AI Security Overview

- **Enterprise vs. public AI:**
  Use **enterprise AI tools** for any confidential or proprietary data. Public/consumer AI tools are suitable only for non-sensitive drafting or research.
- **Permission-based access:**
  Secure enterprise AI should respect existing **user permissions and access controls**—AI can only see what the user is authorized to see.
- **Data ownership & privacy:**
  In enterprise deployments, **customer data stays within the company boundary** and is **not used to train public models** by default.
- **Biggest real risk:**
  Not the AI model itself—but **overshared data, weak permissions, and lack of governance**. AI simply exposes what's already accessible.
- **Tool guardrails:**
  Require SSO, encryption, audit logs, retention controls, and clear policies on what data can be prompted into AI tools.
- **Private data lakes & RAG:**
  A **private data lake** keeps operational data under company control.
  Retrieval-Augmented Generation (RAG) lets AI reference that data without copying it into a model—improving accuracy while maintaining security.
- **Scaling safely:**
  Start with low-risk internal use cases, prove value, then scale only after **governance, labeling, and access controls** are in place.