

---

# StrategyForge

*The Strategic Brainstorm Studio*

## Complete Manual

*Print Edition*

*Your strategy. Your studio. Your rules.*

A single HTML file. No installation. No cloud. No account. No subscription. No telemetry. Everything stays on your machine unless you choose otherwise.

# Contents

## 1. What It Is

## 2. Quick Start

*Install Ollama*

*Open StrategyForge*

*First Synthesis*

## 3. The Source Vault

*Three Folders*

*Uploading Documents*

*Auto-Assignment*

*Adjusting Weights*

*Indexing and Embeddings*

## 4. Lenses

*What a Lens Is*

*The Five Preset Packs*

*Custom Lenses*

*Adjusting Lens Weights*

## 5. The Founder Preset

*Vision*

*Market*

*Model*

*Traction*

*Risk*

## 6. The Go-to-Market Preset

*Positioning*

*Channels*

*Pricing*

*Launch*

*Retention*

## 7. The Investor Preset

*Thesis*

*Diligence*

*Valuation*

*Portfolio*

*Exit*

## 8. The Operator Preset

*Execution*

*Team*

*Process*

*Metrics*

*Culture*

## 9. The Consultant Preset

*Situation*

*Complication*

*Resolution*

*Evidence*

*Action*

## 10. The Synthesis Engine

*How Synthesis Works*

*Writing Good Prompts*

*Reading the Output*

*Copy and Export*

## 11. Deep Dive

*Choosing a Lens and Focus*

*When to Deep Dive*

*Working with Focus Angles*

## 12. Sidebar and Thoughts

*The Scratchpad*

*Search*

## 13. Notes

*Capturing Observations*

*The Value of What the AI Missed*

## 14. History

*Browsing Past Analyses*

*Exporting History*

## 15. LLM Backends

*Ollama*

*Anthropic*

*WebAI and StudioLM*

*Switching Backends*

## 16. Settings

*Display*

---

*Embedding Model*  
*State Management*

## **17. Workflows**

*Pitch Deck Stress Test*  
*Quarterly Strategy Review*  
*Competitive Response*  
*Investment Memo*  
*Go-to-Market Launch*

## **18. Accuracy and Judgement**

## **19. Privacy and Data**

## **20. Keyboard Shortcuts**

## **21. Troubleshooting**

## **22. Glossary**

---

# 1. What It Is

---

StrategyForge is a brainstorm studio for business strategy. You upload your pitch decks, market research, competitive analyses, and working notes into three weighted folders. You choose a set of analytical lenses — perspectives through which an AI reads your material. Then you ask strategic questions, and the engine synthesises answers that draw on your own documents, weighted by how much you trust each source.

It is not a strategy generator. It does not write your pitch deck for you. It does not produce slide templates or fill in business model canvases. It is a thinking tool — a place where you can ask hard questions about your own strategy and get answers that make you reconsider your assumptions.

Think of it as the advisor who has read everything you have written, remembers all of it, and is willing to tell you what does not add up.

*"The synthesis engine is a thinking partner, not an oracle. Apply your own judgement."*

Everything is a single HTML file. No installation. No cloud. No account. No subscription. No telemetry. Open it in a browser and it works. Your data stays on your machine unless you explicitly choose to use a cloud AI backend.

StrategyForge is the second tool in the Forge series, following StoryForge (a brainstorm studio for fiction writers). It shares the same architecture, the same design philosophy, and the same commitment: your work belongs to you, before, during, and after it passes through the engine.

## 2. Quick Start

You can be running your first synthesis in under five minutes. Here is everything you need.

### Install Ollama

Ollama is a free, open-source tool that runs AI models locally on your machine. Nothing leaves your computer. Install it from [ollama.com](https://ollama.com), then open a terminal and run:

TERMINAL

```
ollama serve
```

Then pull the two models StrategyForge needs:

TERMINAL

```
ollama pull llama3.2  
ollama pull nomic-embed-text
```

The first model (llama3.2, roughly 2 GB) handles analysis. The second (nomic-embed-text, roughly 275 MB) handles document indexing. Both run locally and require no API key, no account, and no internet connection once downloaded.

### Open StrategyForge

Open [strategyforge.html](#) in your browser. Chrome, Firefox, Edge, or Safari will all work. Look at the top-right corner for the connection indicator:

- **Green dot** — Connected to Ollama. Ready to go.
- **Yellow dot** — Checking connection.
- **Red dot** — Not connected. Run `ollama serve` in your terminal.

If you see a green dot, you are ready. If not, check the Troubleshooting section at the end of this manual.

### First Synthesis

Here is the fastest path to your first analysis:

- **Step 1.** Go to the **Sources** tab. Upload a .txt or .md file — a pitch deck draft, a strategy memo, anything you are working on. Click **Index All Sources**.
- **Step 2.** Go to the **Lenses** tab. The Founder preset is loaded by default. These five lenses (Vision, Market, Model, Traction, Risk) will all be applied to your question.

