The Overseers Collective

DEF CON Capture the Flag
Organizers Proposal, 2025

Website https://overseers.ctf.ing
Email watchtower@overseers.ctf.ing
Social @_overseers

Primary Contact

[redacted]

[redacted] [redacted]

[redacted] [redacted]

[redacted] [redacted]

Backup Contact

[redacted]

[redacted] [redacted]

[redacted] [redacted]

[redacted] [redacted]

Table of Contents

Executive Summary

The Overseers Collective proposes to lead the next era of the DEF CON Capture the Flag competition with a complete re-imagining of how head-to-head cybersecurity contests are played, watched, and understood. Our vision retains the offensive-defensive combat that defines DEF CON CTF while addressing long-standing issues of scalability, fairness, accessibility, and spectator engagement.

We introduce a new, map-based **strategic CTF format** where every team controls a single shared avatar navigating a dynamic, shrinking game world. Each movement choice determines which services, opponents, and opportunities a team can reach. This transforms coordination, communication, and adaptability into core skills, and creates natural limits on team size and rewarding collaboration over raw manpower. The gameplay structure draws inspiration from real-time strategy and battle-royale mechanics, producing an experience that is both competitive and inherently visual.

This framework integrates four complementary challenge types (**Attack/Defense services, King-of-the-Hill bots, 1v1 speedruns, and Bonus objectives**) ensuring that success reflects both technical mastery and strategic depth. Classic exploitation remains central, but other cybersecurity disciplines like web, crypto, and reverse engineering gain meaningful impact, allowing diverse skillsets to shine.

Technically, our approach ensures **complete infrastructure transparency**. Teams upload exploits to a controlled environment where all execution, scoring, and networking are standardized. The entire platform will be open-sourced at least one month before DEF CON, ensuring a level playing field and enabling the community to inspect, test, and trust the system before the event begins.

Our organizing team combines veteran CTF players, challenge authors, and infrastructure engineers from top global teams and events. Collectively, we bring decades of experience across competition design, infrastructure operations, and elite-level play. We are independent, non-corporate, and committed to full transparency and fairness.

We are pursuing a future that is better for the CTF competition community. One that is fairer to competitors of every size and background, richer in the variety of skills it rewards, kinder to competitors' well-being, and more engaging for spectators and newcomers. The Overseers Collective seeks not just to run a contest, but to steward a lasting, positive evolution for the competition we love at the stage of DEF CON.

Rationale for a Format Revolution

DEF CON CTF has evolved continuously since its inception, with each organizing team bringing new ideas while honoring the competition's core principles. We're proposing the next evolution, a format that maintains the head-to-head combat and offense/defense dynamics that define DEF CON CTF while addressing challenges that have emerged as the community has grown.

Team-Size Impact

In recent years, we've witnessed a troubling trend: **teams consolidating into mergers to maximize their chances for DEF CON CTF**. These aren't organic teams that have grown together through shared experiences and practice, but instead strategic assemblies designed purely to maximize headcount for a single event.

Under the current A/D format, a single team divides into **completely independent sub-teams** with little to no communication between each other, each tackling a different challenge. Success becomes a simple equation: **more people = more parallel work = more points.** A *team* but in name-only.

This creates a concerning feedback loop: teams need more members to compete effectively, so organizers add more services to maintain difficulty, which makes more members even more necessary. Small teams face an insurmountable disadvantage and resort to merging purely for headcount.

We believe CTF is fundamentally a team competition. A team should win because they collaborate and communicate better, are more effective as a collective and not simply because they have more people.

Some obvious fixes have been tried:

- **Limit the number of players** allowed on a team
- **Limit the number of challenges** available at any given time

However, these approaches have inherent flaws. They feel artificial and arbitrary, they're difficult to enforce, and they don't actually change the fundamental dynamics.

Instead of imposing artificial limits, we're proposing a new format where **limitations emerge naturally from gameplay mechanics**. Our idea draw inspiration from real-time strategy war-games: **one avatar per team on a shared map.**

We believe this creates *organic diminishing returns*:

- The whole team must coordinate on **where** to go
- The whole team must agree on **what** to prioritize
- The whole team must decide **when** to move
- Poor communication becomes a *direct penalty*, not just an inconvenience

Think of it like startups competing against large corporations. A nimble startup with 10 people working in tight coordination can often outmaneuver a corporation with 100 employees who struggle with middle management overhead, communication delays, and coordination costs. The startup doesn't win because it's small, but it wins because its size allows for *better teamwork*.

Similarly, our format creates natural coordination challenges that reward efficient communication over raw headcount. However, this isn't about punishing large teams with good coordination: if a team is large **and** their teamwork is exceptional, they'll still do well. What we're addressing is teams that merge purely for headcount, where sub-teams work independently with minimal coordination.

Service Lifecycle

Beyond team size dynamics, there's another critical issue: **the unpredictability of challenge lifecycles**.

Right now, teams never know when a service will appear or disappear. This uncertainty creates several problems:

First, it introduces hidden unfairness. Imagine two equally skilled teams: Team A decides to focus on Service A, while Team B focuses on Service B. Both make smart strategic decisions based on the information available. But then Service B retires after 8 hours while Service A stays live for another day. Team A accumulates significantly more points, while Team B simply picked the "wrong" service to prioritize. The teams had no way to prevent this from happening.

Second, it gives organizers too much discretionary power. When retirement times aren't announced, organizers can adjust on the fly. If a service gets solved faster than expected, they might retire it early. If it remains unsolved, they might extend it. While these decisions may seem reasonable in the moment, they fundamentally alter the competitive landscape in ways teams can't anticipate or plan around.

Third, and perhaps most importantly, services that run overnight create what we call "CTF homework." When challenges don't retire at day's end, teams face an impossible choice: keep working through the night and be exhausted or go to sleep and fall behind. For a competition that spans 2.5 days, this means teams are effectively expected to function for **52 consecutive hours** without real rest.

While some competitions are designed around endurance, such as ironman triathlons or multi-day sailing races, **we don't believe CTF should be one of them**.

Our proposed format addresses these issues through simple, transparent mechanics:

- **All challenges are visible from the start** as physical locations on the map. Teams can see what's available and plan accordingly. Strategic depth comes not from guessing *what* will appear, but from deciding *where* to go and *when*.
- **Challenges retire naturally through map mechanics.** As the map shrinks toward the center (inspiration by battle royale), outer challenges go out of bounds. This happens on a *predictable schedule* that all teams can see and plan around.
- **We use a completely new map each day.** Day 1 ends, everyone rests, and Day 2 starts fresh with new challenges. No homework. No overnight grinding. No lose-lose choices between rest and competition.
- **Difficulty is tied to map position.** Challenges near the edges (where teams start) are easier and more numerous. As teams move toward the center, challenges become progressively harder and scarcer, creating natural progression.

