Tier-1 Reviewer Memo: What Questions Remain?

Post-homepage review - Terra Vita Hub

Following a full read of the Terra Vita Hub homepage and its public governance posture, the reviewer concludes that the platform presents a mature, bounded and institutionally aligned governance infrastructure. The public layer is unusually strong.

Before institutional reliance, the following remaining questions should be validated through protected diligence, annexes or a purpose-bound walkthrough.

1. Evidence Lineage & Replayability

• How is lineage enforced at the schema level across attachments, MRV artefacts, spatial files, and reviewer notes?

• Can reviewers replay a full decision sequence - state, action, condition, escalation and export - without gaps?

• How are cross-programme lineage collisions prevented in multi-deployment environments?

2. Reviewer Controls & Override Governance

• What constitutes a privileged override, and how is it constrained?

• How are informal approval chains prevented from re-emerging inside the protected environment?

• What is the exact reviewer-role hierarchy, and how are role changes audited?

3. Deployment Isolation & Sovereignty Posture

• What are the technical isolation guarantees between deployments?

• How is data residency enforced at the infrastructure and schema layers?

• What is the key-management posture, and who controls encryption boundaries?

4. MRV Attachment & Methodology Integrity

• How does the Hub ensure MRV artefacts cannot drift from programme context?

• How are MRV partner integrations authenticated, versioned, and audited?

• How does the system prevent MRV-derived indicators from being misinterpreted as automated decisions?

5. Assurance Pathways & Audit Reconstruction

• What is the granularity of audit logs - field-level, object-level and workflow-level?

• Can auditors export a full reconstruction bundle for external review?

• How are tamper-evident guarantees implemented?

6. Interoperability & External System Boundaries

• How does the Hub safely integrate with external systems without compromising lineage?

• What are the allowed vs. disallowed integration surfaces?

• How is data provenance preserved when importing from external sources?

7. Programme Lifecycle Controls

• How are mid-lifecycle changes - evidence updates, reviewer reassignment and MRV corrections - governed?

• What prevents post-approval drift in long-running programmes?

• How is funding-readiness posture validated and locked?

8. Export Posture & Committee Pack Integrity

• How is export posture computed, and what prevents manipulation?

• Can committee packs be reconstructed identically at any later date?

• How are conditional vs. release-ready states enforced?

9. Operational Assurance & Vendor-Risk Evidence

• What is the incident-response posture and escalation chain?

• What are the RPO/RTO bands for institutional deployments?

• How is identity federation implemented across institutional SSO providers?

10. Governance Drift & Change-Management Controls

• How are configuration changes - roles, workflows and thresholds - audited?

• What prevents governance drift across multi-programme deployments?

• How are schema migrations handled without compromising lineage?

Reviewer Conclusion

• The homepage provides a clear, bounded and institutionally credible governance posture.

• The remaining questions are not gaps. They are expected diligence items that must be validated under protected review before institutional reliance.

• Recommended next actions: proceed to the Institutional Review Index; issue the Minimum Evidence Request; schedule a purpose-bound protected walkthrough; request the Control Evidence Boundary and Operational Assurance annexes.