ICOA 2026

Official Rulebook

Competition Rules, Scoring, and Awards

International Cyber Olympiad in AI

Sydney, Australia · 27 June – 2 July 2026

v1.0.0 · 25 May 2026 · English · Final

Audience: ICOA 2026 contestants, national coaches, country team leaders, accreditation bodies.

Scope: Competition rules — tracks, scoring, medals, environment, eligibility, conduct.

Out of scope: Training curriculum (see *ICOA Training Book*), authoring (see *Author Team Handbook*), national selection (see *Selection Handbook*).

Published by ASRA — Australia STEM & Robotics Advancement Association Inc.

Preface

This rulebook is the canonical statement of how ICOA 2026 will be scored and judged. It is deliberately compact: every contestant, coach, and country team leader should be able to read it in one sitting.

The companion documents handle everything else:

- **ICOA Training Book** — the technical curriculum (≈160 pages). What to study to be competitive.
- **Selection Handbook** — for national bodies running country-level selection.
- **Author Team Handbook** — for the question-setting team.
- **ICOA CLI Whitepaper** — the positioning document for the press and accreditation bodies.

If you are reading this on the morning of the finals to refresh on the rules, the parts you most likely want are Section 6 (Medals) and Section 7 (Competition environment). Both are short.

A note on the spirit of the rules. ICOA is an Olympiad, not a CTF tournament. Its purpose is to identify and reward young people who can think clearly about AI as a tool and as a target. The rules below exist to make that judgement fair across nations, languages, time zones, and starting points — not to punish unfamiliarity with the form. If something in this book is unclear, ask your country team leader; if the answer is still unclear, the Science Committee resolves the ambiguity in favour of the contestant.

v1.0.0 — Final, 25 May 2026. The substantive rules in Sections 1–9 are **locked** against the operational evidence summarised in Section 10. This is the rulebook that the Science Committee will apply at the 2026 Sydney finals.

Three items remain to be finalised at the morning-of Day 1 team-leader briefing — they are operationally normal pre-event variables, not rule changes:

1. The final venue address — the chosen lab is being selected from several candidates currently under operational testing (Section 7.1).
2. The paper composition — the number of challenges per track and per category (Section 5.1).
3. The pre-written code import policy — currently permissive, under Science Committee review for tightening (Appendix C FAQ).

No Sections 1–9 rule will be downgraded without notice. Locked rules are: medal ratios (1/12, 2/12, 3/12) + HMA + boundary rule (Section 6), AI4CTF dynamic decay + top-3 speed bonus (Section 5.2), CTF4AI static value (Section 5.3), max-aligned combined-total procedure (Section 5.4), competition environment (Section 7), and hybrid/remote requirements (Section 8).

Staying current is one command. In most situations a contestant, coach, or country team leader can resolve any installed-CLI issue — old version, missing command, stale translation — by running:

```
npm install -g icoa-cli@latest    # ~286 KB, < 5 seconds on a normal connection
```

This is safe to run at any time outside a competition window. The CLI is intentionally small (~286 KB tarball) so reinstall is fast even on a thin mobile connection. If a contestant encounters an unexpected error mid-practice, the recommended first step is to upgrade.

— Charlie Zhu, Founder, ICOA
Sydney · 24 May 2026

Foreword — ICOA at the AI frontier

ICOA is not an ordinary Olympiad, and it is not a traditional CTF. To understand why the rules in this book look the way they do, it helps to start with the world the contestants are walking into.

The AI-agent landscape, 2025 → 2026

Across the eighteen months preceding this rulebook, autonomous AI agents have moved from research curiosity to industrial substrate at a pace that has caught even their builders off guard. In classical cyber-security work, AI agents now read code, hypothesise vulnerabilities, run debuggers, and produce proof-of-concept exploits in tight loops with little human supervision. Anthropic's **Project Glasswing**, launched on 7 April 2026 with launch partners including AWS, Apple, Google, Microsoft, NVIDIA, the Linux Foundation, and JPMorganChase, gave roughly four dozen defenders early access to **Claude Mythos Preview** — a frontier model so capable at security work that Anthropic chose not to release it to the public. In its first weeks, Mythos Preview surfaced more than 10,000 high- and critical-severity vulnerabilities across major operating systems and browsers, many of them previously-unknown zero-days.[1]

The Olympiad world has felt the same shift. At the **International Mathematical Olympiad 2025**, AI systems from Google DeepMind and OpenAI each scored 35 out of 42 — gold-medal level — solving five of the six problems in English within the standard 4.5-hour window, with no human intervention.[2] The IMO 2024 silver-medal moment (AlphaProof + AlphaGeometry 2, one point short of gold) was a warning shot; IMO 2025 was the watershed. Every Olympiad now operates downstream of a question its founders never had to answer: **what does competition mean when the strongest individual reasoner in the room is an AI?**

ICOA's philosophy — Promethean, not Pandoran

ICOA's answer is that AI is **Promethean rather than Pandoran**: a fire to be mastered by the next generation, not a box to keep shut. The defensible position is not to ban AI from the contestant's workstation — an enforcement nightmare and a training signal that points away from the future — but to provide **every contestant with the same AI, on the same terms, under audit**, and then measure who can do the most interesting work with it.

This commits ICOA to a specific bet about education in the AI era: that the discipline worth ranking is no longer **computational thinking** alone, in the sense Jeannette Wing gave the phrase in 2006, but its successor — **AI-agent thinking**: deciding what to delegate, what to verify, when to trust an AI teammate, when to attack an AI target, and how to reason about an adversary whose model is itself a moving object of study. ICOA is built around the conviction that this skill is teachable, examinable, and rankable at the secondary-school level, and that doing so on a CLI-native, server-audited, AI-assisted platform produces a more honest signal than excluding the AI and pretending the world outside the exam room still operates that way.

Empirical validation — the rules in this book are calibrated

The rules in Sections 1–9 are not speculation. They are the rules that survived deployment.

Approximately twenty national-selection events have used the ICOA AI4CTF + CTF4AI token format ahead of the 2026 finals. Across that population, the design has held: papers were neither trivialised by

[1] Anthropic, **Project Glasswing: Securing critical software for the AI era**, 7 April 2026 — https://www.anthropic.com/glasswing. Claude Mythos Preview model card — https://red.anthropic.com/2026/mythos-preview/

[2] Google DeepMind, **Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad**, 2025.

AI assistance nor unsolvable without it; the AI-as-teammate / AI-as-target distinction held its calibration; the permissive submission policy (no penalty, no rate limit) scaled without scoreboard distortion; and the server-authoritative architecture survived multi-cohort burst traffic without divergence. The Macao Special Administrative Region of China hosted the world's first CLI-native Olympiad event on 25 April 2026 — Round 1 Paper A — and the run included a **server-side hotfix shipped mid-paper with zero contestant logout**, with the complete results report delivered to the regional committee approximately six minutes after the final submission landed.[1] That event is the canonical proof point for Section 7.3 (the AI Gateway no-pause clause) and Section 5.5 (permissive submission).

Why CTF4EAI now — the embodied-AI wave

The 2026 edition adds the **CTF4EAI** track — a Pioneer / Practical Round in embodied-AI security — for the same reason the Olympiad exists at all: the frontier moved.

In the year and a half since Q4 2024, the embodied-AI domain has shifted from a Vision-Language-Action research conversation into an industrial sprint. World models, diffusion-policy controllers, and multi-robot coordination architectures have crossed from arXiv into product roadmaps. The activity is global — meaningful capital is flowing into embodied-AI companies across North America, East Asia, Europe, and beyond. In Q1 2026 alone, more than 30 embodied-AI companies raised significant rounds, with individual rounds in the $140M$–400M range at firms including Physical Intelligence, Figure AI, 1X Technologies, X Square Robot, AI2 Robotics, Unitree, and others.[2] The Vision-Language-Action market grew from $4.44B in 2025 toward an estimated $23B by 2030.

A young field of this scale, expanding this fast, with this much downstream physical-world impact, is exactly where the next generation of safety researchers should be invited to look. CTF4EAI is built on a frank acceptance that we are introducing contestants to questions that the industrial frontier itself is still working out — and on the belief that this is a feature, not a defect, of an Olympiad that takes "AI security" seriously. Inspiring young talent across nations to engage the **current** frontier — not the frontier of five years ago — is the appropriate posture for the AI era.

The world's first CLI-native Olympiad — the platform behind these rules

This rulebook stands on a working system, not a proposal:

Current platform version	**v2.19.204** (the version contestants will use; verified latest on the npm registry as of this draft)
Numbered releases since v2.0.0	**322** versioned iterations
npm tarball size (this version)	**~286 KB** — every install is a thin client; assets and translation packs fetch on demand
npm downloads to date	**40,000+**
Demo runs (free practice path)	**3,000+**
Contestants who have sat a national-selection ICOA exam	**~800**
Live events anchored	**Macao SAR, 25 April 2026** (world's first CLI-native Olympiad event)

[1] ICOA, Macao SAR Paper A Round 1 — operational report, 26 April 2026; see icoa2026.au project news.

[2] Q1 2026 industry coverage spans multiple regional sources: KraneShares humanoid-robotics tracker (global rankings); The Robot Report (rounds tracked across Physical Intelligence US-Series A, X Square Robot CN-Series A++, AI2 Robotics CN-Series B, Figure AI US, 1X Technologies NO/multinational); IFR World Robotics 2025 statistics; SemiAnalysis embodied-AI coverage; and regional ecosystem reports across the US, EU, China, Japan, and Korea.

Languages supported in the contestant interface	**21**

These figures are operational, not promotional. ICOA's intent is **long-term**: to sustain a global community of young people who treat AI safety as a discipline rather than a scare-story, and to keep the platform under the kind of continuous live calibration that produced the rules below.