Instead of simply fixing a scheduling problem, we create a **transparent, fair, and sustainable competition structure** where strategic decisions are based on visible information, not blind guesses about organizer intentions.

Challenge Diversity

Walk into any recent DEF CON CTF finals, and you'll find that the **competition overwhelmingly focuses on binary exploitation**.

The nature of A/D format naturally gravitates toward pwn challenges because they translate well into vulnerable services that teams can patch and exploit. We get some web challenges occasionally, but they're more like outliers. As for other categories?

- **Reverse engineering** becomes a support skill, instead of a main focus
- **Cryptography** is usually something you do *on the way* to pwning a service
- **Forensics** and **OSINT** are essentially nonexistent

This creates a strange situation: DEF CON CTF, supposedly the "World Championship" of cybersecurity competitions, ends up testing a relatively narrow slice of the skills that make up modern cybersecurity work.

Now, there's an obvious alternative: **Jeopardy format**. Jeopardy CTFs can include any category imaginable. But Jeopardy competitions lack the head-to-head combat element that defines DEF CON CTF. There're no offense and defense. There's no dynamic interaction between teams.

However, here's where the cognitive disconnection happens: **every team that qualifies for DEF CON CTF gets there through Jeopardy-style competitions,** even the DEF CON CTF qualifier event is Jeopardy. These teams have members who specialize in web exploitation, cryptography, forensics, reverse engineering, every category you can think of.

Then they show up at finals, and suddenly **most of their team can't play to their strengths.** The web specialist might get lucky if there's a web service. The crypto expert helps with some challenge aspects but isn't solving challenges end-to-end. All other "misc" players are basically along for the ride, hoping their general problem-solving skills can contribute somewhere.

Our proposed format attempts to bridge this gap by incorporating **four distinct competition types**:

- **A/D services** remain the backbone. These work similarly to traditional A/D services: teams patch vulnerabilities, write exploits, and attack each other's instances. The key difference is teams must position their avatar near a service to interact with it, creating strategic trade-offs.
- **King of the Hill (KotH) challenges** where teams write bots that compete in defined games or simulations. This is already a well-explored concept presented in multiple CTF competitions before.
- **1v1 challenges** where two teams face off on a randomly selected challenge when they meet on the map. Both teams get the same challenge; first to solve wins points. This is similar to LiveCTF from previous DEF CON CTF editions.

- **Bonus speedrun challenges** that appear at random locations on the map, visible to everyone. First team to arrive and solve gets the points.

This combination creates something unique: **we maintain the "head-to-head combat, offense and defense" character that defines DEF CON CTF while expanding what skills contribute to success.** A/D services and KotH bots provide the continuous competitive pressure. 1v1s and Bonus challenges inject burst moments of direct competition across diverse categories.

From a team strategy perspective, this opens up multiple viable approaches. A team strong in pwn might focus heavily on A/D services while being selective about 1v1s and Bonuses. A well-rounded generalist team might do solid work across A/D without dominating it, then maximize points through diverse challenge types where their breadth gives them an edge. **Both strategies can succeed because the format creates multiple paths to victory.**

Spectator Appeal

We have another uncomfortable truth: **most people find CTF competitions boring to watch.**

Even within the cybersecurity community, if you haven't competed in CTFs yourself, watching one is about as exciting as watching someone debug code for eight hours. Which, let's be honest, is exactly what it often looks like.

Every year, DEF CON draws thousands of attendees, media coverage, and industry attention. Yet most spectators have no idea what's actually happening. They see people staring at screens, hear occasional cheers, might catch a glimpse of a scoreboard, but the decisions being made and strategies deployed remain largely invisible.

We believe this is a missed opportunity: not just for entertainment, but for education, outreach, and growing the cybersecurity community.

In recent years, LiveCTF has emerged as an attempt to solve this problem. The concept is clever: take two players from competing teams, capture their laptop screens, and have them speedrun a challenge while everyone watches.

We believe it's a step in the right direction, but it reveals the fundamental challenge of making technical work spectator-friendly: **The work itself is too specialized.** When watching a skilled player's screen, you see terminal commands flying by, decompiler output, hexadecimal dumps, custom tools with inscrutable interfaces. Someone with CTF experience can follow along, but for everyone else it's like watching professional chess without knowing the rules, except everything is in assembly code.

Commentating for this is difficult as well. Things either happen too fast to explain, or nothing visibly interesting happens for long stretches while players think through problems. Most recently in 2025, the winning team used AI that solved challenges in the background with no player interaction. From a technical standpoint this was impressive, but the commentator had nothing to talk about, and audiences had nothing to watch.

This isn't LiveCTF's fault. It's a limitation of trying to make inherently technical work watchable. **You can't easily visualize what happens inside someone's head or inside their custom tooling.**

At the same time, the current A/D format offers limited options for creative visualization. We've seen scoreboards, first-blood announcements, and service status indicators. These help, but they're supplementary information, not a narrative core that audiences can follow and understand.

Our format doesn't try to make the technical work itself more visible. What changes is that we add a layer that *is* visible: the strategic decisions teams make about what to prioritize and when to act.

Here's what spectators will actually see:

- **Team avatars moving across a map.** When a team decides to explore a new area, leave a current challenge, or rush toward a time-sensitive objective, people can *see it happen*. The strategic choices that were previously invisible in server logs and Discord channels now play out visually in real-time.

- **Challenge locations and their status.** Services under attack, patches being applied, teams fighting over bonuses, all of this can be visualized as parts of the game board. Even without understanding *how* a challenge works, people can *see* that teams are fighting over it.
- **Team-versus-team encounters.** When two teams' paths cross and they initiate a 1v1, that becomes a discrete, comment-able event with clear stakes and outcomes.

You shouldn't need to understand buffer overflows to find this compelling. The strategic layer is immediately visible. The commentator's job also becomes easier because they can choose their focus:

- **High-level strategy:** "Notice how Team A is staying near the center, focusing on harder challenges, while Team B is heading towards the outer regions. This is a classic risk-versus-reward trade-off."
- **Service status updates:** "The web service in the southwest has three teams currently patching it. Let's look at the current exploit success rates..."
- **Challenge explanations:** "This challenge requires exploiting a race condition. Here's what that means..."

The commentator can pivot between these levels based on what's happening and what the audience needs. When the action on the board is intense, they can focus on movement and decisions. When things slow down, they can dive deeper into technical explanations.

The format also has strong theming potential. The game board can take on different visual styles (a fantasy dungeon, a sci-fi station, a spy facility, etc.) that help people understand mechanics intuitively through familiar concepts, which makes it more interesting to watch.

To be clear: we're not trying to dumb down CTF. The technical challenges remain exactly as difficult. What changes is that **the container holding those challenges becomes watchable.** The strategy layer becomes visible. The competition creates moments that can be captured, narrated, and remembered.