- **Step 3.** Go to the **Synthesize** tab. Type a question: *"What is the weakest part of our go-to-market strategy?"* or *"If our main competitor launches a free tier, how does our positioning hold up?"*
- **Step 4.** Click **Synthesize**. The engine will read your sources through every active lens and stream its analysis in real time.

That is it. You now have a multi-lens analysis of your strategic question, grounded in your own documents. Read it, argue with it, and let it show you what you were not seeing.

## 3. The Source Vault

The Source Vault is where you upload everything the engine should know. Your pitch deck. Your competitor analysis. Your investor notes. Your quarterly OKRs. The more relevant material you feed it, the more grounded — and surprising — its analysis becomes.

But not all sources are created equal. A polished pitch deck draft represents your current thinking more than a rough brainstorm note from three weeks ago. The Source Vault accounts for this through **weighted folders**.

### Three Folders

Every document you upload lands in one of three folders, each with a default trust weight:

FOLDER	DEFAULT WEIGHT	WHAT GOES HERE
Strategy	90%	Core documents — pitch decks, business plans, financial models, OKRs, roadmaps
Intelligence	60%	Market research, competitor analyses, industry reports, benchmarks
Notes	40%	Working thoughts, meeting notes, brainstorm fragments, loose observations

The weight controls how heavily the synthesis engine trusts material from that folder. A chunk from a 90% folder will outrank an equally relevant chunk from a 40% folder. This means your core strategy documents anchor the analysis, while background research and rough notes add texture without dominating.

### Uploading Documents

StrategyForge accepts **.txt** and **.md** files. Drag them onto the upload zone or click to browse. Multiple files can be uploaded at once.

If you have documents in other formats (PDF, DOCX, Google Docs), convert them to plain text first. Copy-paste into a **.txt** file works perfectly. The engine does not need formatting — it needs your words.

### Auto-Assignment

When you upload a file, StrategyForge assigns it to a folder based on keywords in the filename:

- Files containing **strategy, plan, pitch, deck, model, okr, or roadmap** go to **Strategy**.
- Files containing **research, market, competitor, analysis, report, intel, or benchmark** go to **Intelligence**.
- Everything else goes to **Notes**.

This is a convenience, not a rule. You can always move documents between folders by removing them and re-uploading with a different filename, or by editing the exported state JSON directly.

## Adjusting Weights

Each folder has a weight slider ranging from 0% to 100%. Slide it to match your confidence in that source type. If your research is rock-solid and your strategy draft is still half-baked, raise Intelligence and lower Strategy. The weights are yours to calibrate.

A weight of 0% means the engine still sees the material but gives it no extra trust. It will only surface if nothing else matches the query.

## Indexing and Embeddings

Before the engine can use your sources in analysis, they need to be **indexed**. Indexing breaks each document into overlapping chunks (roughly 800 characters each) and creates a mathematical fingerprint (an *embedding*) for every chunk.

When you ask a question, the engine creates an embedding for your question too, then finds the chunks whose fingerprints are most similar — weighted by their folder trust level. The top matches are injected into the AI prompt as context.

Click **Index All Sources** after uploading. Indexing requires the embedding model (nomic-embed-text) and takes a few seconds per document. Once indexed, sources stay indexed until you clear them.

Embeddings are stored locally in your browser's IndexedDB. They never leave your machine.

## 4. Lenses

A lens is a perspective. It is a way of reading your material and your question that foregrounds certain concerns and backgrounds others. A Market lens asks about customers and demand. A Risk lens asks what kills the business. An Execution lens asks whether the team can actually deliver.

Every synthesis runs your question through *all active lenses simultaneously*. The engine does not just answer your question — it answers it five different ways, then shows you where those answers collide.

### What a Lens Is

Each lens has four components:

- **Name** — What it is called. Vision, Market, Model, etc.
- **Core Question** — The fundamental question this lens asks. For example, the Risk lens asks: *"What kills this? What are we not seeing?"*
- **Weight** — How much emphasis this lens gets in the synthesis (0–100%). Higher-weighted lenses get proportionally more space in the analysis.
- **Focus Angles** — Seven specific sub-angles for Deep Dive mode. The Risk lens, for example, has: competition, execution, regulation, funding, dependency, market-shift, team.

Each lens can be toggled on or off. An inactive lens is ignored during synthesis. This lets you maintain a large palette of perspectives and activate only the ones relevant to your current question.

