Generative AI Integration for Enterprise Workflows

A structured approach to integrating generative AI capabilities into enterprise workflows, selecting appropriate models, establishing data governance, and deploying within security boundaries that protect sensitive information.

PRACTICE

Applied AI and Data

CLIENT PROFILE

Federal organization exploring generative AI capabilities for document analysis, content generation, code assistance, or knowledge management with appropriate security and governance controls

Challenge

- Interest in generative AI capabilities but uncertainty about secure deployment models
- Risk of sensitive data exposure through public AI services
- No framework for evaluating, governing, and monitoring AI model usage

Approach

- Assess use cases and classify by data sensitivity and deployment model requirements
- Deploy private LLM infrastructure or authorized FedRAMP AI services for sensitive workloads
- Implement retrieval-augmented generation (RAG) patterns for domain-specific knowledge
- Establish AI governance policies for acceptable use, monitoring, and responsible AI principles

Typical Outcomes

- Secure, private AI capabilities deployed within authorized boundaries
- Measurable productivity improvements in targeted workflows
- Governance framework enabling responsible AI adoption across the organization

Procurement Paths

- NASA SEWP V and SEWP VI for AI/ML hardware and cloud services
- GSA MAS for AI engineering and advisory services
- CIO-CS (NITAAC) for enterprise AI platforms

Compare vehicles and ordering steps on **How to buy**.

Partner Technology Examples

- NVIDIA
- Microsoft Azure (OpenAI Service)
- AWS GovCloud (Bedrock)
- Elastic

FEATURED PRODUCT

Odin's Edge: the operated site

Norseman runs a secure cloud at your site. You run the mission.

Odin's Edge is Norseman's managed edge platform. It puts compute, storage and network at the site under one controlled architecture, on the hardware you already run, integrated and tested in our lab before it ships. Applications arrive hardened, with the configuration baselines and test reports your security team needs. Norseman operates it around the clock or hands it to your team with the same console. Buy it as one line on your vehicle or as a per site subscription.

The same site runs the AI when the mission is ready.

Frequently Asked Questions

Can federal agencies use generative AI without sending data to commercial cloud providers?

Yes. Norseman specializes in private, air gapped large language model (LLM) deployments that keep all data within the agency boundary. These run on dedicated NVIDIA GPU infrastructure and never route queries or data to commercial AI providers.

How is generative AI procurement handled for federal agencies?

AI infrastructure, particularly GPU compute, is procurable via NASA SEWP V and VI and ITES-4H. Professional services for AI integration are available via GSA MAS and CIO-CS.

What is Odin's Edge and how does it relate to AI?

Odin's Edge is Norseman's managed edge platform. Norseman runs a secure cloud at your site. You run the mission. It delivers the AI a site needs, running at the site on the hardware the agency already runs, for sites where the cloud is not an option.

This is a representative use case. To request a briefing or a quote, contact Norseman Defense Technologies at 410.579.8600 or sales@norseman.com.