SOLVEWORX

T-FLUX ULTRA | ENTERPRISE AI ORCHESTRATION

Security by Architecture

Addressing Security Risks in Enterprise Large Language Model Applications

WHITE PAPER

T-Flux Ultra security architecture, guardrails, compiled execution, data governance and auditability

Secure by architecture. Governed by policy. Controlled by the enterprise.

Solveworx Digital Innovation
T-Flux Ultra — Private, governed enterprise AI

T-Flux Ultra Chief Technology Officer/Data Science

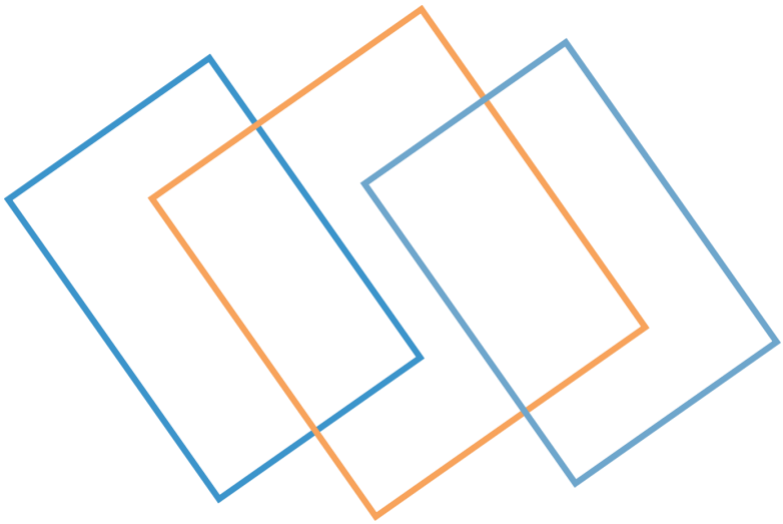

Executive Summary

Enterprise adoption of Large Language Models creates significant opportunities, but it also introduces a new class of security, privacy, governance and operational risks. Security specialists are concerned that LLMs can interpret malicious instructions, interact with sensitive enterprise data, generate unsafe outputs, invoke connected tools and take actions with excessive authority.

T-Flux Ultra is designed to address these concerns by creating a governed enterprise layer around AI. Rather than depending on an LLM to behave securely, T-Flux applies controls around the model and throughout the complete execution lifecycle.

- Private or on-premises deployment and customer-controlled data residency
- Compiled Java and C++ execution architecture with no dependence on Python scripting for orchestration
- Role-based, attribute-based and ACL-aware data access
- Policy and guardrail enforcement before, during and after model execution
- Least-privilege agent and tool permissions
- Controlled API access and output validation
- Data classification, redaction, version control and lifecycle management
- Model provenance, approval, deployment and retirement controls
- Comprehensive logging, tracing and auditability

Core principle: Do not trust the prompt, the model, the retrieved data or the generated output by default. Control the environment in which they operate.

1. Enterprise LLM Security Risks

Security practitioners commonly highlight the following areas of concern when LLMs are connected to enterprise data, applications and automated workflows. T-Flux Ultra addresses these risks using defence-in-depth rather than relying on any single control.

Risk	Enterprise Concern	T-Flux Ultra Response
Prompt injection	Malicious instructions attempt to override intended model behaviour.	Contain impact through least privilege, ACL-aware retrieval and policy enforcement.
Insecure output handling	Generated content is treated as trusted input by downstream systems.	Validate, sanitise and authorise outputs before execution.
Sensitive information disclosure	Confidential information is exposed through prompts, retrieval or responses.	Private deployment, restricted retrieval, redaction and egress control.
Insecure plugins / tools	Connected tools or APIs have weak validation or excessive access.	Authenticated connectors, scoped credentials and controlled capabilities.
Excessive agency	Agents receive more autonomy or privilege than required.	Explicit action boundaries, least privilege and human approval gates.
Overreliance	Users or systems accept inaccurate AI output without adequate verification.	Grounding, citations, traceability, quality controls and human review.
Supply-chain exposure	Third-party models and components introduce unknown risk.	Governed model catalogue, approved dependencies and controlled releases.

2. T-Flux Ultra Security Philosophy

T-Flux Ultra does not treat an AI model as a trusted enterprise application. The model operates inside a controlled environment in which access to data, applications, tools, APIs, users and actions is explicitly governed.

Infrastructure Security: Controls where T-Flux and its models execute.

Data Security: Controls what enterprise information can be accessed and which version is authoritative.

Model Security: Controls which models are permitted to operate and where they may be used.

Policy & Guardrails: Controls model behaviour and decision boundaries.

Agent & Tool Security: Controls which actions can be performed and under whose authority.

Output Security: Controls what happens to model-generated content before it reaches downstream systems.

Monitoring & Audit: Records and explains what occurred across the execution chain.

