

CLOUD AI/ML SERVICES

Avalon builds your AI use case into a secure, working capability inside your own cloud, in weeks, with the documentation your security office needs.

THE PROBLEM WE ADDRESS

Agencies now hold cheap or nearly free AI licenses through GSA OneGov, but lack the engineering capacity to connect those tools to their own data, harden them for government use, and clear the security office.

OMB M-25-21 requires documented minimum risk-management practices for high-impact AI, with compliance reporting due to OMB by September 22, 2026 — non-compliant high-impact AI must be discontinued.

WHAT WE DO

- Use-case feasibility scoring and ranked backlog
- Data readiness assessment and access-rights review
- Extraction and transformation pipeline engineering
- Classification, forecasting, anomaly detection, computer vision
- Retrieval-augmented generation over agency documents
- AWS SageMaker/Bedrock, Azure ML/OpenAI, or Vertex AI development
- Documented evaluation report with known limitations
- Model cards, AI impact-assessment inputs, and monitoring plans

HOW WE DO IT

1 DISCOVERY & DEFINITION

Stakeholder sessions, use-case scoring, and success metrics culminating in a signed-off go/no-go decision.

DELIVERABLE

Go/no-go decision + metrics

3 PILOT BUILD & EVALUATION

Model or GenAI development with iterative evaluation on real agency data against kickoff metrics.

DELIVERABLE

Evaluated pilot

2 DATA READINESS

Access provisioning, secure workspace stand-up in your tenant, and pipeline build until data is flowing.

DELIVERABLE

Data pipeline live

4 PRODUCTION HARDENING

MLOps pipelines, monitoring, and the documentation package your CAIO and security stakeholders accept.

HANDOVER

MLOps pipeline + governance docs

VALUE DELIVERED

- A working capability, not a strategy deck.** A 2-4 person pod delivers a deployed pilot in weeks.
- Evaluated on your real data.** Metrics are agreed at kickoff and tested against actual agency data, not a curated demo.
- Honest evaluation gates.** A pilot can end with a documented "this doesn't work well enough yet" — what M-25-21's testing culture rewards.
- Never one vendor announcement from a dead capability.** Model-portable architecture avoids lock-in as the approved-model roster shifts.

THE ONE THING TO REMEMBER

A working capability in weeks, not a slideware AI strategy.

COMMONLY PROCURED AS

Cloud AI/ML implementation services

AT A GLANCE

Model	Fixed-price, scoped per engagement
Typical	10–16 wks for a pilot; 6–9 months pilot-to-production
NAICS	541512, 541511, 541519, 541611, 541690
UEI	Y9MXMMHY3HA1
CAGE	242T0
Standards	NIST AI RMF 1.0, OMB M-25-21, NIST SP 800-53 Rev 5, FedRAMP-boundary delivery

WHO WE WORK WITH

- CAIO office / program owner scoping an AI use case against M-25-21
- Data/mission owner needing a pipeline connecting AI tools to real data
- ISSO needing AI governance documentation (model cards, monitoring plans)

COMMITMENT TO QUALITY

Avalon is led directly by a PMP-certified founder with a decade of enterprise engineering experience where every engagement gets senior-level attention, never a subcontracted bench. Delivery follows NIST SP 800-53, RMF, and FISMA, evidence-backed and control-mapped by design. ISO 9001/27001 is on the roadmap as the pipeline scales.

Status: Small business federal contractor based in Los Angeles, CA

NIST AI RMF 1.0

OMB M-25-21

NIST SP 800-53 REV 5

FEDRAMP-BOUNDARY DELIVERY