HUMANIS

PRESS RELEASE

Humanis launches Crucible to verify the actions of AI agents on mobile networks

The platform enforces across all agents the operating envelope set by the operator, marking a milestone towards Level 4 automation and safe closed-loop agent optimisation.

BANGKOK, 17 November 2026 – Humanis is launching Crucible, the verification-centric AI platform, today at TM Forum Innovate Asia, giving mobile operators a verification layer they can trust that sits between their closed-loop AI agents and their networks. For every agentic action, Crucible assesses the outcome of proposed changes using a model of the operator's network and checks the results against limits set by the operator before allowing execution.

The launch addresses a key barrier to autonomous network operations: confidence in the decisions that AI agents make. In TM Forum research released ahead of the event, just 2% of Asia-Pacific respondents reported scaling AI agents into core systems for live operational tasks.

"AI offers tremendous potential but also comes with risks. Verification is the key technology needed to bring key AI opportunities ranging from network slicing to agentic energy saving to billions of subscribers globally," said Sy Campos, co-founder and CEO of Humanis. "That means enforcing operating limits across the network, driving decisions across silos in the organisation and retaining evidence that engineers can examine. Crucible provides that control layer between intent and execution."

How Crucible works

Crucible uses a world model, a low-latency digital twin calibrated on the operator's network performance counters that predicts the consequences of changes instantiated by agents and humans alike. It forecasts a range of likely outcomes for the affected cells and services. Changes that pass the checks are deployed on the physical asset. Crucible's gateway blocks changes predicted to breach a limit and holds for human review those for which uncertainty prevents a decision.

The same checks apply to all agents, whether they are built by operators, equipment vendors or Humanis. Crucible runs in the operator's own environment on standard CPUs. No subscriber identifiers are required, and no language model is used in verification, which keeps every decision transparent to the user.

Demonstration at Innovate Asia

Humanis is demonstrating Crucible's gateway in a controlled test on a simulated network. Several agents are taking action on a network: an energy-saving agent, a load balancer and a network-slicing agent. The world model balances the suggestions from all agents to enforce the operator's priorities consistently and gives the operator a view of the trade-offs between its objectives. At one point, an attacker takes control of the network-slicing agent and attempts to cut a slice's radio resources to zero. The world model predicts a breach of the slice's service-level agreement, and the gateway blocks the change immediately. Repeated attempts lead to the agent being quarantined, while the network stays within the limits set by the operator.

Availability

Crucible is available now for mobile operators. An initial assessment estimates the potential gains beyond the operator's existing policies, using four weeks of anonymised historical network data. No integration is required at this stage. During the pilot phase that follows, Crucible runs in shadow, or advisory, mode on a test cluster: it recommends changes, and the operator's engineers decide whether to apply them. As trust increases, the operator can extend Crucible to more areas and more decisions, towards full coverage and higher levels of automation.

Humanis also provides agents for energy saving, coverage and capacity optimisation, and network slicing. Crucible's gateway arbitrates conflicts between agents and protects the network against compromised agents.

To mark the launch, Humanis will credit 50% of pilot fees against the fees for a subsequent operational deployment. The offer applies to the first three operators that register for a pilot during Innovate Asia, 17–19 November 2026, and sign a pilot agreement by 1 March 2027. Subject to contract.

At Innovate Asia

Sy Campos joins the panel "How AI can transform BSS/OSS" on Tuesday 17 November, moderated by TM Forum's Andy Tiller. The panel is part of the 16:00–17:10 session on the Impact Stage.

Humanis will demonstrate Crucible at booth 19 throughout the event. Meetings during the event can be booked at calendly.com/simeon-campos/innovate-asia-tm-forum, and afterwards at crucible.ag/contact.

ENDS

About Humanis

Humanis AI Ltd develops world models and verification technology for AI agents operating critical infrastructure, starting with telecom networks and critical manufacturing. Building on more than $70m of R&D in ARIA's Safeguarded AI programme, whose founding programme director is now Humanis's Chief Scientist, the founding team combines leading research in AI safety, leadership roles and decades of experience in safety-critical industries. Crucible is its platform for mobile operators. Based in London, Humanis operates an information security management system certified to ISO/IEC 27001:2022. More at crucible.ag.

Media contact

Sy Campos, CEO, Humanis
s@humanis.ai | +33 7 82 78 14 82