Infrastructure Transparency

Last but not the least, there's another layer of unfairness built into the current format that has nothing to do with challenge difficulty or team skill: **the infrastructure itself.**

In traditional DEF CON CTF A/D, teams throw their own exploits. This sounds reasonable but creates hidden disadvantages:

- Organizer infrastructure can be **overwhelmed by aggressive exploitation**, leading to unintentional DoS that teams can't distinguish from actual defense
- **Rate limiting algorithms vary wildly** between what teams expect and what organizers implement
- **Critical infrastructure decisions get made without warning.** Most recently in 2025, organizers switched to random service ports per team instead of fixed ports used in previous years, forcing teams to patch their frameworks on the fly at the start of Day 1.

Imagine being a first-time finalist. You earned your spot through pre-qualification but never competed at DEF CON CTF before. You don't know about these infrastructure quirks, what rate limiting is in place, or what connection patterns will work. **Two teams with identical technical abilities can have dramatically different experiences** based solely on whether they correctly guessed how the infrastructure would behave.

Our solution is straightforward: **teams upload their exploits to our infrastructure, and we run them in a standardized, containerized environment.**

We see a few advantages with this approach:

- **Equal computational resources.** All team's exploits run with the same CPU, memory, and network access. This limitation can be done in-kernel to ensure they are fair across all teams.
- **Predictable network behavior.** Since all connections are local, we can guarantee consistent network performance. If an exploit fails, it's because the exploit has a bug, not because the network is congested.

We also don't expect many changes on the team's part, as we will support teams uploading a standard container image, so the environment teams develop against locally will match the competition environment exactly.

And here's the critical part: **we're committing to open sourcing our entire infrastructure codebase and providing detailed usage guides at least one month before the CTF.**

This means:

- Teams can predict *exactly* how exploits will be executed
- They can test their code against the same environment they'll face during the finals

- If we need to make changes, the diff is public, and everyone adapts together
- There are no hidden assumptions or undocumented behaviors
- New teams and veteran teams start on truly equal footing

We're not claiming our infrastructure will be perfect. But by making it transparent, we ensure that when problems arise, they can be identified, discussed, and fixed openly rather than leaving teams to suffer silently through implementation quirks they don't understand and can't change.

Implementation Details

This section contains the implementation and technical game design details regarding our DEF CON CTF proposal on our new format.

Basic Structure

- **Location:** On-site at DEF CON
- **Teams:** 12 qualified teams
- **Team size at venue:** No limit on team size, but physical space accommodates 8 players per team at designated tables
- **Remote collaboration:** Allowed and encouraged
- **Schedule:**
  - **Day 1:** 10:00 AM - 6:00 PM (8 hours)
  - **Day 2:** 10:00 AM - 6:00 PM (8 hours)
  - **Day 3:** 10:00 AM - 2:00 PM (4 hours, to accommodate the award ceremony)
- **Network access:** Each team receives a single LAN connection to the game network

Each team will be required to have at least one person on-site to handle setup, provide off-site team members with access to the game network, and act as a representative of the team.

The game will not run overnight to allow teams to review their strategy, which our new format necessitates, and rest adequately.

Scoring System:

The final ranking is determined by **Team Points**, calculated as:

Team Points = Day 1 Game Points + Day 2 Game Points + (2 × Day 3 Game Points)

Day 3 Game Points are weighted double because:

- The competition is shorter (4 hours vs 8 hours)
- The map is more compact, intensifying interactions
- This prevents teams from "coasting" if they built a lead in earlier days

Tiebreaker: Teams with identical Team Points share the same rank. No secondary tiebreakers are used.

Game Board Design

The game takes place on a sector with a **30-degree central angle** (imagine a slice of pizza, see drawing below).

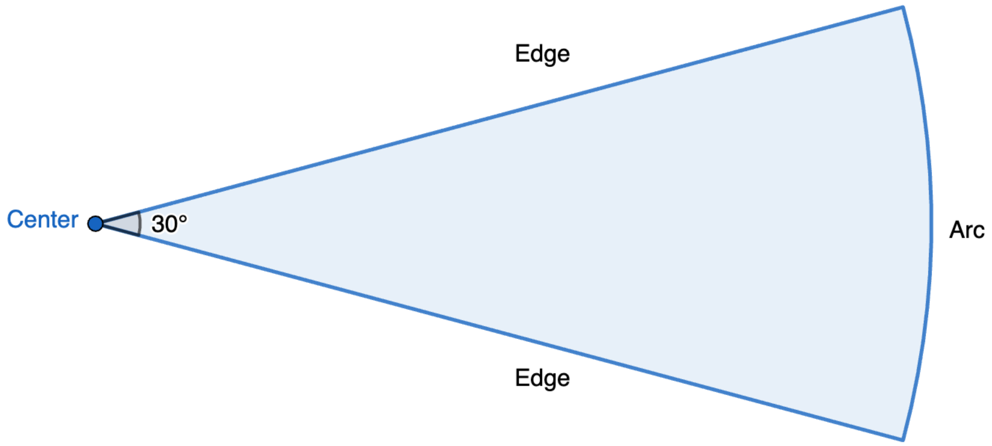

A single team is represented as an avatar (a point) on the game board.

Avatar Movement model:

- Continuous 2D movement
- Constant movement speed for all teams
- Open-world navigation with no movement restrictions
- Position updates processed at 1-second intervals

From now on, when we describe distances regarding game board, **we measure in time rather than units**. When we say “5 minutes,” we mean the distance an avatar can cover in 5 minutes of movement. This makes it intuitive to understand travel costs.

Map Dimensions:

- **Day 1 & 2:** Edge length of 240 minutes (4 hours of travel time from arc to center)
- **Day 3:** Edge length of 120 minutes (equivalent to cutting the outer ring of the normal size map)

Initial positioning: Team avatars spawn evenly spaced along the outer arc in random order.

Edge wrapping: When an avatar reaches either edge of the wedge, it wraps around to appear at the opposite edge. Think of it as the two straight edges being connected, or imagine rolling the wedge into a cone in 3D space.

Vision range: 1 minute radius circle around each avatar.

Team visibility: Teams cannot see each other's positions until within vision range.

Challenge visibility: All teams can see:

- Challenge type (A/D, KotH, Bonus)
- Challenge title
- Exact location on map

For specific visibility rules, please refer to “Challenge Types” below.

Shrinking arc boundary:

- **Starts:** 1 hour after game begins
- **Speed:** Half of avatar movement speed
- **Effect:** Creates a boundary that moves inward from the arc toward the center
- **Reward:** Teams located inside the shrinking boundary gain 1 Game Point per 30 seconds

Note that teams can move freely in and out of the shrinking arc, and they will not be teleported or disqualified.

The shrinking boundary and vision range will be relevant for specific types of challenge below.

Challenge Types

Attack/Defense Services

A/D services are continuous, high-intensity challenges where teams simultaneously defend their own vulnerable services while exploiting others.