3. Prompt Injection

Security concern. Security specialists are concerned that direct or indirect prompt injection can manipulate a model into ignoring intended instructions, retrieving unauthorised information, or attempting prohibited actions.

T-Flux Ultra addresses prompt injection primarily through containment, explicit permissions and least privilege. A successful attempt to influence model behaviour does not automatically grant access to additional data, tools or system authority.

- RBAC and ABAC enforcement
- ACL-aware retrieval and corpus permissions
- Tenant isolation and trust labels
- Jurisdiction and data-classification controls
- Read/write/action restrictions
- Model-specific data permissions
- Tool and API permissions
- Human approval gates for sensitive actions

A model may interpret an instruction, but it does not automatically inherit authority to execute it.

4. Insecure Output Handling

Security concern. A common concern is that model-generated SQL, code, URLs, API payloads or commands may be passed directly to downstream systems and treated as trusted input.

T-Flux Ultra separates model generation from system execution. Outputs can be checked by deterministic application logic and policy before any downstream action occurs.

- Output schema validation
- Sanitisation and type checking
- Allow-listed commands and API operations
- Policy evaluation and business-rule validation
- Action-level authorisation

- Human approval for high-risk actions
- Comprehensive execution logging

Model Output → Validation → Policy Check → Authorised Capability → Execution

5. Sensitive Information Disclosure

Security concern. Security specialists are concerned that confidential data can be exposed through prompts, retrieved content, generated responses, logs or connected external services.

T-Flux Ultra is designed to run on-premises, in private cloud, sovereign cloud or other customer-controlled environments so enterprise information does not need to be sent to uncontrolled public AI services.

- PII masking and redaction
- Data classification and trust labels
- Jurisdiction restrictions
- ACL-aware retrieval
- Tenant separation
- Retention and deletion policies
- Restricted external egress
- Access logging and traceability

The enterprise determines where its data resides, which model may access it, who may request it and how it may be used.

6. Granular Data Control and Version Management

Security and governance requirements do not stop at access control. T-Flux Ultra gives the customer granular control over enterprise information throughout its lifecycle. Data can be version-controlled, classified, approved, superseded, retired, indexed, excluded, retained or deleted according to policy.

- Which version of a document or dataset is authoritative for AI use
- When an updated version becomes active and when older versions are superseded
- Which users, roles, tenants, agents or models may access the information
- Whether content may be indexed, embedded, retrieved or exposed to a model
- Which retention, jurisdiction and classification rules apply
- When human approval is required before a change takes effect
- How changes are logged, traced and audited

Your data. Your versions. Your policies. Your control.

7. Insecure Plugin and Tool Design

Security concern. Enterprise security teams are concerned that LLM-connected tools and plugins can become high-impact attack paths when authentication is weak, inputs are insufficiently validated, or the tool has excessive privilege.

T-Flux Ultra treats tools, connectors and APIs as controlled enterprise capabilities rather than unrestricted extensions of the model.

- Authenticated connectors and scoped credentials
- OAuth or enterprise authentication where appropriate
- Allow-listed APIs and functions
- Function-specific and tenant-specific permissions
- Input and output validation
- Transaction and rate limits
- Approval workflows for sensitive actions
- Complete tool invocation logging

An agent can be given permission to perform a narrow function without inheriting broader system access.

8. Excessive Agency

Security concern. Security specialists are concerned that autonomous agents may be allowed to take actions beyond the minimum authority required, increasing the impact of errors, hallucinations or malicious prompts.

T-Flux Ultra constrains AI agents within explicit authority boundaries and applies the principle of least privilege at the user, tenant, data, tool, workflow and action level.

- Separate read, create, update, approve, delete, transmit and execute rights
- Business-rule and transaction-limit controls
- Tenant and role-specific authority boundaries
- Human-in-the-loop approval gates
- Restricted toolsets for each agent role
- Logging of all agent actions and attempted actions

Bounded agency replaces unrestricted autonomy.

9. Overreliance, Hallucination and Human Oversight

Security concern. A further concern is that users or automated systems may accept AI-generated content as authoritative even when it is incomplete, inaccurate or inappropriate.

T-Flux Ultra reduces this risk by combining grounded retrieval, traceability, quality controls and optional human review.

- Retrieval-grounded responses
- Source citations and evidence traceability
- Response validation
- Multi-model comparison and synthesis
- Model agreement/disagreement analysis
- Confidence and quality indicators where implemented
- Human review for high-risk workflows

- Audit trails linking outputs to data and model versions

The goal is not merely to produce an answer, but to produce an answer that can be traced, verified, challenged and audited.

10. Supply-Chain and Model Provenance Risk

Security practitioners are increasingly concerned about the provenance of third-party models, libraries, datasets and AI components. Unknown or unapproved components can introduce vulnerabilities, licence issues, poisoned data or unexpected model behaviour.

