kept intact.

Photo and genealogy integrations

Planned. Read from the cloud photo libraries a family already keeps, and cite records from the major genealogy platforms without leaving CollabKin.

Community hub

Planned. Questions, discussion and family-history practice, signed in against your existing identity rather than yet another account.

11 Plans

CollabKin is priced per family rather than per seat, because charging for each relative you invite is a penalty on exactly the behaviour the product exists to encourage. Prices below are the Band A list; four regional bands apply.

FREE

\$0 forever

A family starting one tree and rescuing a first batch of photographs. No card required.

- One family tree, 50 people, 5 invited members
- 200 memories and 25 stories, 2 GB of media storage
- Every core capability unlocked — tree, timeline, roles, visibility
- Video memories and recorded stories included, not held back
- Full data, media and standard-format export
- Community support

PRO

\$9/mo, or \$7/mo billed annually

An active family historian building out the tree and the archive properly.

- 3 trees, 1,000 people, 25 invited members
- Unlimited memories and stories, 50 GB of storage
- Uploads up to 2 GB per file
- 30-day record-change history
- Standard-format import and export, and keepsake books
- Email support

BUSINESS (FAMILY)

\$29/mo, or \$24/mo billed annually

A large, multi-branch extended family with many contributors. Fourteen-day trial.

- Unlimited trees and people, up to 150 invited members
- 500 GB of storage, uploads up to 5 GB per file
- One-year record-change history
- Single sign-on in sandbox
- Priority email support
- 99.9% uptime target

ENTERPRISE (HERITAGE)

Custom, from about \$2,400/yr

Family associations, clan and lineage societies, family offices, archives and heritage institutions. Annual contract.

- Custom quotas for members, people and storage
- Single sign-on and SAML
- Dedicated tenancy or self-hosting
- Archival retention commitments and contracted uptime
- Managed onboarding and migration
- Dedicated support

11 Plans, in practice

What can actually be bought today

CollabKin is pre-launch and in waitlist mode. The Free plan is the one that can be provisioned; Pro, Business and Enterprise are fully specified and priced but not yet available for purchase. Treat the paid figures as the published intent for launch rather than a live price list.

Four regional pricing bands

The same plan costs less where incomes are lower. Pro runs \$9 in Band A markets, \$7 in Band B, \$5 in Band C and \$3.50 or the local equivalent in Band D — around ₹399 a month in India, or ₹319 billed annually. Business scales the same way, from \$24 down to about \$10 a month on the annual rate.

Add-ons rather than forced upgrades

Capacity is bought where it is needed: 100 GB of storage at \$4 a month or 1 TB at \$25, a further 1,000 people at \$5, 50 more invited members at \$6, extra trees, extended audit retention and priority support. Service add-ons cover photo restoration credits and assisted oral-history production. Governance capabilities such as production single sign-on, dedicated tenancy and self-hosting are never sold as add-ons below the top tier.

Downgrades never delete anything

Upgrades apply immediately and are prorated; downgrades take effect at the end of the period. If the record then exceeds the lower tier, the over-quota items go read-only rather than being removed, and export stays available on every plan and in every over-quota state.

One graph for the living and the ancestral

The relative you invite, the person in the tree, the face in the photograph and the narrator of the story are one identity. Most tools force a choice between a social app for the living and a genealogy tool for the dead, then leave the family re-keying between them.

Four family roles, not owner and collaborator

Owner, Steward, Contributor and Viewer map to how families actually work: someone runs it, someone curates, most people add, and some just want to look. The historian and memory-keeper role templates further split lineage curation from archive curation, because they are different jobs.

Privacy expressed in the data, not the policy

Roles, per-item visibility, living-versus-ancestor status and guardian rights are schema-level concerns enforced inside the query. That is why the same guarantees hold for a browser, a script and an AI agent without anything being re-implemented three times.

A formal consent lifecycle

Invited, active, declined, removed — with only the invited relative able to accept their own invitation. Joining and leaving a family record is deliberate, which matters most in the branches where relations are strained.

Honesty is a feature of the record

Relationship certainty from confirmed through disputed, and dates that can be marked approximate, mean the tree can distinguish what is known from what is likely. Genealogy that cannot express doubt propagates errors into other families' trees.

Priced per family, in four regional bands

Inviting relatives costs nothing extra because seats are not the unit. Regional bands put the product within reach of diaspora families and markets that dollar-led genealogy platforms price out.

The record is portable by design

Standard-format import and export on every plan including Free, a control to exclude living relatives' sensitive fields on export, and read-only rather than deletion when a family downgrades. There is no version of this where the family's history is held hostage.

Governance inherited, not improvised

Identity, roles, billing, files, conversations, notifications, audit, localisation and accessibility come from shared platform services used across a whole product line, so a small family product ships with governance a standalone application would take years to build.

AI that could not widen access if it tried

The permission boundary predates the intelligence and sits in the retrieval path. Nothing an assistant or agent is told can widen it, and AI is never the party that decides a claim about a person.

You are pre-launch. Why would I put my family's history somewhere with no track record?

That is a fair reservation and we will not argue it away with adoption numbers we do not have — CollabKin has no users, no revenue and no case studies to cite. What we can offer instead is verifiable: standard-format export on every plan including Free, so the record leaves as easily as it arrives; over-quota data going read-only rather than being deleted; and a Free tier substantial enough to test with real relatives before anyone pays anything.

The incumbents have billions of historical records. You have none.