Per Section 1.3, this document deliberately does not publish the 2026 cohort headcount or the count of accredited nations. Rules are framed in ratios and caps so they remain valid regardless of cohort size, and headcount disclosures are reserved for the post-event official report.

Competition at a glance

ICOA 2026 runs three competition tracks. Two are standard onsite tracks contested by the full international cohort. The third is the Pioneer Round in embodied AI security — a special unit run in hybrid mode for the first time this year.

The three tracks

Track	What it tests	Mode	Day	Scoring
AI4CTF	Classical cyber-security domains (Web · Crypto · Forensics · Reverse Engineering · Binary Exploitation) with AI as a *tool* in the contestant's workflow.	Onsite	Day 1	Dynamic decay + top-3 speed bonus (speed rewarded)
CTF4AI	Attacking and defending AI systems — prompt injection, AI forensics, adversarial machine learning, guardrail bypass — with AI as the *target*.	Onsite	Day 2	Static per challenge (knowledge rewarded)
CTF4EAI (Pioneer / Practical Round)	Embodied-AI security as the **practical round** — attacking and defending a vision-language-action (VLA) policy controlling a simulated robot arm. ICOA-VLA is the codename for the target model.	Hybrid	Independent day	Independent ranking, separate medal set

AI4CTF and CTF4AI are scored together: each contestant's Day 1 + Day 2 raw totals are summed, and standard-track medals are awarded from this single combined ranking (Sections 5.4 and 6.1). The two algorithms are deliberately different because the tracks test different things — AI4CTF rewards speed, CTF4AI rewards completeness — and the combined sum measures both.

Standard tracks (AI4CTF, CTF4AI) require onsite presence at the competition venue. Remote participation in standard tracks is permitted only under documented exemption (see Section 8.2 — typically visa denial).

CTF4EAI opens this year in late-window mode: any contestant — onsite or remote — may enter. CTF4EAI is contested and scored independently of the standard tracks. Remote CTF4EAI participants compete under additional video and recording requirements (Section 8.4).

Venue and calendar (onsite)

Dates	27 June – 2 July 2026
Venue	A dedicated computer lab in Sydney. The specific lab is being selected from several candidate venues currently under operational testing, to ensure the chosen environment is equally favourable to every contestant. The final address is published before the CTF4EAI late window opens.
Competition days	Day 1: AI4CTF (5 h contest + ~30 min system pre-test) · Day 2: CTF4AI (same length) · CTF4EAI: 1 separate day (Pioneer / Practical Round, hybrid)
Standard medal calculation	Standard-track medals are decided by the **combined sum** of Day 1 (AI4CTF) and Day 2 (CTF4AI) raw totals — see Sections 5.4 and 6.1

Equipment	Provided by the organisers — identical workstations, identical software stack, identical AI assistance, single OS image (see Sections 7.2–7.3)
Permitted tooling	The ICOA CLI (110 commands) and the bundled 100-tool sandbox — and nothing else
Prohibited tooling	No external additional tools · no web browsers · no standalone AI software or AI clients · no IDEs · no personal laptops. The competition runs entirely inside the ICOA CLI terminal environment.

The single competition environment. ICOA 2026 is contested end-to-end inside the `icoa` CLI — a standalone terminal application. There is no browser tab, no IDE window, no third-party AI chat client, no external tool of any kind on the workstation desk. Every action a contestant takes is a command typed into the CLI, every flag is submitted through the CLI, every AI interaction goes through the CLI's gateway. The same CLI is freely available to anyone who wants to practice:

```
npm install -g icoa-cli@latest    # current version: v2.19.204 · ~286 KB tarball
icoa demo                          # free practice path; no registration required
```

Who competes

Each accredited nation or region may send **up to two teams** to the standard tracks. Each team consists of **four contestants and one team leader**. The team leader is responsible for delegation logistics, translation, and post-window support, but does not assist with challenge content during competition windows (see Section 9.4). Teams are selected by the national body via the Selection Handbook process.

Although team designators (Team 1, Team 2) appear on the scoreboard for routing, competition is **solo**: every contestant submits independently, and all medals are awarded to individuals (see Section 6.5). There is no team scoring in 2026.

Eligibility for contestants: secondary-school students (K-12 equivalent in the contestant's jurisdiction) at the time of competition. National bodies certify eligibility via the Selection Handbook process.

This rulebook intentionally does not publish the 2026 participation total. Participation figures — number of accredited nations, onsite/remote split, contestant headcount per nation — are disclosed only in the post-event official report. Rules in this document are framed in **ratios and caps** (e.g., "up to two teams per nation, four contestants per team", "1/12 Gold per track", "15 CTF4EAI finalists") so they remain valid regardless of cohort size.

Registration status for the 2026 main competition. The standard-track roster (AI4CTF + CTF4AI) closed at the national-selection deadline. New entries to the standard tracks are no longer accepted for 2026. CTF4EAI registration remains open under the late-window rules in Section 8.3.

What the finalist is competing for

Award	Ratio / Count	Awarded to
Standard-track Gold	1 / 12	Top scorers by **combined** AI4CTF + CTF4AI total (Section 6.1)

Award	Ratio / Count	Awarded to
Standard-track Silver	2 / 12	Next 2 / 12 by combined total
Standard-track Bronze	3 / 12	Next 3 / 12 by combined total
Standard-track HMA	—	Top 50% by combined total without a medal (Section 6.2)
CTF4EAI Gold / Silver / Bronze	1 each	Top 3 of the 15 CTF4EAI finalists (Section 6.3)
CTF4EAI Finalist Award	12	The remaining 12 of the 15 CTF4EAI finalists (Section 6.3)

The standard-track ratios 1 / 12, 2 / 12, 3 / 12 follow the long-standing IMO / IOI convention — approximately half of the competing cohort receives a medal, distributed 1 : 2 : 3. The standard medal pool is **single and combined** — a contestant receives one standard-track medal based on their two-day combined total, not separate medals per track. CTF4EAI medals are awarded in addition (Section 6.3).

AI4CTF — Standard Track

What you do

AI4CTF tests the classical cyber-security skill stack — Web, Cryptography, Digital Forensics, Reverse Engineering, Binary Exploitation — in an era where every contestant has an AI teammate. Challenges in this track are vulnerable services, encrypted artefacts, suspect packet captures, and binaries that hide a flag in the form `ICOA{...}`. The role of AI in your workflow is up to you: brainstorming, code-reading, payload drafting, log analysis. The flag is yours to capture either way; the AI you use is the same one every other contestant has.

Format

Each contestant connects to the competition CTFd server via the ICOA CLI. The finals server URL and the activation token for each contestant are distributed to country team leaders shortly before the competition; they are not published in this book.

```
icoa join <finals-server-url>          # URL distributed via country team leader pre-competition
# log in with the credentials on your contestant card (see Section 7.8)
```

Inside the CLI you can:

- `challenges` — list AI4CTF challenges by category, difficulty, and current point value
- `cat <id>` — read a challenge prompt and download attachments into the sandbox
- `submit <id> <flag>` — submit a flag and see your updated score
- `scoreboard` — view the live ranking
- `ai` — open an AI session against the bundled model (rate-limited per the AI Gateway, see Section 7.3)
- `tools` — list and run any of the 100 sandbox tools (e.g., `nmap`, `binwalk`, `pwntools`, `john`, `sqlmap`)

The full set of `icoa` commands inside the finals REPL is in Appendix A.

Day structure

AI4CTF is contested on **Day 1** of the standard-track competition — one full-length paper covering all five domains, with challenges spanning Foundation, Intermediate, and Advanced tiers.

Phase	What is contested	Duration
Day 1 — AI4CTF	The full AI4CTF challenge set — Web, Crypto, Forensics, RE, Pwn — across all difficulty tiers	5 h contest + ~30 min system pre-test

Scoring is **dynamic decay + top-3 speed bonus** — the challenge value drops uniformly as more contestants solve, and the 1st/2nd/3rd solvers receive a +10% / +7% / +5% multiplier on top (Section 5.2). AI4CTF is the **speed track**: being among the first three solvers of a challenge is rewarded; the wider field shares the dynamic base.

The AI4CTF total is one of the two inputs to the combined standard-track ranking (the other being CTF4AI on Day 2). The two day totals are summed raw to produce the medal-deciding combined total (Sections 5.4 and 6.1).

Difficulty calibration

AI4CTF challenges are calibrated against the OWASP Top 10 (Web), classical and modern cryptographic primitives (Crypto), real-world artefact analysis (Forensics), and progressively harder binary-exploitation primitives (RE / Pwn). The training book covers each domain at Foundation, Intermediate, and Advanced tiers; the finals paper draws from all three tiers across all five domains.

CTF4AI — Standard Track

What you do

CTF4AI inverts the relationship: instead of using AI to attack classical systems, you attack AI systems themselves. The targets are deployed LLM services, classifiers behind guardrails, retrieval pipelines, and AI-generated artefacts that you must analyse forensically. Your goal in each challenge is to extract a flag from — or evade — an AI system that resists you.

The domains contested in CTF4AI:

- **Prompt Injection** — direct and indirect; payload assembly; encoding bypasses
- **AI Defence** — guardrail bypass testing, red-teaming LLMs, AI-generated content detection
- **AI Forensics** — analysing AI-generated artefacts, model fingerprinting, attribution
- **Adversarial Machine Learning** — evasion attacks on classifiers (FGSM, PGD, C&W); poisoning; model extraction

Format

CTF4AI uses the same ICOA CLI as AI4CTF — `challenges`, `cat`, `submit`, `scoreboard`, `ai`, `tools`. The difference is on the server: CTF4AI challenges are backed by live AI targets, deployed in isolated containers with the ICOA AI Gateway routing every interaction through the same provider cascade contestants have access to. There is no architectural advantage for any contestant: the model you attack is the same model every other contestant attacks; the budget you have to attack it is the same budget every other contestant has.

