


Dump your thoughts. See what you actually think.



Nodemeadow turns a messy brain-dump into a living spatial map. Paste notes, half ideas, and meeting scraps; it finds the real topics, connects them with reasons, surfaces the open questions, and collapses any cluster into a clean outline. Every node cites the fragments it came from; nothing is invented. Model-agnostic; prices in USD.

LEGAL NAME

Nodemeadow Inc.

ENTITY TYPE

Delaware C-Corporation, US registered

HEADQUARTERS

750 B Street, Suite 2310
San Diego, CA 92101, USA

PHONE

(619) 555-0142

EMAIL

hello@nodemeadow.com

WEBSITE

<https://nodemeadow.com>

FOUNDED

March 2024

STAGE

Pre-Seed, under \$1,000,000 raised

TEAM SIZE

9 people, full time (range 5 to 25)

INDUSTRIES

AI-native productivity and thinking tools

CUSTOMERS

Product people, founders, writers, students, anyone who thinks in fragments. Worldwide.

PRICING

USD only. Free \$0/mo; Thinker \$9/mo; Studio \$24/mo

SUGGESTED DOMAINS

nodemeadow.com (primary) / nodemeadow.net

Availability checked via DNS, verify at registrar before purchase. The .com is canonical in all copy and metadata; the custom domain is attached to the Vercel deployment manually.

LEADERSHIP

Maren Castellanos

CO-FOUNDER AND CEO

Spent eight years as a product lead at two Southern California software companies, mostly untangling roadmaps written on whiteboards. Owns the product and the care rules.

Theo Lindqvist

CO-FOUNDER AND CTO

Built retrieval and clustering systems for a decade before language models made the hard part tractable. Designed the model-agnostic mapping layer and the citation rules.

PRODUCT

A spatial map built from a brain-dump: named clusters, connections with reasons, tensions, one node of open questions, loose thoughts parked at the edge. Edit, merge, add thoughts and re-map (the map morphs), collapse any cluster to an outline or next-steps list, export markdown and image. Free tier enforced server-side; Thinker and Studio in USD via Stripe.

AI-NATIVE, MODEL-AGNOSTIC

Reading unstructured thought and finding its hidden structure was not possible before modern language models, so the model is the product. A thin HTTP layer routes across providers with failover; every map is validated against a strict schema and a citation rule: no source fragment, no node. Thoughts are private by default and never used to train any model.