



CAPABILITY STATEMENT

EBK LLC · EBK PUBLIC SECTOR · AMERICAN-BUILT AI FOR PUBLIC MISSIONS

ENTITY
EBK LLC

UEI / CAGE
JQTMJTJ6F7PL3 / 20FZ2

NAICS
541511 541512 541690 541715
518210

CONTACT
ebk@ebk.tech
(480) 399-3910

OVERVIEW

EBK is a quantitative research firm that studies decision-making under uncertainty and builds governed AI systems. EBK Public Sector, our government practice, works with agencies and teams with prime contractors as their AI engineering partner, on a narrow, well-defined scope beside the prime's program management and service delivery.

CORE COMPETENCIES

- **Auditable decision support:** cited evidence, deterministic scoring, sensitivity analysis and human approval, reproducible end to end
- **Governed AI agents** on an authorized commercial AI product, with human approval at consequential steps
- **MCP connectors**, read-only by default and portable across AI vendors
- **Evaluation and regression testing:** harnesses and suites for field accuracy, fabricated values, escalation, prompt injection and unapproved writes
- **Data engineering with provenance:** a source citation on every record

HOW WE BUILD

- Reuse native product controls (SSO, audit, metering) first
- OpenTelemetry traces for audit and drift detection
- Production only after evaluation gates pass
- Full code delivered to the Government's repository
- House rule: the AI extracts, deterministic code verifies every number

RELEVANT EXPERIENCE (COMMERCIAL)

- **Governed intake agent**, healthcare operations: AI intake on AWS Bedrock; schema-validated, audit-trailed extraction; e-signature; staff board with SSO. In production since May 31, 2026.
- **Executive MCP connector:** OAuth-protected; used daily from Claude and Grok on live data; 91 automated tests.
- **Public-records engine:** 20 live jurisdictions (18 US, 2 Canada) via Socrata, ArcGIS, CKAN, Carto, CSV and a portal scraper; provenance on every record.
- **Survey verification assistant:** deed OCR, traverse-closure checks, drafted legal descriptions; flags numbers for a licensed surveyor; 352 automated tests.

RESEARCH

- ~500-agent evolutionary decision system on a self-hosted 72B open-weights model, with false-discovery-gated selection; calibration error cut from 0.252 to 0.145 on held-out data. Agha, CMC Senior Thesis (2026).