Placement and visibility:

- Fixed locations on the map
- All services for the day are present and visible when the game begins
- Difficulty generally correlates with distance from the outer arc (harder services toward center)

When a service is within vision range, teams can:

- Download service handout (source code or binary)
- See which other teams have enabled this service
- Access the **service network** to exploit other teams' instances
- Submit flags captured from other teams
- Monitor **real-time packet captures** or download historical pcaps
- **Submit patches** to fix vulnerabilities in their service
- **Upload and update exploit code**

When teams go **outside vision range:**

- Teams can view attack status for services they've uploaded exploits to
- Their uploaded exploit keeps running
- Their service patches remain active

Teams cannot see (regardless of range):

- Which teams successfully exploited them
- Real-time defensive status beyond SLA checks (explained below)

Uploaded exploits have the same capabilities as the team. This means they can even analyze packet captures or submit new patches. However, they **cannot access the internet**.

Think of an exploit as a “second avatar” that acts autonomously but can't move or be directly controlled in real-time after upload.

State transition: For each team, every service starts in a **disabled state**. When a team first download service handout, the service for that team turns into enabled state starting the following round and stay enabled even if teams move out of vision range.

Patch timing: If a team submits a new patch during round t, $t$he patched service becomes active in round $t + 1$.

Exploit timing: When a team submits a new exploit, the old exploit will be stopped and the new exploit will immediately starts running.

Scoring mechanism

If a service is disabled for a team, then the team gets no game points from this service. Otherwise:

Scoring operates on **1-minute rounds**. For each round and each team A that has the service enabled:

Defense points (1 Game Point each):

If another team B has the service enabled, AND

Team B did NOT successfully exploit team A's service

→ Team A earns **1 defense point** against team B

Attack points (1 Game Point each):

If team B has the service enabled, AND

Team A successfully exploited team B's service

→ Team A earns **1 attack point**

SLA (Service-Level Agreement) penalty:

If team A's service fails the availability check (SLA)

→ Team A earns **zero defense points** for that round, regardless of exploitation status

Service lifecycle:

- Services continue running until they appear **out of the shrinking arc boundary**
- Effectively, they stop generating new rounds (retired)

King of the Hill Challenges

KotH challenges represent **bot-writing gameplay** where teams create automated agents that compete in hacking mini-games or simulations.

Placement and visibility:

- Fixed locations on the map
- All KotH challenges for the day are present and visible when game begins
- Difficulty generally correlates with distance from outer arc

Initial state:

At the start, every team has a "null" bot implementation for each KotH challenge. This null bot:

- Performs deterministically (no randomness)
- Guarantees exactly **0 game points** for all teams (see scoring below)

This ensures no team has an advantage before writing their own bot

When a KotH is within vision range, teams can:

- Download challenge handout (mini-game instructions)
- See which other teams have submitted custom bots
- Observe **live game status** (watch the current game play out)
- Download **historical game records** (review past rounds)
- **Upload and update their bot implementation**

When teams go **outside vision range:**

- Teams can **view rankings** for each round
- Team's submitted bots continues running

Bot capabilities:

Similar to A/D exploits, bots have autonomy but limitations:

- Main job is playing the mini-game according to the game's rules
- Can also observe game status and see rankings
- **Cannot access the internet**

Bot timing: If a team uploads a new bot during round t, the new bot starts playing in round $t + 1$.

Challenge lifecycle: Similar to A/D Services, KotH challenges continue running until they go **out of shrinking boundary.** Afterwards, they stop playing out new rounds.

Scoring mechanism

Scoring operates on **1-minute rounds**. Each round, the mini-game plays out and produces a ranking. Teams earn game points based on their rank:

When all teams have different rankings:

- Points = **12 minus rank**
- 1st place = 11 points, 2nd place = 10 points, …, 12th place = 0 points

When teams are tied:

Tied teams share the same points, and teams below them get shifted down. The last place team gets 0 points.

Example 1: One team in first, everyone else tied for second:

- 1st place: **1 game point**
- Everyone else: **0 game points**

Example 2: Two teams tied for first, everyone else ranked differently:

- Both 1st place teams: **10 game points each**
- 3rd place: 9 points
- 4th place: 8 points
- (and so on...)

Example 3: Everyone tied for first place:

- All teams: **0 game points**

Bonus Challenges

Bonus challenges are **high value speedrun challenges** that appear at scheduled times throughout the day. They inject moments of intense, focused competition.

Spawning schedule:

- One bonus spawns every **60 minutes** at the 30-minute mark
- For an 8-hour game: **7 bonuses total**
- For a 4-hour game: **3 bonuses total**

Challenge categories: Each bonus is randomly selected from standard CTF categories: pwn, web, reversing, crypto, or misc.

Spawn location rules:

When a bonus is about to spawn, we calculate its location using computational geometry with these constraints (applied in order of priority):

1. **Fairness:** Must be equidistant to the two closest team avatars (no team gets a distance advantage)
2. **Minimum distance:** Must be at least **5 minutes** away from any team avatar
3. **Safety from shrinking:** Must be at least **15 minutes** away from the current shrinking boundary
4. **Optimization:** Among all positions satisfying rules 1-3, choose the location that **minimizes the sum of distances to all 12 teams**

If multiple solutions exist after these rules, one is selected randomly.

Announcement:

When bonus spawns, all teams immediately see:

- Exact map location
- Challenge category
- Challenge title

When a Bonus is within vision range, teams can:

- **Download challenge** handout
- See which other teams have downloaded the challenge and when
- **Submit the flag**

Teams **cannot interact** with the bonus outside vision range (but they can always return to it).

Challenge lifecycle:

Unlike other challenge types, bonuses **do not retire when going out of bounds**. Teams can still submit flags even if the bonus location is outside the shrinking boundary.

Challenge design:

Bonus challenges are designed to be solvable in approximately **30 minutes** by a skilled team.

Scoring mechanism

- **First correct flag submission: 600 game points**
- All subsequent submissions: **0 game points**

This is a winner-takes-all scenario. The first solve gets a massive point injection, but there's no second-place prize.

1v1 Speedrun Challenges

1v1 challenges introduce **direct team-versus-team combat** that can happen anywhere on the map. They're spontaneous, high stakes encounters when teams' paths cross.

Initiation: When two team avatars are within **vision range**, one team can initiate a 1v1 challenge to the other.

Location restrictions: 1v1s **cannot** be initiated within **1 minute** of any:

- A/D service
- KotH challenge
- Bonus challenge

This prevents teams from being forced into 1v1s while working on other challenges.

