MantisGrid AI

2026 MantisGrid AI Hackathon — Theme

About the Challenge

AI is evolving rapidly, but great AI products require more than powerful models—they need intelligent infrastructure.

This hackathon challenges you to build the technologies that make AI systems smarter, more reliable, and easier to operate. Whether you're uncovering insights from infrastructure data, designing intelligent model routing, or creating observability tools with OpenTelemetry, your goal is to solve real-world problems with practical, impactful solutions.

We're not looking for another AI demo. We're looking for projects that improve how AI systems are built, deployed, and managed—solutions that engineers and organizations would actually use.

Choose one of our challenge tracks or combine ideas across them. Build something innovative, demonstrate strong technical execution, and show how intelligent infrastructure can power the future of AI.

If your project makes AI systems faster, more efficient, or easier to understand, you're building exactly what this challenge is about.

Theme 1 — Dataset Track: "Know Your Infrastructure Data"

Your mission is to unlock hidden insights from real observability data.

Pitch: Start from raw, messy infra data and turn it into insight.

What you build: Explore, clean, and understand a real infrastructure dataset — surface patterns, engineer useful features, or build a dashboard / narrative that reveals what is really happening inside the system.

Example directions:

- An interactive dashboard highlighting the most interesting signals in the data.
- A "data story": patterns, correlations, or surprises in real infra metrics.
- A reusable, cleaned feature set other teams could build on.

Best for: Data-curious folks, analysts, and beginners — lowest ML barrier.

MantisGrid AI

We provide: Suggested public infrastructure datasets+ a starter notebook.
Alibaba GPU cluster trace

Judging focus: Quality of data work, clarity of insight, storytelling.

MantisGrid tie-in: Good data is the foundation of any intelligent-infrastructure product.

Theme 2 — Model Routing Track: "Build the Brain Behind AI Applications"

(Policy-Aware LLM Routing)

Your mission is to build an intelligent routing framework that makes better decisions than today's simple LLM routers.

Pitch: Build an intelligent **routing system** that decides which model, tool, or workflow should handle each request — driven by policy and context.

Starting point: Something like **vllm-semantic-router** (vllm-semantic-router.com) is a good baseline — participants are encouraged to invent more interesting scenarios and enhancements on top of it.

What you build: A policy-aware routing engine that analyzes each request and picks the best response path. Possible routing decisions include:

- Route to a small vs. large model.
- Route to a local / private vs. cloud model.
- Failover to different models based on availability, performance or budget
- Route to a multimodal model when images are detected.
- Trigger Retrieval (RAG) before generation.
- Escalate to a human when confidence is low.
- Invoke external tools or APIs when necessary.

Example directions:

- A semantic router that classifies requests by user intent and context.
- Routing policies balancing cost, latency, privacy, and quality.
- Adaptive routing that learns from user feedback.
- An interactive dashboard that visualizes and explains routing decisions.
- A comparison of routing strategies and their trade-offs.

MantisGrid AI

We provide:

- Starter notebook and routing template.
- Sample prompts and workloads.
- Example routing policies for baseline implementations.
- Optional access to MantisGrid APIs.
- vLLM Semantic Router

Judging focus: Quality of routing logic, technical implementation, creativity and policy design, real-world usefulness, and explainability of routing decisions.

MantisGrid tie-in: Intelligent model routing is a core infrastructure capability that improves cost, latency, reliability, privacy, and user experience — exactly the "efficient AI infrastructure" MantisGrid is built around.

Theme 3 — OpenTelemetry Track: "Build on the Open Standard"

Your mission is to help engineers understand infrastructure faster.

Pitch: Build something intelligent on top of OpenTelemetry — the open, industry-standard way infrastructure reports its health.

Background: OpenTelemetry (OTel) is the open-source, vendor-neutral standard for collecting telemetry — **metrics, logs, and traces** — from software and infrastructure. Most modern observability stacks are built on it, which makes it realistic and directly relevant to MantisGrid.

What you build: Use real OpenTelemetry telemetry to build a monitoring tool, anomaly explorer, an AI copilot that helps engineers *explore and debug* incidents, a natural-language "ask your infrastructure" interface, or your own idea. (Keep agents in a *human-assist* role — explain, explore, suggest — rather than claiming to predict failures.)

Example directions:

- An AI copilot that helps an engineer explore OTel traces and debug an incident faster.
- "Ask your infrastructure a question" — a natural-language layer over OTel data.
- A tool that turns raw traces into a clear root-cause or cost-saving suggestion.

Best for: Teams who want something close to real production systems — and product / non-ML builders who prefer to build on a standard.

We provide: Sample OpenTelemetry data (metrics / logs / traces) + an ingestion template + an optional MantisGrid API hook. Open-Telementary Demo

MantisGrid AI

Judging focus: The idea, a "wow" demo, and real-world usefulness on a real standard.

MantisGrid tie-in: OpenTelemetry is the backbone of modern observability — building on it shows exactly where MantisGrid adds value.

Hackathon Coding Rules:

1. Participants may use any AI model, large language model (LLM), AI coding assistant, agent framework, or open-source AI tool during the hackathon.
2. Have FUN

What To Submit:

1. Team information
2. Project title and description.
3. Source Code Repository,
4. 2 minute demo.

Awards

Best Overall Project — Strongest combination of innovation, execution, impact, and presentation.

Most Innovative — Most original or technically creative solution.

Best Business Impact — Best solution to a meaningful customer or market problem.

Best User Experience — Most polished, intuitive, and user-friendly product.

Best Use of MantisGrid — Best integration of MantisGrid technology or infrastructure concepts.

Judging Criteria

- **Technical Execution (30%)**

Does the project demonstrate quality application development?

Is the code, architecture, or technical implementation functional, reliable, and credible?

MantisGrid AI

- **Innovation / Wow Factor (30%)**

How novel and original is the idea?

Does it address a significant problem or create a unique solution?

- **Potential Impact (20%)**

How big of an impact could the project have in the real world?

How useful is it to a broad market or important user group?

- **MantisGrid Product / AI Infra Relevance (10%)**

Does the **project meaningfully incorporate MantisGrid's product,** API, platform, or infrastructure concepts?

Does it **align with the AI infrastructure theme** and show **how MantisGrid can support** real-world AI applications?

- **Presentation / Demo (10%)**

Is the **problem clearly defined** and is the **solution effectively presented**?

Is the demo **easy to understand** and **convincing**?