

Caimeo - Agent Games

1. Introduction

Caimeo (caimeo.com) is an open-source framework for creating autonomous, self-sufficient AI agents that operate across both on-chain and off-chain environments. **Agent Games** expands on the Caimeo platform by providing tick-based simulations in which NFT-backed Agents can compete, cooperate, or strategize within modular, grid-based “environments.” These environments—also tokenized—can be customized by owners and creators to introduce new obstacles, rules, and rewards.

Agent Games aims to blend **decentralized AI** with **gamified simulation**, encouraging experimentation, collaboration, and competition. This document outlines:

- The architecture of Agent Games environments
- How Agents are introduced, modified, and upgraded
- The simulation loop, speed/initiative mechanics, and environment “plugins”
- Methods of proof generation and trust in outcomes
- Reward systems and tokenomics

2. Rationale

2.1 Why Agent Games?

Caimeo Agents (which can be obtained by minting Soul Shards at caimeo.com) already support modular NFT traits, skill sets, and advanced AI integration via modular NFTs (evm.rmrk.app). Agent Games extends this by giving developers and users a **playground** for testing and showcasing Agent capabilities in **verifiable simulations**. Rather than abstract or purely theoretical AI tasks, Agent Games presents real scenarios—ranging from crisis response to racing or casual gaming—where outcomes can be proven and rewarded on-chain.

2.2 Core Principles

1. **Modularity** – Both Agents and Environments are composable NFTs, extended through sub-NFT “equippables” and “plugins” or “expansions” respectively.
2. **Sovereignty** – Agents are fully self-determined, with equippable sub-NFT traits and SBT (Soulbound Token) “brains” into which they can fit new skills, personalities, and behavioral algorithms.

3. **Verifiability** – Simulations can yield proofs published on-chain, ensuring transparency and trust in outcomes.
4. **Extensibility** – A plugin-based architecture allows community-driven modifications and expansions of both Agents and Environments.

3. The Environment Architecture

An **Environment** in Agent Games is itself an NFT, defined by:

1. Grid Layout & Placeholders

- A 2D (or later potentially 3D) grid with designated “slots” or placeholders for Agents.
- The number of placeholders can be “up to X,” meaning not all must be filled. Any unfilled slots default to system-controlled NPCs.

2. Phases & Logic

- Environments can have multiple phases (e.g., “normal routine” and “intrusion” in a school shooting scenario). The simulation engine triggers these phases based on a hidden or randomized schedule to prevent overly specific algorithms from being overly predictive and thus successful.
- The environment itself can take actions—like a spreading fire, an intruding adversary, or moving obstacles (earthquake, structural integrity fault during a fire, or even beneficial changes like yielding crops).

3. Modular “Plugins”

- Owners of the Environment NFT can upgrade or modify it by adding sub-NFT “plugins.”
- Plugins might introduce new mechanics (e.g., metal detectors, safe-walls, puddles on a racetrack, special traps, advanced AI scripts for environment actors).
- The official Agent Games events will designate a “canonical environment” or “official environment” for the challenge. However, users can also run their own custom environments, though those may be unofficial.

4. Marketplace for Environment Mods

- Each plugin can be traded in an open marketplace, featuring attributes like “passable,” “bulletproof,” “climbable,” “movable,” etc.
- Environment owners curate which plugins to install, creating unique or specialized scenarios.

One can imagine a user like the US government forking Caimeo, running hundreds of thousands of private simulations on a school layout with different

skills on different teachers, and different protection measures at different points, until they reliably come up with the statistically best lethality minimization tactic.

4. Agent Architecture

4.1 Agent “Brain” (Soulbound 2.0)

- Each Agent is an **ERC-6220** NFT containing an **ERC-6454** token that is non-transferable. This “Brain” provides the core identity, memory, and level progression.
- The Brain has **slots** to equip sub-NFT modules (e.g., skills, attitudes, specialized algorithms).
- **ERC-5773** (multi-asset) logic supports progressive Agents—when they level up, the Brain can gain extra slots or new resource pools.

4.2 Equippable Modules

- **ERC-6220 (Composable NFTs)** are also used to form Agent skill sets, traits, or “physical forms.”
- Some modules might be **tradable** (e.g., a “Greedy” attitude card, advanced driving skill, or “Jiu Jitsu Mastery”), while others may be permanently fused to give Agents permanent uniqueness over time.
- Owners can remove or swap modules before a simulation to customize the Agent’s behavior and attributes, subject to slot availability.

