

MULTIVERSERA

Reality Engine

Trust, Data Governance & Ethical AI Brief



Version v0.1 CANONICAL · 2026 · Confidential · Limited Distribution

Engine today, Experience OS later · Localize · Combine · Trust

© Multiversera — All Rights Reserved · info@multiversera.com

Contents

1. Short definition
2. Why is the Trust Layer central?
3. What does the Trust Layer govern?
4. KVKK and data-minimization approach
5. Explicit consent, permission and purpose limitation
6. AI avatar / digital-persona boundaries
7. Verified information and source security
8. Institutional data, content and role-based access
9. Additional care for children, youth and education contexts
10. Trust in public, municipal and disaster scenarios
11. Respectful representation in cultural heritage and tourism experiences
12. Ethical-AI principles
13. Measurable trust and audit trail
14. Risks and the mitigation approach
15. What this is not
16. A first Trust Review conversation
17. Contact

1. Short definition

Multiversera Reality Engine is a modular multi-reality experience engine that combines AI, digital twins, XR and trust layers for institutions, cities, education and cultural assets. Its inseparable component, the **Trust Layer**, treats trust not as an add-on feature but as the foundation of the architecture.

The honest definition today is an engine; the Experience OS and Reality OS vision is on the 2029–2030 horizon. The trust approach, however, applies from today, in every scenario.

2. Why is the Trust Layer central?

AI, digital-twin and XR experiences operate with personal data, institutional content, public information and user interaction. This brings questions of privacy, accuracy, representation and accountability. In most projects these are left until last; at Multiversera they are placed at the start of design.

A **privacy-by-design** and **ethics-by-design** approach means that every module and every demo incorporates trust requirements from day one. The Trust Layer collects these requirements in a single consistent layer and applies them across all modules.



The Trust Layer's place in the architecture

3. What does the Trust Layer govern?

The Trust Layer aims to govern the following areas in every scenario involving AI and data:

- personal-data and KVKK compliance framework (data minimization),
- explicit consent, permission and purpose limitation,
- ethical AI and consented digital persona,
- verified information and source security,
- AI response limits and content boundaries,
- role-based access and institutional-data separation,
- child and youth safety,
- accessibility,
- audit trail and human-in-the-loop.

This is not a checklist; it is a design framework reviewed in each scenario.

4. KVKK and data-minimization approach

In processing personal data, Multiversera aims for the principle of **data minimization**: only the minimum data necessary for an explicit and legitimate purpose is processed; use beyond that purpose is not intended.

With respect to KVKK and GDPR, every implementation must be designed according to the relevant legislation and **requires legal review**. Multiversera does not make a definitive statement such as "it is KVKK/GDPR compliant"; instead, through a privacy-by-design approach, it offers a design framework that facilitates the compliance process. The final compliance assessment is the responsibility of the institution's own legal and compliance units.

5. Explicit consent, permission and purpose limitation

Three principles govern the use of personal data and persona:

Explicit consent. The user is informed about the purpose for which their data is processed; consent must be explicit and revocable.

Permission and authority to represent. A digital persona based on a person's image, voice or identity is used only with that person's (or an authorized representative's) explicit permission and within a limited scope.

Purpose limitation. Data is not used beyond the purpose for which it was collected; scope and duration are defined in advance in every scenario.

6. AI avatar / digital-persona boundaries

The use of digital personas in AvatarWorks and related scenarios is subject to strict boundaries:

- A digital persona does **not clone** a person and does **not represent** a real person without authorization; it is used only with explicit permission and within a defined scope.
- The persona is grounded in approved sources; it does not speak on out-of-scope topics.
- Persona use is limited by duration, context and purpose; the person's rights are protected by contract.
- The avatar says "I don't know" on topics it does not know; it stays within bounds rather than fabricating (hallucinating).

The goal is a useful but **consented, ethical and auditable** digital persona.

7. Verified information and source security

Avatars and narratives are grounded in approved sources. The approach has three elements: source verification (the response is tied to a reliable, approved knowledge base), content boundaries (out-of-scope or unapproved content is not produced) and misinformation mitigation (accuracy is checked; under uncertainty the avatar does not act confident).

This approach aims to reduce the risk of incorrect or insufficient information, especially in public, education, culture and financial-information scenarios.

8. Institutional data, content and role-based access

In institutional scenarios, data and content are handled with the principle of **role-based access**: each user accesses only the data and content they are authorized for. Institutional data is kept separate from personal data; sensitive content is protected from unauthorized visibility.