The T-Flux Ultra Model Catalogue provides a governance layer around the AI models available to the enterprise. Model records can include source, architecture, format, version, provenance, approval status, deployment location, permitted tenants and lifecycle state.

Governed lifecycle: Registered → Evaluated → Approved → Deployed → Monitored → Updated → Retired

11. Compiled Java and C++ Architecture

A differentiating characteristic of T-Flux Ultra is its use of compiled Java and C++ software rather than a Python-script-driven execution architecture. This is not a claim that Java or C++ is inherently secure and Python inherently insecure. The security benefit comes from reducing dynamic runtime execution and constraining executable capability.

Controlled runtime: T-Flux can expose a known set of pre-defined, compiled and authorised functions rather than relying on arbitrary runtime scripts.

Reduced dynamic scripting: The architecture can reduce reliance on notebooks, runtime package installation, dynamically generated functions and agent-created scripts.

Deterministic interfaces: Compiled services support strongly defined APIs, interfaces and deployment boundaries.

Release governance: Known binaries, controlled builds and approved versions can be incorporated into conventional enterprise change-management processes

12. Reducing Arbitrary Code Execution

This architectural approach is particularly relevant to agentic AI. A model generating code does not mean that the generated code must be executed. T-Flux can restrict models and agents to approved compiled capabilities.

Controlled T-Flux Pattern	Higher-Risk Dynamic Pattern
Model → T-Flux Orchestrator → Approved Function → Policy Validation → Enterprise System	Model → Generated Script → General-Purpose Runtime → Enterprise System

13. Dependency and Software Supply-Chain Control

Compiled Java and C++ components also support conventional enterprise software controls around the platform itself.

- Dependency management and approved libraries
- Signed builds and controlled releases
- Software bill of materials
- Vulnerability scanning
- Version locking
- Formal CI/CD and deployment pipelines
- Release approval and rollback controls

The objective is to provide customers with a known and auditable software baseline: which version is deployed, which libraries are included, who approved the release, when it was deployed and which capabilities are enabled.

14. Auditability and Traceability

For regulated enterprises, security also requires the ability to reconstruct what happened. T-Flux Ultra can record significant events across the AI execution lifecycle.

- User identity and tenant
- Timestamp and request identifier
- Model and model version
- Retrieved documents and data permissions
- Policy and guardrail decisions
- Tool calls and agent activity
- Output validation
- Final response and downstream action

Auditable chain: User → Request → Data → Retrieval → Model → Validation → Action

15. Monitoring Through the T-Flux Control Planes

Control Plane

- Platform health
- Workload and infrastructure
- Security and audit
- Alerts
- Storage and capacity

Data Plane

- Enterprise data sources

- Ingestion and parsing
- Indexing and vector stores
- Retrieval
- Data quality and access controls

Tenancy Plane

- Tenant-specific workloads
- Model use
- Governance
- Operations
- Capacity and reporting

Models Catalogue Plane

- Available models
- Versions and provenance
- Deployment status
- Performance and routing
- Lifecycle management

16. Defence in Depth

Layer	Control	Purpose
1	**Private Infrastructure**	Keep AI workloads and enterprise data inside the controlled environment.
2	**Identity & Access**	Apply RBAC, ABAC and ACL controls.
3	**Data Governance**	Classify, version, restrict and monitor enterprise knowledge.
4	**Model Governance**	Approve and manage the models available to the organisation.
5	**Guardrails**	Apply enterprise policies to model behaviour.
6	**Agent Controls**	Restrict the actions AI is permitted to perform.
7	**Output Validation**	Treat generated content as untrusted until validated.
8	**Monitoring & Audit**	Record and analyse activity across the full execution chain.

17. Security by Architecture, Not by Prompt

The core distinction in T-Flux Ultra is that security does not depend exclusively on a model following instructions correctly. System prompts and model-level guardrails remain useful, but they should not be the enterprise security boundary.

The surrounding platform enforces access restrictions, policy restrictions, action limits, data boundaries, approval requirements, execution constraints and audit controls. This assumes that AI behaviour may occasionally be unpredictable and designs the system accordingly.

Conclusion

Enterprise adoption of Large Language Models requires a security model that extends beyond conventional application security. LLMs operate across language, data, models, tools and automated actions. The security architecture must therefore govern the interaction between all of these components.

T-Flux Ultra addresses this through an integrated architecture combining private AI infrastructure, governed enterprise data, controlled models, policy guardrails, least-privilege agents, compiled execution, validation and comprehensive auditability.

The objective is not to claim that LLM risk can be completely eliminated. The objective is to ensure that AI operates inside a controlled enterprise environment in which the customer retains authority over the data, models, permissions, workflows, actions and infrastructure.

T-FLUX ULTRA

Secure by architecture. Governed by policy. Controlled by the enterprise.