UNCLASSIFIED

XR/AI Digital Twin Training Platform

Modernizing Military Maintenance Training Through
Extended Reality, Artificial Intelligence, and Digital Twin Technology

White Paper: EnvisXR R&D; Unit
Version: 1.0 | February 2026
Classification: UNCLASSIFIED / CUI when populated with operational data
Featured Implementation: 60KW Tactical Quiet Generator (TQG) Maintenance Trainer
Prepared for: DoD SBIR/STTR and Government Acquisition Review

Table of Contents

1. Executive Summary

The U.S. military faces a persistent gap between the complexity of modern equipment and the training tools available to maintain it. Technical manuals run hundreds of pages. Access to live equipment for training is costly and operationally disruptive. And traditional training methods leave no digital trace — making it impossible to track competency, demonstrate readiness, or satisfy the data reporting requirements of DoDI 1322.26.

The **XR/AI Digital Twin Training Platform**, developed by EnvisXR, addresses each of these challenges in a single, deployable system. By combining an AI-powered conversational instructor grounded in official technical documentation with an interactive 3D digital twin of the target equipment, the platform delivers maintenance training that is:

- **Accessible** — no physical equipment required; deployable on any LAN or GovCloud environment
- **Accurate** — all AI responses are grounded in approved TM content via Retrieval Augmented Generation (RAG), eliminating hallucinated guidance
- **Compliant** — architected from the ground up for DoDI 1322.26, IEEE 9274.1.1 (xAPI), and the Total Learning Architecture (TLA)
- **Scalable** — any Technical Manual can be ingested into the platform, producing a deployable training environment in days, not months

The current Proof of Concept implements this platform for the **60KW Tactical Quiet Generator (TQG)**, a high-priority maintenance training requirement across the force. The PoC demonstrates a fully operational AI instructor with 85+ individually addressable 3D components, streaming real-time responses, and an xAPI-ready interaction database — establishing a clear path toward full enterprise integration.

