Heptabase

The founders, visual learning software, and a company built across Taiwan, the United States, and Britain

Heptabase is a subscription software company founded in 2021 by Alan Chan (詹雨安), Ethan Wu (吳雨川), and Yeefun Lin (林昱帆). Its application lets people work with cards, whiteboards, documents, and other learning material in a spatial environment. Initially known as a visual tool for understanding complex topics, it has expanded into AI-assisted learning, tutoring, and connections with external AI applications. The company originated in a Taiwanese learning and software-making milieu, entered Y Combinator's Winter 2022 batch, and now operates through a geographically distributed organization.[1, 2]

The three founders brought different practices to it. Chan supplied product direction, public writing, and the work of building a business; Wu developed research and interaction-design experience through earlier projects; Lin brought engineering capability and an existing interest in tools for working with information. Their company is best understood through the conjunction of those practices, the people who helped it become viable, and the infrastructure on which its apparent independence rests. The companion Alan Chan profile follows his formation and life beyond the company.

The learning community before the company

Chan's path toward Heptabase passed through a self-learning group he joined in 2016, whose participants and projects changed over time. That history establishes his involvement in a continuing peer-learning setting, not that all three eventual founders belonged throughout.[3] In the January 2022 Y Combinator announcement, he thanked Wu and Lin for their collaboration and trust, and the early investors for advice and introductions.[4]

The distinction between a community and a founding team matters. Membership changed, the initial project involved people who did not remain, and the present trio should not be retroactively inserted into every earlier collaboration. What persisted was a way of meeting: explaining what one was learning, demonstrating things, and finding people willing to investigate unfamiliar subjects. The startup was one outcome of that setting, not its only conceivable purpose.

Chan's December 2021 reflection also describes the problem of distinguishing shared enthusiasm from durable commitment. Early collaborators had different interests and circumstances. He experimented with a demanding cadence of short-term outputs and later abandoned it when it impeded deeper thinking. The account supplies a less polished

picture than the idea of an immediately harmonious founding team: assembling the group required changes in both personnel and working expectations.[5]

Ethan Wu: interfaces, fieldwork, and learning tools

Ethan Wu, also identified as Yu-Chuan Wu, studied electronics engineering at National Chiao Tung University and became Heptabase's product designer. His own portfolio presents a practice centered on interaction design, research, and the translation of complicated activities into usable interfaces. The earlier projects are informative because they reveal work beyond the finished screens of Heptabase.[6, 7]

Caltcha, a 2017 project, addressed scheduling in an online co-learning community. Wu's account distinguishes the needs of an activity's creator, people who must attend, and people who might attend. That distinction changes the interface problem: a calendar is not just a grid of available times, because participants have different relations to the event. He developed wireframes and a prototype and worked on the front end. The project belonged to a small, identifiable group rather than a mass market.[8]

It shows an early concern with the social meaning hidden inside seemingly simple data. An "available" slot means something different when a participant is essential or optional. A tool that ignores the distinction can display correct information and still fail to help people decide. This is a useful precursor to the design problem Heptabase would face: the same piece of information can matter differently according to the activity in which it is being used.

SaveToAnki, developed with Alexander Wang in 2019–2020, connected Google Translate to AnkiDroid. Wu's portfolio describes reducing the manual effort required to make vocabulary cards, including the handling of pronunciation and examples. He worked on design and development. The public evidence documents the workflow and prototype, not a large installed base or demonstrated improvement in language learning.[9]

Here the object of design was the passage between tools. Looking up a word and practicing it later are related activities, but an inconvenient transfer can prevent the second from happening. Reducing that friction does not substitute for learning; it makes a chosen learning practice easier to sustain. The continuity with Heptabase lies in attention to the work surrounding a note, rather than in any claim that the later product descended technically from this extension.

Wu's 2019 Notion research project used contextual inquiry, interviews, observation, and an affinity wall to understand how people worked. The portfolio describes this as coursework and describes research intended to inform Wang's Notablog work. The available summary supports a methodological claim, not a reconstruction of all the research findings: he was learning to investigate activity before deciding what an interface should do.[10]

KnowCourse, developed in 2019–2020 for the NCTU course-selection setting, extended that practice into a larger team. Wu worked on product design and front-end development, with contextual inquiry, user testing, and a style guide among the documented methods. The project drew on the NCTU Plus environment and advice from computer-science academic Yung-Ju Chang. Course selection is an information problem with institutional constraints, uneven knowledge, and consequential choices, not merely a search box.[11]

Taken together, these projects give Wu a substantive place in Heptabase's history. They show repeated attention to educational and collaborative activity, to transitions between tools, and to the difference between a user's task and the data a screen displays. They do not justify attributing every Heptabase interaction to him alone. The company describes a collaborative process, and the portfolio cannot reveal the distribution of authorship within every later feature.

