Simulacrum Labs — Client FAQ

For BD team use when answering client questions. Each question has a Simple answer (explain like talking to a CEO) and a Technical answer (explain like talking to an engineer).

What does Simulacrum do?

Simple: We rent out a network of gaming GPUs to run your simulation training jobs. You upload your project, we run it on our network, and you download the results. Think of it like Uber for GPU compute — instead of buying your own expensive servers, you use ours when you need them.

Technical: Simulacrum is a distributed GPU compute platform for UE5 simulation workloads. Clients submit jobs via a web client or API. The orchestrator dispatches to idle nodes running RTX consumer GPUs with UE5 5.7, Docker with GPU passthrough, and the AirSim/Colosseum stack. We support three job types: rendering with ML annotations (11 types including RGB, depth, segmentation, LiDAR, 2D/3D bounding boxes, COCO format), RL training via PPO/SAC with Stable Baselines3, and hyperparameter sweeps across distributed nodes.

What makes you different from AWS?

Simple: Two things. First, our GPUs are actually better for visual simulation — they have special ray-tracing hardware that Amazon's server GPUs don't. Second, we don't charge for moving data in and out, which can be a huge hidden cost on AWS.

Technical: RTX consumer GPUs (3060/3080/4090/5060 Ti) have dedicated RT cores and rasterization pipelines optimized for real-time rendering. A100/H100 server GPUs are CUDA compute-only with no display outputs and no hardware-accelerated rasterization. For UE5 simulation rendering, RTX hardware is architecturally superior. Additionally, our decentralized network eliminates egress charges — bandwidth is a sunk cost for node operators.

How much does it cost?

Simple: About $0.75 per GPU-hour, which is 75% less than comparable cloud instances. We sell prepaid credit packs: $25 (33 hrs), $100 (133 hrs), $250 (333 hrs). Academic labs get $500 free to try it out.

Technical: $0.75/GPU-hour with real-time credit deduction. Prepaid credit packs via Stripe: $25 (33 hrs), $100 (133 hrs), $250 (333 hrs). 100% beta coupon available. Credits deduct at $0.75/hr during active compute. Billing is per-tenant with support for Clerk Organizations (team billing).

Is my data safe?

Simple: Yes. Your files are encrypted before they touch any GPU in our network. Each job runs in its own sealed container that self-destructs when the job finishes. The person whose GPU is running your job never sees your project — the screen is hidden.

Technical: Per-job AES-256-GCM encryption via Inspect-Then-Encrypt pipeline. Jobs execute in Docker containers with GPU passthrough, seccomp profiles, and ephemeral storage. Hidden rendering via -RenderOffScreen flag. Secure wipe (random byte overwrite + delete) of all job artifacts on completion. Ed25519 PKI authentication between nodes and orchestrator. Scene sanitization scanner rejects malicious zip contents (DLLs, executables, path traversal).

What file formats do you accept?

Simple: You upload a zip file containing your Unreal Engine 5 project. If you're doing AI training, include your Python training script and a requirements.txt file. That's it.

Technical: Zip archive containing a .uproject file, UE5 content, and optional Python scripts + requirements.txt. The web client auto-detects maps and Python scripts from the zip via JSZip. Training scripts must use argparse for hyperparameter sweeps. AirSim/Colosseum projects supported natively — the platform auto-configures SimMode, vehicle names, and GameMode overrides.

What annotations do you generate?

Simple: We generate 11 types of data from your simulation: camera images, depth maps, object outlines, object bounding boxes (2D and 3D), LiDAR point clouds, camera position data, object tracking IDs, and everything formatted for standard ML training tools.

Technical: 11 annotation types via SimulacrumPlugin (UGameInstanceSubsystem): RGB, Depth, Semantic Segmentation (Cityscapes-compatible label map), Instance Segmentation, 2D Bounding Boxes, 3D Bounding Boxes, LiDAR (.pcd), Camera Calibration (intrinsics + extrinsics with 4x4 transform matrices), Tracking IDs, Labels JSON, COCO-format manifest.

How long does a job take?

Simple: A 10-frame annotation job takes about 2 minutes. A 2000-step RL training job takes about 3-5 minutes. Larger jobs scale linearly. Hyperparameter sweeps run in parallel across multiple GPUs, so a 4-variant sweep takes the same time as a single run if you have 4 nodes available.

Technical: Annotation jobs: ~10-15 seconds per frame at 1080p including all 11 annotation types. Training jobs: depends on total_timesteps and environment complexity. PPO with 2000 timesteps on BlocksV2 drone environment completes in ~3 minutes on RTX 3060. Sweep variants dispatch simultaneously to idle nodes — wall-clock time equals the slowest variant.

Can my whole team use one account?

Simple: Yes. You create an organization in our system, invite your team members, and everyone shares the same credit balance and job history. One person manages billing, everyone can submit jobs.

Technical: Clerk Organizations with role-based access. Admin role manages billing and invites. Member role submits jobs. All org members share a single tenant_id for billing, job history, and project storage. Sign in via email, Google OAuth, or GitHub OAuth.

Do you support custom environments or only AirSim?

Simple: We support any UE5 project. AirSim/Colosseum is our most tested integration, but the annotation pipeline works with any UE5 scene. For training, any Python script that connects to UE5 via RPC will work.

Technical: BYOS (Bring Your Own Scene) architecture. Any UE5 5.7 project zipped with a .uproject runs on the platform. SimulacrumPlugin auto-activates for annotation jobs based on the presence of SimulacrumJobConfig.json. For training, any Python script using airsim, gymnasium, stable-baselines3, or custom RL frameworks works inside the Docker container. Dependencies auto-install from requirements.txt.