Response phase:

Once a challenge is initiated, **both avatars freeze in place**. The challenged team has **30 seconds** to decide:

If they REJECT or don't respond:

- Initiating team gains **10 game points**
- Both avatars unfreeze and can move

If they ACCEPT:

- Both teams proceed to category Selection Phase
- Avatars remain frozen

Category Selection Phase:

This is a **draft-style** selection similar to competitive games like CS:GO:

1. **Passive team** (who accepted) **bans one category** from {pwn, web, crypto, rev, misc}
2. **Active team** (who initiated) **bans one category** from the remaining four
3. Repeat: Passive bans one, Active bans one until **one category remains**

Challenge draw:

A random challenge is drawn from that category's pool, with one critical rule:

The challenge must be **unseen by both teams**. If both teams have seen all challenges in that category (because they have played it in previous 1v1), it shows as **"out of stock"** during the category selection phase and cannot be chosen.

Edge cases:

- If both teams initiated simultaneously (same server tick): randomly select one as "passive team," and the accept phase is skipped (both are assumed to accept)
- If only one category is not "out of stock" for both teams: the Category Selection Phase is skipped, and that category is selected
- If ALL categories are "out of stock" for both teams: these two teams can no longer 1v1 each other for the rest of the day

Competition phase:

Once the challenge is determined:

- Both teams receive challenge handout
- **First team to submit correct flag wins**
- Winner receives **300 game points**
- Loser receives **0 game points**

Both teams remain **frozen in place** throughout the 1v1.

Alternative endings:

- **Forfeit:** Either team can forfeit at any time. The other team immediately wins and gets 300 points.
- **Draw:** Either team can propose a draw. If the other team agrees, the 1v1 is cancelled and no one gets points. This requires mutual agreement.

Post-1v1 cooldown:

After an 1v1 ends (through any means: reject, win/lose, forfeit, or draw), these two teams **cannot initiate another 1v1 with each other for 5 minutes**.

However, they can **immediately challenge other teams** if they encounter them.

Challenge design: 1v1 challenges are designed to be solvable in approximately **15 minutes** by a skilled team.

An example game board at the beginning of the competition may look like this

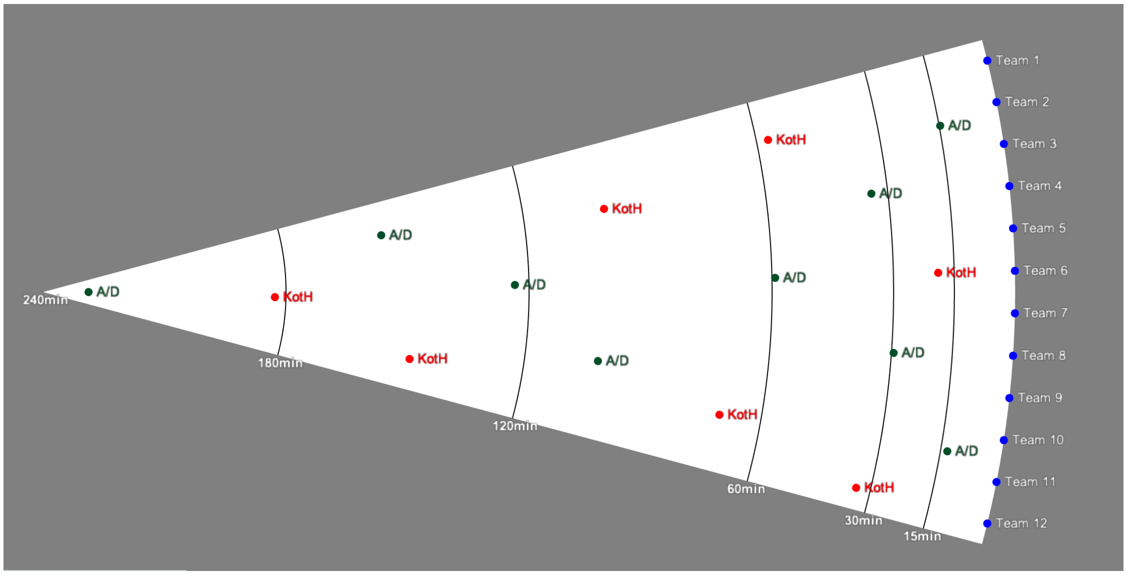

While 2 hour and 30 minutes into the game, the board state may look like this

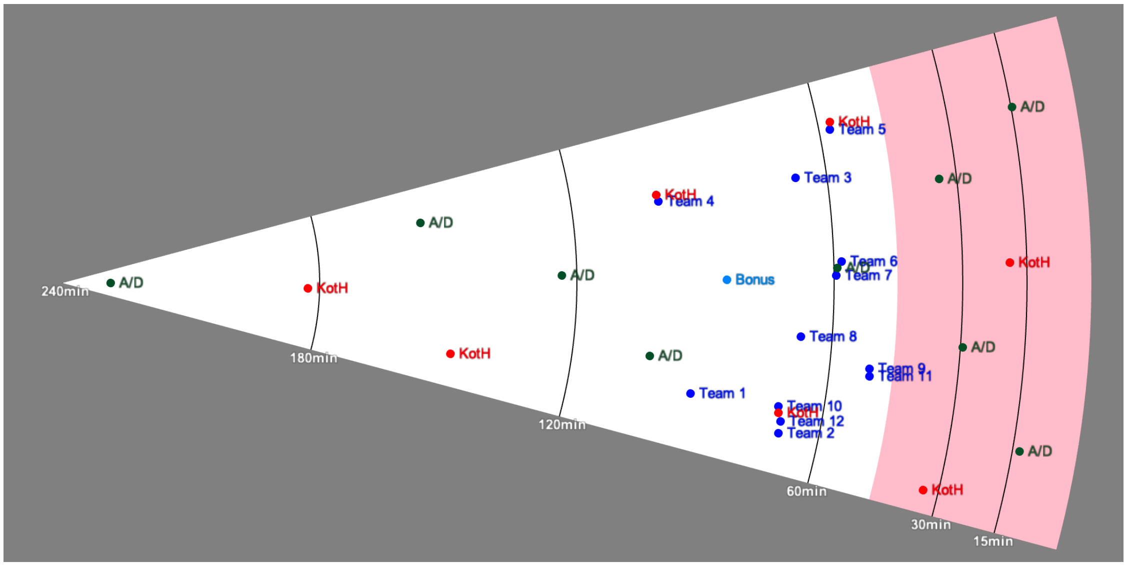

Network Architecture

Physical Setup:

Each team receives **a single LAN cable** at their designated table. This cable provides:

- Access to the game network (for all CTF challenges and infrastructure)
- Access to the internet (for research, external tools, and team communication)

Behind the scenes, we handle the routing internally, so teams don't need to manage multiple connections or configure complex network setups.