### The Five Preset Packs

StrategyForge ships with five preset packs, each designed for a different strategic role:

PRESET	LENSES	DESIGNED FOR
Founder	Vision, Market, Model, Traction, Risk	Founders building and pitching a company
Go-to-Market	Positioning, Channels, Pricing, Launch, Retention	Firms launching or scaling a product
Investor	Thesis, Diligence, Valuation, Portfolio, Exit	Investors evaluating opportunities
Operator	Execution, Team, Process, Metrics, Culture	Leaders managing execution and teams
Consultant	Situation, Complication, Resolution, Evidence, Actions	Advisors structuring strategic analysis

Each preset is a starting point. You can mix presets by loading one, adding custom lenses, or adjusting weights to create your own blend. The presets are covered in detail in Chapters 5–9.

## Custom Lenses

Add your own lens by entering a name and a core question in the Lenses tab. Custom lenses integrate fully into synthesis and deep dive — they are first-class citizens, not afterthoughts.

Some ideas for custom lenses:

- **Regulatory** — "What regulatory forces shape this market, and how might they shift?"
- **Sustainability** — "What are the environmental and social implications of this strategy?"
- **Geopolitical** — "How do geopolitical dynamics affect supply chains, markets, and talent?"
- **Customer Voice** — "What would the customer say about this? What friction are we ignoring?"
- **Second-Order Effects** — "If this works, what breaks? What does success create?"

The best custom lenses are the ones that ask the question nobody on the team is comfortable asking.

## Adjusting Lens Weights

Each lens has a weight slider (0–100%). The weight determines how much emphasis the lens receives in the synthesis prompt. A lens at 100% gets roughly twice the emphasis of a lens at 50%.

Weights are relative. If all five lenses are at 70%, they each get equal emphasis. If you boost Risk to 100% and leave the others at 50%, Risk gets about a third of the synthesis real estate.

Experiment. The same question asked with different lens weights can produce surprisingly different strategic insights.

## 5. The Founder Preset

---

The Founder preset is designed for anyone building a company — from the earliest napkin sketch through Series A and beyond. Its five lenses cover the fundamental questions investors, advisors, and co-founders will ask.

### Vision

**Core question:** *"What future are we building toward? What is the contrarian truth?"*

The Vision lens pushes beyond product features into the worldview that makes the company inevitable. It asks about mission clarity, differentiation that matters, narrative coherence, market timing, and whether the ambition is sized correctly — too small and it is a feature, too large and it is a fantasy.

**Focus angles:** mission, differentiation, narrative, timing, ambition, conviction, clarity

### Market

**Core question:** *"Who is the customer, how big is the opportunity, and what forces shape it?"*

This lens interrogates whether the market is real, reachable, and ready. It probes total addressable market (TAM) assumptions, segmentation logic, secular trends, pain-point severity, willingness to pay, and adoption curves.

**Focus angles:** TAM, segmentation, trends, timing, pain-points, willingness-to-pay, adoption

### Model

**Core question:** *"How does this business make money? What are the unit economics?"*

The Model lens is about mathematics — not just revenue, but the mechanics that make revenue repeatable. It examines revenue streams, gross margins, customer acquisition cost (CAC), lifetime value (LTV), scalability dynamics, pricing structure, and retention.

**Focus angles:** revenue-streams, margins, CAC, LTV, scalability, pricing, retention

### Traction

**Core question:** *"What evidence do we have? What metrics matter right now?"*

Traction is about proof. Not projections — evidence. This lens asks about growth rate, engagement quality, pipeline, milestone velocity, cohort behavior, NPS, and the overall pace of learning.

**Focus angles:** growth-rate, engagement, pipeline, milestones, cohorts, NPS, velocity

## Risk

**Core question:** *"What kills this? What are we not seeing?"*

The Risk lens is the one most founders skip and most investors prioritise. It surfaces competitive threats, execution risks, regulatory exposure, funding dependencies, single points of failure, market-shift scenarios, and team gaps.

**Focus angles:** competition, execution, regulation, funding, dependency, market-shift, team

## 6. The Go-to-Market Preset

---

The GTM preset is built for teams preparing to launch or scale a product. Whether you are building a launch plan, refining your messaging, or rethinking pricing, these five lenses cover the full surface area of going to market.

### Positioning

**Core question:** *"Who are we for, what do we replace, and why is it obvious?"*

Positioning is not taglines — it is the strategic decision about which fight you are picking. This lens examines ideal customer profile (ICP), category design, competitive alternatives, messaging clarity, wedge strategy, proof points, and the narrative that makes the product feel inevitable.

**Focus angles:** ICP, category, alternatives, messaging, wedge, proof-points, narrative

### Channels

**Core question:** *"Where do we find customers? What scales and what does not?"*

Channels are the machinery of distribution. This lens asks about acquisition strategies, distribution partnerships, content-driven growth, paid versus organic dynamics, and whether community can become a channel.

**Focus angles:** acquisition, distribution, partnerships, content, paid, organic, community

### Pricing

**Core question:** *"What is the right price, structure, and packaging?"*

Pricing is strategy expressed in numbers. This lens probes value metrics (what are you charging *per?*), tier architecture, anchoring effects, freemium logic, enterprise packaging, expansion revenue, and competitive pricing dynamics.