True, and we do not plan to compete there. Record archives are enormous, expensive assets built over decades. The honest recommendation is to research where the records are and preserve here — standard-format import brings that research in, and export takes it back out. What the archives do not give you is a private space many relatives build together, with roles, consent and a shared timeline.

Is this real-time, like a shared document?

No. Several relatives can work in the same family space and each other's changes appear, but updates are saved and read rather than streamed live to everyone at once. Real-time co-editing is on the roadmap and is not available today. In practice family history is rarely a race against another editor, but we would rather you knew before you asked.

How do I know AI is not reading things I have marked private?

Because the restriction is not a prompt instruction. Visibility and your role on the tree are combined into the database query that fetches records, so a restricted item is absent from the result rather than filtered out afterwards, and the assistant asks as you rather than as a privileged service. The gap we will name rather than hide: there is no AI-specific read ledger yet, so which answer touched which record is not something you can audit today.

What happens to our record if the company goes away or we stop paying?

Export is available on every plan, in every over-quota state, and includes both the structured tree in a standard format and the media. Cancelling moves the account to the Free quotas at the end of the period; whatever exceeds them goes read-only and stays exportable for the retention window in the terms. Nothing is deleted as a consequence of a downgrade.

We have children in the tree. How is that handled?

Anything filed about a relative whose recorded birth date makes them a child defaults to private, and about any living relative to stewards only — defaults only ever tighten. A verified guardian, or a steward acting for one, can formally demand a correction or removal that an owner or steward must resolve. The honest gap: affirmative consent capture for children is an open product and legal decision, so what exists today is strong defaults plus the guardian request path, not a signed consent workflow.

Do you have SOC 2 or ISO 27001?

No. Practices are modelled on the ISO 27001 control families and designed against SOC 2 criteria, and the underlying platform provides encryption, least-privilege operator access, audit and tenant isolation — but no report and no certificate exists today. Formal certification is committed on the path to general availability. If a certificate is a hard procurement requirement, we are not yet the right answer.

Is there a mobile app?

iOS and Android builds are complete and running, but they are not in the app stores yet, so mobile access today is through the web application, which works on a phone. Store submission is manual work still outstanding. Given that family memory-keeping is camera-first, we consider this a real gap rather than a detail.

Are you going to sell our family data or show us advertisements?

No, and the business model is the reason to believe it. The product is funded by the families who pay for it, which is also why paid tiers matter to us. Family records are not used to train models, there are no public profiles, and there is no advertising surface anywhere in the product.

Can we let an AI agent work on the tree unattended?

We would not recommend it yet. An agent connects as you and is re-checked against your roles on every action, and removals are recoverable and recorded — but destructive tools take no confirmation or dry-run argument today, so a confused agent can delete as fluently as it can create. The confirmation gate is roadmap work, and until it lands, supervise agent write access to any tree you care about.

VS CONSUMER GENEALOGY SUBSCRIPTIONS

They charge per researcher for access to enormous historical-record archives, and relatives end up as passive viewers of one person's tree. CollabKin charges per family and makes every relative a contributor with a role. The honest concession is large: their record archives and DNA networks are assets we do not have and are not building, and for pure record research they remain the better tool.

VS PUBLIC WORLD-TREE PROJECTS

Free, vast and collaborative — and public, with deceased profiles editable or visible to anyone and merge disputes between strangers a routine hazard. CollabKin is members-only with per-item visibility, so the family controls its own record. The concession: they are free at a scale we cannot match, and for connecting into the wider global tree they are genuinely the right place to work.

VS PROMPTED-MEMOIR SERVICES

They are excellent at getting one person to write a book, and the printed hardcover they produce is better than what we generate today. CollabKin captures the whole family — tree, timeline, photographs, audio, video and stories — with interview prompts and a downloadable keepsake book. The concession: print fulfilment is still on our roadmap, so for a physical gift they win.

VS CONSUMER DNA TESTING

DNA tells you where you came from; it tells you nothing about how the family got here. CollabKin deliberately never collects genetic data — no DNA product, no health-records product, no plan for either — so the two pair rather than compete. The concession is straightforward: if genetic ancestry or relative matching is what you want, we do not offer it at all.

VS CLOUD PHOTO DRIVES AND GROUP CHATS

Families already have both, and neither knows that a photograph belongs to a person, that a person has a mother, or that a date matters. CollabKin holds the relationships and the context around the media. The concession: those tools are free, already installed and where everyone's photographs currently live, so the first upload is genuine work.

VS GENERAL-PURPOSE COLLABORATION SUITES

A shared drive plus a spreadsheet plus a chat group can hold a family tree, and costs nothing extra if the family already pays for one. What it cannot do is enforce that a memory about a living child is private, compute how two relatives are connected, or export a standards-compliant family tree. The concession: for a small family with forty relatives, the free improvisation may honestly be enough.

START TODAY

Start the record while there is still someone to ask

CollabKin is in waitlist mode ahead of a 2026 launch. The Free plan is real, the export is real, and everything still on the roadmap is named as such.

- Join the waitlist and get the family space when access opens
- Start one tree with the relatives you can name from memory
- Invite three relatives with the role that fits each of them
- Record one elder this month, using the prompts, before anything else

WAITLIST AND PRODUCT

collabkin.com

HOW FAMILIES USE IT

collabkin.com/use-cases

SECURITY AND PRIVACY POSTURE

collabkin.com/security

THE APPLICATION

app.collabkin.com

TALK TO US

contact@collabkin.com