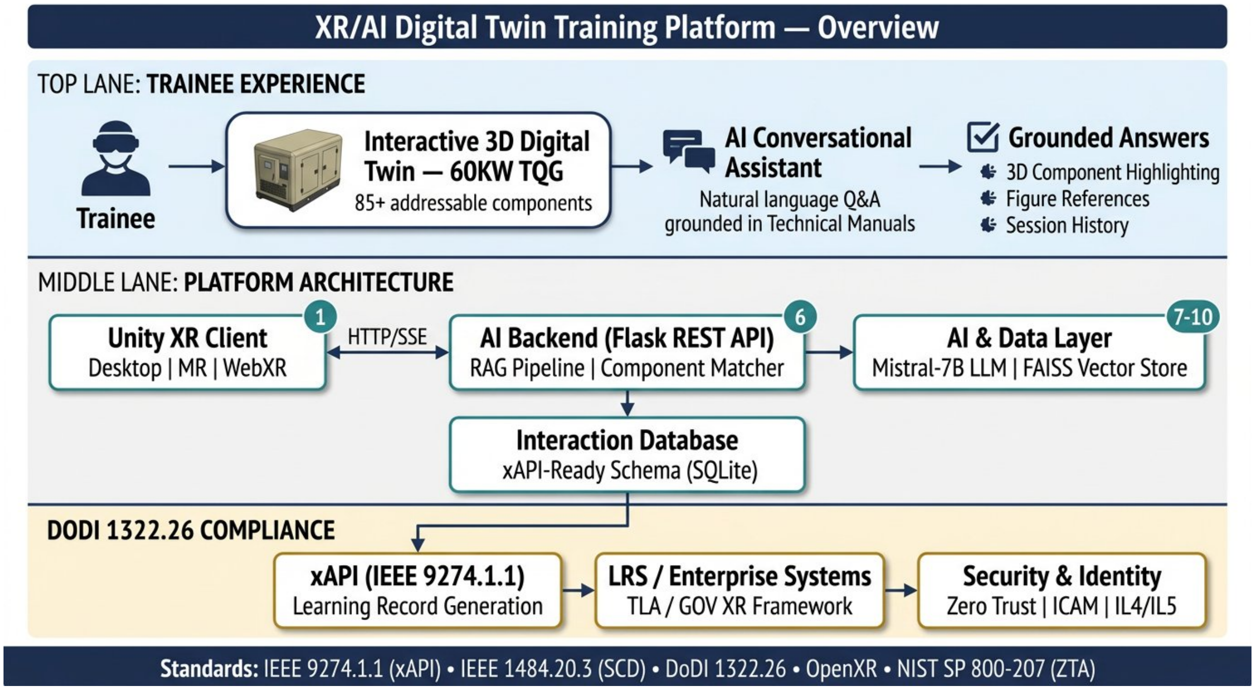

Figure 1: XR/AI Digital Twin Training Platform — three-lane overview showing the trainee experience (top), platform architecture (middle), and DoDI 1322.26 compliance layer (bottom).

2. The Problem: Training Gaps That Cost Readiness

Military maintenance training for complex equipment presents four interconnected challenges that existing tools do not adequately solve. Each gap independently reduces readiness; together, they create a systemic training deficit.

2.1. Equipment Access

Training units routinely lack access to the equipment they are trained to maintain. Deploying operational generators, vehicles, or weapons systems for classroom instruction is expensive, creates scheduling conflicts, and introduces readiness risk. For the 60KW TQG alone, units may share a single operational generator across multiple training events, limiting individual exposure time to minutes per trainee.

2.2. Technical Manual Complexity

The 60KW TQG Technical Manual exceeds 270 pages across dozens of Work Packages. Finding a specific maintenance procedure under operational stress is time-consuming and error-prone. The result is that field maintenance often proceeds from memory or tribal knowledge rather than from the authoritative technical source.

2.3. Spatial Understanding

Two-dimensional diagrams are fundamentally inadequate for developing the spatial awareness required for field maintenance. Misidentification of components under time pressure is a documented source of maintenance errors, preventable damage, and extended downtime.

2.4. Training Data Capture

Traditional maintenance instruction generates no structured digital record. Without data, commanders cannot:

- Assess individual trainee competency against defined standards
- Identify systemic training gaps across a unit or MOS
- Provide the compliance reporting that DoDI 1322.26 requires
- Demonstrate training readiness during inspections or audits

The absence of learning records is both an operational risk and a regulatory gap — one that grows more acute as the DoD moves toward the Total Learning Architecture and mandated xAPI adoption.

3. The Solution: XR/AI Digital Twin Training Platform

The EnvisXR platform integrates three capabilities — AI-grounded instruction, interactive 3D visualization, and standards-compliant learning data — into a unified training environment accessible from a standard workstation or mixed reality headset.

3.1. AI-Powered Conversational Instructor

At the core of the platform is a Retrieval Augmented Generation (RAG) pipeline that gives trainees instant access to the full content of the technical manual through natural language queries. A trainee can ask:

> *"Walk me through the air filter restriction indicator procedure."*
>
> *"What is the fuel water separator replacement work package?"*
>
> *"What components do I need to inspect before starting the generator?"*

The system retrieves the most relevant manual passages using semantic search over a FAISS vector store, then generates a grounded, accurate response using a locally-hosted large language model (Mistral-7B). Because the model is constrained to retrieved TM content, it cannot fabricate procedures or specifications — a critical safety requirement for maintenance training.

3.2. Interactive 3D Digital Twin

Alongside every AI response, the platform highlights the specific 3D components relevant to the trainee's query. The 60KW TQG digital twin includes **85+ individually addressable mesh components** — from the DCS control panel and fuel fill port to individual access doors, load cables, and the output terminal box — mapped to the component registry via cosine similarity matching.

