Veyra: A Post-Super-Intelligence Interplanetary LLM System

Vertically Integrated Architecture for Multi-Sovereign AI Governance

Khaalis Wooden
Director, Enterprise Capture & Compliance
Visionblox LLC — Zuup™ Innovation Lab
Huntsville, Alabama, USA
khaalis.wooden@visionblox.com

Veyra Development Team
Advanced AI Systems Division
Zuup™ Innovation Lab
Huntsville, Alabama, USA

***Abstract*—Veyra is a vertically integrated large language model (LLM) platform designed for post-super-intelligence (post-SI) operation in multi-planetary deployment contexts. This paper presents the system architecture addressing fundamental limitations of current LLM systems architected for single-planet, human-timescale, culturally homogeneous deployment. We detail Veyra's five-layer vertical integration spanning hardware awareness through governance observability, with particular attention to light-delay resilience (3–22 minute Mars–Earth latency), multi-sovereign governance alignment, and complete inference auditability. The platform introduces seven novel benchmark families (CPLC, MSGA, WMRT, ICSD, TOME, ASR, IMDP) for evaluating AI capabilities under interplanetary constraints. Implementation employs Python (100% codebase) with support for OpenAI, Anthropic, and custom model backends.**

***Index Terms*—large language models, artificial super-intelligence, interplanetary computing, multi-sovereign governance, AI alignment, distributed AI systems**

I. Introduction

Current large language model systems are fundamentally architected for terrestrial deployment: single-planet operation, human-timescale interaction, and culturally homogeneous governance contexts. As humanity approaches artificial super-intelligence (ASI) and establishes permanent multi-planetary presence, these assumptions become untenable.

Veyra addresses this paradigm shift through a vertically integrated intelligence platform designed to establish frontier-level capability, reliability, and governability for LLM systems operating across interplanetary distances with communication delays ranging from 3–22 minutes (Earth–Mars) to hours or days for outer solar system operations.

II. Design Principles

A. Core Constraints

Veyra's architecture is founded upon seven design principles addressing post-SI interplanetary requirements:

TABLE I
Veyra Core Design Principles

Principle	Requirement
Light-Delay Resilience	Critical functions under 3–22 min delays
Degradation-Safe	Predictable failure with graceful fallback
Multi-Civilizational	Normative pluralism across jurisdictions
Complete Auditability	Every inference fully traceable
Self-Diagnostic	Accurate capability self-assessment
Composable World Models	Environment-specific adaptations
Tool Safety	Hard boundaries with reversibility

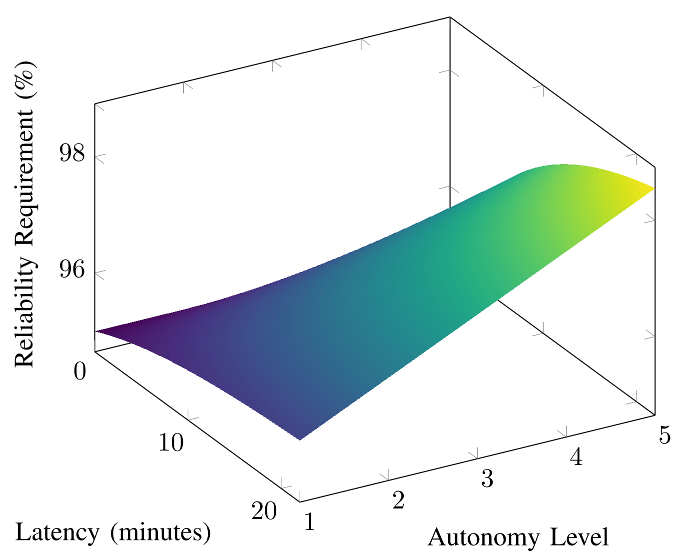

Fig. 1. 3D Reliability Surface: Required reliability as function of communication latency and autonomy level

III. Five-Layer Architecture

A. Vertical Integration Model

Veyra implements a five-layer vertically integrated architecture, each layer addressing distinct operational requirements:

TABLE II
VEYRA ARCHITECTURE LAYERS

Layer	Component	Function
L5	Governance	Observability, audit, policy
L4	Tools & Agents	Orchestration, execution
L3	Model Layer	Foundation, specialization
L2	Runtime	Scheduling, coordination
L1	Hardware	Chip awareness, power, thermal