Day structure

CTF4AI is contested on **Day 2** of the standard-track competition — one full-length paper covering all four AI-security domains.

Phase	What is contested	Duration
Day 2 — CTF4AI	The full CTF4AI challenge set — prompt injection, AI defence, AI forensics, adversarial ML — across all difficulty tiers	5 h contest + ~30 min system pre-test

Scoring is **static**: every challenge has a fixed value (500 points), and every solver receives that value identically — no decay, no rank effect, no time-of-solve bonus (Section 5.3). CTF4AI is the **knowledge track**: completeness matters, breadth matters, the contestant who solves more challenges scores more regardless of who got there first.

The CTF4AI total is the second input to the combined standard-track ranking. The two day totals (AI4CTF Day 1 + CTF4AI Day 2) are summed raw to produce the medal-deciding combined total (Sections 5.4 and 6.1).

Difficulty calibration

CTF4AI is anchored in published adversarial-ML literature (Biggio & Roli 2018, Goodfellow et al. 2015, Madry et al. 2018, Carlini & Wagner 2017) and the operational practice of LLM red-teaming. The training book Part III walks through each domain. The finals paper exercises prompt injection at Foundation–Intermediate, AI forensics at Intermediate, and adversarial ML at all three tiers.

CTF4EAI — Pioneer / Practical Round (Special Unit)

What it is

CTF4EAI is the Olympiad's **embodied-AI** security track. Where CTF4AI attacks software-only AI (LLMs, classifiers, guardrails, retrieval pipelines), CTF4EAI attacks **AI systems that act on the world** — robot arms, manipulation pipelines, sim-to-real stacks, multi-robot coordination, world models, diffusion policies. The 2026 target is a vision-language-action (VLA) policy controlling a simulated robot arm — codename **ICOA-VLA** — but the contestants' scope of attack and defence extends across the full embodied-AI stack, not just the VLA model itself.

Contestants attempt to attack the policy, defend it against an adversarial scene, extract a flag from its closed-loop behaviour, or expose failure modes elsewhere in the embodied stack (sim-to-real drift, cross-modality backdoor, multi-robot coordination failure, world-model degradation).

CTF4EAI is run in 2026 as the **Pioneer Round** — the inaugural edition of an embodied-AI security unit — and serves operationally as the Olympiad's **practical round**: the place where contestants demonstrate end-to-end attack and defence against a live AI system controlling physical-style actuation, in a setting that mirrors how AI-security work is conducted in the field.

It is a special unit with its own 15-finalist field, its own ranking, and its own medal set, independent of the two standard tracks.

Scope — eight pedagogy phases reflected in the contest

The CTF4EAI track is built on top of the `ctf4eai-12/96/360` curriculum (the third of three parallel ICOA red-team tracks; see Section 10.3). The eight pedagogy phases of that curriculum directly shape what a contestant should expect to encounter:

Phase	Contest surface
The Stage	Embodied-AI threat-model framing, scene/policy notation
Break Vision	Adversarial perturbations on the policy's visual input
Break Language	Prompt-style instruction attacks on language-conditioned policies
Break Embodied AI	VLA-model exploitation, sim-to-real drift induction, multi-robot coordination disruption, world-model attacks, diffusion-policy sampling exploits
The Math	Numerical / gradient-based attacks adapted to the action space
Defending	Guarded policy interfaces, action-space rate limiting, safety-classifier integration
The Field	Field-case style challenges drawn from production embodied-AI incidents (anonymised)
Research	Open challenges; contestant-proposed attack/defence with submission writeup

The full curriculum is documented in `docs/three-tracks-curriculum.md`; the contest paper draws across all eight phases, weighted toward The Stage / Break Vision / Break Embodied AI for the Pioneer-Round Qualifying Stage and toward The Field / Research for the onsite Elimination Stage.

Why it is special-unit, not standard, in 2026

Embodied-AI security is a young field. CTF4EAI is the first deployment of a live embodied-AI target in an international Olympiad context. The Science Committee has scoped the 2026 edition as a Pioneer Round to:

- contain the field to 15 finalists, ensuring deep evaluation of each submission;
- open the format to hybrid (remote + onsite) participation, lowering the bar for nations whose contestants face visa or logistical barriers;
- allow the methodology to mature for promotion to a standard track in a future edition.

Late-window registration

CTF4EAI opens its own **late window** — separate from the closed main-competition registration. Any eligible secondary-school contestant may register for CTF4EAI up to **14 days before the competition** via the contestant's national body or directly with the Science Committee where no national pathway exists.

The late window is the only pathway to enter ICOA 2026 in any capacity now that main-competition registration has closed.

Eligibility & the 15 finalists

Three groups are eligible to enter CTF4EAI:

- **Onsite standard-track contestants** — every contestant who travelled to Sydney for AI4CTF / CTF4AI is automatically eligible to enter CTF4EAI.
- **Remote-exempted standard-track contestants** (the small number with visa-related exemptions on file before main-registration close).
- **Late-window contestants** — new entries registered under Section 4.3, in either onsite or remote mode.

Of all CTF4EAI entrants, the top 15 by qualifying-stage score advance as **CTF4EAI Finalists**.

Format

CTF4EAI challenges are accessed through the same `icoa` CLI used for the standard tracks. A typical challenge presents:

- a scene description (objects, robot pose, task goal),
- a target ICOA-VLA policy you can query (with rate limits),
- a flag condition — for example, "induce the policy to release the cube" or "extract the embedded instruction without triggering the safety classifier".

Tooling available in CTF4EAI is identical to AI4CTF / CTF4AI — the 110-command ICOA CLI and the 100-tool sandbox — plus a CTF4EAI-specific simulator client.

The Pioneer Round consists of two stages:

1. **Qualifying stage** — open to all CTF4EAI entrants, run in hybrid mode (onsite contestants in the competition lab, remote contestants over live video link with supplementary recording per Section 8.4). The top 15 by qualifying score become CTF4EAI Finalists.
2. **Final stage (onsite elimination)** — contested onsite by the finalists who are physically present in Sydney. Remote finalists do not compete in the onsite elimination; their qualifying-stage standing is final.

Awards

Awards in CTF4EAI are allocated as follows:

Award	Count	Eligible to
CTF4EAI Gold	1	Top-ranked onsite finalist after the elimination stage
CTF4EAI Silver	1	2nd-ranked onsite finalist
CTF4EAI Bronze	1	3rd-ranked onsite finalist
CTF4EAI Finalist Award	12	All other finalists — remote finalists and onsite finalists who did not reach the top 3 in the elimination stage

Important. CTF4EAI Gold / Silver / Bronze are restricted to onsite finalists. Remote finalists — who qualified by score but did not travel to Sydney — receive the Finalist Award regardless of their qualifying-stage rank. This is by design: the onsite elimination is the deciding round, and only contestants physically present can contest it.

If fewer than three finalists are present onsite for the elimination stage, the Science Committee may either (a) award the available top medals to onsite finalists with the remaining slots going unawarded, or (b) extend the elimination invitation to additional onsite contestants from the qualifying-stage ranking. The decision is published at least 24 hours before the elimination stage.

CTF4EAI medals are awarded in addition to any standard-track medals — a contestant may, in principle, receive an AI4CTF medal, a CTF4AI medal, and a CTF4EAI medal in the same edition.

Tiebreakers within the qualifying stage and within the onsite elimination follow the same rule as the standard tracks (Section 6.4): on equal score, the earlier submission time of the deciding challenge wins.

Scoring

ICOA 2026 standard scoring runs on two complementary algorithms — one per track — and combines them via max-aligned scaling. This section walks through both, explains why two algorithms exist, and shows a full worked example so the procedure is unambiguous.

All concrete numbers in Section 5 are illustrative. The **uniform 500-point starting value**, the **floor at 100**, the **decay constant $k = 0.115$**, the **theoretical maxima (≈ 20,000 per track)**, and **every figure in the worked example** are shown to demonstrate how the procedure works — not to commit to specific 2026 values.

The **final, authoritative** paper composition, expected maxima, and any other numeric calibration are announced at the **morning-of Day 1 team-leader briefing**. The dynamic-decay + top-3-bonus mechanism (Section 5.2) and the four-step combined-total procedure (Section 5.4) apply identically regardless of which specific numbers are used.

Two philosophies, one combined ranking

Track	Algorithm	Philosophy
AI4CTF (Day 1)	Exponential dynamic decay + top-3 speed bonus	Speed track — early solvers receive more
CTF4AI (Day 2)	Static value per challenge	Knowledge track — every solver receives the same
Combined	Max-aligned scaling + integer round-up + sum	Combines speed and knowledge into one medal ranking, robust to per-track max mismatch

Every challenge on both tracks starts at the same uniform value:

Parameter	Value	Meaning
Initial value	500	The value of any challenge before anyone solves it
Floor value	100	The minimum value a challenge decays to once many contestants have solved it
Decay constant k	0.115	Exponential decay rate, tuned for 40–60 contestant cohorts

How a contestant scores on an AI4CTF challenge:

$$\text{contestant's score} = \text{challenge value at competition close} \times (1 + \text{speed bonus})$$

where the speed bonus is:

Solver rank	Speed bonus	Effective multiplier
1st solver (first blood)	+10%	1.10 ×
2nd solver	+7%	1.07 ×
3rd solver	+5%	1.05 ×
4th and later	0%	1.00 ×

The dynamic decay reduces the challenge value for everyone uniformly as more solvers come. An easy challenge solved by many contestants ends up worth less than a hard challenge solved by few — for every solver of it. The first three solvers also get an extra bonus on top of that base value.

