


Proofed: a governed content system for B2B tech companies





Content your engineers will actually trust.

A Claude Project that turns product releases, customer wins, and competitor moves into credible, on-voice content for every reader you sell to.



The problem

B2B technology companies publish for more than one reader at once. The technical user can spot marketing language in a sentence. The buyer funding the purchase needs the same depth in business terms. Most teams write one blurred piece that half-lands with everyone, at a cadence they can't sustain, with a voice that drifts and claims that quietly outrun the proof.

What Proofed is

→ **Not a content generator**

Not a tool that spits out posts. A routing-and-governance layer.

→ **Shared canon first**

Every drafter reads a shared canon before writing a single word.

→ **Signal comes in**

A release, a customer win, a competitor move. An orchestrator classifies it and routes it to a specialized drafter.

→ **QA gate**

A quality check runs on every draft before it ships.

The shared canon

Six governance files every draft reads before writing:

1

Voice and style

Anti-hype rules with a banned-language list

2

Positioning canon

Category, wedge, messaging pillars

3

Audience map

Technical user, operational owner, economic buyer

4

Claims and proof bank

Nothing publishes without a source

5

Competitive matrix

Honest positioning with rules of engagement

6

Search targets

The queries and answer-engine phrasings to win

Five specialized drafters

Release-to-content converter

Turns shipping into a content set

Expert drafter

Technical explainers with admitted limitations

Buyer-narrative drafter

Translates depth into spend, risk, and time-to-value


Comparison builder

Honest competitive pages, no FUD

Customer-story builder

Sourced and approved results only

Plus a repurposing protocol that rebuilds each piece per channel instead of just shortening it.



The proof bank rule

The spine of the product: no number, benchmark, or customer result ships unless it appears in the proof bank with a source.

Claims without sources get flagged [NEEDS SOURCING], not published. The engine refuses to invent statistics.

- ❏ That refusal is what keeps a technical audience trusting the content.

The four-check QA gate

Every draft passes before it is presented:

01

Claims

Every number, capability, and comparison is sourced

03

Answer engine

One primary query, answer-first opening, extractable structure

02

Voice

Plain register, no banned words, active voice, leads with the problem

04

Skeptical-reader pass

The draft survives the three hardest objections a distrustful reader would raise

⊗ Claims and skeptic are hard gates. Nothing publishes if either fails.

Serving readers separately, on purpose

One release, three pieces.

Technical explainer

For the engineer who adopts


Operational angle

For the manager who runs it at
scale


Buyer narrative

For the person who signs off on spend

Each piece stays sharp for its reader instead of gesturing at all three.



Setup and use



The built-in interview configures the canon for your company in about fifteen minutes. Then feed it signals: a release note, a customer win, a competitor launch. Get back drafts, a QA verdict, and a repurposing plan.

Who it's for

Content operations leads, founder-marketers, and DevRel or product-marketing teams at B2B technology companies who publish for both technical and economic buyers.

- ❏ If you only write for one audience, or you don't care whether claims are sourced, this isn't for you.

Content operations leads

Managing multi-audience publishing at scale

Founder-marketers

Wearing every hat, needing a governed system

DevRel and PMM teams

Bridging technical depth and business narrative