

Q3 Roadmap Decision

Three candidate directions for next quarter. This document captures the options, trade-offs, risks, and timeline as structured blocks so an agent can revisit and update each section independently — and so the recommendation can be exported as a prompt for a follow-on review pass.

Options at a glance

◆ A · Docs Platform

Build a hosted publishing target for Noma. Highest revenue ceiling, longest path to value.

◆ B · Research Workflows

Lean into claims/evidence/risk blocks for analyst teams. Narrow ICP, fastest to first paying customer.

◆ C · General Reports

Position Noma as the default format for AI-generated reports across domains. Broadest TAM, weakest wedge.

Decision

DECISION

Start with **Option B — Research Workflows**. Narrowest wedge, fastest signal, keeps the door open to A and C as adjacent expansions.

status proposed

Decision matrix

Dimension	A · Docs	B · Research	C · Reports
Time to first revenue	6–9 mo	6–10 wk	4–6 mo
Wedge sharpness	Medium	High	Low
Existing block fit	Strong	Native	Medium
Defensibility	Network	Workflow	Brand only
18-month revenue cap	High	Medium	High

Claims and evidence

CLAIM

Research and analyst teams are the sharpest wedge for Noma because their existing tools (Word, Notion, Confluence) lack first-class claim/evidence/risk primitives, and they already structure documents this way mentally.

EVIDENCE

Of 11 analyst-team interviews in April, 9 described their current workflow as "copy-paste claims into a doc and hope someone catches stale ones." All 9 said they would pay for a tool that flagged stale evidence automatically.

for [claim-research-wedge](#) · source [user-interviews-apr-2026](#)

CLAIM

A docs platform is the higher revenue ceiling, but time-to-revenue is too long to be the wedge. Better as a follow-on once Noma has format adoption.

EVIDENCE

Comparable docs-platform launches (Mintlify, GitBook) took 12–18 months to reach \$100k ARR; research-tool launches (Mem, Reflect) hit it in 6–9 months with a tighter ICP.

for [claim-docs-too-slow](#) · source [docs-platform-benchmark-2026](#)

Risks

RISK

Research-team ICP is small (~3k orgs globally). Even high conversion caps the business below docs-platform scale. Mitigation: use research wedge to drive format adoption, then expand to docs.

severity medium · **owner** ferax564

RISK

If teams don't keep editing in Noma after first artifact, the workflow value disappears. Mitigation: ship the agent patch protocol in week 3 so updates flow back into source automatically.

severity high · **owner** ferax564

RISK

LLM export quality determines whether agents trust Noma source as canonical. Easy to verify, easy to fix. Tracked separately.

severity low · **owner** ferax564

Timeline

Wk 1 · Format

Parser, AST, frontmatter, JSON export, basic validation.

Wk 2 · Artifact

HTML renderer, default theme, cards/grids/tabs/charts, mobile.

Wk 3 · Agent

LLM export, patch protocol, copy-as-prompt buttons.

Wk 4 · Launch

3 demos, README, spec, comparison page, OSS release.

Open questions

OPEN_QUESTION

Per-seat, per-document, or per-render? Research-team workflows favor per-seat; agent-driven artifacts favor per-render. Decide before week 3.

OPEN_QUESTION

Keep Puppeteer as the report-PDF path or add Typst for longer books? Puppeteer now covers first-class PDFs; Typst may still be worth evaluating for book output.

Agent tasks

agent_task

Re-run the interview tally each month. If `claim-research-wedge` evidence base drops below 8 of 11 supporting interviews (or new interviews contradict), lower the claim's `confidence` attribute and add a `counterevidence` block.

agent_task

Every two weeks, scan `evidence` blocks for `source` attributes older than 60 days. Flag any whose underlying source has changed materially. Do not auto-edit — propose a `replace_block` patch for human approval.

agent_task

On request, package this document's `decision`, top three `claims`, and all `risks` of severity \geq medium into an LLM prompt for a second-opinion review.

Export

Copy decision + risks as a review prompt

Copy summary as Markdown

Copy full document AST

The point of this artifact is not the prose — it's that an agent can re-open it next month, walk the decision/claim/risk graph, and update only the parts that changed. Everything else stays put, and the Git diff stays clean.