**Focus angles:** value-metric, tiers, anchoring, freemium, enterprise, expansion, competitors

### Launch

**Core question:** *"How do we enter the market? What is the sequence?"*

Launch is about choreography. This lens examines timing, beta programs, waitlist mechanics, press strategy, launch partnerships, geographic sequencing, and how to build momentum rather than a single spike.

**Focus angles:** timing, beta, waitlist, PR, partnerships, geography, momentum

## Retention

**Core question:** *"Why do they stay? What drives expansion and loyalty?"*

Retention is the lens most GTM plans underweight. It asks about onboarding quality, activation milestones, habit formation, switching costs, NPS as a signal, churn drivers, and upsell pathways.

**Focus angles:** onboarding, activation, habits, switching-costs, NPS, churn, upsell

## 7. The Investor Preset

---

The Investor preset serves anyone evaluating investment opportunities — venture capitalists, angel investors, corporate development teams, or founders stress-testing how investors will read their company.

### Thesis

**Core question:** *"What is the investment thesis? Why this, why now?"*

The Thesis lens seeks the insight that makes this investment non-obvious. It examines macro trends, proprietary insights, timing arguments, market size, defensibility, founder-market fit, and the contrarian angle.

**Focus angles:** macro-trend, insight, timing, market-size, defensibility, founder-market-fit, contrarian

### Diligence

**Core question:** *"What does the evidence say? Where are the gaps?"*

Diligence is systematic skepticism. This lens probes financials, operating metrics, reference checks, legal risk, IP strength, team assessment, and product maturity.

**Focus angles:** financials, metrics, references, legal, IP, team, product

### Valuation

**Core question:** *"What is it worth? What are the return scenarios?"*

The Valuation lens runs the numbers forward. It considers trading multiples, comparable transactions, DCF scenarios, dilution modeling, term sensitivity, and milestone-based valuation triggers.

**Focus angles:** multiples, comps, DCF, scenarios, dilution, terms, milestones

### Portfolio

**Core question:** *"How does this fit the portfolio? What concentrations emerge?"*

Portfolio is the meta-lens. It asks about overlap with existing investments, diversification gaps, stage concentration, sector weighting, geographic exposure, vintage balance, and reserve allocation.

**Focus angles:** overlap, diversification, stage, sector, geography, vintage, reserves

### Exit

**Core question:** *"How does this return capital? Who acquires, who IPOs?"*

The Exit lens works backward from liquidity. It identifies likely acquirers, M&A comparables, IPO readiness, timeline assumptions, expected multiples, secondary opportunities, and liquidity risk.

**Focus angles:** acquirers, M&A-comps, IPO-path, timeline, multiples, secondaries, liquidity

## 8. The Operator Preset

---

The Operator preset is for leaders who have to make the strategy real. Strategy without execution is decoration. These five lenses cover the machinery of making things actually happen inside an organisation.

### Execution

**Core question:** *"Can we actually do this? What is the critical path?"*

The Execution lens asks the question that separates strategy slides from strategy. It maps the roadmap, surfaces dependencies, identifies bottlenecks, assesses velocity, names the trade-offs, inventories resources, and sequences the work.

**Focus angles:** roadmap, dependencies, bottlenecks, velocity, trade-offs, resources, sequencing

### Team

**Core question:** *"Do we have the right people? Where are the gaps?"*

This lens is about human capital and organisational design. It examines capability coverage, hiring priorities, reporting structure, leadership depth, succession planning, team morale, and strategic alignment across the organisation.

**Focus angles:** capabilities, hiring, structure, leadership, succession, morale, alignment

### Process

**Core question:** *"How does work flow? Where is the friction?"*

Process is the operating system of a company. This lens probes workflows, automation opportunities, handoff points, tool sprawl, documentation quality, technical and organisational debt, and continuous improvement practices.

**Focus angles:** workflows, automation, handoffs, tools, documentation, debt, improvement

### Metrics

**Core question:** *"What should we measure? What drives what?"*

The Metrics lens asks whether you are measuring the right things — and whether you understand the causal relationships between them. It covers KPI selection, leading indicators, dashboard design, review cadence, accountability structures, attribution, and target-setting.

**Focus angles:** KPIs, leading-indicators, dashboards, cadence, accountability, attribution, targets

## Culture

**Core question:** *"What kind of company are we building? What behaviors do we reward?"*

Culture is strategy expressed in daily habits. This lens examines stated values versus lived values, behavioural norms, decision-making patterns, communication style, conflict resolution, speed versus deliberation, and the trust architecture of the team.