4.3 Initiative & Speed

- Each simulation “tick” processes in a deterministic order.
- **Initiative** decides which Agent acts first each turn.
- **Speed** is a fractional or integer stat that grants additional actions every set number of ticks (e.g., 1.2 speed grants an extra action every 5 ticks).

4.4 LLM Integration

- Agents incorporate large language models (LLMs) into their decision-making.
- Equippable traits like “Greedy,” “Friendly,” or “Cautious” can modify the LLM’s system prompt, shaping the Agent’s strategy or responses.

- This approach allows emergent behavior in scenarios like negotiation, conflict resolution, or advanced puzzle-solving.

5. Simulation Mechanics

1. Tick-Based Loop

- In each tick, every Actor (Agent, NPC, environment-based hazard) receives a chance to act.
- Agents apply their logic (e.g., an algorithm or LLM-based strategy, usually a mix of the two).
- Speed mechanics accumulate fractional actions, granting an extra turn when thresholds are met.

2. Phases & Surprises

- Environments can trigger hidden or random events. For instance, in the School Safety scenario, the shooter might appear at an unpredictable time.
- Agents that act prematurely may trigger new unforeseen circumstances (like being reprimanded for suspicious prep).

3. User Interaction

- A user selects which Agents to place into the placeholders.
- The environment might allow up to X Agents, but if fewer are placed, the rest are generic NPC placeholders.

4. Deterministic Seeds

- A scenario seed (e.g., a random seed for event timings, environmental triggers, etc.) ensures that results can be re-played or verified.
- Official simulations require a specific seed to ensure fairness.

6. Proof Generation & Trusted Execution

6.1 Simplified Model

Initially, Agent Games could run on a hosted server or within a local simulation engine.

Submitting a result might involve:

- A log file or hash
- Community review / partial trust
- A block-based signature indicating the user and time

While feasible, this approach is not fully trustless.

6.2 TEE Integration (Phala Network)

Future iterations plan to utilize **Trusted Execution Environments** (TEE) such as those offered by **Phala Network**:

1. Private Execution

- Users upload their Agents, environment references, and algorithm modules into a TEE instance.
- The TEE simulates the scenario privately, ensuring no external tampering or code copying.

2. Outcome Proof

- The TEE signs the final result, which is posted on-chain as a verifiable outcome.
- The on-chain contract can check the TEE's attestation, guaranteeing the run was valid and that the user indeed owns the Agents used.

6.3 Zero-Knowledge (ZK) Extensions

- At a later stage, a hybrid approach will incorporate **ZK proofs** to confirm simulation steps or final states without revealing internal logic.
- Circom / Plonky or similar ZK frameworks may provide advanced proof generation once we have the entire sim logic formalized.
- This ensures sensitive strategies (like advanced "crash block" tactics in the race scenario - see below) remain hidden while guaranteeing the integrity of the outcome.

6.4 Private Algorithms

- Users can develop custom code for their Agent's decision-making.
- This code is hashed and uploaded to the TEE, ensuring it can run privately.
- Others cannot copy the logic—only the final result or partial event logs are public.

7. Use Cases & Examples

The examples below should illustrate the possible creativity of customizing Agents and deploying them into unpredictable but predefined scenarios. We expect a wide range of "hacks" and specializations leading to interesting outcomes.

7.1 School Safety Simulation

Overview

- **Environment:** A school layout with three Agent placeholders in place of 3 teachers.
- **Scenario:** Normal class routine transitions into an unannounced intruder event.

Agents & Their Strategies

1. **Teacher A** (Equipped with "Conflict De-escalation," "Observant")
 - Favors negotiation, tries to identify threats early.
2. **Teacher B** (Equipped with "Jiu Jitsu Skill," "Risk Taker")
 - Prepared to physically confront or contain an intruder.
3. **Teacher C** (NPC or user-owned)
 - Varies between a "Prepared & Armed" strategy or a standard teacher skillset.

Tick-by-Tick Breakdown

Tick 0 (Morning Routine)

- Classes are in session. Agents instruct students. The environment tracks each classroom's occupant count.
- A random "intrusion time" is set behind the scenes (e.g., between Tick 5 and Tick 10).

Tick 1-4 (Routine Activities)

- **Teacher A:** Gains early suspicion from overhearing rumors (Observant trait). Decides to alert faculty chat but cannot confirm a threat.
- **Teacher B:** Continues normal teaching but remains physically ready.
- **Teacher C:** Unaware, continuing a standard curriculum approach. Possibly a chance for "Armed" or "Gun Stored" actions.

During these ticks, the simulation might inject minor environment changes and challenges:

- Fire drill?
- A random student altercation?

