GeoDiver Design Project Final Summary

Group 2: Alex Xu, Andres Popoca, Anthony Arreola, Esteban Garcia-Taquez

Project Overview

GeoDiver is a research-driven simulation game that uses real-time geological data alongside interactive gameplay. The game is set in a world where underground environments are vanishing due to both natural and human-caused activity. Players take on the role of "Divers" who are tasked with documenting, preserving, and defending subterranean ecosystems. The system pulls from live sensor feeds (LIDAR, seismic, water-level data) from real caves and mines around the world, creating a dynamic digital twin of what's actually happening underground. It's built to serve both as a legit scientific research platform and as an engaging game, supporting academic research, environmental monitoring, education, and traditional gaming experiences all at once.

Requirements & Acceptance Tests

GeoDiver manages a pretty wide range of functionality. At its core, the system procedurally generates 3D cave environments from LIDAR point-cloud data with less than 10cm tolerance to real-world coordinates (F01). Players get notified within 5 seconds when real-world seismic events cause terrain shifts in their current mission (F02), and all research data is automatically anonymized before hitting the Global Archive to stay IRB-compliant (F03). Mission instances are generated from the latest sensor data, with 90% loading in under 30 seconds on minimum-spec hardware (F04). Players pick between Standard Extraction (combat included) and Research Mode (non-combat scientific exploration) before deploying (F05), and Research Mode gives access to biological/geological scanning tools that log data with full geospatial metadata (F06).

On the education side, authenticated educators can create classrooms, enroll students, assign cave-system missions, and set learning objectives through a dedicated dashboard (F07), with grade passback to Google Classroom and Canvas via LTI 1.3 happening within 10 minutes of mission completion (F08). Researchers get query access to the Global Archive for searching, filtering, and exporting anonymized records (F09). The gameplay loop also includes Defensive Stand mechanics where players protect data nodes from waves of "Corruption" enemies during timed uploads (F10). Content moderators review player-submitted samples for scientific validity (F11), and system admins have a dashboard for monitoring server health via the Elastic Stack (F12). Offline data caching ensures missions can be completed without internet, syncing within 60 seconds of reconnection (F13).

Data-wise, the system maintains player profiles with role-based differentiation (D-01), enforces a standardized JSON schema for all sensor packets using WGS 84 coordinates (D-02), and stores immutable biological/geological sample records with full metadata (D-03). Mission records track everything from deployment to conclusion in real time (D-04), and cave system master records link real-world locations to their digital twins with version-controlled 3D models (D-05). Classroom records support LMS sync (D-06), and tamper-proof audit logs are retained for at least 3 years through the ELK stack (D-07).

Performance targets include 30 FPS minimum on min-spec hardware and 60 FPS on recommended (P-01), UI response under 200ms (P-02), and mission load times under 30 seconds (P-03). Geospatial accuracy is held to less than 1m horizontal and 2m vertical error (PA-01), cave model dimensions stay within 5% of source LIDAR (PA-02), and all timestamps sync to UTC within 1 second (PA-04). The system targets 10,000 concurrent sessions (C-01), ingestion of 100,000 sensor packets per hour (C-02), and a Global Archive capacity of 50 million records and 10TB (C-03).

Server availability is targeted at 99% monthly uptime (AR-1), with seamless offline-to-online transitions (RT-1) and automatic API fallback to cached data (RT-2). Security is handled through RBAC (ACC-02), mandatory authentication (ACC-01), and zero-PII enforcement on all published research data (PRI-01).

System Design

GeoDiver uses a client-server architecture split across three hardware environments. The Unity game client runs locally and handles rendering, gameplay logic, and offline caching. The backend sits on AWS/Azure with auto-scaling and includes PostgreSQL with PostGIS for geospatial storage, Redis for session caching, and the Elastic Stack for logging. Physical sensors deployed in caves transmit data via cellular, satellite, or store-and-forward methods. Client-server communication uses UDP for real-time gameplay and TCP for reliable data transfers like archive syncs.

The final design breaks into five major subsystems. The Gameplay subsystem handles mission deployment, cave environment generation, HUD presentation, and in-mission logic for both extraction and research modes, using the Strategy pattern to swap combat/non-combat behaviors. The Sensor Integration subsystem receives external data from the

Global Sensor Network and partner APIs (USGS, NOAA), validates and normalizes it through the Adapter pattern, and updates the digital twin while notifying affected clients via the Observer pattern. The Research & Logging subsystem manages biodiversity scanning, species identification, and research log preparation. The Education Management subsystem provides classroom creation, mission assignment, progress tracking, and LMS integration independently from gameplay. The Archive & Synchronization subsystem handles offline caching, conflict-aware sync, PII anonymization, and long-term storage in the Global Archive.

Security is enforced through RBAC with five defined roles (Gamer, Educator, Researcher, Admin, Moderator), using IdentityServer for authentication with SSO support. The system applies the Singleton pattern for session management and offline cache consistency, and the Factory pattern for spawning game entities based on cave data and difficulty. Deployment is standardized through Docker and Kubernetes, and the codebase is organized into nine packages separating domain logic from integration and infrastructure concerns.

Issues

The main open issues center on a few key areas. Formal sensor network partnership agreements with cave and mine operators haven't been finalized yet, so the scope of the initial sensor network at launch is still up in the air. Consultation with indigenous communities about sacred cave sites is ongoing, and the specific access restriction criteria haven't been defined. The offline-to-online data conflict resolution strategy still needs to be nailed down, especially for mission records and sample logs that might contradict server-side data after reconnection. There's also no formal guarantee on long-term USGS/NOAA API stability, which is a dependency risk. The team is still deciding between a December 2027 holiday launch (for the sales boost) versus a January 2028 launch (for a more polished product). Cost-wise, the project is estimated at around $1.2M total, with $750K going toward Year 1 development, and the rest covering infrastructure, licensing, scientific advisory, marketing, and post-launch support. Risks include scaling challenges as the user base grows, hardware durability in harsh cave environments, and the general unpredictability of relying on third-party APIs for core functionality.