

FEDERAL AI / OPTIMIZE

Ask what will break before the schedule does.

Training platforms manage learning. Optimize manages the mission outcome. It is the mission intelligence layer for a readiness domain — sensing the systems of record, forecasting slip and risk, simulating courses of action, and staging decisions for the humans who own them. Current Air Force delivery.

THE OPERATING LOOP

Sense

Authoritative systems ingested continuously — scheduling, maintenance, personnel, weather.

Understand

One reconciled source of truth where there were many conflicting ones.

Forecast

Flying-hour delivery risk and slip risk surfaced days early, with drivers named.

Simulate

Courses of action scored — including the minimum-change recovery.

Decide + Learn

COAs staged for human review. No operational write without authority.

DOCUMENTED ENGAGEMENT OUTCOMES

12–15 hrs/wk

ANALYST TIME RECLAIMED

1

SOURCE OF TRUTH, FROM MANY

Same-day

DATA FIXES, WAS WEEKS

\$0/seat

MISSION-PRICED, NOT PER-USER

Outcomes reported from a delivered engagement; verification available in a briefing.

CLAIM DISCIPLINE

- **Foundation** — capabilities running in delivery today.
- **Configurable module** — built, adapted to your domain during onboarding.
- **Research option** — clearly labeled; scoped only with eyes open.

Every capability on the public page carries one of these three labels. We do not sell research as foundation.

THE 90-DAY PROOF

Kill criteria, published up front.

Optimize deploys against one readiness domain with agreed measures and explicit kill criteria. If the forecast doesn't beat the status quo inside 90 days, you stop — and we say so first. A decision ledger traces every recommendation back to its evidence.

DISCLAIMERS

No U.S. Air Force endorsement implied. The public demonstration at federal.ai/optimize runs on synthetic data and performs no operational writes. Named-unit references available only with unit consent, in an appropriate setting.