



Zindi's First Agentic AI Challenge

AI for Good Global Summit 2026 | Paul Kennedy



Zindi is the leading data science and AI competition platform focused on Africa and other developing and emerging markets. We build skills and spark innovation to solve real-world problems.

120K+
users

28%
women

190+
countries



ZINDI AT A GLANCE



500+

AI problems
solved



120K+

AI builders
served



28%

female
participation



\$1.5M

paid to Zindians



450+

universities
supported



200+

Zindi
ambassadors

ZINDI ITU PARTNERSHIP



31

AI challenges
(and counting!)



15,000+

participants



\$200K+

invested in
winners



Telco Troubleshooting Agentic Challenge by ITU

€40 000 EUR

Completed (~2 months ago)

1402 joined 354 active

Skills you will learn

- Agentic AI
- Fine-tuning
- Large Language Models

Info Data Leaderboard

Accept terms

Start ⓘ

17 Apr 26

Close ⓘ

19 May 26

Reveal ⓘ

03 Jun 26



Can you build an AI agent that understands and optimizes telecom networks?

Note Phase 2 has begun! Find extra information about participating in Phase 2 [here](#)



Telco Troubleshooting Agentic Challenge – Global Launch
AI for Good

Beyond prediction: we asked teams to build agents that go beyond predicting faults: they needed to investigate it, check the troubleshooting rules, and choose the fix.

Find the likely fault

Apply the correct rules

Fix with suggested actions



How the challenge was designed



shared base
model

- Qwen3.5-35B-A3B
- Fine-tuning is allowed
- Emphasises agent design



network
tracks

- Wireless/5G networks
- IP networks



challenge
phases

- agent design
- testing on new data
- final evaluation

Two tracks

	Track A: Wireless	Track B: IP
Focus	Diagnose faults in wireless / 5G cells	Troubleshoot multi-vendor IP networks
Tasks	Single- and multi-answer questions	Open-ended answers
Data	2,000 train + 500 test questions	50 → 100 → 70 questions
Tools	Agent Tool Server simulation APIs	Simulated CLI: Huawei / Cisco / H3C
Scoring	Intersection-over-union	Exact match

A three-phase format



Design



Generalise



Final run

The agent's toolbox



Simulation

query a network
sandbox



Memory

recall past cases



RAG

fetch the right
SOP



Skills

tools for real metrics



Fine-tuning

adapt the base model



CoT

reason step by step

The data & environment

Element	Detail
Base model	Qwen3.5-35B-A3B — fine-tuning allowed, no architecture or scale swaps
Track A data	2,000 labelled cases + 500-question test sets
Track B data	Open-ended tasks, released in batches
Environment	Agent Tool Server (server.py) exposes the network simulation APIs
Public dataset	Hosted on Hugging Face under CC-BY-SA 4.0
Signals	KPIs, cell configuration, signalling and handover logs

Scoring & efficiency

Answering time

under 5 min

5 – 10 min

10 – 15 min

over 15 min

Score kept

100%

80%

60%

0%

Phase 3 multiplies each correct answer by a time discount, so fast, correct agents win.

Prizes

€12.5K

€5K

€2.5K



AI for Good



HUAWEI

GSMA™



The challenge by the numbers



1,400

users engaged



349

active builders



109

countries
reached



3,949

submissions



12K+


hours of DS work



42

African countries

Top 10 countries by number of participants



The winners

	Track A: Wireless
1st	tel_number_one
2nd	WILL
3rd	TrainyMcTrainFace

Track B: IP
SLAI BOYS
Cisco_Team
bamps53

What we learned



Design over size



Built to generalise



Agents can
do it



Lessons learnt

Design over size: a shared base model moved the competition to agent design rather than model scale.

Built to generalise: the multi-phase format plus the Phase-3 time penalty rewarded agents that stay accurate on unseen data and run efficiently

Agents can do it: the top solutions show that agentic AI can carry out genuine telecom troubleshooting, not just classify examples.

Prize power: The high prize pool attracted a large competitor community and brought new competitors to Zindi.

ZiND!



Thank you!



Antonio de Domenico
Huawei



Megan Yates
Zindi