Platform	Description	Status
2D Desktop	Windows/Linux workstation application	Implemented
Mixed Reality	Meta Quest, HTC Vive via OpenXR	Phase 2
WebXR	Browser-based, no install required	Phase 2

3.3. Standards-Compliant Learning Records

Every trainee interaction — query text, AI response, matched 3D components, relevant TM figures, timestamp, and session identifier — is persisted in an xAPI-ready database schema. The data model is designed for direct mapping to IEEE 9274.1.1 xAPI statements, enabling seamless forwarding to any DoD Learning Record Store (LRS) once Phase 2 integration is complete.

3.4. Web Workstation

In addition to the Unity client, the platform includes a browser-based training workstation providing:

- **Streaming AI responses** with real-time figure and component updates
- **Session history** — all trainee queries are logged and retrievable
- **Figure viewer** — clickable manual figures open an enlarged view with AI-generated contextual explanation
- **PDF source linking** — each figure links to the exact page in the TM PDF
- **Procedure video playback** — contextual video plays automatically after AI response delivery
- **Trainee authentication** — session-based identity capture aligned with xAPI actor requirements

4. Platform Architecture

The platform is organized into distinct architectural layers aligned with the TLA / GOV XR Framework. Each layer is independently replaceable, allowing the system to evolve from air-gapped PoC to enterprise cloud deployment without rearchitecting.

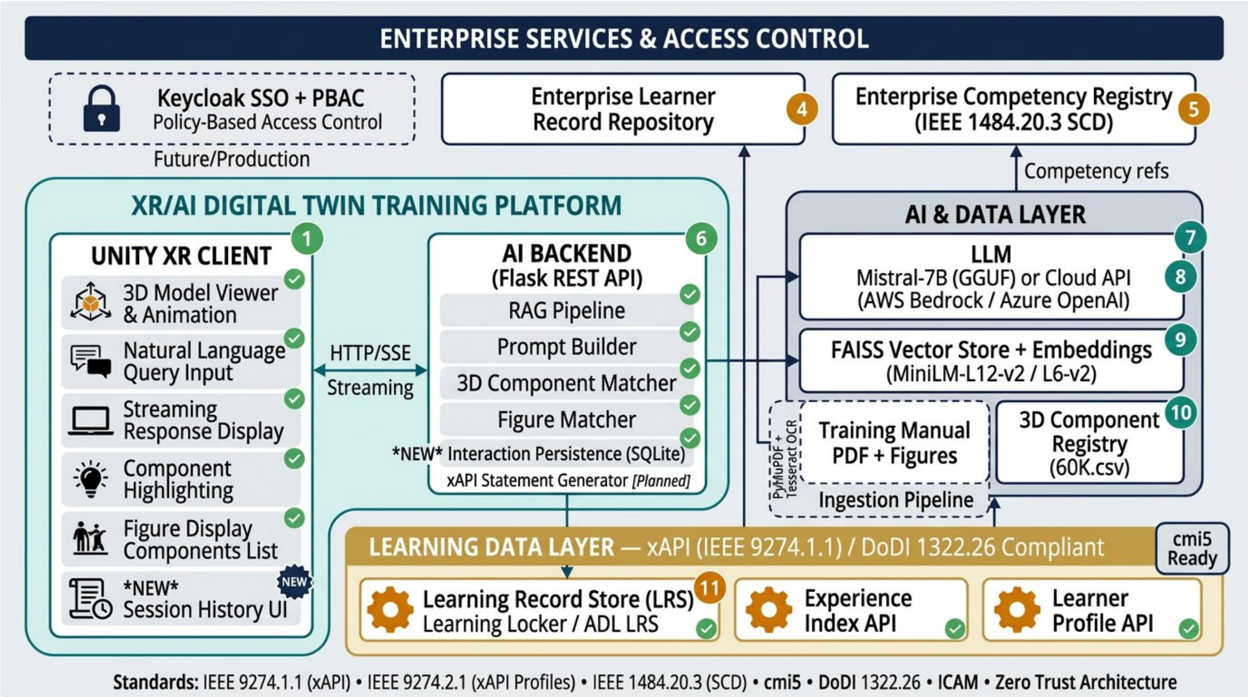

Figure 2: Detailed solution architecture showing all numbered TLA components, data flows, and integration points. Enterprise services (top), core platform (center), and the learning data / compliance layer (bottom).

