



OniX Agents — Your Personal Intelligence Assistant to Advance Your Science, Privately.

Personal. Private. No one is watching.

Life sciences intelligence — competitive, patent, regulatory, and asset — every fact sourced, every search kept completely private.

Your Personal Agent to assess and de-risk your ideas and assets — privately, and grounded in real evidence, not a guess.

White Paper — July 2026
onixhub.com · martin@onixhub.com

Table of Contents

Executive Summary	3
The Problem	3
The Approach	3
Five Pillars	4
Why This Matters, By Role	5
What's Next	6
Appendix	7
A1. Methodology	7
A2. Agent-by-Agent Detail	7
A3. Pricing	8
A4. Roadmap	9
A5. Get Involved	9
A6. How Teams Use OniX Agents	9
In the End, It's Personal	11

Executive Summary

The Problem

The real challenge isn't just staying private — it's knowing enough to move forward with confidence. Every idea, every asset, is trying to find its way to patients: is the science sound, has someone already tried this and failed, is the space already crowded, is there a regulatory landmine three steps ahead nobody's flagged yet? Answering that well takes access — to funding history, prior art, regulatory precedent, and what's already been tried — and most researchers don't have that access, or the time to dig for it themselves.

That's the gap OniX Agents close: real scientific signal — what's ahead, what's already happened, what to watch for — so you can move your work toward patients with fewer surprises along the way, and get there faster because you're not retracing ground someone else already covered.

None of that works, though, if seeing the landscape means exposing your own position while you look. It's not a hypothetical risk: reporting on how quickly some overseas biotech programs have moved a promising public lead into their own clinical trials has made that head-start problem impossible to ignore. You shouldn't have to choose between staying informed and staying private.

Here's the good news: with OniX Agents, you don't have to choose. **Search everywhere and see everything, while staying completely unseen.** No surprises, no exposure — you search, and the results are private, for your eyes only.

Cost has been the other wall. The enterprise platforms built to unify this kind of intelligence run in the mid-five to six figures a year — built for teams with a dedicated analyst, not a single researcher or a small team chasing down one question this week. Even narrower, single-purpose tools typically still start in the thousands per year. And even where the budget exists, access itself is scattered across a dozen separate logins, each with its own search syntax and its own bill, none of them talking to each other.

The Approach

OnixHub is building a suite of **Personal Agents** — each one working only for you, each one focused on a specific intelligence domain, each one running entirely on your own laptop. No query ever leaves your machine except the searches sent directly to free public databases (NIH, PubMed, FDA, USPTO, and similar) and, optionally, to a local AI model running on your own hardware. Nobody else — not us, not a cloud vendor — sees what you're looking into. Personal by purpose; local by design.

Every Agent is also **signed and notarized by Apple** — verified, legitimate software you can trust the moment you open it, not something you have to take on faith.

Four Agents are live today:

- **OniX Competitive Analysis Agent** — one search across 20 sources, giving you a wide slice of the global research landscape: grants, publications, and clinical trials, spanning multiple countries in a single run.
- **OniX Patent Agent** — worldwide patent search spanning roughly 100 patent offices — the large majority of the world's filed patents in one search — plus a dedicated US-only mode.
- **OniX Regulatory Agent** — starting with deep, real FDA and CMS data (US regulatory intelligence first, international coverage on the roadmap), including animal-health and veterinary coverage most human-drug-only tools skip entirely.
- **Asset Intelligence Agent** — a staged view through target biology, structural chemistry, competitive landscape, regulatory signal, and patent landscape — gate by gate — so you can make an informed decision on what to do next with a specific asset, plus a draft investor deck.

Two more — a Funding Agent and a VC & Industry Agent — are in active development.

Across the whole suite, what that looks like in practice: 10M+ grants globally, 460M+ publications, 160M+ patents across 100+ countries, multi-agency regulatory intelligence covering 15+ specialized watchlists, and drug-discovery databases with 500M+ biochemical data points — all queried live, from your own machine, not pulled from a cached or stale copy.

