

Bring AI into the agency. Keep the data in the agency.

Enterprise AI substrate for federal civilian missions — deployable in AWS GovCloud, Azure Government, or on-prem, with FedRAMP-ready architecture and BYO LLM.

Every agency CIO wants to deploy AI. Every agency security officer requires FedRAMP. Every commercial AI product is years from FedRAMP-High — or doesn't pursue it at all. Sparcle is architected to be deployable in your existing sovereign cloud authorization boundary. You don't wait for our paperwork to start your pilot.

FedRAMP boundary

"Can we deploy this inside our existing ATO boundary?"

Sparcle: Yes. Bolt + Aeira deploy via Helm into your existing GovCloud or Azure Government Kubernetes cluster. The full software stack is open to your security team's review (source escrow at pilot tier). **Your existing ATO covers Sparcle; we do not require a separate FedRAMP authorization to begin operating.**

CUI + NIST 800-171

"How does this comply with NIST SP 800-171 for our contractor base?"

Sparcle: The 4-axis ACL with `country` · `office` · `role-band` · `employment-type` maps to the 800-171 access-control family. Documents pre-filter at index time; access decisions audit-log in a tamper-evident Merkle chain. CMMC-aware contractors adopt the same architecture in their enclave.

FIPS 140 cryptography

"Do we have to wait for FIPS 140-3 module validation?"

Sparcle: No, because Sparcle does not ship its own cryptographic module. KMS operations route to **your** validated module — Vault FIPS, AWS KMS FIPS endpoints, Azure Key Vault Premium, or PKCS#11 HSM. Your FIPS posture covers the crypto path; Sparcle is the orchestration layer.

What's shipped today vs what ships in a 90-day federal pilot

GovCloud / Azure Government deployable Helm chart	SHIPPED
AWS KMS + HashiCorp Vault + PKCS#11 HSM support	SHIPPED
Air-gap mode (no outbound traffic)	SHIPPED
4-axis ACL with <code>country</code> + <code>office</code> for CUI handling	SHIPPED
Sealed Merkle audit chain + standalone verifier	SHIPPED
BYO LLM — on-prem or sovereign-cloud inference	SHIPPED
Webhook-driven real-time ACL invalidation	PILOT
End-to-end ingest saga	PILOT
FedRAMP Moderate authorization (~12-18 mo with sponsor)	NOT HELD
FIPS 140-3 module validation (uses your validated KMS)	NOT HELD

Why federal civilian buyers choose us over Glean / Copilot

- Microsoft Copilot is in GCC-High but lock-in is total.** Microsoft-only sources, Microsoft-only inference, Microsoft-only roadmap. If your agency runs Workday + ServiceNow + Salesforce + Confluence on the side, Copilot doesn't help.
- Glean / Hebbia / Harvey are not in GovCloud.** They cannot authorize on your timeline. Earliest realistic GovCloud date is 2027+; you can pilot Sparcle in your enclave in 60 days.
- No vendor lock-in on the LLM.** Switch between Anthropic Claude on Bedrock, Llama 3.3 on AWS GovCloud, Azure OpenAI in Government, or on-prem — without changing your Sparcle deployment.

Recommended deployment: any self-hosted Bolt tier — **Authorize, Backbone, or Certify** — customer-hosted in GovCloud / Azure Government inside your existing ATO boundary. All are fully customer-hosted and send zero outbound traffic to Sparcle. There is no Sparcle-managed hosting tier, so no Sparcle-managed sub-processor is added in any tier.

Brief us in 30 minutes. Demo in your sovereign cloud in 60 days.

Sovereign enterprise AI — inside your perimeter, with your LLM, on your audit trail.

sparcle.app/federal · bolt@sparcle.app · v2 · 2026-07