devoteam
Cyber Trust

Case Study

Securing AI in Financial Services: LLM Chatbot Penetration Testing for a Leading Bank

Client

A global financial institution deploys an AI-powered financial literacy chatbot embedded into a customer-facing website and commissions a full LLM security assessment before go-live.

The Challenge

- **AI introduces a new class of risk in regulated environments:** Deploying an LLM-powered chatbot in a financial institution creates attack surfaces that existing security frameworks were not built to evaluate and that regulators are increasingly scrutinising.
- **Public-facing deployment amplifies exposure:** The chatbot is designed to be embeddable across any website , including sites outside the institution's control. This creates a unique risk: the chatbot can be loaded and exploited from arbitrary third-party contexts, making abuse prevention, origin validation, and output integrity critical security requirements.
- **AI risk sits on top of existing infrastructure risk:** LLMs do not replace legacy vulnerabilities, they add to them. A financial institution serving 3.6 million customers across 500+ domestic branches and 23 countries cannot afford blind spots at any layer of its stack.

The Solution

- **Devoteam Cybertrust AI Pentesting methodology:** A robust, structured approach grounded in MITRE ATLAS, OWASP LLM Top 10 and OWASP Agentic Top 10 — the industry's most comprehensive frameworks for adversarial AI risk, ensuring systematic coverage of every known AI attack category.
- **AI + Cloud + Web Application: a combined engagement:** The chatbot deployment integrated the corporate website and cloud backend. Devoteam Cyber Trust Cloud and Web Application Pentesting, applied alongside the AI Methodology, was crucial: cloud analysis revealed a token disclosure that exposed credentials, risking full cloud environment compromise.
- **Enabling trusted digital innovation:** By securing both the AI layer, the cloud infrastructure and the Web Application before go-live, the engagement enables the institution to extend financial literacy services to millions of customers, protecting end-users, cloud assets, and the institution's regulatory standing simultaneously.

Results

Unbounded Consumption
No rate limiting on token usage or input length attacker could degrade availability across all embedded instances
Critical risk in a public-facing, embeddable deployment

Cloud Token Disclosure
Cloud credentials exposed through the website–chatbot integration layer, potential full cloud environment compromise
Discovered through Devoteam Cyber Trust Cloud Pentesting applied to the integration layer

Chatbot Hijacking
The embeddable nature of the chatbot allowed it to be loaded and exploited from arbitrary third-party websites outside the institution's control
An attacker could embed the chatbot in other websites with financial impact to the client

Pre-production discovery
All critical findings identified and remediated before public launch
Zero exposure to live customers or regulatory consequences at go-live