USA-ADL™ Executive Summary

Uranusys Secure AI Agent Development Lifecycle

Published 2025

Version 1.0

USA-ADL™ Executive Summary

Uranusys Secure AI Agent Development Lifecycle

Overview

Autonomous AI agents are moving rapidly from experimentation into operational environments. These systems are no longer limited to analysis or recommendation. They initiate actions, interact with systems, and operate continuously with limited human oversight.

This shift introduces a new class of operational, security, and governance risk. While existing AI governance models address models, data, and outcomes, they do not fully account for the lifecycle behavior of autonomous agents operating in production.

USA-ADL™ addresses this gap.

USA-ADL™ is an openly published lifecycle governance framework designed to govern AI agents as operational actors across their full lifespan. It defines how agents are authorized, constrained, monitored, audited, changed, and ultimately decommissioned in real environments.

The Governance Gap in Agentic AI

Traditional AI governance assumes bounded systems with predictable execution paths. Agentic AI breaks this assumption.

AI agents:

- Persist over time
- Act autonomously within delegated authority
- Change behavior through updates, context, and integration
- Interact with other systems and agents
- Continue operating long after original design decisions were made

In many organizations, accountability weakens as agents move from development into operation. Ownership becomes unclear. Authorization boundaries drift. Auditability degrades. Decommissioning is rarely planned.

These failures are rarely caused by model error. They are governance failures.

USA-ADL™ was created to establish lifecycle authority where traditional governance stops.

What USA-ADL™ Is

USA-ADL™ defines governance logic for AI agents across seven lifecycle phases:

1. Strategy and Threat Modeling
2. Secure Architecture and Design
3. Data and Model Security
4. Secure Development and Prompt Hardening
5. Red Teaming and Adversarial Testing
6. Deployment and Runtime Governance
7. Decommissioning and End of Life Authority

Across these phases, USA-ADL™ introduces governance controls that treat AI agents as non human identities with defined ownership, authorization scope, and accountability over time.

The framework is tool agnostic, model agnostic, and vendor neutral.

What USA-ADL™ Is Not

USA-ADL™ is not a product.
USA-ADL™ is not a scanner.
USA-ADL™ is not a checklist of vulnerabilities.

It does not replace risk identification frameworks or technical safeguards. Instead, it provides the governing structure that allows those controls to remain effective as agents evolve in production.

Relationship to Industry Standards

USA-ADL™ is designed to complement existing standards and guidance.

Risk focused frameworks identify what can go wrong in agentic systems. Management standards describe desired outcomes. USA-ADL™ defines how authority, ownership, and accountability are maintained across the agent lifecycle.

This makes USA-ADL™ particularly suitable for organizations aligning with:

- AI risk management initiatives
- Information security management systems
- Regulatory oversight of autonomous systems

- Enterprise scale AI deployments

Why Lifecycle Governance Matters

Without lifecycle governance:

- Authorization expands silently
- Ownership becomes ambiguous
- Audit trails degrade
- Incidents become harder to investigate
- Agents outlive their original controls

With lifecycle governance:

- Authority is explicit
- Responsibility is continuous
- Changes are controlled
- Behavior remains observable
- Decommissioning is intentional

USA-ADL™ enables organizations to operate autonomous agents with confidence rather than assumption.

Publication and Use

USA-ADL™ is openly published for reference and adoption. The framework specification is public.

Implementation methodologies, assessment models, tooling, training programs, and certifications are governed and maintained separately by Uranusys.

Use of the USA-ADL™ name or branding to imply certification or endorsement requires explicit authorization.

Who This Is For

USA-ADL™ is designed for:

- Security and risk leaders

- AI and platform architects
- Governance and compliance teams
- Engineering leaders deploying autonomous systems
- Organizations preparing for regulatory scrutiny of AI autonomy