28 September 2026 · Singapore

Call for Papers

OASec 2026

OASec
2026 · SINGAPORE

Overview

OASec 2026 is an open forum on AI safety and security in Asia Pacific, taking place on **Monday, 28 September 2026 in Singapore**.

We invite researchers, CISOs, policymakers, and practitioners to submit talks on how AI systems are built, deployed, governed, attacked, and defended. Topics may span offensive research, defensive techniques, enterprise deployment, governance, and emerging risks across foundation models, copilots, autonomous workflows, and agentic systems.

We welcome original research, technical deep dives, field reports, live demos, and open-source tooling, with a particular interest in talks grounded in hands-on experience, practical lessons, and real-world safety and security outcomes. Sales pitches, vendor marketing, and solution showcases are out of scope for CFP talks. Whether you have a new attack technique, a production case study, a governance insight, or a tool to share, we want to hear from you.

What We're Looking For

- **Practical and technical depth** — real-world experience, working demos, and actionable takeaways are strongly preferred over high-level overviews
- **Original work** — novel research, new tools, fresh perspectives, or unique case studies
- **Diversity of perspective** — we encourage submissions from academia, industry, government, and independent researchers
- **Regional relevance** — regional insights, regulatory developments, or deployment stories are a plus (but not a requirement)

Submission Tracks

We are accepting submissions across four tracks. Session-to-stage mapping will be determined after the CFP closes based on the submissions received.

Talk format: 20 minutes total, with 15 minutes for presentation followed by 5 minutes of Q&A.

Enterprise: Real-world adoption of AI in production. We welcome talks on enterprise use cases, architecture patterns, deployment strategies, integration challenges, evaluation and reliability practices, lessons learned, and case studies from operating AI systems at scale. Topics may include internal copilots, AI assistants, SOC workflows, developer platforms, and agentic AI in business-critical environments.

Governance: The policies, controls, standards, and organizational frameworks needed to deploy AI responsibly and securely. We welcome talks on AI regulation, compliance, risk management, assurance, supply chain risk, access control, data governance, and operating models for GenAI and agentic AI. This track also includes discussion of AI safety versus AI security, along with standards and risk frameworks such as the OWASP Top 10 for LLMs and Agentic Applications, which help organizations identify, prioritize, communicate, and mitigate risk in AI systems.

Offense: Research, vulnerability disclosures, and demonstrations showing how AI systems can be attacked, manipulated, or made to fail. Relevant topics include AI red teaming, jailbreaks, prompt injection, tool abuse, goal hijacking, vulnerability discovery, model misuse, and attack paths targeting both AI models and agentic applications.

Defense: Techniques, solutions, and operational practices for securing AI systems. We welcome blue team talks and research on guardrails, policy enforcement, evaluation and testing, monitoring, detection, incident response, SOC agents, threat hunting, secure-by-design patterns, and defensive automation for both AI models and agentic applications.

Please note: CFP talks must be educational and experience-based. Vendor marketing, sales pitches, and solution-led presentations are not accepted.

Submission Guidelines

Each submission should include:

- **Title**
- **Abstract** (max 300 words)
- **Proposal Details** (max 1,000 words total, Markdown or PDF) covering:
  1. **Detailed Outline** — Major topics/sections of the talk with approximate time allocation across the 15-min presentation slot
  2. **Key Takeaways** — 3 actionable takeaways the audience can apply
  3. **Novelty & Relevance** — What is new or original about this work? Why is it relevant to AI safety and security practitioners now?
  4. **Target Audience & Prerequisites** — Who benefits most? Any assumed background knowledge?
  5. **Disclosure & Prior Presentation** *(if applicable)* — Has this content been presented elsewhere? If disclosing a vulnerability, what is the current disclosure status?
  6. **Demo Details** *(if applicable)* — What will be demoed? Any special AV or network requirements?
  7. **Vendor Neutrality Statement** *(if applicable)* — If your employer sells products/services in this space, briefly explain how the talk will remain educational and vendor-neutral
- **Speaker bio** (up to two speakers)
  - LinkedIn (required)
  - GitHub (optional)
  - Any other URLs (open-source projects, hacking competitions, etc.)
- **Track** (Enterprise / Governance / Offense / Defense)
- **Demo included?** (yes/no)

Multiple submissions are welcome — each speaker may submit up to **3 proposals**. Each proposal must be submitted separately.

Important Dates

Milestone	Date
CFP opens	1 May 2026
CFP closes	1 July 2026
Review period	1–15 July 2026
Acceptance notifications	16 July 2026
Speaker confirmation deadline	23 July 2026
Final titles & abstracts due	20 August 2026
Agenda published	1 September 2026

How to Submit

Submit your proposal at: `https://oasec.org/`

For questions, contact: Alex Leung <alex.leung@aift.io> and Kentaroh Toyoda <kentaroh.toyoda@aift.io>

Organizers

VULCAN

AI.f(T)