**Focus angles:** values, norms, decision-making, communication, conflict, speed, trust

## 9. The Consultant Preset

---

The Consultant preset follows the classic structured problem-solving framework: situation, complication, resolution, evidence, action. It is designed for anyone who needs to produce rigorous, defensible strategic analyses — whether for clients, boards, or internal leadership.

### Situation

**Core question:** *"What is the current state? What is the context?"*

The Situation lens establishes the baseline. What does the landscape look like? What is the relevant history? Who are the stakeholders? What constraints bind the analysis? What assets exist? Where is the company positioned today? What trends are in play?

**Focus angles:** landscape, history, stakeholders, constraints, assets, position, trends

### Complication

**Core question:** *"What changed? What is the problem or opportunity?"*

The Complication lens identifies the disruption — the thing that makes the current state unstable. It probes urgency, root causes, impact scope, downstream consequences, timeline pressure, and severity.

**Focus angles:** disruption, urgency, root-cause, impact, consequences, timeline, severity

### Resolution

**Core question:** *"What should we do? What are the options and trade-offs?"*

This lens generates and evaluates alternatives. It considers evaluation criteria, specific recommendations, feasibility assessments, risk profiles, sequencing logic, and how to build buy-in for the chosen path.

**Focus angles:** alternatives, criteria, recommendation, feasibility, risk, sequence, buy-in

### Evidence

**Core question:** *"What supports this? What data, analogies, or precedents?"*

The Evidence lens asks what backs up the recommendation. It examines quantitative data, qualitative signals, external benchmarks, case studies, experimental results, survey data, and analogies from adjacent domains.

**Focus angles:** quantitative, qualitative, benchmarks, case-studies, experiments, surveys, analogies

### Action

**Core question:** *"What happens Monday? Who does what by when?"*

The Action lens turns analysis into movement. It assigns owners, establishes timelines, sets milestones, identifies resource requirements, plans communication, defines governance, and builds escalation paths.

**Focus angles:** owners, timeline, milestones, resources, communication, governance, escalation

## 10. The Synthesis Engine

---

The Synthesis Engine is the heart of StrategyForge. It takes your question, retrieves the most relevant chunks from your sources (weighted by folder trust), and passes them to the AI alongside instructions generated from your active lenses.

### How Synthesis Works

When you click **Synthesize**, this is what happens behind the scenes:

- **1. Retrieval.** The engine creates an embedding of your question and compares it against every indexed chunk in your Source Vault. It selects the top six most relevant chunks, weighted by their folder trust level.
- **2. Prompt Construction.** The engine builds a system prompt that lists every active lens with its core question and relative weight. Higher-weighted lenses receive more emphasis.
- **3. Context Injection.** The retrieved source chunks are appended to your question, so the AI can ground its analysis in your specific material.
- **4. Streaming.** The response streams in token by token, appearing in real time. When complete, basic Markdown formatting is applied.

### Writing Good Prompts

The synthesis engine responds to the quality of your question. Here are some principles:

**Be specific.** "What should our strategy be?" is vague. "Should we enter the enterprise segment before or after achieving 1,000 SMB customers?" is actionable.

**Include tension.** The best prompts contain a dilemma: "Our retention data says customers love the product, but our NPS is declining. What is happening?"

**Name what you fear.** "What happens to our unit economics if customer acquisition cost doubles in the next twelve months?" forces the engine to confront specific scenarios.

**Ask what-if questions.** "If our primary competitor acquires our biggest integration partner, how does our positioning change?" generates insight that planning documents rarely contain.

**Challenge your own assumptions.** "Our pitch says we have no competition. What are we wrong about?" is the kind of question StrategyForge was built for.

## Reading the Output

A synthesis output has a consistent structure:

- **Per-lens analysis.** Each active lens provides a focused paragraph applying its perspective to your question.
- **Tensions.** Where do the lenses conflict? Where does optimising for one dimension hurt another?
- **Connections.** Where do the lenses reinforce each other? Where is there leverage?
- **Blind spots.** What has the analysis not considered? What questions remain unanswered?
- **Surprising insight.** The single observation the engine considers most unexpected or non-obvious.

The tensions section is often the most valuable. Strategy is the art of navigating trade-offs, and the multi-lens approach is designed to make those trade-offs visible.

## Copy and Export

Every synthesis can be copied to your clipboard or exported as a Markdown (.md) file with a full metadata header including the prompt, timestamp, model, and active lenses. Use this to build a library of strategic analyses over time.

# 11. Deep Dive

---

Where Synthesis gives you breadth across all lenses, Deep Dive gives you depth through a single one. Choose a lens, choose a focus angle within that lens, and ask your question. The engine goes narrow and deep.

## Choosing a Lens and Focus

The Deep Dive tab shows a dropdown of all active lenses and a row of focus pills — the specific sub-angles within that lens. For example, selecting the **Risk** lens reveals focus pills for competition, execution, regulation, funding, dependency, market-shift, and team.

