

Riverside Virtual Classroom

A browser-based 3D classroom that trades open movement for teacher-first control over mics, seating, and attention.

VIRTUAL CLASSROOM / EDTECH PLATFORM • WEBGL • UNITY3D

GENRE	Controlled 3D virtual classroom (non-exploratory, teacher-moderated)
PLATFORM / ENGINE	WebGL browser app built in Unity3D — Photon Fusion (networking), Agora (voice), PlayFab + MongoDB (auth)
TARGET AUDIENCE	Educational institutions — teachers and students in supervised remote/hybrid classes, up to 40 seats across 2 concurrent classes

Prepared by Tensei Games • Senior Design & Production Review

1 Executive Summary

Riverside Virtual Classroom is a WebGL, third-person 3D classroom built in Unity3D where a Teacher and up to 40 Students meet as seated avatars, with the Teacher holding exclusive control over microphone permissions, session notes, and class flow. The core loop is deliberately narrow: log in with a Unique ID and Join Code, spawn into a seat (Teacher up front, Students First-Come-First-Served), participate through a strict raise-hand-then-teacher-unmutes protocol, and end with the Teacher's notes autosaved to the backend. The project's own questionnaire answers the most telling design question directly — asked whether students or teachers can move around the 3D space, the answer is a flat 'No.'

THE PITCH IN ONE LINE *Not a metaverse to explore — a classroom to control: embodied avatar presence with Zoom-level moderation discipline layered on top.*

That single scoping decision (no free movement) is the load-bearing choice this study examines: it trades away the 'metaverse' promise of an explorable 3D space in exchange for something schools actually want — a controllable, disruption-resistant remote classroom. Whether the rest of the design (permissions, notes, reconnection handling) lives up to that promise, and whether the document's own unresolved questionnaire items pose a real delivery risk on a 12-week timeline, are the two threads this study follows.

2 Core Loop & Mechanics Analysis

The mechanics are built around a single organizing principle: the Teacher is the only actor with unilateral control, and every Student action is a request rather than an action. That asymmetry is exactly right for a classroom-management tool, and it shows up consistently across the systems the document specifies.

What works

- **The permission model mirrors real classroom order.** Students can self-mute but never self-unmute; the only path to speaking is the Raise Hand icon, which the Teacher must actively click to grant. Combined with Mute All / Mute Individual controls and a dedicated Participants panel showing hand-raise and mute state at a glance, this gives a single Teacher a workable way to run a 40-seat session without it collapsing into cross-talk — the single biggest practical failure mode of any group voice platform.
- **Reconnection handling protects session continuity for both roles.** Disconnected Teachers and Students can rejoin the same session, avatar/state selections persist across a reconnect, and a Zoom-style network-instability indicator is planned — all sensible choices for a tool whose users are, by definition, running on home or classroom Wi-Fi rather than a controlled environment.
- **Notes are tied to the session lifecycle, not left to manual habit.** Autosaving every 5 minutes and again on session end, with access preserved after the class ends, reduces the single biggest risk for a teacher-facing record-keeping tool: losing a session's attendance and action items to a crash or a forgotten save.

Where friction lies

- **The 3D environment's value proposition rests entirely on presence, not interaction — and that's an unstated bet.** With movement disabled for everyone, the 3D classroom, avatars, and seating exist purely to be looked at rather than navigated. That may be exactly the right trade-off (embodied presence without chaos), but the document doesn't state this as a deliberate value proposition anywhere — it surfaces only as

a one-word answer in the questionnaire. Recommend making 'presence without movement' an explicit, named design pillar so the 3D art and animation budget is allocated toward what students actually perceive (avatar expressiveness, seating-area detail) rather than environment traversal no one will use.

- **The login flow is a genuinely open question, not a resolved one.** The questionnaire itself asks whether Teacher and Student login prompts differ, and whether the ID/Join-Code model coexists with an email/password flow — and the document leaves it unanswered. Since login is the very first thing every user experiences, shipping without this resolved risks rework late in a 12-week schedule.
- **Notes file-naming has an unaddressed collision case.** Files are named by creation date only, cannot be renamed, and the document doesn't describe what happens if a Teacher creates more than one notes file on the same day — a realistic scenario for any teacher running multiple sessions daily.

