



LevelUp

AI Proposal Evaluation System

Your data. Your model. Your call. ▶ levelup.falconlogic.ai

Source-Selection-Grade Proposal Evaluation

LevelUp reads every offeror volume against your source selection criteria and returns auditable, citation-backed findings and ratings — in hours, not months. Built for the organizations that can't put competition-sensitive or classified proposals in a commercial AI vendor's hands.

AIR-GAPPED & PRIVATE BY DEFAULT

- Deploy the full stack — API, UI, and inference — in your own AWS account or isolated enclave, IL4–IL6 through JWICS/TS-SCI.
- Inference runs in-VPC on your own GPUs; proposals stay in your S3, encrypted and CloudTrail-logged.
- Commercial cloud models (Bedrock, OpenAI) are opt-in only — never the default.

OPEN-SOURCE MODELS YOU CONTROL

- Runs on open-weight US models — Llama, Gemma, Nemotron, GPT-OSS. No black-box vendor model in the loop.
- You own the weights — fine-tune on your own historical evaluations without handing data to a third party.

DOCUMENT INTAKE & SOLICITATION BUILD

- Volume-tagged, multi-document uploads scored as one body of evidence, with paragraph-indexed DOCX/PDF parsing and OCR for scans.
- Build the solicitation straight from the package: Volumes, Items, Factors, and Standards pulled from Section L/M and the PWS/SOW.
- Pre-evaluation format and page-count checks, plus optional AI pre-review that flags front matter and non-substantive pages.

AUDITED & PROTEST-DEFENSIBLE

- Verbatim, paragraph-indexed citations on every finding; a quote validator rejects anything it can't match to the source.
- Full run history — steps, prompts, payloads — retained for protest-defensible recordkeeping.
- Deterministic rollups: the rubric sets the rating, the LLM only writes the rationale.

SPEED & CROSS-CHECK

- Every offeror, every volume, evaluated in minutes to hours — not days or months.
- Run the same evaluation across multiple LLMs, each acting as an independent board member. Keep the consensus, discard the rest.

HOW TO DEPLOY

- Self-host: your AWS account, up to TS/SCI. No third-party API calls unless you opt in — the security-first path.
- Hosted: our managed levelup.falconlogic.ai environment. Lower entry cost, no infrastructure to run — for non-classified work.
- Industry pricing from \$999 per proposal, no commitment. Start with one draft, scale when you need it.

LevelUp is the proposal evaluation system that is both AI-native and self-hostable to TS/SCI — competing workflow tools have no AI and competing AI tools are cloud-only.

See the features, request access, and sign in to try it at levelup.falconlogic.ai