Click a focus pill to select it. The engine will concentrate its analysis on that specific angle, while still drawing context from your sources via RAG retrieval.

## When to Deep Dive

Use Deep Dive when a synthesis has flagged something worth exploring further. If the Tensions section reveals a conflict between your Market and Model lenses, deep-dive into the Model lens with a focus on pricing or margins. If the Blind Spots section mentions regulatory risk, deep-dive into a custom Regulatory lens.

Deep Dive is also useful for preparing specific sections of a pitch deck, strategy document, or board presentation. Need to stress-test your competitive moat? Deep-dive into Thesis with a focus on defensibility.

## Working with Focus Angles

Each preset lens comes with seven focus angles, chosen to cover the most useful sub-dimensions. Custom lenses get four default angles (overview, detail, risk, opportunity), but you can always add more by editing the exported state JSON.

Focus angles are not constraints — they are starting points. The AI will follow the angle but also draw on anything relevant from your sources. Think of them as the question behind the question.

## 12. Sidebar and Thoughts

---

The left sidebar contains the **Thoughts** panel — a simple scratchpad for capturing ideas, questions, and fragments that occur to you while working.

### The Scratchpad

Click **+ new thought** to create a new entry. Each thought has a title and a body. Use it for:

- Questions you want to synthesise later
- Hypotheses triggered by an analysis
- Fragments of strategy language you want to refine
- Disagreements with the engine's output
- Connections you notice between different analyses

Thoughts are stored locally and included in state exports. They are not indexed by the AI — they are for you alone.

### Search

The search bar at the top of the sidebar filters thoughts by title and content. Type a keyword and the list narrows in real time.

On mobile screens (under 600px), the sidebar slides in from the left. Tap the hamburger menu to open it, tap the overlay to close.

## 13. Notes

---

The Notes tab is a timestamped notebook for your own observations and insights. Unlike Thoughts (which are sidebar fragments), Notes are meant to capture your reactions to the engine's output.

### Capturing Observations

After running a synthesis or deep dive, switch to Notes and write down what struck you. Not what the engine said — what *you* noticed. What assumption did it make that you disagree with? What connection did it miss? What question did it raise that you had not considered?

### The Value of What the AI Missed

The best strategic insights often come not from what the AI says, but from the gap between what it says and what you know to be true. The engine does not know your team's morale, your board's risk appetite, or the conversation you had with a customer last Tuesday. Those gaps are where your judgement adds the most value.

Notes are persistent across sessions and included in state exports. Individual notes can be deleted.

## 14. History

---

The History tab stores all syntheses and deep dives (up to 50 of each). Each entry shows the prompt, type (Synthesis or Deep Dive), timestamp, and model used.

### Browsing Past Analyses

Click any history item to expand it and see the full response. This is useful for reviewing past analyses, comparing how your thinking has evolved, or finding an insight you remember but did not save.

### Exporting History

Each history item can be copied to clipboard or exported as a Markdown (.md) file. The export includes a full metadata header with the original prompt, timestamp, model, and (for syntheses) the active lenses and their weights.

## 15. LLM Backends

---

StrategyForge supports four AI backends. You can switch between them at any time in the Settings tab. Your sources, lenses, notes, and history are all preserved when you switch.

### Ollama

Ollama is the default and recommended backend. It runs AI models locally on your machine. Nothing leaves your computer.

- **Cost:** Free
- **Privacy:** Complete — all inference happens on localhost
- **Speed:** 5–30 seconds per response, depending on hardware
- **Default model:** llama3.2 (roughly 2 GB)
- **Requirements:** 8 GB RAM minimum, 16 GB recommended

### Anthropic

The Anthropic backend uses Claude via the Anthropic API. It is faster and more capable than local models but requires an internet connection and an API key with billing configured.

- **Cost:** Pay-per-use (see [anthropic.com](https://anthropic.com) for current pricing)
- **Privacy:** Prompts and retrieved source chunks are sent to Anthropic's servers. Full documents are never transmitted.
- **Speed:** 2–5 seconds per response
- **Default model:** claude-sonnet-4-20250514

To use the Anthropic backend: go to Settings, select Anthropic, and paste your API key. The connection indicator will turn green immediately.

### WebAI and StudioLM

These are Ollama-compatible backends for custom server setups. Enter the server URL in Settings. They use the same API format as Ollama but on a different address.

### Switching Backends

Backend switching preserves all state. Your documents, lenses, notes, thoughts, and history remain exactly as they are. Only the AI connection changes. The model name auto-updates to a sensible default for the new backend.

## 16. Settings

---

### Display

**Vignette** — Adjustable from 0% (none) to 100% (heavy). The vignette is a radial gradient that darkens the edges of the screen, focusing attention on the centre. It is a visual preference, not a functional setting.

**Theme** — Toggle between dark and light mode. Both themes include a subtle paper noise texture for visual warmth. The toggle is also available in the top bar.