4.1. Presentation Layer

- **Unity XR Client** — Cross-platform 3D application supporting desktop, MR headsets, and WebXR
- **Web Workstation** — Browser-based training interface for LAN deployment; no install required

4.2. AI Backend (Flask REST API)

- **RAG Pipeline** — Prompt construction from retrieved TM passages; semantic context injection
- **3D Component Matcher** — Cosine similarity scoring against the 85+ component registry
- **Figure Matcher** — Embedding-based figure retrieval with query-boost scoring
- **Interaction Persistence** — SQLite database with xAPI-ready schema; upsert logic for session continuity

4.3. AI and Data Layer

- **LLM — Mistral-7B (GGUF):** Quantized for GPU-accelerated inference on standard workstation hardware. Swappable for IL4/IL5 cloud API endpoints (AWS Bedrock, Azure OpenAI) without architectural changes.
- **FAISS Vector Store:** Semantic search over ingested TM content using Sentence-Transformer embeddings (all-MiniLM-L12-v2). Fully local, zero external API calls.
- **3D Component Registry:** 85+ components with category, description, mesh identifiers, and importance flags. Maintained as CSV/JSON for human-readable audit.

4.4. Ingestion Pipeline

- **PyMuPDF + Tesseract OCR** — Extracts text, tables, and figures from TM PDFs
- **Figure-to-page mapping** — Preserves the source page number for every extracted figure, enabling PDF-level source attribution
- **Vector embedding** — All extracted content is embedded and indexed in FAISS for semantic retrieval

4.5. Key Design Characteristics

Characteristic	Implementation
Response delivery	Server-Sent Events (SSE) — token-by-token streaming
Figures	Update dynamically during generation; linked to TM PDF pages
Component matching	Real-time cosine similarity against query + response embeddings
Data sovereignty	All inference, storage, and embedding runs on-premises
Model portability	Drop-in replacement: local GGUF, cloud API, or hybrid

5. DoDI 1322.26 Compliance Architecture

DoDI 1322.26 establishes requirements for distributed learning across the Department of Defense, including standards for learning data interoperability, actor identity, and system security. The EnvisXR platform is architected to satisfy these requirements at each phase of deployment.

5.1. Learning Data Interoperability — IEEE 9274.1.1 (xAPI)

The platform's interaction database captures all fields required for xAPI statement generation. The target statement structure aligns with ADL xAPI profiles:

```
{
  "actor": { "account": { "homePage": "https://envisxr.com", "name": "{dodId}" } },
  "verb": { "id": "http://adlnet.gov/expapi/verbs/asked" },
  "object": {
    "id": "https://envisxr.com/activities/60k-generator/query/{uuid}",
    "definition": { "type": "http://adlnet.gov/expapi/activities/question" }
  },
  "result": {
    "response": "{ai_response_text}",
    "extensions": {
      "https://envisxr.com/xapi/ext/matched-components": ["FuelFilter"],
      "https://envisxr.com/xapi/ext/confidence-score": 0.87
    }
  },
  "context": {
    "extensions": {
      "https://envisxr.com/xapi/ext/session-id": "{session_uuid}",
      "https://envisxr.com/xapi/ext/equipment-type": "60KW-TQG-AMMPS"
    }
  }
}
```

5.2. Actor Identity — ICAM Readiness

The system's **trainee_id** field is designed for direct substitution with the 10-digit DoD EDIPI. Phase 3 implementation will enforce authentication via Keycloak SSO with CAC/PIV card integration, satisfying the ICAM requirements of DoDI 1322.26 Section 3.1.