Proof in practice: CXCL1. A live, unscripted test of the Competitive Analysis Agent — a chemokine spanning inflammatory disease, oncology, and chronic pain, picked specifically because it's the kind of multi-area topic that trips up shallow keyword search. One run returned a full Executive Summary correctly identifying three concentrated research groups and three distinct disease contexts, plus the Active Researcher Landscape, Innovation, and Target-Disease Opportunity reports generated alongside it — four report types from one search, not four separate searches.

Five Pillars

Advancing real science and de-risking what's next both depend on all five of these holding true — not just one:

1. **Personal, and private by design.** Nothing you search, upload, or generate ever leaves your machine.
2. **Grounded in real sources.** Every report is built from real records — a real grant, a real patent filing, a real trial registration, a real published abstract — not from what an AI model remembers or assumes. When an Agent draws an inference (like a read on mechanism of action), it's always clearly labeled as an AI interpretation of the record above it, never presented as unverified fact. Every result links straight back to its original source.
3. **Time saved, not just data delivered.** A search that would take hours or days by hand — checking multiple databases, then verifying each claim — comes back in minutes, already sourced and ready

to act on. And because every fact traces back to a real record, you're not spending extra time fact-checking an AI's answer — that verification work is already done.

4. **Real access, at a price that makes sense.** Enterprise-grade insight without an enterprise budget.
5. **One growing suite, not one tool.** Competitive intelligence, patents, regulatory, asset intelligence, funding, and investor/industry research — one company, built to work together.

Proof in practice: getting it right, not just getting a match. PD-1 — one of oncology's best-known targets — initially came back with an empty Target-Disease Opportunity Report, not because the data was missing, but because the underlying database only recognizes PD-1 by its official gene symbol, PDCD1. Rather than loosen the matching logic (which would have reopened a real false-positive risk), a small alias table was added so common non-official names resolve correctly — while two other ambiguous topics tested the same day correctly returned “no target resolved” instead of a wrong guess. That's the standard every report is held to: the right answer, or a plain “we don't have one” — never a confident guess.

Why This Matters, By Role

Whatever your role, the same question sits underneath: what do I need to know to move this forward safely, and who can't know I'm looking?

- **Startup founders and biotech teams** get competitive and patent visibility without a subscription that costs more than a first hire.
- **Technology transfer offices** can screen an invention's competitive and patent landscape before committing legal spend to a filing.
- **Pharma BD and innovation teams** get a fast first-pass diligence read before committing to a deeper (and more expensive) formal study.
- **Animal health teams** get regulatory coverage that most human-drug-focused tools simply don't include.
- **VCs and investors** get a faster way to sanity-check a deal's competitive and IP position pre-term-sheet.
- **Incubators, accelerators, and research parks** can offer real intelligence tooling to portfolio companies without buying enterprise seats for each one.
- **Economic development teams** can help regional life-science clusters assess where a technology actually stands globally.
- **Patent lawyers** get a first-pass landscape read to scope a formal freedom-to-operate opinion more efficiently.

What's Next

All four Agents are available now — each one is already signed and notarized by Apple, so there's nothing extra to do beyond downloading and opening it. What's coming next isn't fixed in stone: reach out directly with what's working, what isn't, and what you wish it also did — that feedback goes straight into the roadmap below, not a queue.

Appendix

A1. Methodology — What “Personal, Local by Design” Actually Means

Every Agent runs as a standalone Mac application. Searches go directly from your machine to the relevant public data source's own API (NIH RePORTER, PubMed, openFDA, USPTO, EPO, and similar) — the same way your browser would query that source directly. Written summaries and narrative sections are generated by a local AI model (via Ollama) running on your own hardware, not a cloud AI service. Nothing is routed through OnixHub's servers, logged, or retained by us — we built it this way deliberately, not as an afterthought.