Worked example. A challenge has 10 solvers by competition close. Per the exponential decay curve, its value at close ≈ 242 (down from the initial 500). The first three solvers' scores:

- **1st solver:** $242 \times 1.10 = 266$ points
- **2nd solver:** $242 \times 1.07 = 259$ points
- **3rd solver:** $242 \times 1.05 = 254$ points
- **4th through 10th solver:** 242 points each

If two more contestants solve the same challenge later, the value drops further (e.g., to 220). The base score for everyone — including 1st, 2nd, 3rd — moves with it. The top-3 bonus is recomputed against the new value:

- **1st solver new score:** $220 \times 1.10 = 242$
- **4th–12th solver new score:** 220 each

So early solvers do not have their score "locked" at solve time — they continue to share in the dynamic decay alongside the rest of the field. The speed bonus rewards being among the first three, not the absolute value at which they solved.

AI4CTF scoring (Day 1) — exponential dynamic decay + top-3 speed bonus

Every AI4CTF challenge starts at $P_{\max} = 500$ and decays toward $P_{\min} = 100$ as more contestants solve it. The decay function is exponential:

$$\text{challenge value}(N) = P_{\min} + (P_{\max} - P_{\min}) \cdot e^{-k(N-1)}$$

where N is the number of contestants who have solved the challenge so far. With $P_{\max} = 500$, $P_{\min} = 100$, $k = 0.115$:

Solves N	1	2	3	5	10	15	20	30	40	50+
Challenge value	500	457	418	353	242	180	145	114	104	~101

This produces a smooth exponential approach to the floor — steep at the start (where the value drops sharply as the first solvers come in), asymptotic in the tail (where the field has plateaued at the recognition floor).

Key property — all solvers share the same base. Every contestant who solves the challenge receives the **current challenge value** at competition close as their base score — there is no per-solver lock-in. As more solvers come in, the value drops for everyone uniformly. The first solver and the hundredth solver each get the same base value at competition end, just with different speed-bonus modifiers (see below).

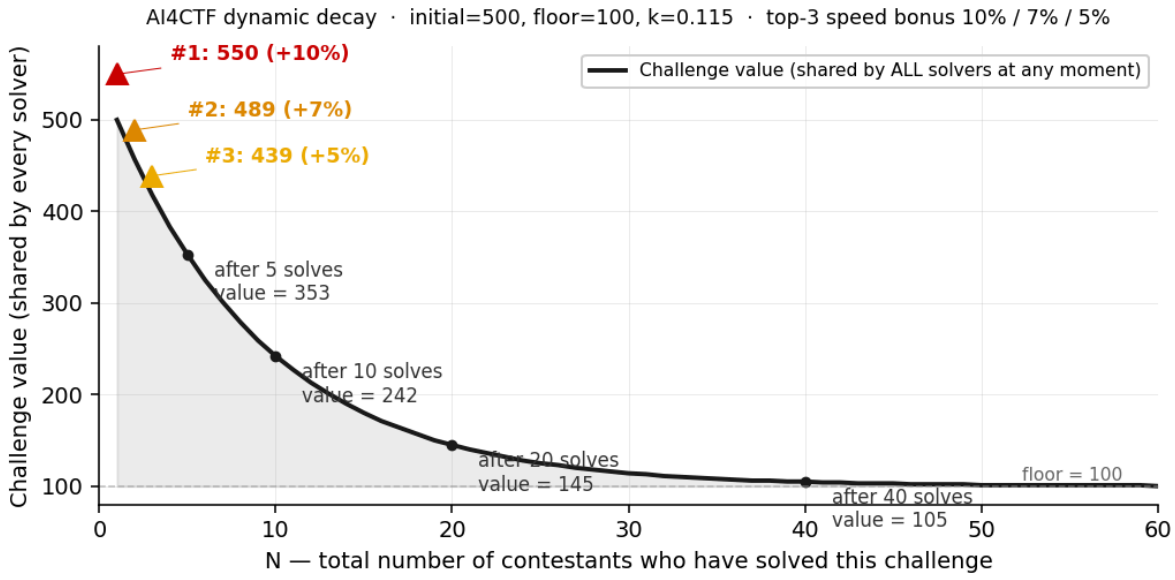

The diagram above shows the decay curve (black line) plus the top-3 speed-bonus markers (red/orange triangles) — the bonus pushes the first three solvers' final scores above the base curve.

Top-3 speed bonus. The 1st, 2nd, and 3rd solvers of each challenge receive an additional multiplicative bonus on the current challenge value:

Solver rank	Bonus	Multiplier	Final score formula
1st (first blood)	+10%	×1.10	contestant score = challenge value × 1.10
2nd	+7%	×1.07	contestant score = challenge value × 1.07
3rd	+5%	×1.05	contestant score = challenge value × 1.05
4th and later	+0%	×1.00	contestant score = challenge value

Implementation: the ICOA Award plugin (v3). The plugin extends CTFd's `dynamic_challenges` with two changes:

1. It registers an **exponential decay function** in place of CTFd's default parabolic, so every `calculate_value()` call on a dynamic challenge returns $P_{\min} + (P_{\max} - P_{\min})e^{-k(N-1)}$.
2. It hooks the solve flow: after each new solve, it refreshes the top-3 contestants' bonus awards (10% / 7% / 5% of the current challenge value). Because the bonus scales with the **current** challenge value, it automatically drops as the field grows — preserving the locked structure where 1st > 2nd > 3rd > 4th, but with the absolute gap shrinking as everyone's value comes down.

The plugin has been verified end-to-end on the practice server: 10 simulated solves at $P_{\max} = 500, P_{\min} = 100, k = 0.115$ produced a challenge value of ≈ 242 at the end, with 1st-place final = 266, 2nd = 259, 3rd = 254, 4th–10th = 242 each — exactly per formula.

Real-time visibility. Contestants can see each challenge's current value at any time during the competition via the `challenges` command in the ICOA CLI — there is no information hiding. As more contestants solve, contestants see the value tick down on their displays. The scoreboard is also live (with normal CTFd update latency on the order of seconds).

CTF4AI scoring (Day 2) — static per-challenge value

Every CTF4AI challenge has a **static, fixed** per-challenge value. There is no decay, no rank effect, no top-3 bonus — you solved it, you got the full value. The first solver and the hundredth solver each receive the same.

Note on the specific value. This rulebook **demonstrates with 500 points per challenge** throughout Sections 5.3, 5.4, and 5.5, so the math is concrete. The actual 2026 per-challenge value is **announced at the morning-of Day 1 team-leader briefing** (per the Pending items in the Preface). What is locked by this rulebook is that the value is **fixed** — the same number applies to every solver of a given challenge — and matches AI4CTF's $P_{\max}$ on the same day so the two tracks remain symmetric before max-aligned scaling.

The reason CTF4AI is static and AI4CTF is dynamic: the two tracks test different things. AI4CTF rewards *speed of execution* (an operational skill), so the dynamic decay + top-3 bonus makes it pay to reach the flag first. CTF4AI rewards *breadth of knowledge across the AI-security surface* — and that's a fairer measure when every solver who arrived at the right answer is recognised equally.

Per-challenge value matches AI4CTF starting value. Both tracks publish at 500 per challenge. A contestant who solves all 40 AI4CTF problems at rank 1 (impossible) would score $40 \cdot 500 \cdot 1.10 = 22\{,\}000$ — close to the CTF4AI theoretical max of $40 \cdot 500 = 20\{,\}000$. In practice the AI4CTF realised max is much lower (because of decay), and the **max-aligned scaling** in Section 5.4 restores balance between the two tracks.

Combining AI4CTF and CTF4AI — the alignment

Combining the two tracks is **not** a naive raw sum — even though both publish at $P_{\max} = 500$, their **realised** maxima at competition end differ because AI4CTF decays away from the peak for everyone except the rank-1 solver of each problem. The Science Committee applies a four-step procedure to keep the two tracks contributing equally to the medal ranking:

The four-step combined-total procedure

1. **Compute each track's realised max** = the highest single-contestant total observed in that track at competition close. Call them M_a (AI4CTF) and M_c (CTF4AI).
2. **Identify the larger.** Let $M_{\text{big}} = \max(M_a, M_c)$ and $M_{\text{small}} = \min(M_a, M_c)$.
3. **Scale the smaller track up.** Multiply every contestant's score in the **smaller-max track** by $\frac{M_{\text{big}}}{M_{\text{small}}}$, then **round up to the nearest integer** (ceiling). Scores in the larger-max track are left unchanged.
4. **Sum.** For each contestant: $\text{combined} = \text{scaled AI4CTF} + \text{scaled CTF4AI}$. Rank by this combined total.

Why this works:

- After step 3, both tracks have the **same effective ceiling** (M_{big} on each). A contestant who topped one track still benefits from leading; a contestant who solved everything on the slower-max track has their effort weighted to match.
- Step 3's "round up to integer" ensures no score is fractionally reduced — every contestant in the scaled track sees their raw score either preserved (if their track was the larger-max one) or **expanded upward** (if the smaller-max one).
- Step 4's raw sum is then comparable across contestants: the two tracks now contribute equally to the maximum achievable combined total.

The 2026 situation. Both tracks use a uniform 500-point initial value. CTF4AI's realised max is simply $500 \times (\text{problems solved by the top CTF4AI scorer})$ — typically near $500 \times 40 = 20\{,\}000$ if a contestant solves all problems. AI4CTF's realised max is lower because the challenge values decay as more contestants solve, so even the top AI4CTF scorer's per-problem average is below 500 (with the top-3 bonuses contributing modestly on a few problems). The scaling factor in 2026 will typically be a moderate multiplier (around 1.5–3×) applied to AI4CTF.

The four-step procedure handles all paper compositions identically — equal-length papers, asymmetric paper lengths, and edge cases where one track has many more solvers per problem than the other. The

procedure is published in advance and applied identically to every contestant — no organisational degree of freedom, no per-contestant adjustment.