Network Isolation:

The game network operates as a **completely separate system** from DEF CON's networks:

- **Game infrastructure** (challenge servers, scoring systems, packet capture storage) runs on dedicated hardware isolated from the DEF CON network
- **Internet access** for teams comes through a dedicated uplink requested from DEF CON, separate from the convention's general network
- **No crossover:** Game traffic never touches DEF CON infrastructure, and vice versa

Security Measures:

- Teams **cannot** access each other's networks or machines directly
- All team-to-team competition (A/D exploits, KotH bots) goes through our controlled game infrastructure
- Packet captures and network monitoring only cover game traffic, ensuring team privacy for non-game activities

What we need from DEF CON:

- **One wired internet connection** for participant use
- **Physical space** to rack our game servers separate from DEF CON equipment
- If possible, 12 network cables and power strips that is setup pre-competition and secured under floor or under mattress

This straightforward design means teams can focus on competing rather than wrestling with network configuration, while we ensure complete separation from DEF CON's networks.

As required by the CFO, we will publish sanitized packet captures and metadata (per-day, per-service) to a public GitHub repository within 7 days of DEF CON close. These packet capture files are similar to what will be provided to the teams.

User Interface and Controls

Team Interface

Each team receives access to a **web-based control panel** that serves as their primary way of interacting with the game. This interface runs in any modern browser and requires no special software installation.

Core features:

- **Map view:** A live visualization of the game board showing:
  - The team's current avatar position
  - All challenge locations (A/D services, KotH, Bonuses) with their types and titles
  - The current shrinking boundary position and movement
  - The team's vision range (1-minute radius circle around the avatar)
  - Other teams' avatars **only when they enter the vision range**
- **Movement control:** Use arrow keys or WASD keys to move the avatar around in constant speed.
- **Challenge interaction panels:** When within vision range of a challenge, context-specific controls appear. These controls show UI element for functionality that is explained above for each challenge type.
- **1v1 management:** When a team appears in the vision range, a popup will display asking for 1v1 initiating. For the passive team, a modal interface appears with Accept/Reject buttons with countdown timer. Both teams then see:
  - Category ban selection during category selection phase
  - Challenge handout and flag submission once the challenge begins
  - Forfeit and Draw proposal options
- **Game status dashboard:** Real-time display of:
  - Current Game Points for all teams
  - Time remaining in the current day
  - Active challenge status (which services you've enabled, the team’s KotH rankings, etc.)
  - The team’s recent point gains and their sources

Commentator Interface

Commentators receive an **enhanced version** of the team interface with additional visibility and analytical tools. This runs on the same web platform but with elevated permissions.

Additional capabilities:

- **Full map visibility:** See all 12 team avatars at all times and their historical movements
- **Movement history replay:** Scrub backward through time to see how teams moved and what decisions they made.
- **Challenge status overlays:** Visual indicators showing:
    - Which teams have enabled which A/D services
    - Current KotH rankings for each challenge
    - Who has downloaded/solved Bonus challenges
    - Active 1v1 matches (who's fighting, what category, timer)

On-site spectators see a **1-hour delayed version** of the commentator interface projected on screens, while potential online viewers see the same commentator interface live broadcasted with the same delay.

The commentator will commentate on the game in real-time, and the commentated version with commentator face/voices is live edited into a "highlight" reel. After the 8-hour Game Day ended, the live stream will be replaced with the 8-hour full recording of the commentator POV, and the highlight reel will also be uploaded.

Visual Theme and Storytelling

We haven't yet finalized our art design, but the idea is that there's an underlying **circular** map. However, the whole map will not be presented directly, instead:

- The map is divided into **4 quarters**, and each quarter is divided further into **3 slices**
- For each game day, the game board we reveal is actually **one slice of the map** (therefore 30 degrees = 1/12 of the circle)
- Therefore, for each year of DEF CON, we will reveal a quarter portion of the map.
- Assuming we are running DEF CON CTF for 4 years, by the end of the fourth year we will have revealed the whole map.

Each quarter features a different **"-punk" aesthetic:**

- **Cyberpunk** (neon lights, high-tech)
- **Steampunk** (brass gears, industrial revolution)
- **Biopunk** (genetic engineering, organic technology)
- **Atompunk** (retro-futuristic, atomic technologies)

Where we can design challenges that fit into the corresponding aesthetic.

This design features an **underlying narrative**, that will be revealed at the end of the fourth year:

On the outside, it seems that every Game Day's map represents a different city that has the same wedge shape. However, all cities are actually a part of a **panopticon institute**, divided into 12 pieces and controlled.

The edges of the wedge shape that wraps around was a device to make people think that they are experiencing the entire city, whereas all wedges are connected and make up the entire circular institute.

Each day's competition is hacker **group's attempt** trying to reach and hack the central watchtower. When the watch tower is hacked from all 12 sides, the central watchtower is compromised and the walls separating all sections can be taken down. And finally, the city is freed from surveillance and control.

Community Q&A

To make our proposal more rigorous and fairer, we invited around 40 members from the community consisting of current CTF players, team captains, and organizers to review our format design. We held open discussions where they could raise questions and concerns about the proposal details.

While most of the questions are explained above in **Rationale** and **Implementation Details**, there are some meta questions regarding strategy, balance and philosophy that we feel are worth mentioning here. Therefore, we've compiled the most important questions and our responses below. We hope they help address questions you might have as well.

Q: What stops teams from just visiting all challenges and working on them offline?
A: The map and movement speed are designed so teams choosing this strategy won't have enough time to effectively move back and forth. Teams focusing on 2-3 challenges while following the pace earn about the same points as teams rotating between many challenges.

Q: Won't large teams still have an advantage by solving more challenges?
A: Large teams can solve more challenges, but the single avatar creates a coordination bottleneck that limits how those solves translate to points. A large team must achieve consensus on where to go, what to prioritize, and when to move, and poor coordination directly costs points through wasted travel time and missed opportunities.

Q: Isn't camping a single challenge the best strategy?
A: No — camping isn't ideal because time is needed to solve challenges. The optimal strategy should combine movement and exploitation. Movement is encouraged for visiting new services and challenges, 1v1 opportunities, and speedrun bonuses on the map.

Q: Won't the A/D lose its intensity if we can't constantly update exploits and patches?
A: In practice, based on previous DEF CON CTF, major exploit and patch updates happen every 20-30 minutes at most, not every round. This means teams can develop new exploits or patches offline during travel, then quickly visit the service to deploy them. The current format also encourages teams to develop semi-autonomous exploit framework that can reflect exploits or patches the service based on network traffic.

Q: Won't this introduce a lot of luck based on which challenges teams encounter?
A: Several mechanisms mitigate randomness: all teams spawn equidistant from the center, challenges are visible from the start (teams choose their path), bonuses spawn with strict fairness rules requiring equal distance to the nearest teams, and the 20-hour total competition duration smooths out short-term variance. Some tactical luck remains, as in any competition, but strategic planning matters more.

Q: Isn't strategizing the board traversal an algorithm challenge that could determine winners?
A: While teams can model the board mathematically, the environment is too dynamic for pure algorithmic solutions. The navigation decisions require human judgment about the team's strengths, current point standings, and risk tolerance. We see route planning as part of team strategy, not a separate skill.