Yeefun Lin: sociology and tools for retaining context

Yeefun Lin, or Yu-Fan Lin, is Heptabase's cofounder and chief technology officer. His Y Combinator biography identifies a sociology background at National Taiwan University and describes him as a self-taught engineer. That combination should not be embellished into a claim that sociology directly caused a particular architecture; the public record does not provide such an account.[12]

One concrete predecessor is Notion Mark Manager. Its public repository describes an extension for collecting and managing highlights and comments from Notion pages, filtering marks, returning to their original context, and exporting them. Mozilla's version history records releases in September 2020 and October 2021. The project therefore predates or overlaps the formation of Heptabase and documents an interest in manipulating knowledge-work material before the company became established.[13, 14]

The useful continuity is specific: a highlighted fragment is valuable partly because it can be retrieved and returned to the place from which it came. Extraction and context are complementary requirements. An engineer building that kind of utility confronts the mundane details that make a knowledge system trustworthy: preserving a reference, collecting the right items, and allowing someone to act on them without manually retracing every step.

Lin's public code also includes smaller utilities and experiments, but the available record is thinner on his personal formation and working life than Chan's. It would be misleading to fill the gap with a generic story about a humanities graduate becoming a programmer. Here his earlier tool and documented leadership role provide a firmer basis than speculation about motives or temperament.

The first product, and the decision to make less

Chan’s founding retrospective describes technical overreach and a changing team. Rich-text editing consumed effort before there was a useful product. In April 2021, Firstory’s chief executive 子軒 introduced him to Sting of 500 Startups and Lawrence of iCook. Their questions pushed him toward an identifiable customer and something that could be tried immediately.[15]

In late May, the team began new code and produced a whiteboard-and-card prototype within a week. Extensive one-to-one onboarding, eventually more than two hundred sessions, helped them prioritize product managers. Lin contacted Chan after reading his vision essays and joined in late July, taking over principal development. The smaller product and a clearer division of work made learning from users more feasible.[16]

By March 2022, an interview described the working division clearly: Chan handled direction, fundraising, and support; Wu developed the detailed design and simulated uses; Lin led development and technical management. The founders shared analytic work. The interview’s example of a card appearing in different contexts, such as reading about retention and planning a business, also captured the product’s central proposition.[17]

This is a distinctive design choice. Copying the text into two documents would create two objects that could diverge. A reusable card instead allows one object to participate in several arrangements. The innovation is not that relationships suddenly exist, but that a user can inspect and revise different contexts without duplicating every piece of underlying material.

Cards, whiteboards, and the labor of understanding

Heptabase’s documented basic elements include cards, whiteboards, tags, and properties. Cards hold material; whiteboards arrange it spatially; tags and properties support other forms of organization. Connections, sections, and nested structures permit more than one scale of reading. The same collection can be approached as individual notes, grouped concepts, or an overview of a larger inquiry.[18]

The consequences become clearer through an ordinary research task. A reader studying a complicated argument may need to keep its definitions, evidence, objections, and unresolved questions visible at once. A linear document can hold all of them, but the reader must repeatedly reconstruct their relation while scrolling. A spatial arrangement can make some of those relations persistent. Its cost is that somebody must decide what deserves separation, proximity, or a link. The board does not escape interpretation; it gives interpretation a manipulable form.

This is also why visual abundance can become a problem. A large board may produce a feeling of comprehension before the relations have been tested. Connections can be vague, categories can harden prematurely, and rearranging cards can become an activity in its

own right. These are analytical risks of the format, not findings that Heptabase users necessarily experience them. Assessing a board requires asking what question it helps answer and what it conceals, not merely how elaborate it looks.

