ProCleanLakes

BUSINESS COACHING REPORT
LEAN MVP DEVELOPMENT FOR A VR/AR BRATES LAKE RECOVERY SIMULATOR

1. Coaching context

The proposed business concept is a VR/AR digital simulation tool that would help users visualise how Brates Lake could evolve from its current degraded condition toward healthier ecological states under different restoration scenarios. Potential users could compare present conditions with proposed interventions and use the visualisation to support restoration planning, stakeholder communication, public consultation, education, or investment discussions.

The principal business risk is not whether an advanced immersive platform can be built, but whether identifiable customers will use and pay for it. A complex VR/AR product could require substantial expenditure on environmental data, modelling, software development, three-dimensional assets, interfaces, hardware compatibility, validation, and maintenance before the value proposition has been demonstrated.

The participant therefore requested support in applying Lean Startup principles and defining a Minimum Viable Product before committing to full development. The coaching assignment focused on selecting one use case, defining what should and should not be built, choosing early testers, designing moderated tests, measuring behaviour rather than relying only on positive comments, and establishing evidence-based investment criteria.

2. Coaching objectives

The overall objective was to design a testable MVP that can establish whether the proposed simulation creates sufficient decision or communication value to justify further investment.

The first objective was to define the smallest credible product experience. The second was to identify the organisations and professional roles that should participate in early testing. The third was to establish a practical testing protocol, including tasks, interview questions, observations, and behavioural indicators. The fourth was to define go, revise, and stop criteria connected to customer commitment, data access, repeat use, and willingness to support a paid pilot.

3. Lean startup methodology

The recommended methodology follows the Build Measure Learn cycle. *Build* means creating only the minimum experience required to test the central hypothesis. *Measure* means observing whether target users understand and use the prototype for a real task. *Learn* means deciding whether to continue, change the use case, modify the customer segment, or stop before additional development costs are incurred.

The central hypothesis is responsible organisations, and restoration stakeholders will find a visual comparison between the current state of Brates Lake and one evidence-based restoration scenario useful enough to reuse, recommend, support with data, or fund through a pilot.

The MVP should test value before technical scale. It does not need to prove that every future hydrological or ecological process can be simulated. It must prove that a clearly identified user can understand one scenario better, communicate it more effectively, or make a better-informed planning decision because of the tool.

4. Recommended MVP scope

The first MVP should present one defined area or viewpoint of Brates Lake in two comparable states: the current condition and one proposed restoration scenario. The restoration scenario must be specified with ecological experts and based on transparent assumptions. The prototype may be a lightweight interactive web-based 2D or 3D experience, a VR headset or advanced AR implementation is not required to test the initial value proposition.

Essential MVP functions should include:

- a short, guided introduction explaining the user task and the limits of the prototype;
- a current-state visualisation using available and appropriately referenced information;
- one restoration scenario with visible, labelled interventions and an indicative future state;
- a simple before-and-after comparison, slider, toggle, or guided transition;
- brief explanations of the scenario assumptions, expected effects, time horizon, and uncertainty;
- a mechanism for recording task completion, comments, usefulness ratings, and requests for follow-up;
- an administrator function for correcting text, assumptions, and data sources without rebuilding the experience.

The following functions should be postponed: multiple restoration scenarios, real-time sensor integration, predictive risk forecasting, multiplayer collaboration, complex avatars, gamification, multi-lake libraries, headset-specific development, photorealistic rendering, and automated ecological recommendations. These features should be added only when evidence identifies a paying use case that requires them.

5. Selection of early testers

The first test group should contain approximately 8 to 12 participants and represent at least three organisational perspectives. A small, deliberately varied group provides more useful evidence than a large group of general users who do not influence planning, budgets, data access, or procurement.

Tester group	Suggested number	Value to test	Relevant commitment
Lake administration and public or water authorities	3-4	Planning, consultation and institutional communication	Pilot discussion, internal referral or procurement guidance
Ecologists, restoration planners and technical consultants	2-3	Scientific credibility, assumptions and professional usability	Data review, scenario validation or expert participation

NGOs and community engagement specialists	2-3	Public understanding, participation and educational value	Use in a workshop, consultation or campaign
Funded project or communication representatives	1-2	Dissemination, reporting and funder communication	Budget feedback or inclusion in a project proposal

Early testers should be selected because they perform a relevant task, influence a decision, control or provide useful data, communicate with stakeholders, or can sponsor a pilot. Friends, students, or general technology enthusiasts may provide interface feedback but should not be treated as proof of market demand.

6. Testing protocol and interview questions

Each session should last approximately 30 to 45 minutes and should be moderated individually where possible. The participant should explain the task without selling the concept, allow the user to explore independently, observe difficulties and questions, and conduct the interview only after the task has been completed.

Recommended test task: *Use the prototype to compare the current condition of Brates Lake with the proposed restoration scenario. Identify the main intervention, explain what you believe would change, and describe whether this comparison would help you in your professional role.*

During observation, the moderator should record whether the user starts without assistance, finds the comparison function, identifies the intervention, understands that the future state is a scenario rather than a guaranteed prediction, asks for missing information, and chooses to share or revisit the result.