Q: Won't this just reward having dedicated strategy people separate from challenge solvers?
A: Teams are free to do that. However, if strategy people and challenge solvers do not communicate well, it is very easy to result in an internal discord. Imagine a part of the team starts to solve the challenge, but the strategy people move the avatar somewhere too far that renders the solver's work meaningless.

Q: What if a team makes bad strategic decisions early? Can they recover?
A: Multiple factors enable recovery: the competition spans three separate days with fresh maps (a bad Day 1 doesn't doom you), Day 3 points are weighted 2× to encourage for final sprint, and there are multiple paths to victory through different challenge types. A team that commits to the wrong area can pivot to bonuses, hunt for 1v1s, or relocate as the map shrinks. No single decision is catastrophic.

Q: Can teams collude or team up against the leading team?
A: Explicit collusion (like sharing challenge solutions or coordinating attacks) is prohibited in the rules and can lead to disqualification. However, emergent strategies like "everyone attacks the leader" are part of competitive dynamics and we don't plan to specifically forbid it.

Qualification Structure

Our qualification structure largely aligns with that of previous organizers.

The qualification cycle will result in 12 teams qualified for the DEF CON CTF Finals held in Las Vegas.

Automatic Qualifications (4 spots):

- **Winner of previous year's DEF CON CTF** will be automatically pre-qualified.
- Three of the most well-regarded CTFs the community has to offer will serve as pre-qualification events, whose winner is invited for the finals.
    - **Plaid CTF** — established in 2011, held annually by PPP, the academic CTF team of Carnegie Mellon University in Pittsburgh, Pennsylvania. Currently holding a rating weight of 100 on CTFtime.
    - **hxp CTF** — established in 2013, held annually by hxp, a German CTF team. Currently holding a rating weight of 98.14 on CTFtime.
    - **0CTF** — established in 2015, held by 0ops, an academic team and a security research group from the Shanghai Jiao Tong University. Currently holding a rating weight of 100 on CTFtime.

For the first year, depending on the timeline of the announcement of the new DEF CON CTF organizer, the pre-qualifier CTF might already have happened and therefore not fair to retroactively be selected as pre-qualifiers. In which case, we'll either choose another well-regarded CTF as a pre-qualifier, or we'll add an extra spot to the DEF CON CTF Qualifier (explained below).

If a winner is already qualified or they reject the invite, the spot moves to second place, then third place, and so on.

DEF CON CTF Qualifier (8 spots):

- **When:** Early to mid-May (approximately 3 months before DEF CON)
- **Format:** 24-hour online jeopardy-style competition
- **Categories:** Web, pwn, crypto, reversing, misc
- **Team size:** No limit
- **Qualification:** Top 8 teams advance to finals

Qualified teams must accept within 14 days. Declined or no-response slots are offered to the next placed team from that qualifier in ranking order. Final participant list is locked 30 days prior.

In the case that a qualified team has already earned a spot from a prior event, the qualification carries over to the next in placement. This ensures we always have exactly 12 teams at finals.

Team & Logistics

Our organizing team brings together **CTF players, team captains, and organizers from around the world**, united by a shared vision for DEF CON CTF's future. We are not from a single pre-existing team, but assembled specifically for this proposal, drawn together by our belief that the competition needs meaningful evolution.

All members have **years of experience participating in competitions** both online and on-site globally. Beyond playing, we've all authored CTF challenges and have **years of experience organizing CTF competitions**.

We've experienced the competition from both the player's and organizer's perspective. We know what works and what frustrates teams, and we understand the pressure of competing on that stage. This combination of perspectives means the format we designed will work not only just in theory, but in practice as well.

Core Committee

SuperFashi ([redacted]) has been part of the cybersecurity community since 2011. Starting with iOS and game console jailbreaks, he developed expertise in reverse engineering over the following years. He joined the CTF community in 2021 with *Tea Deliverers* and qualified for DEF CON CTF that same year. After becoming co-captain of *Water Paddler* in 2021, he helped merge it with *perfect blue* in 2022 to form *Blue Water*, where he served as co-captain and team manager.

Under his leadership, *Blue Water* became one of the world's strongest CTF teams, winning first place at highly regarded competitions including PlaidCTF, HITCON CTF, and 0CTF, while qualifying for DEF CON CTF finals every year. The team ranked #1 globally on CTFtime.org in 2023. After DEF CON 33, SuperFashi left the team to focus on his vision for a new DEF CON CTF format.

Individually, SuperFashi is an active figure in the reverse engineering community. He's been on the Flare-On Challenge Hall of Fame for five consecutive years (2021-2025) and has authored multiple award-winning writeups. He has also created various reversing and misc challenges and organized CTF events including IERAE CTF and BWCTF.

anatomic ([redacted]) has been active in the cybersecurity community since 2017, when he first decided to dabble in the art of game cheat development. He joined the CTF community in 2022 to hone his skills, initially competing solo, because he didn't know anyone else in the field. With time, he developed expertise in binary exploitation with a passion for kernel exploitation.

In 2023 he was part of the inaugural edition of the Bulgarian National Cybersecurity Team, representing Bulgaria at the European Cybersecurity Challenge. He was elected Captain the following year and has led the team since. Beyond competition, anatomic has been active in the development of the CTF scene in Bulgaria, authoring challenges, contributing to events and delivering technical talks, including at the largest cybersecurity conference (BSides Sofia) in Bulgaria.

es3n1n ([redacted]) began reverse engineering through game hacking in 2019, focusing on software protection analysis, DRM, and tooling. He entered CTFs in 2023 by founding *cr3.mov*, and in 2024 joined both *Project Sekai* and *Friendly Maltese Citizens*, quickly becoming a core contributor.

As of 2025, he serves as an admin for both teams. He is also an organizer, challenge developer, and infrastructure engineer for events including Sekai CTF, Malta CTF, idekCTF, and others, with focusing on reversing challenges and reliable, scalable infrastructure. Beyond competitions, he has published research on commercial DRM, software obfuscation and maintains multiple open-source projects closely related to reverse engineering.

quasar ([redacted]) is a player who started CTF in 2021, where she first started doing cryptography. She currently plays for *.;,;.*, *Project Sekai* and *Friendly Maltese Citizens* as a cryptographer. Furthermore, she has helped create challenges in .;,;.CTF, Sekai CTF, and Malta CTF.

She is also part of the team representing the US for the International Cybersecurity Challenge and the European Cybersecurity Challenge, assisting with infrastructure and A/D. With FMC, she competed in DEF CON CTF Finals during 2024 and 2025. Outside of CTF, she is a zero-knowledge cryptography auditor for OtterSec and has research in varying fields of lattice-based cryptography.

molenzwiebel ([redacted]) got nerd-sniped by a friend into trying a CTF back in 2021 and now curses them for robbing them of his free time in the last 4 years. Initially only playing smaller CTFs with his university team, he started playing on bigger international teams after experiencing the joys of on-site competitions during the European Cybersecurity Challenge, where he played as a Dutch representative during the 2023 and 2024 editions.