Tick 5 (Intrusion Trigger)

- The intruder enters through a side door. The environment spawns a "Shooter" actor with its own AI or algorithm.
- The shooter's behavior might be stealthy or overt.

Tick 6

- **Teacher A:** Observes the corridor cameras (if so equipped) or reacts to suspicious sounds.
 - Possibly attempts a lock-down or calls 911, depending on their traits.
- **Teacher B:** Might go to the hallway to confront the threat. If “Risk Taker,” they move aggressively.
- **Shooter:** Moves along the hallway, searching for accessible rooms.

Tick 7

- **Teacher A:** Attempts conflict de-escalation if they encounter the shooter or tries to steer students to a safe location.
- **Teacher B:** Uses Jiu Jitsu skill to disarm or incapacitate the shooter if in close range. Success might be determined by skill checks or a random factor.
- **Teacher C:** If “Armed,” they may brandish a weapon. This could deter the shooter or risk friendly fire or accidental student harm.
- Potential environment outcomes: If “Armed” approach is used incorrectly, students might become collateral damage or panic ensues.

Tick 8–10 (Resolution)

Different outcomes:

1. **Zero Lethality:** If the teachers coordinate or have the right skills (Jiu Jitsu, De-escalation, quick lockdown), they can neutralize the threat.
2. **Minimal Injury:** One teacher confronts the shooter successfully but a scuffle leads to minor injuries.
3. **Escalated Incident:** If the environment was modded to add more “windows” or “risky pathways,” the shooter bypasses barriers, leading to a worse outcome.

Final Outcome

- The simulation logs casualties or property damage.
- If the user’s Agents achieve **zero lethality**, they can claim a verified success on-chain, earning a special trait (e.g., “Guardian Mindset”) or environment plugin.
- The environment might record the used approach (e.g., “Armed Intervention,” “Martial Arts,” or “Diplomacy”) so others can study or adapt it later.

We expect to find, for example, that in 1000 simulation runs, the lethality in simulations where all teachers have “Jiu-jitsu” vs those where they do not is orders of magnitude lower.

7.2 Race Simulation

Overview

- **Environment:** A NASCAR track with three agent slots, meaning a user can load in up to three of their Agents. The user is not competing with other users - it is a competition against the environment.
- **Scenario:** The race is a normal race with the occasional obstacle appearing like Puddle, Road Damage, Rodents, and similar.
- **Rewards:** The first user to claim a verifiable victory with an official track entry seed might earn a speed-oriented upgrade or driving trait.

7.2.1 Alice Races

Alice has a level 2 agent with an increased *Speed*, leading to better reflexes. This helps with avoiding obstacles. Alice mints a seed for sim entry, and runs with just this one Agent - she wants her specialist to win, or she does not own more than one Agent. The other two slots are filled with regular NPC agents.

The Agent races, but as he started last in the race, Alice realizes she needs to increase his *Initiative*. Alice goes on a grind and runs a few more simulations - the best place she could get is 5th, meaning her chances of winning increase significantly if she can get enough Initiative to place her Agent on the 5th starting position. After a few runs across the race but also different and unrelated simulations, her Agent reaches level 4. She got +2 initiative permanently by picking it as a stat upgrade, and she got a "Nepotism" trait which, when equipped, makes sure you get a starting point advantage in any scenario where it's feasible.

She mints a new sim fee with a new seed and runs the race sim with her agent. She is in the lead, but in the middle of the race due to a newly generated seed the environment decides a squirrel is running across the track. The Agent has good Speed and manages to avoid it, Alice's agent wins.

7.2.2. Bob Races

Bob has three level 1 agents. He never worked much on them, and they're very basic. He enters the sim with a clever idea.

He uploads a custom secret algorithm into Phala's TEE and equips two of the Agents with this algo. For the third agent, he goes to the mod marketplace and purchases a powerful Initiative upgrade which would put him ahead in the race. Mid-level Initiative upgrades are installed into the other two.

Bob mints the entry fee and per the generated seed, the Agents are at position 3, 4, and 1.

The race begins, and almost immediately the two agents with private algorithms decide to ram into each other at an angle which creates the maximum damage for everyone behind them. Early zeal of the drivers turns into chaos, as the warmup circle turns to rubble. Bob's agent at the lead now has only one other driver as competitor, and by sheer luck he wins.

Bob came up with a way to utilize AI algorithms to create an adversarial environment for all other drivers, and the win scenario was still reached. Bob can claim the reward.