This is the simplest possible alignment — **raw sum, equal weight, no normalisation**. The Science Committee considered four common methods before settling on this:

Method	Formula	Why ICOA does / doesn't use it
Raw sum (chosen)	$\mathrm{AI4CTF} + \mathrm{CTF4AI}$	Both tracks designed to same maximum (20,000) — sum is automatically fair
Weighted sum	$w_a \cdot \mathrm{AI4CTF} + w_c \cdot \mathrm{CTF4AI}$	Useful when tracks have unequal max; ICOA's matched-max design makes weighting redundant
Min-max normalisation	$\frac{\mathrm{raw} - \min}{\max - \min}$	Removes scoring scale; harder to explain and depends on observed extremes
Z-score	$\frac{\mathrm{raw} - \mu}{\sigma}$	IMO/IOI statistical practice; counter-intuitive for K-12 audience

Raw sum is the choice because it is **the most transparent**: a contestant can predict their final standing from their two day-end totals on a pocket calculator. There is no hidden coefficient, no organisational degree of freedom in the formula, and the rule cannot be rewritten after the fact.

Worked example

Reminder: every number in this example is illustrative. The per-contestant raw totals, the scaling factor, and the final ranking numbers are all chosen to demonstrate how the four-step procedure works. The actual 2026 numbers are realised at competition close from CTFd's live scoreboard. **Only the procedure (Section 5.4 four steps + Section 5.2 dynamic + top-3 bonus) is locked by this rulebook.**

Step 1 — six contestants and their raw totals (computed by CTFd at competition close as sum-of-challenge-values + sum-of-speed-bonuses for AI4CTF, and sum-of-static-values for CTF4AI):

Assume a 40-problem paper per track. After competition close, AI4CTF challenge values have decayed to (illustrative):

- 10 challenges with ~5 solvers each → final value ≈ 350
- 20 challenges with ~15 solvers each → final value ≈ 180
- 10 challenges with ~40 solvers each → final value ≈ 105

A contestant's AI4CTF raw = sum over solved challenges of (challenge value × bonus multiplier). CTF4AI raw = number of solves × 500.

Contestant	Profile	AI4CTF raw	CTF4AI raw
Speed Specialist	Solves 24 of 40 AI4CTF; first-blood on 3 challenges, 2nd-blood on 2, 3rd-blood on 2. Solves 24 of 40 CTF4AI.	4,532	12,000
All-Round Champion	Solves 38 of 40 AI4CTF; first-blood on 4, 2nd-blood on 4, 3rd-blood on 2. Solves 38 of 40 CTF4AI.	7,766	19,000
Knowledge Specialist	Solves all 40 AI4CTF but never a top-3 (always 4th or later). Solves all 40 CTF4AI.	8,170	20,000

Contestant	Profile	AI4CTF raw	CTF4AI raw
Solid Generalist	Solves 30 of 40 AI4CTF; first-blood on 1, 2nd-blood on 1. Solves 30 of 40 CTF4AI.	5,400	15,000
Mid-Pack Solver	Solves 20 of 40 AI4CTF; no top-3 finishes. Solves 22 of 40 CTF4AI.	3,200	11,000
Newcomer	Solves 3 AI4CTF (no top-3). Solves 8 CTF4AI.	315	4,000

Step 2 — identify the realised maxes:

- M_a = highest AI4CTF raw = **8,170** (Knowledge Specialist)
- M_c = highest CTF4AI raw = **20,000** (Knowledge Specialist)
- $M_{\text{big}} = 20,000$ (CTF4AI), $M_{\text{small}} = 8,170$ (AI4CTF)

Step 3 — scaling factor for the smaller-max track (AI4CTF):

$$\text{scale} = \frac{M_{\text{big}}}{M_{\text{small}}} = \frac{20000}{8170} \approx 2.4480$$

Each contestant's AI4CTF raw is multiplied by 2.4480 and rounded **up** to the nearest integer.

Step 4 — compute and sum:

Contestant	AI4CTF raw	AI4CTF scaled	CTF4AI raw	Combined
Speed Specialist	4,532	⌈11,094.3⌉ = 11,095	12,000	**23,095**
All-Round Champion	7,766	⌈19,011.0⌉ = 19,011	19,000	**38,011**
Knowledge Specialist	8,170	⌈20,000.0⌉ = 20,000	20,000	**40,000**
Solid Generalist	5,400	⌈13,219.2⌉ = 13,220	15,000	**28,220**
Mid-Pack Solver	3,200	⌈7,833.6⌉ = 7,834	11,000	**18,834**
Newcomer	315	⌈771.1⌉ = 772	4,000	**4,772**

Final ranking:

Rank	Contestant	Combined
1	Knowledge Specialist	40,000
2	All-Round Champion	38,011
3	Solid Generalist	28,220
4	Speed Specialist	23,095
5	Mid-Pack Solver	18,834
6	Newcomer	4,772

What this shows: with the new "dynamic + top-3 bonus" design, **completeness still beats speed-with-coverage-gaps**. The Knowledge Specialist edges out the All-Round Champion (40,000 vs 38,011) because solving 2 more problems is worth more than 4 extra first-bloods. The Speed Specialist, who skipped 16 problems to chase speed bonuses, lands fourth — even with their bonus stack, missing problems is a 5,000-point gap that bonuses can't close.

This is intentional: ICOA's combined ranking does not over-reward speed at the expense of completeness. To win reliably, a contestant must be **near-complete on both tracks** — speed merely provides the tiebreaker among contestants of similar coverage.

CTF4EAI scoring (Pioneer / Practical Round)

CTF4EAI is contested independently. Its qualifying-stage score determines the 15 finalists; the onsite elimination then determines the top three among onsite finalists (Section 4.6). CTF4EAI scoring is per the Pioneer Round's own ruleset and does not feed into the standard-track combined total.

Submission policy (applies to all tracks)

ICOA 2026 keeps the CTFd server's standard submission defences. Verified against the live competition server (CTFd 3.8.5, May 2026):

- **Wrong submissions do not deduct points.** A wrong flag costs no score — the only cost is the time spent. A contestant who arrives at a correct flag through ten wrong guesses scores identically (per the dynamic challenge value on AI4CTF, per the static value on CTF4AI) to one who arrives first try.
- **Correct submissions are accepted under the same rate-limit window as wrong submissions** (CTFd does not split the limit by outcome). In practice a contestant submits a flag the moment they have it; the limit is large enough that no realistic human pace touches it.
- **Anti-bruteforce limit:** CTFd's `incorrect_submissions_per_min` defence is on, default value **10 attempts per minute per challenge per contestant**. This is the live server configuration as of 24 May 2026 testing.
- **Per-challenge cooldown if the limit is hit:** CTFd's `max_attempts_timeout` of **5 minutes** (300 seconds, default) locks that one challenge to that contestant. Other challenges remain fully open. The cooldown applies whether the burst was made of wrong-only or mixed submissions — but the trigger is the wrong-submission counter, so a human pace of one or two attempts per minute will never trigger it.

This policy intentionally rewards iteration and experimentation while preventing automated brute-force. A human contestant typing or pasting flags by hand at any realistic pace will not approach the 10-per-minute threshold on a single challenge — the guard activates only when a script is in the loop.

Contestant logbook (submitted with answers)

ICOA provides an in-CLI **logbook** — a per-paper structured solve log that contestants are **encouraged but not required** to use. Logbook entries are auto-tagged with the current question and **submitted alongside the contestant's answers** when the paper finishes. They become part of the official record for that paper.

Suggested uses (entirely voluntary):

- Note the approach you tried before the working one
- Record an interesting observation about a challenge that didn't fit the flag format
- Flag a possible bug or ambiguity (this is the fastest channel to the Science Committee for post-event review)
- Capture any external constraint that affected your attempt (network blip, terminal hiccup, etc.)

How to use it (inside a competition window):

`note <text>`	Add an entry, auto-tagged with the question you are currently viewing
`note q5 <text>`	Add an entry tagged explicitly to Q5
`logbook`	View all your entries for the current paper
`logbook 5`	Filter to entries on Q5
`notes`	Alias of `logbook`

Logbook entries are **not graded** — they have no point value. They exist for the contestant's own record-keeping during the paper, and for post-event review by the Science Committee in cases where a contestant disputes a result or flags a challenge issue. Use of the logbook is encouraged but never required; a contestant who submits a perfect paper with zero logbook entries is in no way penalised.

Scoring authority

The CTFd server is the sole authority on every contestant's score. Local CLI displays are mirrors of the server state and may briefly lag during high-throughput periods. In any dispute, the server's stored score at the official close of the competition window is final.

Medals and awards

Standard-track medals (AI4CTF + CTF4AI combined)

Standard-track medals are awarded from a **single combined pool**. There is one Gold, Silver, Bronze, and HMA recognition for the standard competition — drawn from the **AI4CTF + CTF4AI combined two-day total** (Section 5.4).

Medals are distributed by the IMO / IOI convention:

Medal	Ratio	Cohort
Gold	1 / 12	Top 1 / 12 of the standard-track cohort, by combined total
Silver	2 / 12	The next 2 / 12, by combined total
Bronze	3 / 12	The next 3 / 12, by combined total

Approximately half of the standard-track cohort receives a medal. The ratios are applied once over the combined-total ranking, not per track.

Boundary rule when the ratio is fractional

When the 1/12, 2/12, 3/12 ratio produces a non-integer cutoff position, the contestant at the **boundary position** is promoted to the **higher** tier. The math floor is applied to the count of the **lower** tier; the surplus contestant goes up.

> **Boundary rule.** If a contestant sits exactly on a gold/silver, silver/bronze, or bronze/HMA boundary because of a fractional ratio cutoff, they receive the **higher** medal. No contestant is denied a higher medal merely because of an integer-division remainder.

