

From commissioned experience to operational capability: Designing virtual worlds for evolution and industrial value

Ricard Gras^{1,2,*}

¹University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany

²Co-Founder, Captic GmbH, Himmelmutterweg 5/7/8, 1170 Vienna, Austria

Abstract

XR has become highly effective at producing compelling demonstrations, simulations and bespoke experiences. Yet challenging environments expose a weakness in the conventional production model: the conditions, systems and requirements surrounding a virtual worlds deployment rarely remain static. Connectivity changes, assets evolve, data sources are replaced, AI capabilities advance, hardware becomes obsolete and operational users develop new requirements. A system designed as a finished experience can therefore become obsolete soon after deployment. We argue that virtual worlds and XR should instead be designed as an operational capability: modular, interoperable, secure, deployable across heterogeneous environments and capable of continuous evolution. We propose four principles—deployment before demonstration, interoperability before lock-in, modularity before bespoke production, and evolution before completion—and complement them with a fifth: empower the operator organisation to own the lifecycle. Drawing on Captic’s experience supporting industrial metaverse, we argue that deployment success depends not only on technology but also on onboarding, operational support, transparent due diligence, realistic treatment of supplier constraints, and rigorous security and privacy practices. The central proposition is simple: the value of this sector should be measured not only by what can be demonstrated at delivery, but by what can continue to be changed, integrated and used afterwards.

Keywords

Extended Reality, XR, operational deployment, industrial metaverse, interoperability, digital twins, AI, resilience, lifecycle, human-centred deployment, cultural heritage

1. The deployment gap

Many virtual worlds projects still follow a production logic inherited from agency and media work: define a brief, commission a bespoke experience, build around a selected technology stack, demonstrate it, and deliver a finished artefact. This model is effective when the objective is a bounded experience. It is less suited to operational environments in which the system must remain useful for years. Industrial organisations do not operate in finished environments. Equipment changes, procedures change, data changes, suppliers change and AI capabilities evolve. A digital representation that is accurate at commissioning can therefore become incomplete without any failure in its original design. The relevant question is not simply whether virtual worlds and/or XR works, but whether it can absorb change without requiring a new project each time. [1, 2]

2. Design for operational deployment

2.1. Operational design principles

We propose that deployment constraints become design inputs from the beginning. The development sequence should move from “design → build → demonstrate → deploy” towards “operational requirement → deployable capability → validate → adapt → scale”. This changes what good engineering means.

2nd International XR4CE Workshop, November 3, 2026, Valencia, Spain

*Corresponding author.

EMAIL: r.gras@captic.io (R. Gras)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Deployment before demonstration. Prototype against real data, devices, workflows, connectivity, security constraints and user responsibilities. A demonstration that cannot survive its operating environment should not be treated as an early version of the operational system.

Interoperability before lock-in. Design to exchange information with digital twins, enterprise systems, 3D assets, AI services and future platforms wherever practical. The objective is not to predict which technology will dominate, but to reduce dependence on any single one [2].

Modularity before bespoke production. Separate data, spatial representation, interaction, AI, collaboration and deployment capabilities so individual elements can evolve independently. The target should be lower total cost of evolution, not merely lower cost of first delivery.

Evolution before completion. Treat deployment as the start of a lifecycle: deploy → observe → adapt → extend → redeploy. New AI models, devices, data sources and workflows should be incorporable without fundamental redevelopment [1, 2].

Resilience by design. Test degraded communications, uncertain data, unavailable services, changing physical conditions and human error, including graceful degradation and recovery.

3. Empowering the organisation, not creating dependency

Operational deployment also changes the relationship between supplier and client. The objective should not be to leave an organisation dependent on the team that produced its original experience. It should be to give the organisation the capability, confidence and processes required to maximise its investment in the industrial metaverse. In practice, this requires more than software: structured onboarding; documentation and operational support; reusable tools and processes; transparent technical and supplier due diligence; clear identification of platform and AI-service limitations; security and privacy by design [3]; and an explicit understanding of what data, models and infrastructure can or cannot be controlled by each supplier. Where organisations face restrictions on particular AI providers, models or cloud services, the architecture should accommodate those constraints rather than conceal them or make them an afterthought.

3.1. Beyond industrial applications –Cultural Heritage

This principle also applies beyond industrial operations. Cultural heritage organisations may have equally ambitious aspirations for immersive representation, access and interpretation, but face different constraints around preservation, provenance, ownership, accessibility and long-term stewardship [4, 5]. The implementation changes; the underlying requirement for an evolvable and well-governed XR capability does not. Cultural heritage organisations can benefit from the same principles, particularly where immersive technologies are used to reconstruct sites, interpret collections, widen access or create shared educational and visitor experiences. Many of the underlying user-experience requirements are similar: intuitive access, cross-device participation, reusable assets, interoperability and the ability to evolve experiences without rebuilding them from scratch. In most cases, security requirements may be less demanding than in industrial environments, but provenance, intellectual property, accessibility and long-term preservation become correspondingly important. Though a generalisation, most cultural institutions may operate with more limited budgets or technical capacity. This makes collaboration particularly valuable: shared frameworks, reusable components and partnerships between cultural, research and technology organisations can reduce duplication while building sustainable expertise. [4, 5, 6, 7]

4. The legacy-engine question

Much of our sector remains organised around established legacy engines and the production workflows that have grown around them. These technologies have played an important role in making immersive experiences possible, but their strengths as production environments do not necessarily make them

the right foundation for long-lived operational use. The issue is not the engine itself; it is the level of dependency on an obsolete paradigm that can prevent innovation, modularity and adaptation. Conventional engine-led production can require highly specialised skills, proprietary workflows and substantial teams operating between the organisation with the operational need and the technology itself. Production is consequently expensive and time-consuming, and iteration after delivery can require returning to the same specialist production chain. What appears efficient for creating a finished experience can become inefficient when the requirement is to continuously adapt that experience.