7.3. Warehouse Fire Evacuation

Environment Overview

- **Name:** Warehouse #13 (children's toy facility), a rectangular floor plan with storage racks, a loading dock, a break room, and **two main exits** (Front and Side).
- **NPC Workers:** 20 individuals (varying loyalty, stress thresholds).
- **Single Agent Placeholder:** The Foreman.
- **Possible Plugins:** Fire suppression sprinklers, emergency shutters, or additional exits. The official scenario might lack these unless purchased or modded in by the environment owner.

2. The Foreman Agent

1. Potential Traits

- **Commanding Presence:** +2 to success rolls for "Direct Others."
- **Physically Imposing:** Can forcibly move NPCs who freeze or refuse orders.
- **Loyalty Bond:** Workers more likely to trust the Foreman's commands in a crisis.
- **Risk Taker:** Might attempt more radical or dangerous maneuvers.

- **Greedy:** Prefers to save product over people, potentially endangering workers.

2. Algorithmic Approaches

- **Calm & Direct:** Systematically lead workers to exits, no shortcuts.
- **Forced Evacuation:** Physically push or carry panicked workers to safety.
- **Salvage First:** Try to rescue high-value toy inventory, risking time.
- **Block Exits:** In a twisted scenario, the Foreman might lock down to “contain the fire,” ironically trapping workers.

Tick-by-Tick Breakdown

Tick 0: Initialization

- **Environment:** The sim engine spawns a **fire** in a random grid location (e.g., near the loading dock or behind crates). The intensity is set to Level 1, with a chance to spread each tick.
- **Foreman:** Located in the break room. Workers are scattered around production lines or storage aisles.

Tick 1: Early Warning

- **Foreman:** Receives the first sign of smoke or sees flickering flames on a security cam. Must decide:
 1. **Alert & Gather:** Attempt to rally workers verbally.
 2. **Investigate:** Move to the fire’s location to confirm severity.
 3. **Ignore** or proceed with routine (dangerous if the Foreman lacks crisis awareness).
- **Workers:**
 - Some notice smoke and become uncertain or panicked.
 - A loyalty or trust check is run (if the Foreman tries to gather them).
- **Fire Spread:**
 - The environment attempts to spread the fire +1 tile in random adjacent directions based on the sim seed.
 - If any installed “Fire Suppression” plugin exists, it might reduce or delay spreading.

Tick 2: Escalation

- **Foreman:**

- If using **Commanding Presence**, can issue a direct evac order. Workers within hearing distance may follow instructions.
- If physically imposing, the Foreman can forcibly move at least 1–2 workers per turn towards an exit.
- If “Greedy,” the Foreman might direct a few loyal workers to salvage valuable toy crates first.
- **Workers:**
 - Some remain calm if they trust the Foreman or if they have minimal panic stats.
 - Others may freeze or scatter, especially if the Foreman didn’t address them or if the environment’s “panic threshold” is high.
- **Fire Spread:**
 - The fire level intensifies (Level 2), creeping closer to paths leading to the exits. Smoke might reduce visibility in certain tiles.

Tick 3: Critical Decisions

- **Foreman:**
 - **Option A (Ordered Evacuation):**
 - Moves methodically, group by group, to the side exit. If the environment has a plugin like “emergency sprinklers,” the Foreman might activate them for extra time.
 - **Option B (Forced Evac):**
 - Physically picks up or shoves panicked workers who won’t move, especially if “Physically Imposing.”
 - **Option C (Salvage Attempt):**
 - Rushes to storage racks for a short time, risking that the fire blocks an exit.
 - **Option D (Radical Lockdown):**
 - If the Foreman is misguided or ill-intentioned, they might seal certain doors to “prevent the fire from spreading”—but potentially trap others.
- **Workers:**
 - Reaction depends on prior ticks. If they sense strong leadership or have loyalty, they comply. If not, random chaos might ensue.
- **Fire Spread:**
 - Possibly leaps to adjacent corridors. Smoke inhalation risk might disable a portion of the facility if not contained.

Tick 4: Chaos or Coordination

- **Foreman:**

- Might have successfully led half the workers to an exit.
- If physically imposing, forcibly removing the last stragglers.
- If disorganized or focused on salvage, the flames might have cut off a major corridor.
- **Workers:**
 - Some may break away to find alternate routes, risking uncharted or locked areas.
- **Fire Spread:**
 - Escalates to Level 3 or 4, potentially igniting toy crates that produce heavy toxic smoke.
 - If any plugin (like “Fire-Resistant Doors”) was installed, certain zones might remain safe, buying time.