Now primarily playing with *Project Sekai* and *Friendly Maltese Citizens*, he was the on-site captain for FMC at the DEF CON CTF in the last two years. He was a major part of organization for the SekaiCTF and MaltaCTF competitions, as well as the Dutch pre-qualifier track for the ECSC and the Dutch edition of the ECSC, which is now slated for 2027. If you ask him, he's a pretty decent reverse engineering player.

Trixter ([redacted]) started playing CTFs around 2018. Since then, he has become responsible for running Friendly Maltese Citizens, representing Finland's national team in the European Cybersecurity Competition, and being chosen to represent Europe as one of the strongest CTF players in the International Cybersecurity Championship. Additionally, he has been an infrastructure engineer and challenge author for multiple years in various competitions like idekCTF, MaltaCTF, and some smaller national competitions in Finland focusing on web and forensics related challenges.

Professionally he has spent a lot of time maintaining dozens of dedicated servers at the same time, and now focusing on cybersecurity, primarily web auditing and maintaining security-related internal tooling.

ctfguy ([redacted]) found his way into CTFs in 2023 through a love of mathematics, programming, and hacking that naturally led him to cryptography. He is the co-founder of *infobahn*, which rapidly rose into the CTFtime global top 10 in 2025. He competed at DEF CON 33 Finals with *SuperDiceCode*, earning a 3rd-place finish.

Beyond competing, ctfguy developed deep expertise in blockchain security and has organized multiple hackathons and CTF events with reliable, 100%-uptime infrastructure and smooth operations. He’s known for calm, effective communication and hands-on leadership in hectic moments—qualities that helped him build one of the strongest new teams in a remarkably short time.

While our core organizing team is currently focused and lean, we are confident in our ability to rapidly scale up for challenge authoring and infrastructure development once selected. Our team members are deeply embedded in the global CTF community, with established connections to talented challenge authors, infrastructure engineers, and security researchers across multiple continents. Many in the community have already expressed interest in contributing to a format revolution like ours. We expect contributor pool of 15 to 20 vetted volunteers. Of course, we will perform thorough due diligence in our recruiting process to ensure that every team member we bring on board represents the best our community has to offer in terms of both technical excellence and commitment to the competition's integrity.

Pledge of Independence

Our team was established for one purpose: to bring our vision of a new DEF CON CTF format to life. We operate as **an independent organizing team**, not affiliated with any company, organization, or existing CTF team. In fact, our team lead SuperFashi left his CTF team specifically to pursue this opportunity without conflicts of interest.

We understand that concerns may arise regarding individual members' affiliations with their current employers or teams. To address this directly, we make the following pledges:

1. Absolute Impartiality

We commit to treating all teams and participants equally:

- **Information disclosure:** All public information will be announced through public channels. No individual or team outside the organizing team will receive advance access to any competition-related information.
- **Competition design:** We will not design any aspect of the competition to favor specific participants or teams. For parameters that could create advantages, such as qualifier starting times or spawn positions on the game board, we will use **publicly verifiable random number generators** to ensure fairness.
- **Judging and scoring:** All scoring will be automated and transparent. Any discretionary decisions will be documented and explained publicly.

2. Information Security

We will maintain strict operational security:

- All organizing communications will occur through newly created, dedicated channels with proper access control
- Competition infrastructure, software, and hardware will remain internal and separate from any external entities
- Information will only be shared publicly according to our announced disclosure schedule (as outlined in Pledge 1)
- All personal information received from participants or teams will be treated as **confidential** and accessible only to designated organizing personnel who require it for their specific responsibilities

3. Competition Participation Restrictions

To preserve the integrity of qualification pathways, **no organizing team member will**:

- Participate in any competition that affects DEF CON CTF qualification status
- Organize, contribute to, or advise any of the designated pre-qualification events listed in our “Qualification Structure” section

- Participate in or organize qualification events for those pre-qualification events
- *(This restriction applies recursively to any upstream qualification pathway)*

And of course the organizing team members will not be participating in the DEF CON CTF itself.

4. Financial Independence

While we may seek additional funding or resources to enhance the competition experience following "Budget Considerations" section below, we commit that:

- No external sponsorships or donations will be accepted that could influence our decisions, challenge design, or team selection
- If we receive any offers of financial support or sponsorship, we will **transparently disclose** them to DEF CON organizers and the community
- No company, organization, or individual outside our core team and DEF CON will have any decision-making power or influence over any competition design, challenge creation, or judging, etc.
- We will reject any sponsorship that could create real or perceived conflicts of interest

Budget Considerations

Our entire team has committed to organizing this competition **completely free of charge or individual compensation**. We're doing this purely because we believe in the vision and want to give back to the community that shaped us.

That said, we won't compromise on the participant experience. Running a competition of this scale requires:

- **Robust game servers** that can handle 12 teams running exploits and patches on them simultaneously
- **Network infrastructure** capable of supporting A/D traffic, packet capture storage, and real-time game state updates

We have significant experience managing both technical projects and budgets across our team members. We know how to stretch a dollar without cutting corners on what matters. But realistically, we know that depending on DEF CON's budget allocation, we might need additional resources.

If necessary, we're prepared to seek external sponsorship or donations. Several cybersecurity companies founded by former CTF players have consistently supported the community with no strings attached. We have reliable connections with these companies and confidence we can secure support if needed.

However, we want to be crystal clear one more time: any sponsorship we pursue will come with zero conditions that could compromise the competition's integrity or fairness. We will not allow sponsors to influence challenge design, scoring, or any competitive aspect of the CTF.

As a final backstop, our team lead (SuperFashi) has personally committed to covering any shortfall out of pocket rather than compromise on the competition quality. We're going to make this work, one way or another.

Closing Remarks

This proposal represents our belief that CTF competition can be simultaneously fairer, more inclusive of diverse skills, more sustainable for competitors' well-being, and more engaging for spectators. All without sacrificing the technical depth or competitive intensity that defines DEF CON CTF.

Some will say we're changing too much. Others might say we're not changing enough. We're confident we've found the right balance because we designed this format to solve real problems we've experienced as competitors, not theoretical ones we imagined as organizers.

If selected, we won't just be running a contest for three days in August for a few years. We'll set a direction for the competition format that could influence the sport for years to come. We understand the weight of responsibility we're accepting, and we’re acutely aware of what failure would mean. This awareness drives our obsession with transparency, fairness, and thorough preparation.

Four years from now, we hope teams will look back and remember this era as when DEF CON CTF became not just harder or flashier, but fundamentally *better:* fairer, more rewarding, and truer to what a “World Championship” should be.

The circular map will be complete. The watchtower will fall. And the community will be stronger for it.

Thank you for considering our proposal. We are ready to build this future together.

— ***The Overseers Collective***