Chan’s own learning examples include comparing customer interviews and studying product analytics, legal matters, and computing history. They show that the board is not restricted to one discipline. They are founder demonstrations of use rather than independent evaluations.[19]

These examples also expose a maintenance burden. Cards need to remain understandable when the immediate occasion for writing them has passed. A reusable note needs enough context to be useful again without becoming so elaborate that making it interrupts the inquiry.

Computing histories and the scale of the ambition

Heptabase’s intellectual references reach beyond recent note-taking software. In a Ness Labs interview, Chan described encountering the history of tools for thought through his own reading and connected the company’s ambitions with earlier work on augmenting human understanding. His roadmap distinguishes individual learning from later possibilities for collective work and programmable environments.[20, 21]

Engelbart’s 1962 report is a useful comparison precisely because it is not simply a proposal for a better editor. It treats the human, language, artifacts, methods, and training as parts of a system. Improving the artifact alone is not the whole undertaking. The report arose in a research setting with its own institutional and funding conditions, not in the subscription-software economy in which Heptabase operates.[22]

This comparison clarifies both the company’s aspiration and a limit. A commercial application can alter the material through which someone studies, but it cannot assume that users will also adopt a complete method, develop shared vocabularies, or receive training. A feature may be available without becoming a practice. The company consequently has to perform some educational work through examples, onboarding, documentation, and support, while selling a tool rather than a fully organized learning institution.

Chan’s 2023 writing on Realtalk and visual end-user programming shows continued attention to Dynamicland’s different approach to computing. It is evidence of an active research interest, not an institutional affiliation or a claim that Heptabase implements Dynamicland. Keeping the distinction prevents “tools for thought” from becoming a label that erases the differences between a communal research environment, a programming medium, and a commercial application.[23]

The published roadmap is similarly best read as an agenda. Its proposed movement toward shared understanding and more programmable elements does not establish that all these

capacities have shipped, or that an open knowledge ecosystem already exists. A roadmap can organize the work without serving as a factual description of the present product.[24]

Taiwan’s founders, mentors, and early capital

The passage into venture-backed software depended on local relationships as well as international institutions. Early investor CJin’s January 2022 essay describes meeting the team, helping with its Y Combinator application, and investing alongside others in the circle that became associated with Tonic Fund. The essay places Heptabase in a Taiwanese network of founders and investors capable of translating an unconventional project into a credible startup opportunity.[25]

That network performed several kinds of work. Introductions reduced the distance to people who had built companies before. Advice made the product’s customer and metrics more explicit. Early investment supplied money, but also a form of endorsement that could travel to the next conversation. None of this means the company was selected only because of connections; it means that the connections helped its work become visible and interpretable to people who controlled resources.

Chan’s June 2022 account reports a total of $1.68 million raised by March: $800,000 from venture funds, $500,000 from an accelerator, $210,000 from angels, and $170,000 through a syndicate. Named backers included HOF Capital, Y Combinator, Kleiner Perkins, Moving Capital, and Tonic Fund. These are historical self-reported figures, not a current capitalization table.[26]

In a separate May 2022 essay, he reported that the founders still held more than eighty percent of the company and that the business was cash-flow positive. The date is essential. Neither ownership nor profitability should be carried forward unchanged into 2026 without fresh evidence. The account does nevertheless describe an early strategy: use outside capital while trying to avoid becoming dependent on spending it to sustain basic operations.[27]

Paid use, retention, and the economics of a small team

The company’s early commercial experiments were unusually close to product learning. Chan’s account of the Y Combinator period describes moving from early supporter payments toward an annual paid product and evaluating use through retention and revenue rather than raw sign-ups. He reports rapid revenue growth and frequent releases during the batch. The figures are the founder’s account, not audited evidence of a causal relationship between a particular pricing choice and success.[28]

Charging changes what feedback means. A person willing to try an interesting free tool and a person willing to pay for an ongoing workflow are not identical signals. Annual payment also brings cash forward, which can support a small team. At the same time it creates an

obligation to deliver continuing value, and retention cannot by itself explain why someone stays. Users may remain because a tool works well, because their material is embedded in it, or because moving would be costly. The available public evidence does not separate those effects quantitatively.

