

CASE STUDY — DELIVERED ENGAGEMENT

From five conflicting systems to one readiness picture.

How an Air Force training organization reclaimed 12–15 analyst hours a week by replacing manual readiness reconciliation with a governed mission intelligence layer. Details below are stated at the level the customer has cleared for public use; named-unit verification is available in a briefing, with unit consent.

THE SITUATION

Readiness questions — *which aircraft is most likely to miss tomorrow's schedule, and why?* — were answered by analysts hand-reconciling scheduling, maintenance, and personnel records across multiple systems that routinely disagreed. Each weekly readiness picture consumed days of skilled time, was stale on arrival, and carried no consistent evidence trail. Data errors discovered downstream took weeks to trace and fix.

THE CONSTRAINT SET

- Authoritative systems could not be replaced — only read and reconciled.
- No operational writes by any automated system; humans own every decision.
- Deployment inside the organization's approved environment, not a vendor cloud.
- Every recommendation must show its sources and its reasoning.

WHAT WAS DELIVERED

Sense & reconcile

Continuous ingest from the systems of record into one reconciled source of truth. Conflicts surfaced with provenance instead of silently resolved.

Forecast & simulate

Flying-hour delivery risk and schedule-slip risk forecast days ahead, drivers named; courses of action scored, including the minimum-change recovery.

Decide, governed

COAs staged for human review with a decision ledger. No operational write performed by the system — by design, not by policy memo.

OUTCOMES REPORTED BY THE ORGANIZATION

12–15 hrs/wk

ANALYST TIME RECLAIMED

1

SOURCE OF TRUTH, WAS
5+

Same-day

DATA FIXES, WAS WEEKS

Days early

SLIP RISK SURFACED

WHY IT HELD UP

The engagement was run to Federal AI's operating doctrine: acceptance measures and kill criteria agreed before work began; a Cleared FDE embedded with the analysts who own the readiness picture rather than a remote delivery team; and a fail-closed contract — the system proposes and explains, the mission authority decides. The same doctrine is published at federal.ai and applies to every engagement.

SEE IT AGAINST YOUR DOMAIN

A 30-minute briefing maps this pattern to your readiness domain and states whether the 90-day proof applies.

federal.ai → Start an engagement

Claims boundary: figures above are engagement outcomes reported by the customer organization for its own operation; they are not a promise of identical results elsewhere. No U.S. Air Force endorsement implied. The public product demonstration runs on synthetic data.