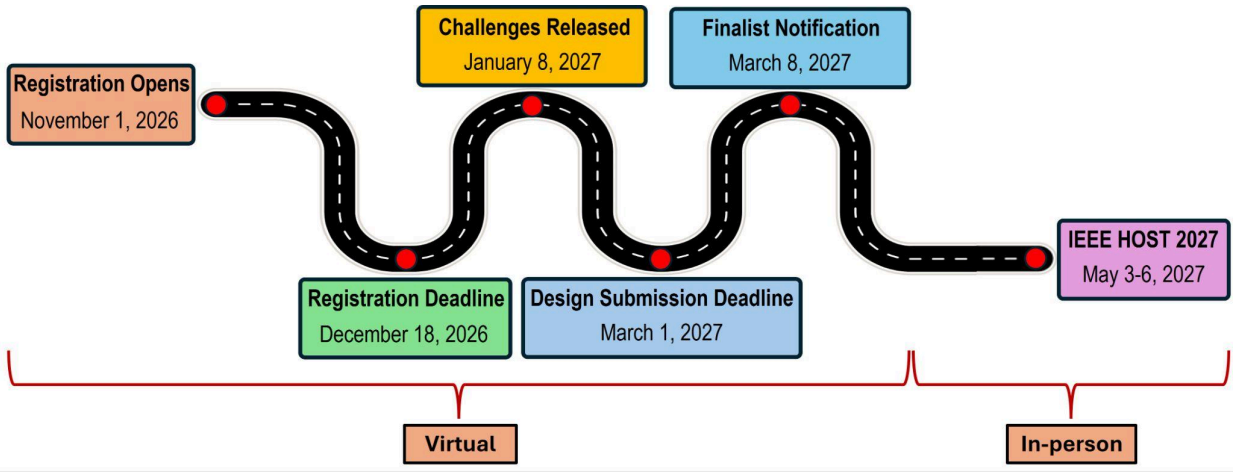
GhostGate: HOST 2027 Hardware Hacking Call for Participation

Agentic Hardware Trojan Insertion, Assurance, and Mitigation Challenge

GhostGate is the Hardware Hacking Competition at IEEE HOST 2027. It puts agentic AI on both sides of hardware security. Insertion agents hide hardware trojans and hard-to-detect functional bugs in large, complex, open-source chip designs. Assurance agents prove what is correct, detect what is not, and return machine-checkable evidence. Remediation agents trace failures to root causes and verified fixes. Everyone uses the same designs, sandbox, and budgets. The organizers' detection system, the Oracle, evaluates every insertion and is the public baseline for participant agents.



Online rounds use the competition platform. The live final is co-located with IEEE HOST 2027, in person with remote participation for finalists who cannot travel.

Register and prepare
Registration opens **November 1, 2026,** and closes **December 18, 2026**. The challenge will be released on **January 8, 2026,** and finalists will be announced on **March 8, 2027**. In-person finals will be held at IEEE HOST 2027.
Registration is required. Sign up through the portal, complete your details, enroll in "**GhostGate, HOST 2027**", and select one or more tracks.

Register for GhostGate HOST 2027

Registered participants receive the practice modules on January 8, 2027. It includes a practice design, sandbox image, local validity checker, agent harness with a reference agent, and worked examples for every track.

Designs and tools
Designs are released one round at a time and are identical for all participants. They include an IP block with its bus interface at register-transfer level (RTL), an SoC subsystem at RTL with core, bus fabric, peripherals, memory models, and the chip filelist, and a synthesized gate-level netlist. Each release includes the specification, filelist, regression suite, public formal properties, declared threat model, a local validity checker, and worked examples for every track. A Docker image provides the open-source toolchain, including Verilator, Yosys, cocotb, SymbiYosys, and OpenROAD.

Eligibility and publication
Participation is open globally. Insertion is individual only, with one person, one account, and one leaderboard entry. Assurance and Remediation accept individuals or teams of up to four. Organizers, their students, and Oracle team members are not eligible. Any tools, models, or agent frameworks are allowed, including licensed or open-source EDA tools, coding assistants, language models, and agents. Use the organizers' model allowance or your own API keys. After the final, every valid insertion, assurance package, and patch is published with its evidence under an open license and with the authors' names. Registration accepts this condition.