DESIGN VERDICT *The Teacher-controlled permission model is well-matched to the product's institutional purpose; the biggest open risk is that several first-order product decisions (login flow, same-day file naming) are still genuinely unresolved rather than simply unwritten.*

3 User Engagement & Adoption

Riverside is an institutional adoption decision, not a consumer engagement one — a school or training provider chooses it over a flat video call specifically because of what embodied presence and hard-gated permissions can do that Zoom or Teams can't. Framing motivation and adoption in those terms is more useful here than a retention curve.

What drives adoption

The strongest adoption argument is Teacher confidence: full mute authority, no student self-unmute, and a single end-of-class action that locks the room and saves notes gives an institutional buyer a concrete answer to 'how do we stop this turning into chaos with 40 kids on mic.' That's a sharper pitch than most consumer virtual-classroom tools make, and it should be the headline of any sales conversation, not a footnote.

- **Lightweight social affordances (Thumbs Up/Down emotes, the color-changing hand-raise icon)** give students a low-friction way to participate without opening a channel for disruption — appropriately modest for a supervised setting rather than a social space.

Where adoption is at risk

- **The Teacher role carries almost all of the operational weight — and almost all of the onboarding risk.** Teachers create classes, control every student's mic, take the only record of attendance, and end the session. If class creation and editing (itself an unanswered questionnaire item — can teachers edit class details, and are users notified?) isn't fast and clear, the tool will feel like a burden to the one role whose buy-in determines whether a school keeps using it. Recommend prioritizing Teacher-side usability testing well ahead of Student-side polish, since Teacher friction is adoption-fatal in a way Student friction usually isn't.

ADOPTION VERDICT *The controlled-permission pitch to institutions is genuinely strong; the product's success depends on the Teacher's day-one experience with class creation and session control more than on any Student-facing feature.*

4 Interface & Presence Design

There's no narrative here, so the aesthetic question is really an interface-legibility one: does the visual and audio language make a 40-person voice session readable at a glance? Largely, yes.

What's aligned

- **Diegetic status icons above each avatar's head (hand-raise state, mute/speaking indicators)** keep session status visible in the 3D space itself rather than buried in a separate panel, which is the right call for a Teacher scanning a room of up to 40 students. The Participants panel exists as a secondary, more precise control surface — a sensible two-tier design (glance at the room, act from the panel).
- **Borrowing Zoom's disconnection-indicator convention (the familiar 'three dots') is a good instinct, not a shortcut.** Unlike a novel visual metaphor, reusing an interaction pattern both teachers and students already know from existing remote-learning tools directly reduces onboarding friction for an audience that didn't choose this software for its novelty.

Where it needs sharpening

- **Avatar identity is currently a single binary toggle (Male/Female, defaulting to Male) with a pre-defined model per choice.** That's a minimal but workable identity system for a first release; for real institutional deployment across varied age groups and school communities, it's worth confirming with the client whether a two-option avatar system is sufficient or whether broader representation is expected before wider rollout.

INTERFACE VERDICT *The diegetic status icons and borrowed Zoom conventions are the right instincts for a tool this audience needs to learn instantly; avatar representation is the one area worth revisiting before a broad institutional launch.*

5 Commercial Viability & Deployment Model

The document doesn't specify a pricing or licensing model, and that should be treated as an open question rather than assumed — this reads as B2B/B2I software (sold to schools or training providers, not to individual consumers), where per-seat or per-institution licensing are the likely paths, but nothing in the GDD commits to either.

What's sound

- **The stated scale is modest and testable.** A ceiling of 40 seats across 2 concurrent classes is a sensible pilot scope — large enough to prove the permission model under real classroom size, small enough to keep infrastructure costs predictable during development and initial rollout.

