

How a major regulated financial institution built confidence in generative AI

Hands-on demonstrations and a secure implementation model accelerated software delivery while maintaining governance and a consistent user experience.

In financial services, generative AI adoption requires more than faster code generation. Security, data protection, governance, operational reliability, and professional accountability must be addressed before AI becomes part of everyday software delivery.

A major wholesale financial institution initially limited access because stakeholders were concerned about how proprietary code, business requirements, and sensitive information might be processed. OGI replaced abstract discussions with hands-on demonstrations based on real modernization challenges.

CHALLENGE

Building trust in a highly controlled environment

Stakeholders needed clear answers about data processing, protection of proprietary code and business data, access controls, code review, and accountability for the final implementation.

Development teams also needed to modernize aging frameworks, upgrade back-end and front-end components, address technical debt, and recover time invested in a portal that had not delivered the expected functionality. The challenge was to demonstrate measurable value within the controls required of a financial institution.

SOLUTION

Hands-on demonstrations in a secure AI environment

OGI took a hands-on adoption approach, using cases from the institution's active technology backlog. The first demonstration showed how generative AI could help modernize applications built on an older enterprise development framework and related components.

Using existing requirements and interface designs, a second demonstration rebuilt a representative portion of a portal that had not delivered the expected functionality. The result included business rules, application behavior, back-end functionality, database connectivity, and the same consistent user experience.

Stakeholders observed the full workflow, from analysis and prompt preparation through code generation, review, testing, and refinement by experienced developers. OGI also showed how an approved enterprise AI environment could limit model context, set boundaries around data processing, manage access, and support monitoring and governance.

RESULT

From experimentation to responsible adoption

An internal comparison indicated that work resembling an earlier effort of several months could be recreated as a functional implementation in approximately two weeks. In one modernization initiative, AI-assisted development reduced the estimated timeline from eight weeks to four, an estimated 50% reduction.

The 50% figure is a project-level estimate, not an institution-wide benchmark. Access expanded through established approval processes, and AI became part of daily delivery. Some teams requested more capacity after reaching their allocated usage limits. A data project that had assumed AI was prohibited later incorporated the approved environment.