This matters because challenging environments are inherently iterative. An industrial organisation may need to introduce a new asset, data source, AI capability, device, workflow or operating procedure without commissioning another substantial production cycle. If every meaningful change requires specialist redevelopment, the system becomes a static snapshot of an operational environment rather than a capability that evolves with it. The architectural consequences extend further. Engine-dependent experiences can constrain interoperability, cross-platform deployment and the ability to combine technologies from different suppliers. They can also make innovation dependent on the roadmap and commercial priorities of an external technology ecosystem. In strategic sectors, this creates a less visible but increasingly important question of technological autonomy: organisations should understand where critical dependencies sit, what happens if a supplier changes its terms, capabilities, availability or strategic direction, and how readily the organisation can move to alternative technologies. [1, 2]

This is particularly relevant as AI accelerates the pace of technological change. New models, agents, spatial interfaces, devices and computational approaches may emerge faster than traditional XR production cycles can accommodate. An architecture designed around replaceable components and open interfaces can incorporate these advances; an architecture tightly coupled to a legacy production stack may make adopting them disproportionately expensive. The alternative is not to reject engines, but to manage their architectural authority. Engines can remain useful tools within a broader technology-agnostic architecture, rather than becoming the environment within which the entire operational capability is constructed. Data, assets, AI services, interaction models and deployment targets should be able to evolve independently wherever practical. The result is a different proposition for virtual worlds: not a finished experience delivered through a technology stack, but an interoperable capability capable of absorbing new technologies as they arrive. For organisations investing in the industrial meta-verse, that distinction may determine whether today's deployment becomes tomorrow's platform—or tomorrow's technical debt.

5. A different measure of success

Operational XR should be evaluated against a broader set of measures: interoperability, adaptability, portability, resilience, modularity, maintainability, security, privacy and lifecycle cost, alongside conventional usability and performance measures. One particularly useful metric is the cost and time required to introduce a meaningful change after deployment. The provocative proposition is that the best XR system may not be the one that produces the most impressive demonstration. It may be the one that is still useful five years later—and can incorporate what the next five years bring.

6. Position and discussion

Captic's position is that virtual worlds should increasingly be treated as operational infrastructure rather than commissioned media. This does not reduce the importance of experience design; it places experience design inside a longer lifecycle of integration, adoption, governance and evolution. We propose that the XR4CE community examine how deployment architecture, supplier independence, interoperability, lifecycle support and organisational empowerment can become first-class research and design concerns. The practical challenge is not simply to bring XR into the field. It is to ensure that, once there, it can continue to change with the field.

Declaration on Generative AI

During the preparation of this work, the authors used generative AI tools for drafting and language refinement. The authors reviewed and edited the resulting content and take full responsibility for the publication's content.

References

- [1] J. Guo, J. Leng, J. L. Zhao, X. Zhou, Y. Yuan, Y. Lu, D. Mourtzis, Q. Qi, S. Huang, X. Song, Q. Liu, L. Wang, Industrial metaverse towards Industry 5.0: Connotation, architecture, enablers, and challenges, *Journal of Manufacturing Systems* 76 (2024) 25–42. doi:10.1016/j.jmsy.2024.07.007.
- [2] X. Tu, R. Ala-Laurinaho, C. Yang, J. Autiosalo, K. Tammi, Architecture for data-centric and semantic-enhanced industrial metaverse: Bridging physical factories and virtual landscape, *Journal of Manufacturing Systems* 74 (2024) 965–979. doi:10.1016/j.jmsy.2024.05.016.
- [3] D. Cayir, A. Acar, R. Lazzeretti, M. Angelini, M. Conti, S. Uluagac, Augmenting security and privacy in the virtual realm: An analysis of extended reality devices, *IEEE Security & Privacy* 22 (2024) 10–23. doi:10.1109/MSEC.2023.3332004.
- [4] M. S. Anwar, J. Yang, J. Frnda, A. Choi, N. Baghaei, M. Ali, Metaverse and XR for cultural heritage education: Applications, standards, architecture, and technological insights for enhanced immersive experience, *Virtual Reality* 29 (2025) 51. doi:10.1007/s10055-025-01126-z.
- [5] C. Lin, G. Xia, F. Nickpour, Y. Chen, A review of emotional design in extended reality for the preservation of culture heritage, *npj Heritage Science* 13 (2025) 86. doi:10.1038/s40494-025-01625-x.
- [6] D. Buragohain, Y. Meng, C. Deng, Q. Li, S. Chaudhary, Digitalizing cultural heritage through metaverse applications: Challenges, opportunities, and strategies, *Heritage Science* 12 (2024) 295. doi:10.1186/s40494-024-01403-1.
- [7] N. Partarakis, M. N. Katsantonis, E. Zidianakis, A systematic review of extended reality (XR) applications in cultural heritage, *Heritage* 9 (2026) 215. doi:10.3390/heritage9060215.