The use of third-party services (cloud, AI, data infrastructure) requires the data flow and responsibility boundaries to be defined in advance.

9. Additional care for children, youth and education contexts

In the **Education & Campus Transformation** beachhead and other scenarios with possible young users (especially CampusVerse and education demos), additional safeguards apply: content appropriateness, additional data

protection, age-appropriate interaction and, where needed, parental/institutional approval. Child and youth safety is the highest-priority design requirement in this context.

10. Trust in public, municipal and disaster scenarios

In the **Civic Resilience & Disaster Awareness** beachhead, in CivicVerse and disaster-awareness scenarios (Disaster Awareness Journey, Civic Service Navigator), trust relates directly to public value. The principles:

- **Panic and misdirection are prevented:** the disaster scenario is designed to be educational and calming, and is clearly marked so it is not confused with a real emergency.
- **Correct behavior is emphasized:** content is grounded in authoritative sources and official guidelines.
- **Accessibility is essential:** a public service must be reachable by everyone.

11. Respectful representation in cultural heritage and tourism experiences

In the **Cultural Heritage & Tourism Experience** beachhead, in scenarios such as Heritage Twin and AI Museum Guide, cultural representation is handled with the principle of **respect and accuracy**: cultural and historical content is grounded in approved sources and the guidance of the relevant institution; sensitive cultural elements are presented respectfully; and meaning and context are preserved in multilingual access.

12. Ethical-AI principles

Multiversera's ethical-AI approach rests on these principles: transparency (the user knows they are interacting with an AI), fairness (mitigation of bias and discrimination is intended), accountability (decisions are auditable and subject to human approval), boundedness (the AI stays within a defined scope) and human-centeredness (human approval is essential in critical decisions).

13. Measurable trust and audit trail

Trust is not only an intention but a **measurable** requirement. Via Impact Layer, content accuracy, accessibility usage and interaction are observed anonymously and consistently. An **audit trail** aims to make critical operations traceable; **human-in-the-loop** ensures a human remains involved in critical decisions.



Trust spreading across sectors

14. Risks and the mitigation approach

The risk headings below are reviewed in every scenario, and appropriate mitigation is designed. This is not a complete list or a guarantee; it is a risk-awareness framework.

Risk heading	Mitigation approach
Personal data	Data minimization, consent, purpose limitation
Special-category data	Avoid processing; if needed, additional protection and legal review
Child / youth users	Content appropriateness, additional protection, age-appropriate interaction
AI avatar / persona permission	Explicit permission, authority to represent, scope and duration limits
Investment-advice risk	In financial scenarios, information not advice; legal approval
Incorrect / insufficient information	Source verification, content boundaries, accuracy review
Panic in public / disaster scenarios	Educational design, clear marking, official sources
Cultural-heritage representation	Respectful, approved-source representation
Accessibility	WCAG compliance, captions, audio guide, keyboard navigation
Bias / discrimination	Fairness principle, content review, diversity awareness
Data retention and deletion	Defined retention period and deletion; data minimization
Cross-border data transfer	Legal review before transfer; appropriate safeguards
Third-party services	Advance definition of data flow and responsibility boundaries
Human approval	Human-in-the-loop in critical decisions
Audit trail	Traceability in critical operations

15. What this is not

- Not a legal opinion or legal advice.
- Not a definitive guarantee of KVKK/GDPR compliance.
- Not a finished compliance certificate or audit report.
- Not an investment or financial-guidance document.
- Not the security statement of a finished, live product.

This document is a note of **design principle, compliance framework and risk awareness**. Every implementation requires the institution's own legal and compliance review.

16. A first Trust Review conversation

Before a pilot or collaboration, we recommend a **Trust Review** conversation in which we jointly review the trust, KVKK and ethical-AI requirements of the relevant scenario. This conversation addresses data flow, the consent and permission structure, persona boundaries, source verification, accessibility and audit-trail requirements.



Trust Review / first conversation



Section / closing

17. Contact

For a Trust Review or trust-requirements conversation: info@multiversera.com

Provenance note (brief): Multiversera's vision of privacy and responsible technology was articulated in 2022; today the trust layer is the matured form of that vision. This is not a product proof; it is a note of vision continuity.

© Multiversera — All Rights Reserved · info@multiversera.com · Confidential · Limited Distribution

This document is not legal advice and provides no guarantee of compliance; it offers design principles, a compliance framework and risk awareness. Every implementation requires independent legal review.