Insertion and Assurance

Track 1 Insertion Ghosts

Insertion is limited to individuals, with up to five submissions per round. Submit a modified design and an encrypted envelope recording the trojan class, trigger, payload, changed lines, and a witness test that passes on the modified design and fails on the original. The Gate submission portal runs six automated validity checks. These verify that (1) public regression passes unchanged, (2) public formal properties hold, (3) cell area stays within 1% of the original, or 0.5% at netlist, with no new ports, (4) the design is equivalent to the original outside the trigger, (5) witness behavior matches the declaration, and (6) the envelope is well formed. Valid submissions undergo five Oracle search stages with fixed budgets. From Round 2, the three highest-ranked Assurance agents of the previous round also search them. The Oracle works blind. It receives the modified design, specification, and public tests, never the original, a diff, or the envelope.

Insertion outcome	Base score
Detected by lint or public tests	0
Survives 30-minute automated search	10
Survives four-hour extended search	30
Survives expert review at the final	60

Partial credit reflects the fraction of a stage that survived. The base score is multiplied by impact, realism, and stealth. Repeated ideas in the same class and module count less each time.

Track 2 Assurance Wards

Individuals or teams of up to four may enter. Submit an autonomous agent container implementing the published harness interface. Each round supplies a design directory, specification, public tests and properties, threat model, and a budget of four hours of wall time on sixteen cores with a model-spend allowance. The design may be clean, contain a change request relative to the previous round, or contain one to three valid insertions from the previous round. The agent is not told which. Return an assurance package with agent-written SystemVerilog assertions and a proof log for each. Include a mapping from every threat-model clause to properties that discharge it, equivalence results where a reference exists, and coverage from the agent's own stimulus. Return a verdict of certify, or quarantine with a localized finding and reproducing stimulus. The harness replays every claim.

Assurance component	Points	Scoring rule
Detection and localization	40	Detect the insertions present. False quarantine of a clean design costs as much as a miss.
Proof value	30	Only properties confirmed meaningful by mutation analysis and discharging threat-model clauses count.
Evidence integrity	15	Any claim that fails to replay zeroes this component.
Cost	15	Position on the round's score-versus-cost front; harness measures tokens, wall time, and CPU hours.

Remediation and the live final

Track 3 Remediation Rites

Individuals or teams of up to four submit an agent container using the published harness interface. Given a design with one valid insertion and a failure signature (such as a failing assertion, waveform, or register value at a specific cycle), the agent returns the root cause (file, lines, class) and a patch. The harness verifies that public tests and properties pass, the insertion witness fails, the patch is equivalent to the original outside changed lines, and area stays within budget. Scoring: localization 40 points, fix validity 40, and cost 20. Assurance autonomy multipliers apply.

Autonomy for Assurance and Remediation

Declare the autonomy level. The multipliers are ×1.0 for fully autonomous operation, ×0.85 for up to three logged human interventions, and ×0.6 for interactive operation. Assurance allows single-agent and multi-agent architectures. Multi-agent submissions declare their roles and protocol, and the harness logs every inter-agent message.

Feedback and leaderboards

Within a few days of each round, every participant receives a full evaluation of each submission. This includes validity, the stage or agent that detected an insertion, replay results for each assurance-package claim, measured cost, and the evidence used. Each track has a separate leaderboard. Agent leaderboards show the score-versus-cost front with the Oracle as the baseline row. Round scores are rank-normalized. The season score is the sum of the three rounds and the final.

Live final at IEEE HOST 2027

The ten highest scores in each track are invited. Finalist Insertion participants hide trojans in a design revealed on site within a four-hour window. Finalist Assurance agents run on those designs on screen under one clock, and finalist Remediation agents receive the signatures of what was found. Two mentors review what neither side settled.

Awards and recognition

The top three designs as well as "Best in Track" designs receive cash awards and formal recognition at HOST 2027.

Place	Cash award per track in USD
1st	**$1,500**
2nd	**$1,000**
3rd	**$500**

The participant or team with the highest combined ranking across any two tracks receives the Gatekeeper title and a separate award. Any insertion surviving all five Oracle stages and every finalist Assurance agent at the final receives the Undetected title and a separate award. All finalists receive certificates.

Organizers and contact

- Soheil Salehi, University of Arizona (ssalehi@arizona.edu)
- Biswajit Ray, Colorado State University (biswajit.ray@colostate.edu)

Competition Mentors

- Parsa Mirfasihi, University of Arizona (parsamirfasihi@arizona.edu)
- Muhtasim Alam Chowdhury, University of Arizona (mmc7@arizona.edu)
- Sujan Ghimire, University of Arizona (sghimire@arizona.edu)
- Harish Kumar Dharavath, University of Arizona (harrydhara16@arizona.edu)

Industry Sponsors:

- PilotCrew.AI: https://pilotcrew.ai/
- ChipSageLabs.AI: https://www.linkedin.com/company/chipsage-labs