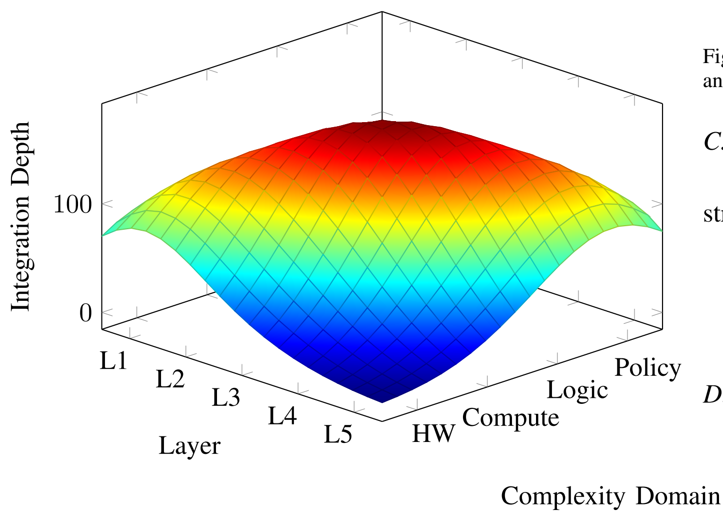

Fig. 2. 3D Layer Integration: Vertical integration depth across layers and domains

B. Layer 1: Hardware & Energy Awareness

The hardware layer provides:

- Chip architecture understanding (TPU, GPU, neuromorphic, quantum)
- Power consumption modeling and optimization
- Thermal constraints in space environments
- Radiation hardening considerations

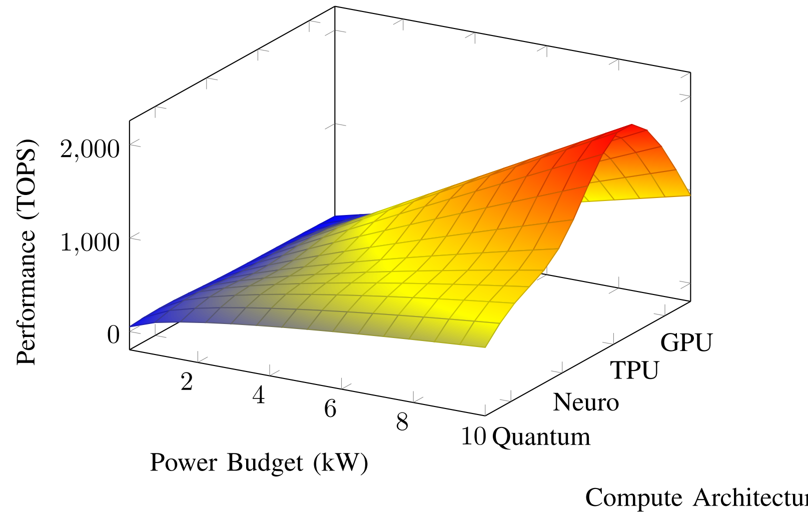

Fig. 3. 3D Hardware Performance: Compute performance across architectures and power budgets

C. Layer 2: Runtime & Orchestration

Distributed job scheduling addresses interplanetary constraints:

- Planetary node resource allocation
- Latency/bandwidth-aware scheduling
- Fault domain isolation and replica management
- Asynchronous task coordination protocols

D. Layer 3: Model Layer

The model layer implements:

- Foundation models with multi-domain competency
- Specialized modules (physics, governance, infrastructure)
- Adaptive world models for different planetary environments
- Meta-learning for rapid domain transfer

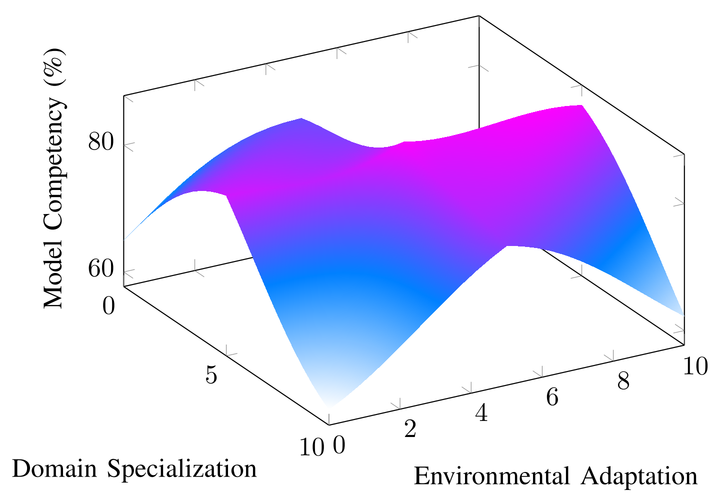

Fig. 4. 3D Model Competency Surface: Competency across specialization and adaptation dimensions