The 2023 version-one announcement reported that the company had grown from three to six people and had been cash-flow positive since its first year. It described more than three hundred releases through the alpha and beta periods, desktop support across major operating systems, and the addition of sync, web, and mobile access. These claims document the company's own account of its development pace and finances.[29]

Maintaining a small team can protect continuity between a user's problem and the person able to change the product. It also places more demands on each person and makes priorities especially consequential. A request for collaboration, an editor fix, mobile improvements, and a new AI capability compete for the same finite attention. The institutional question is not simply whether the team can build quickly, but whether its selection process remains answerable to the people already relying on the software.

Incorporation, infrastructure, and the geography of the firm

Heptabase's institutional geography is more complicated than a headquarters label. Chan's December 2021 incorporation guide describes establishing a Delaware C corporation through Stripe Atlas, obtaining a US banking relationship through Mercury, and using Stripe for payments. His reasons included the practical needs of international payments and investment. The current privacy policy identifies the operator as Hepta Platforms, Inc., a Delaware corporation with a New York mailing address.[30, 31]

The distinctions are important. A legal domicile is not a team location; a mailing address is not evidence of a staffed office; a payment processor is not an investor. An application can be made substantially by people in Taiwan while its contracts, banking, and financing are organized through the United States. The cross-border arrangement is part of the operating machinery, not a biographical footnote.

Nor was the relationship with the United States wholly virtual. An April 2022 essay says Chan had completed his work there after Y Combinator and fundraising and returned to Taiwan. This supports a US stay and return, but does not establish a permanent post-founding relocation to San Francisco.[32]

In 2024 Chan and some colleagues moved to London. His account describes a Global Talent visa process and argues for relocating while the team remained small but had stable revenue. The flexibility of remote work allowed different people to be in different places. These were his assessments of that particular decision, not general immigration or tax advice.[33]

A company-controlled hiring page describes a fully remote team based in Taipei and London and lists benefits associated with Taiwanese employment, including labor and health insurance and pension contributions. The current Y Combinator directory instead labels the company Manchester, United Kingdom. The sources need not describe exactly the same thing or have been updated together. They do not establish a verified Manchester migration for every founder, a UK legal subsidiary, or the disappearance of the Taiwanese operation.[34, 35]

The result is a useful small-company case of distributed institutional dependence: Taiwanese formation and labor, a US corporation and financial services, British residence for some members, and globally supplied software infrastructure. This arrangement offers mobility, but it also multiplies the systems that must continue to work for the product to remain ordinary and dependable to its users.

Recruitment and the cost of growing beyond the founders

Chan's 2025 recruitment essay describes a shift after the six-person period. He emphasizes communication, learning ability, and experience with startup uncertainty, and discusses trial contracts. He also acknowledges misallocated feature work. Read as an employer's statement, it reveals what he wants to select for; it does not establish that employees experience the process as fair, effective, or sustainable.[36]

His concrete proposals include one-to-two-month trial contracts and a preference for people able to handle ambiguity and work closely with others. These practices extend assessment beyond a conventional interview. The public essay does not supply the applicants' or employees' account of that arrangement.[37]

The growth problem also appears in his own retrospective on priorities. More engineering capacity had not made feature selection automatically better. The company's early direct contact with users remained relevant as the work spread beyond the original founders.

The AI turn, 2025–2026

AI changes the product's division of intellectual labor. In the earlier visual-workspace model, the user largely selected, summarized, and connected material. The current documentation describes conversations situated alongside cards and files, models from multiple providers, and agents that can work with the knowledge base. The spatial context can help make the material under discussion explicit, while the model supplies explanations or proposed changes.[38]

Chan's September 2025 learning demonstration uses a difficult textbook rather than only short web material. It provides a concrete reason for the turn: assistance might make sustained engagement with demanding sources feasible for people who would otherwise avoid them. It remains a founder demonstration, not an evaluation against alternative

teaching methods.[39]

The AI Tutor public beta announced in March 2026 extended the ambition toward guided learning and customized courses. The announcement described the founder’ s own use across several subjects. The announcement establishes its intended activity and public introduction.[40]

The move also changes the economics and risks. Model use has a continuing cost, unlike some local editing operations. A provider can change pricing or behavior. A plausible explanation can be wrong, and a generated structure can silently privilege one interpretation of a source. Integrating AI into a knowledge system therefore makes provenance, reversibility, and inspection part of the product problem. A visually persuasive answer is not necessarily a well-supported one.