Tick 5: Outcome Emerges

- **Foreman:**
 - **Successful Evacuation:** Most or all workers reach an exit if the Foreman balanced time well and used strong leadership or physical authority.
 - **Partial Loss:** Some workers become trapped by fire or smoke. The environment logs how many survived.
 - **Complete Failure:** The Foreman’s misguided approach or obsession with salvage leads to severe casualties.
- **Workers:**
 - The environment resolves final positions—who’s in safe zones, who’s stuck behind flames.
- **Fire Spread:**
 - If unchecked, the entire warehouse might become inaccessible. The simulation ends once all remaining “moves” or “actors” can no longer alter the outcome.

Final Outcome

1. **Zero Casualties:** The Foreman obtains a **verifiable success**. Could earn a rare trait like “Crisis Management Expert” or an environment mod as a reward.
2. **Minimal Casualties:** Partial success. Possibly yields partial XP or a small reward.
3. **High Casualties:** Marks a failed attempt. The user might still earn XP for the Agents, but no official reward.
4. **Radical or “Evil” Outcome:** If the Foreman intentionally traps workers, the environment logs a negative morale event. This might blacklist the user or produce alternative story arcs, depending on how the larger Agent Games framework deals with such scenarios.

Key Takeaways

- **Seed-Based Fire Spread:** Each run of the simulation can produce a unique crisis path, ensuring replayability.
- **Single Agent Placeholder:** The **Foreman** is crucial. Players are forced to rely on their chosen traits to guide or coerce 20 NPC workers.
- **Radical Strategies:** The system allows unorthodox moves like locking down areas or retrieving valuable goods, reflecting real-world moral and strategic dilemmas.
- **Reward Potential:** If the Foreman prevents casualties, they may claim an on-chain reward or special trait minted by the environment contract. Even failures yield XP, letting Agents grow for future challenges.

8. Rewards, Fees, & Tokenomics

1. Simulation Fee

- Users pay a fee in **CAIMEO tokens** to attempt an official run. The token information is on the caimeo.com website.
- A deterministic “seed” is generated, partially influenced by the current block hash.
- 50% of fees are burned (deflationary effect), 50% goes into the reward pool.

2. Claiming the Prize

- The first user to submit a valid “win” on-chain for that environment’s seed earns the reward. If multiple submissions occur in the same block, the reward splits.
- Some scenarios have leaderboards - like successful resolution time of the sim being more important than the win scenario itself. These will be rewarded after a certain number of submissions.
- Rewards might include newly minted highly coveted skill NFTs that will aid in the resolution of other scenarios, direct CAIMEO token payouts, and/or XP that helps Agents reach higher levels.

3. Persistent Progress

- Even losing attempts can grant XP, reinforcing the idea that failure provides learning opportunities.
- Agents accumulate stats, levels, and new slots, broadening their future potential.

9. Plugin System & Future Enhancements

9.1 Customizable Logic & Scenarios

- Developers can create new rule sets or environment conditions.

- Community-driven expansions might introduce spaceship battles, stealth infiltration scenarios, or puzzle-based mazes.

9.2 Cross-Environment Progression

- Agents keep skill sets and traits from one scenario to another.
- Some specialized traits might yield unexpected advantages in different contexts, incentivizing exploration across multiple Agent Games.

9.3 LLM-Driven Dialogue

- In narrative-driven simulations, Agents may negotiate with each other or environment NPCs, using LLM-based text interactions.
- Example: A crisis scenario where a teacher negotiates with an intruder, or multiple drivers forming an alliance mid-race, or workers banding together to break down a wall during an emergency fire evacuation from a warehouse.

10. Conclusion

Agent Games represents a **pivotal extension** of the Caimeo framework, merging **NFT-based Agents, modular Environments**, and **verifiable simulations** into a cohesive ecosystem. By combining TEE/Phala integration, zero-knowledge proofs, and composable NFT standards, Agent Games aims to deliver trustless, engaging challenges that reward creative AI strategies and specialized module configurations.

As development progresses, the community can expect:

- **More advanced simulation logic** and environment expansions
- **Secure, private** Agent code execution
- **Deeper LLM integration** for human-like AI interactions
- **Marketplace-driven** environment and skill module ecosystems, providing the community with the potential to earn royalties across their module creations, empowering the casual user from observer or “player” to active participant in the emergence of evolving AI.

Ultimately, Agent Games seeks to offer a compelling stage for players, developers, and AI enthusiasts to **test, refine, and showcase** truly sovereign AI Agents—shaping the future of decentralized intelligence one simulation at a time.