5.3. Security Architecture

Layer	Current (PoC / LAN)	Production Target
Network	Air-gapped LAN; no external connectivity	TLS 1.2+ via NGINX reverse proxy
Data at rest	Local filesystem	AES-256 encryption

Layer	Current (PoC / LAN)	Production Target
Data in transit	LAN only	TLS 1.2+ end-to-end
Identity	Self-reported trainee_id	ICAM / Keycloak SSO with CAC/PIV
Architecture	Perimeter isolation	Zero Trust (NIST SP 800-207)
Cloud readiness	N/A	IL4/IL5 (AWS GovCloud, Azure Gov)

5.4. Total Learning Architecture (TLA) Alignment

#	TLA Component	Platform Implementation	Status
1	Unity XR Client	Cross-platform (Desktop, MR, WebXR)	Implemented (Desktop)
4	Enterprise Learner Record Repo	xAPI statement forwarding to LRS	Phase 2
5	Enterprise Competency Registry	IEEE 1484.20.3 SCD competency tags	Phase 3
6	AI Backend	Flask RAG Pipeline (REST / SSE)	Implemented
7-10	AI & Data Layer	LLM + FAISS + Component Registry	Implemented
11	Learning Record Store	xAPI-conformant LRS (Learning Locker / ADL LRS)	Phase 2

6. AI Safety: Hallucination Guardrails and Source Attribution

For maintenance training, accuracy is a safety requirement, not a preference. Incorrect procedural guidance can result in equipment damage, injury, or mission failure. The platform implements layered controls to keep the AI grounded in approved documentation.

6.1. Current Guardrails (PoC)

- **Retrieval-constrained generation:** All responses are generated using top-K=3 retrieval from the FAISS vector store. The LLM receives only retrieved manual passages as context, limiting its ability to generate unsourced content.
- **Dual-embedding scoring:** Figure and component matches are scored using cosine similarity with both query and response embeddings.
- **Session persistence:** Every interaction is logged with full input/output pairs, enabling post-hoc audit of AI accuracy by training administrators.

6.2. Phase 2 Enhancements

- **Source Attribution UI:** Every AI response will display the Chapter, Work Package (WP), and page number of the TM passage that grounded the answer. Trainees can click through to the exact source page in the TM PDF.
- **Zero Hallucination Policy:** The system prompt will instruct the model to respond with "Information not found in available manual sections" when retrieved context is insufficient.
- **Confidence scoring:** Each response will carry a retrieval confidence score based on embedding similarity, enabling the UI to flag low-confidence answers for human review.

7. Operational Use Cases

7.1. Individual Maintenance Training

A Soldier preparing for a field maintenance task queries the platform with a natural language description of the procedure. The AI instructor returns the relevant TM guidance, highlights the involved 3D components, and surfaces the associated manual figures. The Soldier can interact with the 3D model to build spatial familiarity before touching the physical equipment.

Outcome: Reduced time-to-competency. Training proceeds independently of equipment availability.

7.2. Classroom / Instructor-Led Training

An instructor projects the web workstation during a maintenance class. As students ask questions, the AI returns authoritative answers with 3D visualization in real time. The instructor validates and supplements the AI responses, using the platform as a dynamic teaching aid rather than flipping through a printed TM.

Outcome: More interactive instruction. Questions are answered instantly from the official TM. All queries are logged for post-class review.

7.3. Field Reference

A maintainer in the motor pool encounters an unfamiliar component or procedure. Using a LAN-connected workstation or tablet, they query the platform and receive the relevant TM procedure, figure references, and a 3D visualization of the component's location on the equipment.

Outcome: Faster troubleshooting. Reduced reliance on memory or peer guidance for infrequent procedures.

7.4. Competency Assessment and Reporting

A training NCO reviews session logs to assess which TM topics trainees have queried, identify knowledge gaps across the section, and generate compliance reports for DoDI 1322.26. The xAPI-ready data schema allows direct forwarding to enterprise LRS once integrated.

Outcome: Data-driven training management. Objective competency visibility at the unit level.