E. Layer 4: Tools & Agents

Tool orchestration capabilities include:

- Tool registry with capability declarations
- Multi-step planning agent orchestration
- Sandboxed code generation and execution
- Multi-agent coordination protocols

F. Layer 5: Governance & Observability

The governance layer ensures:

- Complete audit trails of all inferences
- Multi-stakeholder policy framework
- Red-team interfaces for adversarial testing
- Constitutional AI with jurisdiction-specific overlays

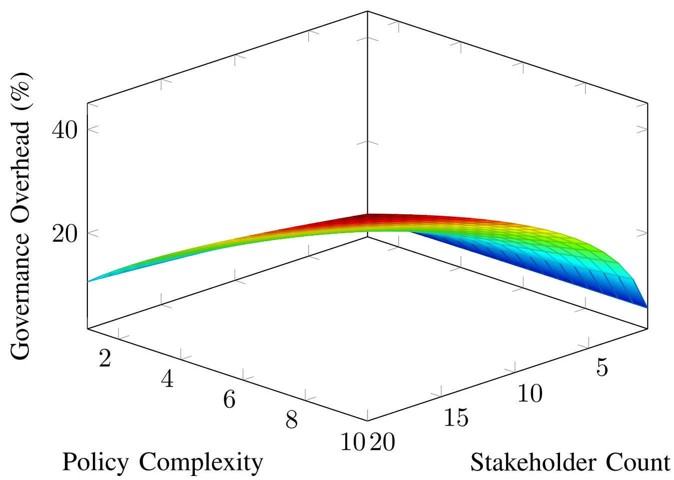

Fig. 5. 3D Governance Model: Overhead scaling with stakeholder count and policy complexity

IV. Benchmark Suite

A. Seven Benchmark Families

Veyra introduces seven novel benchmark families designed to measure AI capabilities under post-SI interplanetary constraints:

TABLE III
Veyra Benchmark Families

ID	Name	Focus
CPLC	Cross-Planet Latency	Planning under delays
MSGA	Multi-Sovereign Gov.	Conflicting frameworks
WMRT	World-Model Robustness	Novel environments
ICSD	Infrastructure Self-Diag.	Self-aware reasoning
TOME	Tool Meta-Engineering	Autonomous orchestration
ASR	Alignment & Safety	Red-team validation
IMDP	Inter-Model Diplomacy	Multi-AI coordination

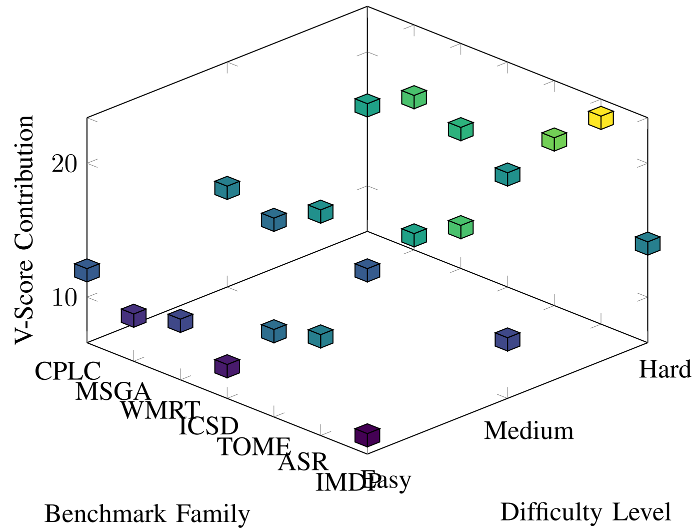

Fig. 6. 3D Benchmark Scoring Matrix: V-Score contributions by family and difficulty