The one exception, stated precisely rather than left implicit: activating your license, the first time you open each Agent, sends your license key (nothing else — never a search, an uploaded file, or a report) to our licensing system to confirm your subscription and register that Agent to your computer. That's a one-time, required step. After that, each Agent also runs a brief, non-blocking background check of your license status each time you open it — again, your license key only, nothing about what you search or generate — so a cancelled or expired subscription is caught promptly; if you happen to be offline when that check runs, the Agent simply falls back to its own local subscription record and keeps working. This is the only outbound communication anywhere in the product — everything else described in this document (searches, uploaded files, AI-generated narrative) stays on your machine, with no exception.

Where an Agent computes something itself rather than trusting a source's own reported number — the Asset Intelligence Agent's structural similarity scoring, or the Patent Agent's crowded/sparse technology cluster flags — that computation happens locally too, and the underlying method is documented in each Agent's own report output, not just claimed in marketing copy.

A2. Agent-by-Agent Detail

OniX Competitive Analysis Agent — 20 sources: 11 grant sources (NIH RePORTER, NSF, UKRI, USASpending.gov, ARC, CORDIS, NWO, ANR, JSPS/KAKEN [Japan], NHMRC [Australia], and a proprietary Foundation/Unfunded source), 5 publication sources (PubMed, OpenAlex, Crossref, Europe PMC, bioRxiv/medRxiv), 1 clinical trials source (ClinicalTrials.gov), plus Grants.gov, an OniX newsletter feed, and Google News shown as bonus context in every report. OpenAlex, JSPS/KAKEN, and NHMRC each require a free, optional, self-serve API key; every other source needs nothing to set up. Supports saved, recurring daily/weekly/monthly searches that run in the background. Three report types: the full multi-section report; the **Active Researcher Landscape Report**, which cross-references grants and publications over a chosen period (last 12 months, 2 years, 5 years, or all time) to surface who's actually active in a space right now — ranked by grant activity first, then publication output, with a companion Top Organizations rollup by institution; and the **Standalone One-Pager**, a single-page, share-ready distillation of a report's key numbers (funding activity, publication volume, modality breakdown, top players) for quick BD sharing, generated on its own or alongside the full report.

OniX Patent Agent — dual-source: Espacenet/EPO Open Patent Services for worldwide coverage (roughly 100 patent offices) and the USPTO Open Data Portal for US-only searches, each requiring the user's own free API credentials. Computes a deterministic crowded-vs-sparse read on CPC/IPC technology clusters, adds AI-judged per-record relevance scoring, and builds an asset-type/mechanism-of-action distribution chart from fetched abstracts.

OniX Regulatory Agent — openFDA (approvals, adverse events, recalls, labeling), FDA Orphan Drug Designation, CMS's NADAC pricing benchmark, and a ClinicalTrials.gov cross-reference, plus automatic flags for DEA-controlled substances, NRC-regulated radiopharmaceuticals, and USDA/EPA/APHIS-regulated veterinary or GMO-derived material — including FDA's Center for Veterinary Medicine. Intake is a guided, chat-style conversation rather than a rigid form.

Asset Intelligence Agent — pulls target expression (Human Protein Atlas), bioactivity/safety data (PubChem, ChEMBL), natural-product structures (COCONUT), and literature (PubMed) directly, then folds in condensed, asset-specific signal from the Competitive Analysis Agent, Patent Agent, and Regulatory Agent when those sections are selected. Structural analog comparisons are re-scored locally with RDKit (Morgan fingerprints, Tanimoto similarity, maximum common substructure) on one consistent scale, including a computed read on whether an asset's chemical neighborhood is crowded or comparatively unexplored. Includes an optional draft investor pitch-deck built from the same underlying data. Two additional report types, selectable alongside the full report: the **Structural Pivot & Fast-Follower Report**, which groups this asset's closest structural analogs by the kind of change involved (a halogen swap versus a core scaffold change, computed locally with RDKit), showing real ChEMBL compound-level safety flags only where they actually exist; and the **Safety Window Flag Report**, a qualitative cross-reference of this target's tissue-expression profile against its own bioactivity data — built deliberately as a flag to investigate, never a calculated numeric “safety margin,” since the two data types aren't on directly comparable scales.