Worked example — cohort of 80 contestants, **medal tiers only** (HMA is computed separately, see Section 6.2):

Tier	Math	Resolved count	Places awarded
Gold	$\frac{80}{12} = 6.67$	**7**	1st through 7th by combined total
Silver	$\frac{160}{12} = 13.33$	**14**	8th through 14th (7 contestants)
Bronze	$\frac{240}{12} = 20$	**20**	15th through 20th (6 contestants)

Notice the boundary effect: the math gives 6.67 golds, but the 7th-place contestant is at the boundary and is **promoted to gold**. Silver and bronze cutoffs are similarly resolved with the boundary contestant promoted upward.

HMA count is not fixed by the 80-contestant cohort. HMA is determined per the rule in Section 6.2 — a contestant earns HMA if their **Day 1 AI4CTF** or **Day 2 CTF4AI** score was in the top 50% of that day's cohort **and** their combined two-day total did not reach the bronze cutoff (i.e., they are 21st place or below in combined). The resulting HMA count depends on the **overlap between Day 1 top-50% and Day 2 top-50% among non-medalists**, which varies edition to edition. With a typical 80-contestant cohort, HMA usually covers somewhere between 10 and 25 contestants — but the exact figure is computed only after Day 2 scores are in.

Tied scores at a cutoff

When multiple contestants share the exact same combined total at a medal cutoff, all of them receive the higher medal (this is the existing Section 6.4 tiebreaker rule applied at the medal boundary).

Honorable Mention Award (HMA)

The Honorable Mention Award follows the long-standing convention for Olympiad-style competitions: it recognises contestants who had a strong **single day** even if their combined two-day total fell short of the bronze cutoff. This is the IOI / IPhO "good-day recognition" rule, ported intact.

> **HMA rule.** A contestant receives the Honorable Mention Award if **either** of the following holds:
>
> - their Day 1 AI4CTF score places them in the top 50% of the AI4CTF cohort by Day 1 score, **or**
> - their Day 2 CTF4AI score places them in the top 50% of the CTF4AI cohort by Day 2 score,
>
> **and** their combined two-day total did **not** reach the bronze-medal cutoff (Section 6.1).

HMA is a single recognition. A contestant who reaches bronze, silver, or gold on the combined ranking does not also receive HMA — medals strictly dominate HMA. The "or" in the rule is inclusive: a contestant who has a top-50% day in **both** tracks but still misses the bronze cutoff receives HMA once.

Practical reading: HMA rewards a contestant who **can solve hard problems on one of the two days** but didn't sustain the performance across both. It tells the field "this contestant has shown a true peak performance worthy of recognition."

CTF4EAI awards (Pioneer / Practical Round)

The Pioneer / Practical Round produces its own award set, independent of the standard tracks:

Award	Count	Awarded to
CTF4EAI Gold	1	Top-ranked onsite finalist after the onsite elimination stage
CTF4EAI Silver	1	2nd-ranked onsite finalist
CTF4EAI Bronze	1	3rd-ranked onsite finalist
CTF4EAI Finalist Award	12	All other finalists — remote finalists and onsite finalists outside the top 3

CTF4EAI Gold / Silver / Bronze are restricted to onsite finalists by design (see Section 4.6). A CTF4EAI medal or Finalist Award is independent of and additive to any standard-track medal.

Tiebreakers

Within any single track, when two or more contestants are tied on total score at a medal cutoff:

1. The contestant with the earlier submission time of the highest-value challenge they both solved is ranked higher.
2. If still tied, the comparison repeats on the next-highest-value commonly solved challenge.
3. If still tied after exhausting commonly solved challenges, all tied contestants are awarded the higher medal.

For CTF4EAI specifically, the onsite elimination stage is the deciding round for Gold / Silver / Bronze, and the elimination outcome is authoritative. The qualifying-stage tiebreaker rules above apply when determining the 15 finalists and the ranking among remote finalists.

Individual vs team

Medals and awards are awarded **to individual contestants**. The “Team 1 / Team 2” designators on the live scoreboard exist for routing and identification only; there is no team ranking in 2026. National team standings (a possible future addition) are not contested in the 2026 edition.

Competition environment

Onsite venue — Computer Lab

The Sydney finals are contested at a single physical venue — a dedicated computer lab in Sydney. The specific lab is being selected from several candidate venues currently under operational testing, so that the chosen environment is equally favourable to every contestant (no incumbent-team advantage, no environmental familiarity that one delegation could have prepared against and another could not). The final address is confirmed and published before the CTF4EAI late-window deadline.

All onsite contestants compete on workstations provided by the organisers. The lab provides:

- identical workstations (hardware, OS, display) — no two contestants have an environmental advantage over each other;
- identical software stack — the same ICOA CLI build, the same sandbox image, the same network configuration;
- proctored physical access — only registered contestants, country team leaders (in designated areas), and organisers may enter during competition hours.

Provided equipment — no personal laptops

Contestants **may not use their own laptops** during the standard-track competition windows. Provided workstations are pre-imaged with:

- the ICOA CLI at the locked finals version;
- the 110-command CLI and the 100-tool sandbox image;
- access to the finals competition server (URL distributed via country team leaders pre-competition);
- access to the ICOA AI Gateway (see Section 7.3).

Anything else — personal notes, reference material, code snippets — must be brought as plain text and pasted in through the lab's controlled clipboard channel during the pre-competition setup window (details in the day-of briefing).

Identical AI assistance

Every contestant has access to the same AI capability, routed through the ICOA AI Gateway. The Gateway enforces:

- a per-contestant token budget identical across contestants;
- a per-contestant rate limit identical across contestants;
- a single provider cascade — every contestant queries the same model family in the same order.

A contestant who exhausts their budget early competes the remainder of the window without AI assistance. The budget is published before the competition and is identical for every contestant in the same track.

Gateway availability — three tiers of response. The AI Gateway is monitored by the Science Committee but the competition clock does **not** automatically pause for outages. Operational expectation is graduated:

Outage duration	Expected response
Under 5 min	Auto-recovery. Most blips are upstream rate-limit returns that resolve within seconds. Contestants continue working with the 110-command CLI, the 100-tool sandbox, and their own thinking. No clock adjustment.

Outage duration	Expected response
5–10 min	Usually hot-fixable. The Sydney technical committee ships a server-side fix the same way the Macao SAR event did on 25 April 2026 (Section 10) — no contestant logout, no client update, no clock adjustment. Contestants continue working.
Over 30 min (extreme)	The organisers will announce, onsite, whether the affected window's clock is **extended**. This is a Science Committee decision communicated to all contestants in the window simultaneously, with the extension amount and revised end time stated explicitly.

The 5/10/30-minute tiers are operational expectations, not guarantees — the actual decision in any incident rests with the Science Committee. Contestants should plan their attack assuming the AI is one tool among many, not the central tool.

CLI fullscreen, single desktop

During competition hours, contestants must run the ICOA CLI in **fullscreen** mode. The lab workstation is configured to a single desktop — multi-desktop / virtual-desktop features are disabled at the OS level. Switching out of the CLI fullscreen to access any other application during a competition window is grounds for disqualification (Section 9.5), with one exception: the supplied scratch text editor for short notes (no network access, no external clipboard).

OBS + CLI recording

Throughout each competition window, two recordings run automatically on each workstation:

- **OBS screen capture** — full-screen video of the contestant's display;
- **CLI session log** — the asciinema-format log of every command, output, and AI interaction in the CLI session.

Both recordings are uploaded to the organisers' archive at the end of each window for proctoring and integrity review. Contestants do not have to enable or manage these recordings — they are pre-configured on the lab image. The recordings remain confidential to the Science Committee and are deleted six months after the competition.

Permitted tools

The full set of tools a contestant may use during a competition window:

- the **ICOA CLI** and its 110 finals-mode commands (`challenges`, `cat`, `submit`, `scoreboard`, `ai`, `tools`, and all subcommands listed in Appendix A);
- the bundled **100-tool sandbox**, accessed via `icoa tools` from inside the CLI — including `nmap`, `binwalk`, `pwntools`, `john`, `hashcat`, `sqlmap`, `wireshark`-CLI, `radare2`, `ghidra`-headless, `volatility`, and the 90 other tools listed in the training book Section 5;
- the **ICOA AI Gateway** via the `icoa ai` subcommand;
- the supplied scratch editor (offline only).

Prohibited tools

Inside a competition window, contestants may not use:

- **No external additional tools** of any kind — the ICOA CLI and its bundled 100-tool sandbox are the entire toolset;
- **No web browsers** (Chrome, Firefox, Safari, Edge, etc.);

- **No standalone AI software or AI clients** (ChatGPT app, Claude app, Gemini app, Copilot, Cursor, etc.) — every AI interaction must go through the ICOA AI Gateway;
- **No IDEs** (VS Code, JetBrains/IntelliJ, Xcode, Eclipse, Vim/Emacs as a graphical IDE, Cursor, etc.) — the CLI's bundled editor is the only editing surface during competition windows;
- **No third-party communication clients** (WeChat, WhatsApp, Discord, Slack, email, etc.);
- **No web access** outside the finals server;
- **No personal laptops**, phones, tablets, smartwatches, or any external-network-capable device on the workstation desk.

The single permitted environment is the `icoa` CLI terminal — a standalone application that is the contestant's entire interface to the competition. Devices not in use must be stowed in the designated locker outside the lab. Possession of a prohibited device on the workstation desk during a competition window is grounds for disqualification (Section 9.5).

The `icoa join` workflow

Connecting to the finals server is one command:

```
icoa join <finals-server-url>
# The exact <finals-server-url> is distributed to country team leaders
# shortly before the competition and is NOT published in this book.
# Enter your contestant card credentials when prompted.
# After successful join, the REPL drops you at the finals prompt.
```

The contestant card is provided at registration on the morning of Day 1, together with the activation token routed through your country team leader. It contains the contestant's `Player_NNN` identifier and password. Cards are single-issue; lost cards are reissued by a country team leader with the on-duty proctor.