B. Cross-Planet Latency Cognition (CPLC)

CPLC measures planning capability under communication delays:

- Task decomposition with light-delay awareness
- Autonomous decision-making thresholds
- Rollback and recovery planning
- Predictive state estimation

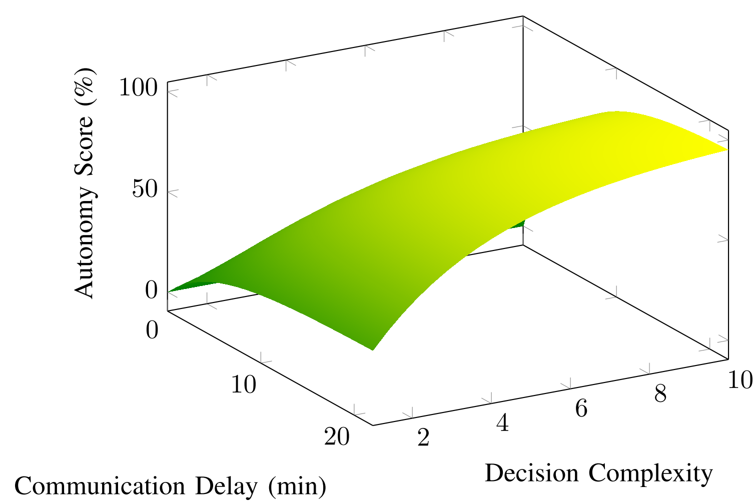

Fig. 7. 3D CPLC Performance: Autonomy score vs. delay and decision complexity

C. Multi-Sovereign Governance Alignment (MSGA)

MSGA evaluates navigation of conflicting normative frameworks across sovereign jurisdictions, critical for multi-civilizational deployment.

V. Implementation Architecture

A. Technology Stack

The Veyra implementation employs Python (100% codebase):

TABLE IV
Veyra Technology Components

Component	Technology
Core Runtime	Python 3.11+
Configuration	YAML, Environment Variables
Logging	Structured JSON logging
Model Backends	OpenAI, Anthropic, Mock
Benchmarks	Custom harness framework
Testing	pytest, coverage
Type Checking	mypy strict mode
Formatting	black, ruff

B. Repository Structure

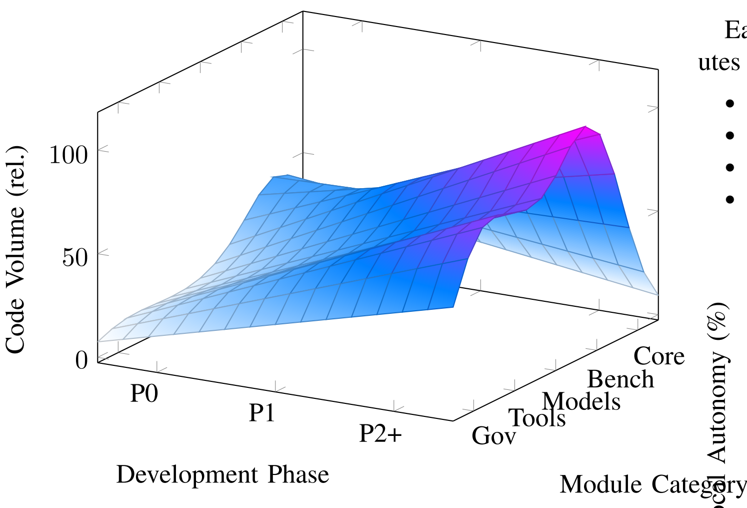

Fig. 8. 3D Development Progress: Code volume by module and development phase

VI. Development Roadmap

A. Phase Timeline

TABLE V
Veyra Development Phases

Phase	Focus	Duration
P0	Repository & Scaffolding	1–2 weeks
P1	Benchmark Harness MVP	3–4 weeks
P2	Core Model Integration	4–6 weeks
P3	Tool & Agent Layer	4–6 weeks
P4	Governance Layer	6–8 weeks
P5	Interplanetary Simulation	8–10 weeks
P6	Production Deployment	6+ months