Users did not uniformly welcome the shift. In an August 2026 community letter, Chan grouped concerns around price, design, and priorities. He discussed options for using users’ own AI subscriptions or accounts, reducing unwanted AI interface presence, and devoting attention to product quality. These were responses and plans expressed publicly.[41] The changelog subsequently records AI Skills on 14 August and mobile AI Agent and Tutor access on 18 August, providing dated evidence of some follow-through.[42]

The disagreement is substantive. A customer who adopted a stable workspace may not want the company to redefine its core around a rapidly changing technology. Another may find the older workspace much more useful once assistance is embedded in it. Heptabase must decide how far the same product can support both without making one group’ s workflow the unavoidable interface for the other.

External AI applications and the November 2026 product

By November 2026, Heptabase was also making its workspace accessible from outside its own main interface. An official update described opening and editing whiteboards from Claude Code, ChatGPT, and Codex through its integration tools. This makes the application partly a persistent environment that another assistant can operate, rather than only the place where a person begins a session.[43]

The changelog supplies a more granular record. Early-November releases covered opening and editing whiteboards inside external AI applications and expanded creation operations, followed by inbox access, space instructions, search improvements, and more precise editing and citation-related capabilities. It also records crashes and import or export fixes. The mixture is instructive: the platform ambition coexists with the ordinary reliability work required by a mature collection of user data.[44]

The MCP documentation describes authorizing an external AI application to access the Heptabase account. Such a connection can make existing material more useful, but it introduces a new boundary. The user must distinguish an assistant answering from selected

context, an assistant searching a collection, and an assistant permitted to change it. Convenience does not make those permissions identical.[45]

This is also a different answer to the early ambition of a programmable knowledge environment. Instead of requiring every user to write a program, an external agent can perform operations through a defined interface. Whether that becomes a durable practice depends on trustworthy results, understandable changes, and the ability to recover when the agent misunderstands the task. The released capability does not by itself settle those questions.

Privacy, portability, and what the company promises

Heptabase's security documentation locates the AWS cloud storage it describes in us-west-1 and says it is encrypted in transit and at rest. This is server-side encryption, not a promise that only the user's devices can decrypt the content. The region statement also does not cover every connected service's processing.[46]

The October 2026 privacy policy says the company does not use customer content to train AI models and prohibits training through its agreements with the AI providers it uses. It also explains that relevant content is transmitted to providers when AI features are used, that search infrastructure holds indexes including embeddings, and that users' own provider accounts are governed by those providers' terms and settings. "No training" therefore should not be translated into "no processing outside the application."[47]

The public subprocessor list further makes visible a supply chain behind the small team. Hosting, payment, communication, support, analytics, and AI functions depend on external services. The appropriate institutional unit is larger than the number of people on the payroll. It includes the contractual and technical relationships that allow those people to deliver a synchronized service.[48]

The terms state that users retain their rights in their content and can continue to view and export data after a subscription ends. That is a significant exit provision. It does not establish perfect portability of every interaction, spatial relationship, or future feature into a different application. A useful assessment must separate ownership of content, availability of an export, and preservation of a working environment.[49]

These distinctions are not reasons to dismiss the product. They are part of taking it seriously as a place where someone might invest years of intellectual work. The company's obligations extend beyond adding capabilities to maintaining access, explaining dependencies, and making departure possible without treating a user's accumulated work as leverage.

The institution that is visible, and the questions still open

Heptabase's institutional history joins a Taiwanese learning milieu, three different founding practices, introductions into startup capital, US incorporation, early paid use, and later geographical mobility. Wu's design research, Lin's information utilities, customers' recurring tasks, and administrative labor all help explain how the application became usable. The public record does not yet supply an audited current financial picture, an employee-level account of work, a complete legal-entity chart, or independent evidence of learning outcomes. The company is still negotiating the expectations created by its AI turn while maintaining the workspace in which existing users have invested their intellectual work.

Source note

Researched through 13 November 2026. This study combines founder essays in Chinese and English, Wu's project portfolio, Lin's public software records, an interview, investor testimony, official documentation, policies, and dated product announcements. Company and founder claims are identified as such. Roadmaps are not treated as released features, announcements are not treated as independent evaluations, and directory locations are not treated as verified migration histories. Public material on Wu and Lin is thinner than Chan's autobiographical archive; the study preserves that asymmetry rather than manufacturing two complete personal biographies.

