

Doordash for Counter Terrorist — brief

Concept draft. Same dispatch pattern as a delivery app: parse a messy human request, validate it against rules, geofence it, confirm twice, log it immutably — only the payload differs.

The idea

The JTAC talks like a human. The AI does the paperwork. MCP tools do the math. Humans make every fire decision.

Flow: JTAC speaks naturally → AI parses the four-input brief (target location, description, mark type, friendly positions) → AI fans out to MCP tools (grid validation, no-strike check, danger-close math, geometry cross-check, terrain, threat rings, readback diffing) → AI reads back a validated brief → three human gates in series (JTAC confirms, pilot confirms, cleared hot) → everything logged immutably.

The split

- **LLM:** fuzzy parsing of speech, discrepancy flagging, readback composition.
- **MCP tools (deterministic code):** every calculation. A language model never computes danger-close.
- **Humans:** the only trigger authority. The AI drafts, validates, flags — never clears hot.

Charges

- **Safety path — free, unlimited, flat platform rate.** All validation tools. Cost must never influence whether a check runs.
- **Premium — metered (pay-per-use or credits).** Weapon-tuning, forecasting, multi-mission planning, 3D terrain, debrief analytics.
- **Enterprise — annual contract.** Dedicated hosting, custom tools, on-prem, SLA.

Full architecture with charges table: see </drafts/ai-strike-mcp-architecture.html>.