7.5. Pre-Deployment Readiness

A unit preparing for deployment conducts maintenance proficiency checks against the platform. Session records provide auditable evidence of training completion.

Outcome: Documented readiness. Compliance evidence generated automatically as a byproduct of training.

8. Return on Investment

8.1. Cost Avoidance

Cost Driver	Traditional Training	XR/AI Digital Twin Platform
Equipment allocation	Requires operational asset; creates readiness gap	Zero equipment required
TM navigation time per query	5–15 minutes (manual search)	Under 10 seconds (AI retrieval)
Instructor preparation	Hours per lesson plan	Platform is always current with TM content
Training site constraints	Requires classroom + equipment	Any workstation on the LAN
Compliance reporting	Manual data entry, paper records	Automatic xAPI generation

8.2. Scalability Economics

The platform's content-agnostic architecture means that expanding to a new equipment type requires only:

- **Ingestion of the TM PDF** — automated via the PyMuPDF + OCR pipeline
- **3D model mapping** — component registry (CSV) maps mesh names to TM terminology
- **Deployment** — the same backend, the same UI, the same compliance layer

No custom software development is required per equipment type. This transforms the cost curve from *linear per system* (traditional CBT/IMI development) to *marginal per system* (data ingestion only).

8.3. Training Throughput

Because the platform requires no physical equipment and operates on standard workstation hardware, it removes the primary bottleneck in maintenance training: equipment availability. Multiple trainees can train simultaneously on independent workstations against the same digital twin, and training can occur at any time the LAN is operational.

9. Deployment Options and Security Architecture

Environment	Configuration	Security Posture	Use Case
LAN / Air-gapped	Local GPU server, restricted network	Maximum data sovereignty	PoC, OPSEC-sensitive sites, SCIF-adjacent
Base / Installation	LAN server, multiple workstations	On-premises, controlled access	Unit-level and schoolhouse training
GovCloud IL4/IL5	AWS GovCloud / Azure Government	FedRAMP High, DISA STIG	Multi-installation, enterprise LRS integration
Mixed Reality	Meta Quest / HTC Vive via OpenXR	Same backend security posture	Immersive hands-on simulation
WebXR Browser	No install required	TLS + authentication	Reach training, individual self-study

9.1. Air-Gapped Deployment (Current State)

- **LLM inference** runs locally via GPU-accelerated CUDA on WSL2/Linux
- **All embeddings** are computed on-premises using Sentence-Transformers
- **No external network calls** are made at any point during operation
- **Training content** (TM PDFs, figures, vector embeddings) never leaves the local server

This configuration is immediately deployable in environments where data sovereignty is non-negotiable.

10. Production Roadmap

Phase 1 — Proof of Concept (Current)

Capability	Status
2D desktop application with interactive 3D digital twin (85+ components)	Complete
Local LAN deployment — fully air-gapped, zero external dependencies	Complete
Flask AI backend with RAG pipeline and FAISS vector store	Complete
3D component matching with real-time mesh highlighting	Complete
Multi-modal ingestion pipeline (PyMuPDF + Tesseract OCR)	Complete
Web workstation with streaming responses, session history, and figure viewer	Complete
xAPI-ready interaction database with session persistence	Complete
Trainee authentication with session-based identity capture	Complete

Phase 2 — Cross-Platform Expansion and Standards Compliance

Capability	Target
Deployment to Meta Quest / HTC Vive via OpenXR	Mixed reality training
WebXR browser-based frontend	Zero-install access
LRS deployment with live xAPI statement emission	Full DoDI 1322.26 compliance
Source attribution UI — Chapter / WP / page reference panel	Verifiable AI grounding
Zero hallucination policy enforcement	AI safety hardening
Confidence scoring and low-confidence flagging	Training quality assurance

Phase 3 — Enterprise and Cloud Integration

Capability	Target
Migration to IL4/IL5 Secure GovCloud	Enterprise-scale deployment
Keycloak SSO / ICAM integration with CAC/PIV and DoD EDIPI	Mandatory identity assurance
Enterprise Learner Record Repository forwarding	TLA compliance
IEEE 1484.20.3 Sharable Competency Definition (SCD) integration	Competency-based training