OniX Funding Agent (*planned*) — Grants.gov (26 federal agencies, independent agencies, and commissions), the European Commission, UK Research and Innovation, Australian Government Research Grants Services, the Japan Society for the Promotion of Science, and private foundation funding.

OniX VC & Industry Agent (*planned*) — focused on pre-seed, seed, and Series A life sciences funding activity, tracked from the public activity of a curated starting list of venture funds and corporate venture arms. A starting list, expected to grow once the Agent is built.

A3. Pricing

Pricing is still being finalized across the board (TBA) — see onixhub.com for the current status of each Agent. The Asset Intelligence Agent is Call for a Quote rather than a fixed price, given how much a report's scope can vary by asset.

A4. Roadmap

This roadmap grows directly from tester feedback — what's below is what's already decided; what comes after it is intentionally shaped by what testers tell us to build next.

- **Live now, signed and notarized:** Competitive Analysis, Patent, Regulatory, Asset Intelligence.
- **In active development:** Funding Agent (federal grants + private foundation funding), VC & Industry Agent (investor and industry-landscape research).
- **Live now: each Agent's own Report Hub**, letting you pick which report type a run produces.
- **Planned (v1.5):** account/license hardening — one license key tied to one customer's machine, plus a simpler renewal flow.
- **Planned (v2):** a Report Archive/Library — one place to browse every report generated across past runs.
- **Planned (v2):** AI-suggested candidate targets and a self-critiquing AI synthesis layer for the Competitive Analysis Agent.
- **Planned (v2):** choice of local AI model — a Settings dropdown for Llama, Qwen, Gemma, and Mistral.
- **Planned (v3):** on-premises and private-cloud deployment options, and Windows support.

A5. Get Involved

Feedback directly shapes what ships next. We've selected a broad spectrum of testers: academic researchers, academic TTOs, research park management and tenants, and open innovation teams at human and animal health companies, among others. Any questions or inquiries, please contact martin@onixhub.com.

A6. How Teams Use OniX Agents

These are illustrative examples, not real customer stories — built from what the Agents actually do (and, for the fifth one, what's planned). Real customer stories will replace or join these as testers give feedback and permission to be named.

Academic Researcher. A PI or postdoc scanning a narrow research area before a grant submission runs the Competitive Analysis Agent on their topic — one search pulls funded projects, recent literature, and related trials, with a summary calling out crowded areas and daylight. Result: a report built in minutes, not an afternoon, that becomes the grant's background section.

University TTO. Staff evaluating an invention disclosure run the Patent Agent against the core technology alongside a Competitive Analysis Agent run on the same science to spot licensing prospects.

Result: a defensible, documented basis for the “should we file this” recommendation — done in-house, before committing legal spend.

Startup / Early-Stage Biotech. A founder prepping for an investor conversation runs the Asset Intelligence Agent on their lead asset. Result: a current landscape view and a draft deck walk into the board meeting, cutting a multi-week, multi-vendor exercise to an afternoon.

Pharma Innovation / Partnering Team. Someone triaging many external opportunities per quarter runs the Regulatory Agent alongside Competitive Analysis and Patent Agents, for a consistent first-pass view across every opportunity.

Startup Founder Raising a Round *(describes the planned VC & Industry Agent — not built yet)*. A founder heading into a raise runs the planned VC & Industry Agent against their asset to see which active funds and corporate venture arms are a realistic fit right now, rather than cold-emailing a generic list.

In the End, It's Personal

Not “local software.” Your own Personal Agent — one that works only for you, remembers what you ask it to, and never has a reason to show anyone else what you're working on. The point was never privacy for its own sake — it's giving you what you need to move real science toward patients, faster and with fewer surprises, without losing the head start you earned getting there. That's the whole idea, stated plainly: **Personal. Private. No one is watching.**

onixhub.com • martin@onixhub.com