Mashreq Bank Bounty Starter Pack

This document outlines topics, tools, and general tips and tricks that participants are encouraged to familiarise themselves with before the competition, should they choose to do this bounty.

This document does not describe the bounty. You're going to have to wait till the event for that 😀.

Difficulty Level

Difficulty: Intermediate (Conceptual and design-focused)

This bounty is approachable for participants with basic programming knowledge but requires participants to have an understanding with concepts like AI reasoning, explainability, and governance by the end of the bounty.

Recommended Team Composition

Teams with at least one or two members with basic programming knowledge are recommended.

Multidisciplinary teams are welcome to take on this challenge, provided the team includes at least one technically familiar member to support implementation and prototyping.

Participants from backgrounds such as Computer Science, Engineering, Business, and Finance may all contribute meaningfully to different aspects of the solution.

Topics encouraged to look at

1. Responsible AI and Governance

Participants are encouraged to explore established frameworks for managing AI risks, specifically focusing on transparency, accountability, and fairness.

Resources:

- Microsoft Responsible AI: https://www.microsoft.com/en-us/ai/responsible-ai
- Google Cloud Responsible AI: Responsible AI Best Practices

2. Human-in-the-Loop Design

Participants are encouraged to explore design patterns where AI acts as a signal-triaging assistant rather than a fully autonomous decision-maker.

They should understand how to design interfaces that allow human operators to validate, override, or contextualize AI insights during a crisis.

Resources:

- LangChain Documentation: Human-in-the-Loop
- Stanford HAI: Humans in the Loop: Design of Interactive AI

3. Synthetic Data Generation

Participants are encouraged to explore technical methods for procedurally generating datasets if data sources are for some reason not available.

Resources:

- OpenAI Cookbook (Data Generation): Synthetic data generation with OpenAI
- LangChain Documentation: Data Generation Use Case
- Faker (Python Library): Faker Documentation
- Hugging Face: Synthetic Dataset Generation Guide

4. Agentic Workflows

Participants are encouraged to explore how to design AI systems that can plan, reason, and execute multi-step tasks, rather than just generating a single text response.

Resources:

- DeepLearning.AI: What are Agentic Workflows?
- LangChain Documentation: Agents & Tool Use
- OpenAI API: Function Calling Guide

5. Operational Resilience in Banking

Participants are encouraged to explore what "resilience" means for banks.

They should understand the difference between "noise" (minor complaints) and "signals" (systemic risks like liquidity rumors or service outages)

Resources:

- Bank of England: Operational Resilience Policy
- Basel Committee: Principles for Operational Resilience

General Tips and Tricks

- This bounty will reward participants who design for safety and explainability. A solution that handles 50 inputs with perfect ethical guardrails is better than one that processes 1,000 inputs with no explanation. So think of ways to test programs.

- AI-assisted development tools can be useful for ideation, exploration, and implementation.

- If using AI powered tools that have access to your repository, always use plan mode and reiterate the plan before you let the agent execute the commands.

- If you have a team from beforehand, set up the repository and look into git commands in case you are using version control. https://education.github.com/git-cheat-sheet-education.pdf