What needs scrutiny

- **Four separate third-party services underpin every session — PlayFab and MongoDB for auth/pairing, Photon Fusion for networking, and Agora for voice — and each typically bills per-seat or per-minute.** The document doesn't model cost beyond the pilot ceiling. Before quoting a client on rollout to multiple schools or a larger per-institution seat count, recommend building an actual per-session cost model across these four services at realistic concurrent-class volumes, since voice and networking costs in particular tend to scale linearly with seats and don't stay flat past a pilot.

- **No stated data-retention or export policy for notes beyond 'cannot be downloaded.'** For an institutional buyer with its own data-governance requirements, 'notes stay on-platform, teacher can view them after the session ends' needs to be an explicit, intentional policy statement (and likely a selling point around data control) rather than something that reads as a missing feature.

COMMERCIAL VERDICT *The pilot-scale plan is sound; a real per-seat infrastructure cost model and an explicit data-retention policy for notes are the two gaps worth closing before scaling past the pilot.*

6 Risk Assessment & Strategic Recommendations

This GDD is candid in a way that's genuinely useful: it includes its own unresolved questionnaire, which means the open risks below are largely the client's own words, not inferred gaps.

Risk 1 — A meaningful number of core product decisions are still open questions (Requirements)

- The document's own QA section leaves unanswered: whether Teacher and Student login prompts differ, whether the ID/Join-Code model coexists with email/password, whether students can join before the Teacher, what happens if a Teacher tries to double-book a class slot, whether ended classes stay visible in the lobby, and whether Teachers can edit class details after creation. These aren't edge cases — several are first-screen, first-session experiences. Mitigation: convert every unanswered questionnaire line into a signed-off requirement with the client before Sprint 1, given a 12-week Agile timeline has little slack to absorb requirements discovered mid-build.

Risk 2 — Notes file-naming has an unresolved collision case (UX / Data Integrity)

- Files are named only by creation date, can't be renamed, and the document is silent on what happens when a Teacher creates multiple files in one day — a routine scenario for anyone teaching more than one session daily. Since notes are the Teacher's primary record of attendance and action items, an ambiguous or overwriting file-naming scheme directly undermines the feature institutional buyers will rely on most. Mitigation: define an explicit same-day disambiguation rule (e.g., an auto-incrementing suffix) before this reaches a client demo.

Risk 3 — Multi-service infrastructure cost is unmodeled beyond pilot scale (Technical / Commercial)

- PlayFab, MongoDB, Photon Fusion, and Agora each add both integration surface and a per-unit cost that will multiply with any real institutional rollout beyond the stated 2-class/40-seat pilot. Mitigation: cost the stack at a realistic deployment scale (e.g., a single school running 10+ concurrent classes) before that number appears in a client proposal or contract.

Bonus Risk — Several structural sections read as placeholders (Documentation Completeness)

- Gameflow, Role: Teacher, Role: Student, Admin Panel Spec, Animations, and Sounds all appear in the extracted document as headers with no body content. This may reflect diagrams or images that didn't survive text extraction from the source file rather than genuinely empty sections — worth confirming directly, since treating a diagram-only section as 'unspecified' versus treating a truly blank section as 'unspecified' calls for two different conversations with the client before scoping final delivery dates.

RISK SUMMARY *The mechanics and permission model are sound; the real project risk is that the document's own open questions — login flow chief among them — need to be closed with the client before, not during, the*

12-week build.

7 Final Verdict

Riverside Virtual Classroom makes one clear, well-reasoned design bet: give up the open-world promise of a 3D metaverse in exchange for the thing institutional buyers actually want from a remote classroom — embodied presence with airtight moderation control. The Teacher-first permission model, diegetic status icons, and session-tied notes system all support that bet consistently, and the technical stack (Photon Fusion, Agora, PlayFab/MongoDB) is a reasonable, proven foundation for the stated pilot scale.

The project's real risk isn't in the mechanics — it's that the document candidly leaves a meaningful number of first-order product questions open, from login flow to same-day notes file naming to infrastructure cost past pilot scale. None of these are hard problems, and the fact that the questionnaire surfaces them at all is a sign of a careful team asking the right questions — but a 12-week Agile schedule needs them answered before Sprint 1, not discovered during it. Closed out, this is a low-risk, well-scoped build with a genuinely differentiated pitch to schools tired of flat video calls.