### Embedding Model

The embedding model is used for source indexing (RAG). The default is **nomic-embed-text**, a lightweight model that produces high-quality embeddings. You can change this to any Ollama-compatible embedding model.

If you change the embedding model, you should clear existing embeddings and re-index your sources, as different models produce incompatible embeddings.

### State Management

**Export State (JSON)** — Saves everything: documents, lenses, custom lenses, notes, thoughts, history, settings, and backend configuration. This is your complete workspace as a single file.

**Import State** — Loads a previously exported JSON file. The import validates the file structure before applying it. Malformed files produce an error message without corrupting your current state.

**Reset Everything** — Clears all state after confirmation. This cannot be undone. Export first if you have work you want to preserve.

## 17. Workflows

---

StrategyForge is flexible by design, but here are five concrete workflows that show how the pieces fit together in practice.

### Pitch Deck Stress Test

You have a pitch deck and you want to know where it breaks.

- **Upload** your pitch deck text to the Strategy folder.
- **Load** the Founder preset.
- **Add** a custom lens: **Skeptic** — "If this were a bad investment, why? What is the most generous reading a passing investor would give?"
- **Synthesize**: "Read this pitch and tell me where the narrative is weakest, where the evidence is thinnest, and where a sophisticated investor would push back."
- **Deep Dive** into Risk with a focus on competition. Then into Model with a focus on LTV.
- **Note** every objection you cannot immediately answer. Those are your revision priorities.

### Quarterly Strategy Review

It is the start of a new quarter and you need to assess where the company stands.

- **Upload** last quarter's OKR results, this quarter's plan, and any relevant market updates.
- **Load** the Operator preset.
- **Synthesize**: "Based on last quarter's results and this quarter's plan, what are the three most important things to get right in the next 90 days?"
- **Deep Dive** into Metrics with a focus on leading-indicators.
- **Switch** to the Founder preset and re-run: "Does this quarter's plan still serve the two-year vision?"

### Competitive Response

A competitor just made a significant move and you need to respond.

- **Upload** the competitor's announcement, your current positioning doc, and any market research.
- **Load** the Go-to-Market preset.

- **Synthesize:** "Our competitor just [describe the move]. How does this change our positioning, pricing, and channel strategy?"
- **Deep Dive** into Positioning with a focus on alternatives.
- **Add** a custom lens: **Customer Impact** — "How does this change what the customer sees, feels, and chooses?"

## Investment Memo

You are writing an investment memo for a portfolio company or fund LP.

- **Upload** the company's pitch deck, financial model, and your diligence notes.
- **Load** the Investor preset.
- **Synthesize:** "What is the strongest investment thesis for this company, and what are the three most significant risks?"
- **Deep Dive** into Valuation with a focus on scenarios.
- **Deep Dive** into Exit with a focus on acquirers.
- **Export** both analyses as Markdown. Use them as raw material for your memo.

## Go-to-Market Launch

You are planning a product launch and need to align the team.

- **Upload** your positioning document, competitor pricing, beta feedback, and launch plan.
- **Load** the Go-to-Market preset.
- **Synthesize:** "What is the most important thing to get right in the first 30 days of launch?"
- **Deep Dive** into Launch with a focus on momentum.
- **Deep Dive** into Retention with a focus on onboarding.
- **Switch** to the Founder preset and synthesize: "Does this launch plan match the company's positioning story?"

## 18. Accuracy and Judgement

---

StrategyForge uses AI language models, and AI language models are not perfect. They can be confidently wrong. They can miss obvious things. They can hallucinate connections that do not exist.

This is not a bug — it is a feature of how these tools work. The synthesis engine is a provocation partner, not a decision-maker. It is designed to surface perspectives, tensions, and blind spots that you might not have considered. Whether those perspectives are correct is for you to judge.

Some guidelines:

- **Treat every output as a draft.** The engine is showing you a possibility, not a truth.
- **Challenge what feels right.** Confirmation bias is real. The most dangerous outputs are the ones that tell you exactly what you wanted to hear.
- **Verify specific claims.** If the engine cites a trend, a statistic, or a market dynamic, verify it independently. Language models do not have access to real-time data.
- **Value the tensions over the conclusions.** The most useful part of a synthesis is often the Tensions section — where different lenses disagree. That is where the real strategic decisions hide.
- **Use it to think, not to decide.** StrategyForge should make you a better thinker, not replace your thinking.

## 19. Privacy and Data

---

StrategyForge stores all data locally in your browser's localStorage (application state) and IndexedDB (vector embeddings). No data is transmitted to any server except the AI backend you configure.

- **Ollama:** All calls go to localhost. Nothing leaves your machine.
- **Anthropic:** Prompts and retrieved source chunks are sent to Anthropic's API. Full documents are never transmitted — only the relevant chunks selected by RAG retrieval.
- **WebAI / StudioLM:** Calls go to the server address you configure.