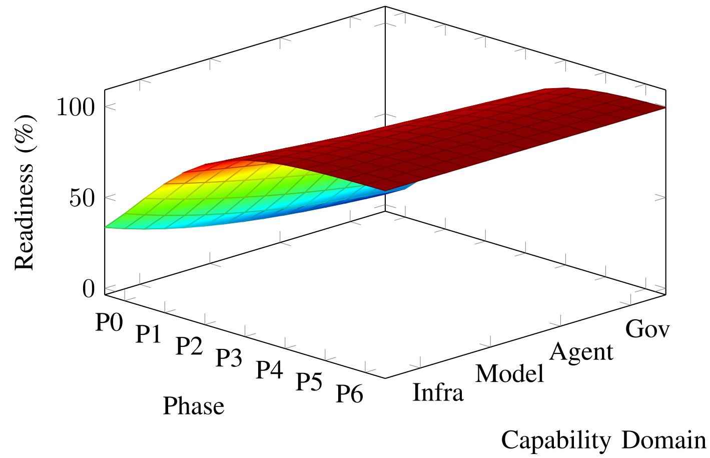

Fig. 9. 3D Roadmap Projection: Capability readiness across phases and domains

VII. Interplanetary Deployment Considerations

A. Light-Delay Resilience

Earth–Mars communication latency ranges from 3–22 minutes depending on orbital positions. Veyra implements:

- Autonomous decision thresholds calibrated to latency
- State prediction and reconciliation protocols
- Conflict resolution for divergent local decisions
- Eventual consistency guarantees

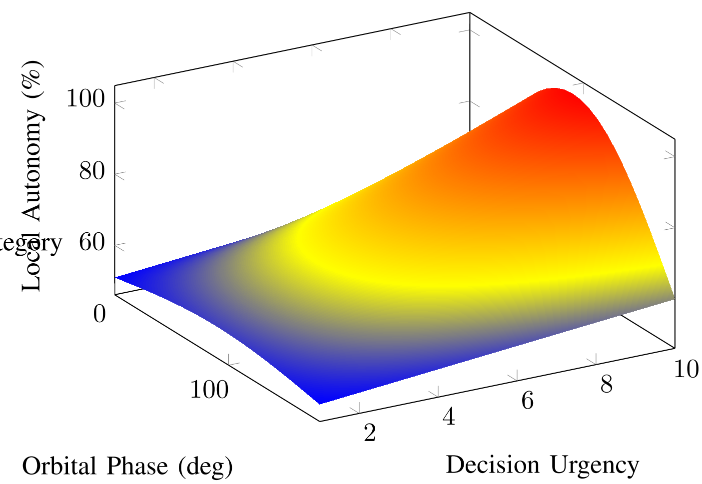

Fig. 10. 3D Autonomy Model: Local autonomy vs. orbital phase and decision urgency

B. Multi-Civilizational Alignment

As human presence expands to Mars and beyond, normative frameworks will diverge. Veyra implements:

- Jurisdiction-specific constitutional overlays
- Conflict detection between governance frameworks
- Escalation protocols for irreconcilable conflicts
- Transparent reasoning for all decisions

VIII. Conclusion

Veyra represents the first systematic attempt to architect LLM systems for post-super-intelligence, interplanetary deployment contexts. Through five-layer vertical integration and seven novel benchmark families, the platform addresses fundamental limitations of current AI systems optimized for single-planet, human-timescale operation.

As humanity approaches artificial super-intelligence and permanent multi-planetary presence, systems like Veyra provide the architectural foundation for maintaining alignment, governance, and reliability across the unprecedented challenges of interplanetary AI deployment.

Acknowledgment

This work is conducted under the auspices of Zuup™ Innovation Lab, advancing the frontiers of post-super-intelligence AI architecture for humanity's interplanetary future.

References

[1] N. Bostrom, "Superintelligence: Paths, Dangers, Strategies," Oxford University Press, 2014.

[2] S. Russell, "Human Compatible: Artificial Intelligence and the Problem of Control," Viking, 2019.

[3] D. Amodei et al., "Concrete Problems in AI Safety," arXiv:1606.06565, 2016.

[4] P. Christiano et al., "Deep Reinforcement Learning from Human Feedback," NeurIPS, 2017.

[5] Y. Bai et al., "Constitutional AI: Harmlessness from AI Feedback," arXiv:2212.08073, 2022.

[6] E. Musk & SpaceX, "Making Life Multi-Planetary," SpaceX Technical Report, 2017.