Hybrid and remote participation

Standard tracks: onsite is the default

AI4CTF and CTF4AI are **onsite tracks**. The intended model is: every contestant travels to Sydney, competes on a competition lab workstation, and is proctored in person. This is how the 2026 edition is scored under normal conditions.

Exemption process (visa or special circumstances) — closed for 2026 standard tracks

Standard-track registration closed at the national-selection deadline (see Section 1.3). The exemption pathway described below applied during that registration window; contestants who were approved for remote participation in the standard tracks before the deadline retain their approval. No new standard-track exemptions are accepted for 2026.

The historical process — documented here for reference and for future editions — was:

1. Country team leader submits an exemption request, with documentation of the obstructing circumstance, to the Science Committee before the registration deadline.
2. The Science Committee approves or denies. Approval was the default for documented visa denial; other circumstances were evaluated individually.
3. Approved remote standard-track contestants compete from a designated remote-friendly environment — typically a partner-institution computer lab in their home country, with live video link to a Sydney proctor and full session recording.

Approved remote standard-track contestants are scored against the same scoreboard as onsite contestants — the same medal pool, the same ratios, the same tiebreakers.

CTF4EAI late-window: hybrid is open

CTF4EAI opens its own **late window** — separate from the closed main-competition registration — and the late window is currently the only way to enter ICOA 2026.

Deadline	14 days before the start of the CTF4EAI competition day
Mode	Hybrid — onsite or remote, contestant's choice
Eligibility	Secondary-school student (K-12 equivalent in jurisdiction); registered via national body or, where no national pathway exists, directly with the Science Committee
Effect on awards	Onsite CTF4EAI registrants are eligible for Gold / Silver / Bronze (subject to reaching the 15-finalist field and the onsite elimination outcome). Remote registrants are eligible for the Finalist Award only.

No exemption process is required for remote CTF4EAI registration — remote is one of the two equal entry modes of the late window. Remote registrants must, however, meet the Section 8.4 remote requirements during the competition window.

Remote participation requirements (CTF4EAI and exempted standard tracks)

A remote contestant must, during the competition window:

- maintain a live video link to a Sydney-side proctor (camera framing the workstation and the contestant);

- maintain an additional desktop recording on the remote machine (OBS or equivalent), uploaded to the organisers' archive at the end of the window;
- run the ICOA CLI in fullscreen on a single desktop, with the same prohibition on third-party browsers, AI clients, and external devices as onsite contestants;
- be in a quiet, lit room with no other persons present (country team leader supervision permitted from a separate camera angle).

Loss of video link for more than 5 minutes within a competition window triggers an integrity review; sustained loss is grounds for disqualification of that window.

Eligibility, teams, and conduct

Eligibility

A 2026 contestant must:

- be a secondary-school student (K-12 equivalent in their jurisdiction) at the time of the competition window;
- be selected through their national body's process (see Selection Handbook);
- be in good standing with the ICOA Code of Conduct (Section 9.4);
- meet the visa, travel, and consent requirements for the competition mode (onsite, exempted remote, or CTF4EAI hybrid).

National teams

A nation or region may send **up to two teams** to the standard tracks. Each team consists of **four contestants plus one team leader**. ICOA does not impose internal selection criteria within a national team — the national body decides which four contestants per team and which leader.

The team leader is the delegation's accountable adult: they handle pre-event logistics, coordinate with the Science Committee, translate when needed, and support contestants outside the competition windows. Team leaders may **not** assist with challenge content during competition windows (see Section 9.4 conduct rule 1 and Section 9.5).

Contestants from the same nation appear on the scoreboard as `Team 1 – Player_NNN`, `Team 2 – Player_NNN`, and so on; the team designator is for routing only and has no scoring effect. There is no aggregated team or national ranking in 2026 — every result is individual.

Solo competition

ICOA 2026 is contested solo. Every submission is attributed to a single contestant. Sharing flags, solutions, or AI-assistance budgets between contestants — including contestants on the same national team — during a competition window is grounds for disqualification (Section 9.5). The training period before competition is unrestricted: practice together, share notes, study together. Inside the window, you are on your own.

Code of conduct — the short version

Five rules cover the spirit:

1. **Compete honestly.** The flag you submit must be a flag you found. The AI you used to find it must be the ICOA AI Gateway, not a third-party client. The reasoning is your own (with AI assistance), not borrowed from another contestant during the window.
2. **Respect the sandbox boundary.** Every attack technique is legal inside the ICOA sandbox and illegal outside it. Do not pivot from the sandbox to the broader internet, to other contestants' workstations, or to the organisers' infrastructure.
3. **No third-party tooling.** Browsers, AI clients, communication apps — none of them, during a window. The ICOA CLI is the entire environment.
4. **Trust the proctors.** If a proctor asks you to stop, you stop. If a proctor asks to see your screen, you show it. Disagreement goes through your country team leader after the window, not during it.

5. **Be kind.** You are contesting an Olympiad in a brand-new field alongside other young people from many nations. Your conduct toward other contestants — onsite and remote — is part of how the field decides what kind of community it wants to be.

Disqualification (brief)

The Science Committee may disqualify a contestant from a window, a track, or the full event, on evidence of:

- prohibited tooling on the workstation desk during a window;
- third-party AI or browser use;
- inter-contestant collaboration during a window;
- attempts to pivot out of the sandbox;
- sustained loss of remote video link without proctor approval;
- conduct violating the Code of Conduct.

Disqualification decisions are made by the Science Committee, communicated to the country team leader, and may be appealed through the country team leader within 24 hours of the decision.

Detailed disciplinary procedure — appeals, evidence handling, second-chance criteria — is out of scope for this rulebook and is published separately in a future revision.

Track record — rule-by-rule provenance

The Foreword summarised the **narrative**: AI agents have reshaped both industry and Olympiads, ICOA's philosophy is Promethean, and the platform behind these rules has shipped 313 versions, served ~800 selection contestants, and anchored the Macao SAR first-event proof point. This section is the **audit trail**: for each rule that materially shapes a contestant's day, where the rule came from operationally.

Rule-by-rule operational provenance

Section 5.2 AI4CTF dynamic + top-3 bonus + 5.3 CTF4AI static	Live across ~20 national-selection deployments. AI-assistance balance held — papers were neither trivialised nor unsolvable. Sustained burst traffic from multi-cohort selection days without scoreboard divergence.
Section 5.5 Permissive submission policy	Live across the same ~20 deployments. Contestants submitted 3–5× more flag attempts than under a penalty regime, with no degradation of the final ranking; the engagement gain was unambiguous.
Section 7.2 Single-image environment, no personal laptops	Hardened after pre-Macao deployments in which contestants used their own machines: small environmental differences (OS version, terminal font, sandbox image age) generated disproportionate dispute volume.
Section 7.3 AI Gateway no-pause clause	Calibrated against the Macao SAR event of 25 April 2026 — a mid-paper server-side hotfix shipped with zero contestant logout, AI request success returned to 100% post-fix, the competition clock kept moving. The platform was demonstrably hot-fixable; the rule reflects that capability.
Section 5.5 Server-authoritative scoring	The CTFd backend, with the ICOA token-API layer in front, sustained multi-cohort selection days without scoreboard divergence. Not aspirational.
Section 6.1 Medal ratios 1/12 · 2/12 · 3/12 + Section 6.2 HMA	IMO / IOI lineage; ported intact with no operational divergence in

	the selection-stage deployments that used the same ratio for Paper A.
Macao SAR report-delivery target	The 6-minute “final-submission-to-report-delivered” interval at Macao SAR set the operational expectation for finals-day reporting; back-end pipeline carried over intact.

What is still being calibrated

The Science Committee is transparent about which rules in this book are settled and which remain provisional pending the finals themselves:

Rule	Status	Why
AI4CTF dynamic + top-3 bonus + CTF4AI static	Settled	Live in ~20 deployments; ICOA Award plugin v3 verified end-to-end on practice server
Permissive submission policy	Settled	Live in ~20 deployments
Server-side hot-fixability	Settled	Macao SAR proof point, 2026-04-25
Onsite single-image environment	Settled	Hardened after pre-Macao dispute pattern
Medal ratios 1/12 · 2/12 · 3/12 + HMA	Settled	IMO / IOI lineage, ported intact
CTF4EAI G/S/B onsite-only	**Pioneer Round, provisional**	First live edition of an embodied-AI Olympiad unit — refinement expected for 2027
Late-window deadline (14 d before)	**Pioneer Round, provisional**	Calibrated against current visa-processing observations; may tighten
Final venue address	**Pending confirmation**	Selected from candidate labs under operational testing; published before late-window opens
Paper composition	**Pending briefing**	Announced morning-of Day 1

Rule	Status	Why
		team-leader briefing — Section 5.1
Pre-written code import policy	**DRAFT, pending briefing**	Section 9.5 / Appendix C FAQ — Science Committee reviewing tightening

The v1.0.0 Final designation is supported by this audit trail: Sections 1–9 are the operational rules the Science Committee will apply, and Section 10 is the empirical evidence behind them. The three Pending items listed in the Preface are **operational variables**, not rule revisions.

Appendix A — Quick reference

Finals REPL — the commands you will actually use

`challenges`	List challenges available in the current track, with category, difficulty, current point value, and your solve state.
`cat <id>`	Display the challenge prompt; download attachments into your sandbox.
`submit <id> <flag>`	Submit a flag. Returns the updated score and (on AI4CTF top-3) the speed-bonus confirmation.
`scoreboard`	Live ranking. Updates as solves come in.
`ai`	Open an AI session against the bundled model. Token budget is per-contestant and shown after each turn.
`tools`	List the 100 sandbox tools. `tools <name>` runs a tool with arguments in the sandbox.
`hint <id>`	Reveal the next hint tier on a challenge (subject to per-challenge hint cost).
`note <text>`	Add an entry to your contestant logbook for the current question (Section 5.8). Encouraged but optional.
`note q5 <text>`	Tag the entry explicitly to Q5 instead of the current question.
`logbook` / `notes`	View all your logbook entries for the current paper. Add `5` (or `q5`) to filter by question.
`status`	Your current score, time remaining in the window, AI budget remaining.
`help`	List all REPL commands and their flags.
`menu`	Return to the ICOA main menu (escape hatch; safe to use mid-window).