StrategyForge has no telemetry, no analytics, no tracking, no cookies, and no network calls other than those listed above. Your strategy documents never leave your machine unless you explicitly choose a cloud backend.

## 20. Keyboard Shortcuts

---

---

SHORTCUT	ACTION
Ctrl/Cmd + S	Save state to localStorage
Ctrl/Cmd + E	Export state as JSON

---

State is also auto-saved with debouncing (100ms delay) after every change. The keyboard shortcuts are primarily useful for forcing an immediate save before closing the browser.

## 21. Troubleshooting

---

### Status shows "offline"

Run `ollama serve` in your terminal. Check for the llama icon in your menu bar (Mac) or system tray (Windows). If Ollama is running but StrategyForge still shows offline, check that the URL in Settings matches (default: `http://localhost:11434`).

### "Model not found" error

Run `ollama pull llama3.2` in your terminal. The model needs to be downloaded before first use.

### Indexing fails

Run `ollama pull nomic-embed-text` to download the embedding model. Indexing requires this model specifically.

### Responses are slow

This is normal for local AI. Ollama runs models on your CPU (and GPU if available). Close memory-intensive applications. If speed is critical, switch to the Anthropic backend for cloud-based inference.

### Import fails

Ensure the file is valid JSON exported from StrategyForge. The import validates the file structure before applying it. If the file is corrupted, the error message will indicate the issue.

### Browser storage full

Browser localStorage has a size limit (typically 5–10 MB depending on your browser). If you hit this limit, export your state, clear the browser storage, and re-import with fewer documents. Consider archiving older sources that are no longer actively relevant.

### Anthropic returns an error

Check that your API key is valid and has billing configured. The Anthropic backend requires an active API key with available credits. Error 401 means invalid key. Error 429 means rate limited — wait a moment and try again.

### Fonts do not load

StrategyForge loads Crimson Text and IBM Plex Mono from Google Fonts on first use. This requires an internet connection. After the first load, fonts are cached by your browser and available offline in most cases.

## 22. Glossary

---

**CAC** — Customer Acquisition Cost. The total cost of acquiring a new customer, including marketing, sales, and onboarding expenses.

**Chunk** — A segment of a source document (roughly 800 characters) created during indexing. Chunks overlap at sentence boundaries to preserve context.

**DCF** — Discounted Cash Flow. A valuation method that estimates the present value of future cash flows.

**Deep Dive** — A single-lens analysis with a specific focus angle. Narrower than synthesis, designed for depth.

**Embedding** — A mathematical fingerprint (vector) representing the semantic meaning of a text chunk. Used for similarity search in RAG retrieval.

**Focus Angle** — A specific sub-dimension within a lens. Each lens has seven focus angles for Deep Dive mode.

**GTM** — Go-to-Market. The strategy for bringing a product to customers, covering positioning, channels, pricing, and launch.

**ICP** — Ideal Customer Profile. A detailed description of the customer most likely to buy and succeed with your product.

**IndexedDB** — A browser-based database used by StrategyForge to store vector embeddings locally.

**Lens** — A perspective through which the synthesis engine reads your material. Each lens has a core question, weight, and focus angles.

**localStorage** — A browser-based storage mechanism used by StrategyForge to persist application state between sessions.

**LTV** — Lifetime Value. The total revenue expected from a customer over the entire duration of their relationship with the business.

**NPS** — Net Promoter Score. A measure of customer satisfaction based on the likelihood of recommending a product.

**OKR** — Objectives and Key Results. A goal-setting framework that links ambitious objectives to measurable outcomes.

**Ollama** — A free, open-source tool for running AI models locally. The default and recommended backend for StrategyForge.

**Preset** — A pre-configured set of five lenses designed for a specific strategic role (Founder, GTM, Investor, Operator, Consultant).

**RAG** — Retrieval-Augmented Generation. The technique of retrieving relevant source material and injecting it into AI prompts for grounded analysis.

**Source Vault** — The document management system in StrategyForge. Three weighted folders containing your uploaded source material.

**Synthesis** — A multi-lens analysis that examines a question through all active lenses simultaneously, revealing tensions and connections.

**TAM** — Total Addressable Market. The total revenue opportunity available for a product or service.

**Vignette** — A visual effect that darkens the edges of the screen, focusing attention on the centre of the workspace.

**Weight** — A numerical value (0–100%) that controls the relative emphasis of a folder or lens in the analysis.

# StrategyForge

*Version 1.0 · April 2026*

*Your strategy. Your studio. Your rules.*

---

Everything is a single HTML file. No installation. No cloud. No account. No subscription. No telemetry.

You own your data. Every synthesis, deep dive, note, and exported state file you create with StrategyForge belongs entirely to you. No rights are claimed over your content. No license is imposed on your output.