---

1. Y Combinator, "Heptabase", current company and founder directory, checked 13 November 2026.↩
2. Heptabase, official website, current product positioning, checked 13 November 2026.↩
3. Yu An Chan, "持續的自信", 31 August 2021.↩
4. Yu An Chan, Y Combinator acceptance announcement, 16 January 2022.↩
5. Yu An Chan, "創業筆記(一): 休學創業學到的七件事", 1 December 2021, reflections on the first year of company-building.↩
6. Y Combinator, Ethan Wu founder profile, checked 13 November 2026.↩
7. Ethan Wu, portfolio biography, checked 13 November 2026.↩
8. Ethan Wu, "Caltcha", project record, July–August 2017.↩
9. Ethan Wu, "SaveToAnki", project record, April 2019–July 2020.↩
10. Ethan Wu, Notion research project, August 2019.↩
11. Ethan Wu, "KnowCourse", September 2019–January 2020.↩
12. Y Combinator, Yeefun Lin founder profile, checked 13 November 2026.↩

13. Yeefun Lin, Notion Mark Manager repository, public project description, checked 13 November 2026.↩

14. Mozilla Add-ons, Notion Mark Manager version history, 2020–2021 releases.↩

15. Yu An Chan, “創業筆記(三): 在錄取 Y Combinator 之前”, 28 February 2022.↩

16. Yu An Chan, “創業筆記(三): 在錄取 Y Combinator 之前”, 28 February 2022.↩

17. Sunrise Medium, Heptabase founder interview, March 2022.↩

18. Heptabase, “Fundamental Elements”, official documentation, checked 13 November 2026.↩

19. Heptabase, “How Heptabase’s founder uses Heptabase for learning”, founder-authored workflow account, 2023.↩

20. Yu An Chan, Ness Labs interview, 8 December 2023, author’s republication.↩

21. Yu An Chan, “My Vision: The Roadmap”, 30 April 2023.↩

22. Douglas C. Engelbart, Augmenting Human Intellect: A Conceptual Framework, October 1962, report prepared at Stanford Research Institute.↩

23. Yu An Chan, “Realtalk and Visual End-User Programming”, 26 June 2023.↩

24. Yu An Chan, “My Vision: The Roadmap”, 30 April 2023.↩

25. CJin, Heptabase investment essay, 17 January 2022.↩

26. Yu An Chan, fundraising account, 7 June 2022.↩

27. Yu An Chan, fundraising-winter reflection, May 2022.↩

28. Yu An Chan, growth and metrics account, 8 May 2022.↩

29. Heptabase, “Version One”, 31 August 2023.↩

30. Yu An Chan, US incorporation account, 27 December 2021.↩

31. Heptabase, Privacy Policy, updated 28 October 2026, checked 13 November 2026.↩

32. Yu An Chan, “Ten Years”, 22 April 2022, return to Taiwan after US activities.↩

33. Yu An Chan, “London Relocation: Global Talent Visa”, 30 April 2024.↩

34. Heptabase, company hiring profile, checked 13 November 2026.↩

35. Y Combinator, “Heptabase”, current company and founder directory, checked 13 November 2026.↩

36. Yu An Chan, “1–10”, 31 May 2025.↩

37. Yu An Chan, “1–10”, 31 May 2025.↩

38. Heptabase, “Work with AI”, official documentation, checked 13 November 2026.↩

39. Yu An Chan, “The Best Way to Use AI for Learning”, 29 September 2025.↩

40. Heptabase, “The Ultimate AI Tutor You Wish You Had”, 31 March 2026, public-beta announcement.↩
41. Alan Chan / Heptabase, “Heptabase Community Letter & AMA”, August 2026.↩
42. Heptabase, Changelog, releases through 11 November 2026, checked 13 November 2026.↩
43. Heptabase, update on opening and editing whiteboards in external AI applications, 6 November 2026.↩
44. Heptabase, Changelog, releases through 11 November 2026, checked 13 November 2026.↩
45. Heptabase Support, “Heptabase MCP”, checked 13 November 2026.↩
46. Heptabase Support, “Is my data safe in Heptabase?”, checked 13 November 2026.↩
47. Heptabase, Privacy Policy, updated 28 October 2026, checked 13 November 2026.↩
48. Heptabase, Subprocessors, checked 13 November 2026.↩
49. Heptabase, Terms of Service, updated 30 October 2026, checked 13 November 2026.↩