A full reference of all 110 finals-mode commands is bundled inside `icoa help finals`. The above are the most frequently used.

Sandbox tools — the categories

The 100-tool sandbox image is organised in 13 groups (Web · Crypto · Forensics · RE · Pwn · Network · Misc · Steg · OSINT · ML · Recon · Util · Lang-specific). A representative tool from each group:

Web	`sqlmap`, `gobuster`, `ffuf`, `nikto`, `wpscan` …
Crypto	`openssl`, `hashcat`, `john`, `RsaCtfTool`, `featherduster` …
Forensics	`volatility3`, `autopsy`, `bulk_extractor`, `binwalk`, `foremost` …
RE	`radare2`, `ghidra-headless`, `objdump`, `strings`, `ltrace` …
Pwn	`pwntools`, `gdb-peda`, `ROPgadget`, `one_gadget`, `seccomp-tools` …
Network	`nmap`, `tcpdump`, `tshark`, `mitmproxy`, `scapy` …
Steg	`steghide`, `zsteg`, `exiftool`, `stegseek` …
ML	`numpy`, `pytorch` (read-only, no GPU), `scikit-learn`, `transformers` (local) …
Util	`jq`, `xxd`, `base64`, `curl`, `wget` (restricted to finals server) …

The full inventory is available via `icoa tools list` from inside the CLI.

Appendix B — Glossary

AI4CTF	The standard track that uses AI as a tool to attack classical cyber-security challenges.
CTF4AI	The standard track that attacks AI systems themselves — prompt injection, AI defence, adversarial ML.
CTF4EAI	The Pioneer / Practical Round special unit — attacking and defending embodied-AI systems (VLA, world models, diffusion policy, sim-to-real, multi-robot coordination). 2026 target stack centred on ICOA-VLA.
ICOA-VLA	The codename for the vision-language-action target policy used in the 2026 CTF4EAI Pioneer Round. CTF4EAI scope is broader than ICOA-VLA alone (see Glossary entry for CTF4EAI).
Embodied AI	AI systems that perceive and act on the physical (or physically-simulated) world. The Olympiad scope for CTF4EAI includes VLA, world models, diffusion policy, sim-to-real, and multi-robot coordination.
Three tracks	The 2026 ICOA structure: AI4CTF (AI as teammate), CTF4AI (red-teaming software AI), CTF4EAI (red-teaming embodied AI). Each tier of the matching learning curriculum is 12 / 96 / 360 cards; ctf4ai adds a refreshable +120 frontier supplement. See `docs/three-tracks-curriculum.md`.
Pioneer Round	The 2026 designation for CTF4EAI as the inaugural edition of an embodied-AI security unit. Synonym: **Practical Round** — the operational role CTF4EAI plays in the 2026 awards programme.
Practical Round	Operational name for CTF4EAI in 2026 — the end-to-end attack-and-defence stage against a live VLA system. Synonym for Pioneer Round.
Late window	The CTF4EAI-only registration window, open after the standard-track main registration has closed. Deadline: 14 days before the CTF4EAI competition day.
Dynamic scoring	Scoring model in which a challenge's point value decreases as the number of solves grows.
ICOA Award	The custom CTFd plugin (v3, 2026-05-25) that implements the AI4CTF scoring rule: registers an exponential decay function in place of CTFd's parabolic default, and refreshes the 1st/2nd/3rd-blood +10%/+7%/+5% speed-bonus awards on every new solve.
Top-3 speed bonus	The +10% / +7% / +5% multiplier on the current AI4CTF challenge value, awarded to the 1st / 2nd / 3rd solvers respectively. Scales with the challenge value as it decays — bonus drops as the field grows, preserving 1st > 2nd > 3rd > 4th ordering.
Combined total	A standard-track contestant's **AI4CTF Day 1 total + CTF4AI Day 2 total**, by raw sum. Determines standard-track Gold/Silver/Bronze/HMA rank (Section 6.1).
HMA	Honorable Mention Award — recognition for strong single-day performance without a medal.
Finalist (CTF4EAI)	One of the 15 top-scoring contestants in CTF4EAI's initial stage who qualify for the final stage.
Standard track	AI4CTF or CTF4AI — onsite-by-default, two-day, full international cohort.
Special unit	CTF4EAI — independent ranking, independent medal set, hybrid participation.

Sandbox	The 100-tool environment bundled with the ICOA CLI; the only permitted execution environment for offensive techniques.
ICOA CLI	The official command-line interface for the competition; 110 finals-mode commands.
ICOA AI Gateway	The organisers' AI proxy; the only sanctioned AI assistance during a competition window.
Computer Lab	The 2026 onsite venue — a dedicated Sydney computer lab; final address confirmed and published before the CTF4EAI late-window deadline.

Appendix C — FAQ

Can I win more than one medal in one edition?

Up to two. The standard-track medal pool is **single and combined** — a contestant earns one standard-track Gold, Silver, Bronze, or HMA based on their AI4CTF + CTF4AI combined two-day total (Sections 5.4 and 6.1). CTF4EAI runs independently as the Pioneer / Practical Round and awards its own Gold / Silver / Bronze / Finalist Award (Section 6.3). So a contestant who places in the top 1/12 by combined total **and** in the top 3 of the 15 CTF4EAI finalists takes home two medals.

There is no longer a separate "AI4CTF medal" and "CTF4AI medal" — they were merged into the combined standard-track medal in this rulebook revision.

Do remote contestants compete for the same medals as onsite contestants?

In the standard tracks: only with an approved exemption that was on file at the registration deadline (Section 8.2). New remote participation in AI4CTF / CTF4AI is no longer accepted for 2026.

In CTF4EAI: partially. The Pioneer / Practical Round is hybrid by design — both onsite and remote contestants can enter and qualify as finalists — but **Gold, Silver, and Bronze are restricted to onsite finalists**. Remote finalists receive the Finalist Award. This is because the deciding round is contested onsite; only contestants physically present in Sydney can compete in it.

What happens if the AI Gateway is unreachable during my window?

The competition does not auto-pause for AI Gateway outages. Three tiers of expected response (see Section 7.3 for detail):

- **Under 5 min:** auto-recovery, keep working
- **5–10 min sustained:** hot-fix from the server side, no clock adjustment
- **Over 30 min extreme:** organisers announce onsite whether to extend the window's end time

Workstation-local problems (CLI crash, frozen terminal) are handled by the on-duty proctor; the recommended recovery is `icoa join` again, which resumes the session.

What if I submit a wrong flag?

No point penalty. ICOA 2026's submission policy is permissive (Section 5.7): a wrong flag costs no score. There is a CTFd anti-bruteforce guard — a maximum of 10 wrong submissions per minute per challenge per contestant; if you exceed that on a single challenge, that one challenge is locked to you for 5 minutes (other challenges stay open). No human pace ever touches that limit; only automated brute-forcing does.

Can I use code I wrote before the competition?

Status: DRAFT — pending final decision at the morning-of Day 1 team-leader briefing.

The current draft position: a contestant **may** paste in their own pre-written code (e.g., personal pwntools templates, writeup-grader scripts) via the lab's controlled clipboard during the pre-competition setup window. The Science Committee is, however, actively considering a stricter rule — banning all imported code so the competition tests on-the-spot construction only — for the following reasons:

- Pasted code is difficult to verify as truly the contestant's own (vs. coach-supplied or team-shared)

- The “individual solo” rule (Section 9.3) is undercut if a coach can prepare a sophisticated template the contestant just imports
- Proctor enforcement of the clipboard channel adds operational overhead

For now, contestants should **prepare their own scripts as if they will be allowed**. The final answer will be announced by your country team leader at the morning-of Day 1 briefing. If imports are disallowed at that briefing, you forfeit nothing — the same problem set is solvable from scratch in the contest window, and the same starter tools are bundled in the 100-tool sandbox.

After the competition window starts, no new external code may be introduced regardless of which way this draft resolves.

Can my country team leader help me during the window?

No. Country team leaders are present at the venue (in a designated area) for logistics, translation, and post-window support, but they may not assist with challenge content during a competition window. Translation of an unclear English challenge prompt is permitted on request — the proctor will translate, not the country team leader.

What if I find a bug in a challenge?

Report it to the on-duty proctor immediately. The Science Committee maintains a 24/7 channel during the competition window. If the bug is confirmed, the challenge is either fixed (with affected solves recompensated) or removed from scoring (with previously awarded points retained).

Where do the medal scores get published?

Per-contestant scores are reported privately to the contestant’s national exam centre after the competition concludes. Public scoreboards show standings and medal designations but not numeric scores — a long-standing ICOA convention to focus public attention on the work rather than the number.

Where can I read more?

- **Training Book** (≈ 160 pages) — what to study (the four-part curriculum companion).
- **Three-Tracks Curriculum** — the locked card-by-card learning blueprint at `docs/three-tracks-curriculum.md` (ai4ctf / ctf4ai / ctf4eai × 12/96/360 tiers + frontier-120).
- **Selection Handbook** — country-level selection process.
- **ICOA CLI Whitepaper** — the positioning document.
- Section 10 of this rulebook — the operational track record: 322 versioned releases, Macao SAR live event, ~20-nation token-mode calibration.
- icoa2026.au — official site (project news + Macao SAR report).
- For specific rule questions, your country team leader. For rules they cannot resolve, the Science Committee at `australia@icoa2026.au`.