

AI-NATIVE FEDERAL CAPTURE PLATFORM

Find the opportunity. Predict the winner. Write the draft.

"Legacy tools tell you an opportunity exists. PAGIntel tells you whether you'll win it, who you're up against, what language wins it, and writes the first draft."

HOW IT WORKS Six stages, one continuous loop

01 / SYNC

Continuous opportunity feed

Every SAM.gov notice from the last 90 days, all notice types including awards. Refreshed every few minutes. Full descriptions and attachments enriched on a rolling cycle.

02 / FINGERPRINT

Capability fingerprinting

Paste a capability statement or website copy. Extracts NAICS, set-asides, target agencies, keywords, win themes, past performance, certifications, and clearances held.

03 / MATCH

Semantic matching + filters

Profiles and opportunities become vector embeddings. Hard filters for NAICS, set-aside, deadline, agency, price, and active security clearance — toggling clearance refreshes the list instantly.

04 / PRICE

Contract value intelligence

Award amount when published; otherwise an estimate from comparable NAICS x agency awards. Placeholder \$0/sub-\$1K awards excluded, winsized p5/p95 bands, minimum sample (5+ at NAICS-4, 10+ at NAICS-2).

05 / PATTERNS

Win-pattern intelligence

Three years of comparable awards from USASpending.gov per opportunity. Identifies winners, language and structural cues, plus a Go/No-Go panel with HHI, top-vendor share, and \$/day burn.

06 / DRAFT + INTEL

Proposal generation, dossiers, alerts

AI drafts grounded in the opportunity, your profile, and winning patterns — section by section, export-ready. Per-agency dossiers, saved-search alerts, and a nightly backtest that keeps accuracy honest.

HOW IT DIFFERS FROM THE EXISTING MARKET

CAPABILITY	PAGINTEL	LEGACY CAPTURE TOOLS
Semantic AI matching	Yes. Continuously refreshed vector matching on both sides.	Keyword and NAICS filters only.
Contract value estimation	Yes. Median + winsized p5/p95 from comparable NAICS x agency awards.	Listed award amount only, or nothing.
Active security clearance filtering	Yes. Toggle excludes clearance/secret/TS-SCI notices instantly.	Manual keyword filtering.
Win-pattern + Go/No-Go intelligence	Yes. HHI, top-vendor share, \$/day burn per NAICS x agency bucket.	Generic templates.
AI proposal drafting grounded in opportunity, firm, and winners	Yes. Section by section, placeholders preserved, export-ready.	None, or a thin AI wrapper.
Nightly accuracy calibration against real awards	Yes. Backtests the matcher against itself.	No.
Price point	Yes. A fraction of market rate. Modern serverless stack.	\$15k to \$60k per seat, per year.

WHO IT'S BUILT FOR

Small and mid-size primes

Including 8(a), WOSB, and SDVOSB firms. Replaces a \$30k seat tool plus a part-time capture analyst.

Firms entering new agencies

Agency dossiers and win-pattern extraction compress years of relationship learning into a usable brief.

Teaming-heavy shops

Top-recipient and HHI analysis surface the obvious teaming targets for any given opportunity.

High-volume bidders

For anyone bidding more than ten opportunities a year, proposal drafting alone pays for the platform on the first bid.