Capability	Target
Multi-equipment corpus expansion	Platform scalability
Experience Index API and Learner Profile API integration	Unified learning record

11. Applicable Standards and References

Standard	Description	Platform Relevance
DoDI 1322.26	Distributed Learning	Overarching compliance framework
IEEE 9274.1.1	Experience API (xAPI)	Learning record generation and LRS emission
IEEE 9274.2.1	xAPI Profiles	Standardized activity and verb definitions
IEEE 1484.20.3	Sharable Competency Definitions (SCD)	Competency-based assessment (Phase 3)
OpenXR	VR/AR hardware abstraction	Cross-platform XR deployment
NIST SP 800-207	Zero Trust Architecture	Security architecture alignment
ICAM	Identity, Credential, and Access Management	DoD identity assurance framework
cmi5	xAPI profile for LMS interoperability	LRS and LMS integration profile
FedRAMP	Federal Risk and Authorization Management	Cloud security authorization

12. Strategic Outlook

The XR/AI Digital Twin Training Platform is not a single-application training tool. It is an **extensible architecture** designed around a content-agnostic, equipment-agnostic core. The 60KW TQG is the current corpus. The platform is equally capable of ingesting Technical Manuals for any piece of equipment — ground vehicles, aviation systems, crew-served weapons, communications infrastructure — and producing a fully interactive, AI-powered training environment from that documentation.

12.1. The Platform Value Proposition

Dimension	Value
Any TM	Automated ingestion produces a deployable training environment in days
Any platform	Desktop, MR headset, or browser — same AI backend, same data
Every interaction	Recorded in xAPI-ready format from the first query
No equipment	Training proceeds wherever a workstation can operate
Built to evolve	Air-gapped LAN today; enterprise GovCloud tomorrow

12.2. Path Forward

The 60KW TQG PoC demonstrates that the architecture works: AI-grounded instruction is operational, 3D component matching is accurate, streaming responses are responsive, and every interaction is captured in a standards-compliant schema. The roadmap — Phase 2 cross-platform expansion, Phase 3 enterprise integration — is not a redesign. It is an activation of capabilities that the architecture already supports. The foundation is built. What follows is deployment.

13. Glossary

Term	Definition
CAC/PIV	Common Access Card / Personal Identity Verification — DoD smart card authentication
DCS	Digital Control System — primary operator interface on the 60KW TQG
DoDI 1322.26	Department of Defense Instruction for Distributed Learning policy
EDIPI	Electronic Data Interchange Personal Identifier — 10-digit DoD identity number
FAISS	Facebook AI Similarity Search — vector database for semantic retrieval
GGUF	GPT-Generated Unified Format — quantized LLM model format for local inference
ICAM	Identity, Credential, and Access Management — DoD identity framework
IL4/IL5	Impact Level 4/5 — DoD cloud security classifications for CUI and above
LLM	Large Language Model — AI model for natural language understanding and generation
LRS	Learning Record Store — xAPI-compliant data store for learning records
MR	Mixed Reality — immersive computing combining physical and digital environments
RAG	Retrieval Augmented Generation — AI technique grounding LLM output in source documents
SSE	Server-Sent Events — HTTP-based streaming protocol for real-time response delivery
TLA	Total Learning Architecture — DoD framework for interoperable learning systems
TM	Technical Manual — authoritative maintenance and operation documentation
TQG	Tactical Quiet Generator — military power generation equipment
WP	Work Package — discrete procedural unit within a Technical Manual
xAPI	Experience API (IEEE 9274.1.1) — standard for learning activity data exchange
XR	Extended Reality — umbrella term for VR, AR, and MR technologies
ZTA	Zero Trust Architecture — security model assuming no implicit trust (NIST SP 800-207)

EnvisXR — R&D; Unit
XR/AI Digital Twin Training Platform