Post-test questions should include:

- What decision, discussion, or communication task could this prototype support?
- Which part was most useful, and which part was unclear or unnecessary?
- What information or evidence would you need before using it professionally?
- Who inside your organisation would use, approve, fund, or provide data for this tool?
- How do you currently communicate restoration scenarios, and what does that process cost or require?
- Would you request access again for a meeting, consultation, proposal, or planning exercise?
- Would your organisation provide data, expert time, a letter of interest, or funding for a pilot?
- What would prevent the organisation from adopting or purchasing this solution?

Questions such as “Do you like it?” or “Would you use it?” should not be used as primary validation because they encourage polite and hypothetical answers. Behaviour, follow-up requests, data provision, introductions, scheduled meetings, and budget discussions are stronger signals.

7. MVP validation framework

The following criteria provide an initial decision framework. Thresholds should be treated as practical hypotheses and reviewed after the first testing round rather than as universal rules.

Dimension	Indicator	Initial threshold	Evidence source
Recruitment	Relevant users complete the test	8-12 users from at least 3 organisation types	Attendance and role record
Usability	Core comparison completed without major assistance	At least 75% of testers	Observation checklist
Comprehension	User correctly explains intervention, change and uncertainty	At least 70% of testers	Task explanation
Usefulness	Prototype supports a real planning or communication task	At least 60% rate 4 or 5 out of 5	Post-test rating and example
Repeat behaviour	User requests reuse, sharing or a second session	At least half of testers	Recorded follow-up action
Data commitment	Organisation offers relevant data or expert review	At least 2 organisations	Named data or review commitment
Commercial signal	Organisation discusses a paid or co-financed pilot	At least 1 credible institution	Meeting, budget path or proposal request

8. Commercial hypothesis

The product should be positioned as a decision-support and stakeholder-communication service, not primarily as a VR/AR novelty. The customer pays for clearer understanding of restoration alternatives, more effective consultation, stronger project communication, or better presentation of evidence. Immersion is a delivery mechanism rather than the value proposition itself.

The most plausible initial revenue model combines a paid prototype or institutional pilot, scenario-development and data-preparation fees, workshop or consultation delivery, and an annual licence or service agreement for updates and reuse. Custom development should be priced separately from recurring access so that the cost of new data, modelling, and visual assets is not absorbed by the base licence.

The likely economic buyer is a lake administrator, public authority, water-management organisation, restoration consortium, environmental consultancy, or funded project. NGOs, schools, and community groups may be users or partners, but should be considered primary customers only when they control a relevant communication or project budget.

9. Risks and safeguards

The most important risk is presenting a visually persuasive scenario as if it were a scientific prediction. Every visualisation should distinguish observed data, modelled

assumptions, expert judgement, uncertainty, and illustrative content. Technical claims must be reviewed by appropriate ecological, hydrological, and restoration specialists.

Other risks include insufficient or incompatible data, excessive customisation for each institution, inaccessible interfaces, hardware dependence, costly updates, and unclear ownership of supplied data or visual assets. Pilot agreements should define data rights, review responsibilities, permitted uses, update scope, confidentiality where relevant, and the limitation that the tool supports rather than replaces statutory or scientific assessment.

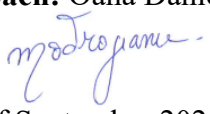
10. Recommended action plan

- Weeks 1 and 2 should be used to select the exact user role, professional task, lake area, restoration scenario, and evidence required. At least five discovery interviews should be completed before development begins.
- Weeks 3 and 4 should produce a clickable storyboard or low-fidelity prototype. Two expert reviews should verify the scenario logic and the communication of uncertainty. The prototype should be revised before formal user testing.
- Weeks 5 and 6 should cover 8 to 12 moderated tests using the same task, observation sheet, and interview guide. Results should separate stated opinions from observed behaviour and concrete commitments.
- Weeks 7 and 8 should compare results with the validation framework, identify one institutional pilot candidate, prepare a priced pilot proposal, and decide whether to continue, revise, or stop. Additional VR/AR functions should enter the roadmap only when a validated customer task requires them.

11. Conclusion

The concept is suitable for Lean Startup validation because its value can be tested before a complete immersive platform is built. The recommended MVP is a lightweight comparison of the current condition of one defined area of Brates Lake with one transparent restoration scenario. Its purpose is to test understanding, professional usefulness, and institutional commitment rather than technological sophistication.

Further investment is justified when relevant users complete and understand the task, identify a recurring professional application, request reuse, and convert interest into data access, expert participation, or a paid pilot. If these signals do not emerge, the participant should revise the use case or customer segment before adding new VR/AR functionality.

Coaching requested by: Bogdan George Dragan (Romanian demosite)
Business coach: Oana Daniela Modrogeanu
Signature: 

Date: 15th of September 2026