

SOL ENGINEERING SERVICES

The Daniel Group LLC

SUPPORTING DOCUMENTATION PACKAGE

ALICE: Adaptive Labor Intelligence and Capacity Engine

SBIR Topic DLA26BZ03-NV011 | Defense Logistics Agency

ARTIFACT 2 OF 11

Architecture Decision Record Register

Document Version 1.0 | July 2026

Status: DRAFT — Pending Gate-2 Stakeholder Review

Classification: Unclassified

Document Control

Field	Value
Artifact ID	SDP-ART-02
Artifact Title	Architecture Decision Record Register
Artifact 1 Dependency	SDP-ART-01 (Research Log) — Approved Gate-1
Package	ALICE Supporting Documentation Package (SDP)
Source Documents	ALICE TNA v1.0 (July 2026); ALICE PRD v1.0 (July 2026); ALICE CONOPS v2 (July 2026); SDP-ART-01 (Research Log)
Prepared By	SOL Engineering Services | The Daniel Group LLC
Prepared For	Defense Logistics Agency | SBIR DLA26BZ03-NV011
Version	1.0 (Draft)
Date	July 2026
Classification	Unclassified
Status	Draft — Pending Gate-2 stakeholder approval before Artifact 3 is produced
ADR Count	7 fully specified ADRs (ADR-001 through ADR-007) + 3 candidate ADRs for stakeholder determination
ID Convention	All IDs preserved from TNA v1.0. No IDs modified.

1. Purpose and Scope

This Architecture Decision Record (ADR) Register is Artifact 2 of the ALICE Supporting Documentation Package. It expands the seven ADRs documented in TNA Section 21 into fully specified governance records, adds three candidate ADRs identified during cross-artifact analysis that were not formally captured in the TNA, and establishes the register as the authoritative governance reference for all ALICE architecture decisions across Phase I and Phase II.

Each ADR in this register includes all fields required for governance-quality decision documentation: a unique stable identifier, title, current status, decision version history, the decision context (forces and constraints that drove the decision), the decision statement, full rationale, all alternatives considered with documented rejection reasons, consequences (both intended and unintended), risks, related architecture principles, related research log entries, related PRD requirements, related CONOPS activities, implementation guidance, and open questions specific to the decision.

The register is a living document. ADR status may change as Phase I execution resolves open questions. All status changes require stakeholder review and must be reflected in this register before dependent implementation decisions are finalized. Artifact IDs used in this register are consistent with TNA v1.0 and SDP-ART-01 and will be preserved across all subsequent SDP artifacts.

2. ADR Register Overview

2.1 Complete ADR Index

ADR ID	Title	Status	Phase	Category	Primary Architecture Principle	Key Risk
ADR-001	Modular Service-Oriented Architecture	Accepted	I and II	Structural	AP-5 (Modular Service Decomposition)	Architectural complexity overhead managed by interface contracts
ADR-002	Four-Zone Data Architecture	Accepted	I and II	Privacy and Data	AP-4 (Data Zone Separation)	Zone boundary enforcement must be technical, not reliant on application-layer controls
ADR-003	Integration Abstraction Layer for DCPDS/DCHRMS Migration	Accepted	I and II	Integration	AP-9 (Open Standards); AP-7 (Configuration Over Customization)	DCHRMS migration timeline for pilot sub-component is an unresolved open question
ADR-004	Identity-Centric Security with ZTA Principles	Accepted	I (principles); II (full)	Security	AP-3 (Defense-in-Depth with Identity-Centric Controls)	CAC integration feasibility in Phase I environment is an assumption pending confirmation
ADR-005	Immutable Audit Logging as a Dedicated Service	Accepted	I and II	Security and Compliance	AP-6 (Immutable Audit by Design)	Log retention storage provisioning must be planned for 7-year minimum at full enterprise scale
ADR-006	Phase-Gated Security Posture	Accepted	I and II	Security and Compliance	AP-10 (Phase-Gated Security Posture)	Phase I data handling must be confirmed as permissible by DLA legal counsel before execution
ADR-007	Synthetic Population Treated as CUI Pending Formal PIA	Accepted	I and II	Privacy and Compliance	AP-1 (Privacy as Architectural Invariant)	Synthetic data classification ambiguity may affect Phase III commercializatio

ADR ID	Title	Status	Phase	Category	Primary Architecture Principle	Key Risk
						n
ADR-008 (Candidate)	Simulation Engine Access Restricted to Synthetic Zone Only	Proposed — Awaiting Stakeholder Confirmation	I and II	Privacy and Security	AP-1; AP-4	If the simulation engine were to access de-identified records, Privacy Act risk would increase substantially
ADR-009 (Candidate)	Read-Only External Integration Enforced at Service Account Level	Proposed — Awaiting Stakeholder Confirmation	I and II	Integration and Security	AP-2 (Read-Only External Integration)	Service account misconfiguration could grant write access to source systems
ADR-010 (Candidate)	Configuration-Driven Operational Parameters (No Code Changes for Threshold Adjustments)	Proposed — Awaiting Stakeholder Confirmation	I and II	Maintainability	AP-7 (Configuration Over Customization)	Configuration data must itself be version-controlled and change-logged

2.2 ADR Status Definitions

Status	Definition	Action Required
Accepted	The decision has been made, rationale is documented, and the architecture reflects this decision. Implementation teams shall follow this decision.	No stakeholder action required unless a superseding decision is proposed.
Proposed	A candidate decision has been identified and documented but has not yet received explicit stakeholder acceptance. The architecture currently reflects this decision implicitly.	Stakeholder review and explicit acceptance or rejection required at Gate-2.
Deferred	The decision is recognized as necessary but cannot be made until one or more open questions are resolved.	Identify resolution path and responsible decision owner. Revisit at Phase I completion or earlier if the blocking question resolves.
Superseded	A prior decision has been replaced by a newer decision. The superseded record is retained for historical reference.	Document the superseding ADR ID in the Superseded By field.

3. ADR Detail Records

ADR-001: Modular Service-Oriented Architecture

ADR-001	Modular Service-Oriented Architecture
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Structural Architecture
Decision Owner	SOL Lead Architect
Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0 approval

Context

ALICE must be delivered in two phases with materially different scope and constraints. Phase I ($100,000, six months) produces a prototype covering organizational ontology, synthetic data generation, scenario simulation, and a decision support interface prototype. Phase II expands to enterprise scale, live system integration, and an Authority to Operate. A monolithic architecture that bundles all capabilities into a single deployable unit would make independent testing and validation of individual capabilities impractical, would prevent horizontal scaling of computationally intensive components (the simulation engine under concurrent user load), and would create a deployment dependency chain that could block Phase II scope expansion.

The CONOPS OV-1 defines five distinct architectural layers (Data Ingestion, Organizational Ontology, Synthetic Workforce Engine, Agent-Based Simulation, Decision Support) that have different performance, security, and scalability characteristics. These layers map naturally to independent service components.

Decision

Decompose ALICE into twelve logical service components organized in six service layers. Each component operates within a defined service boundary and communicates with other components exclusively through documented interface contracts. No component accesses the internal state of another component directly. The twelve components are: LC-01 (Interface Control), LC-02 (Data Ingestion and De-Identification), LC-03 (Schema Abstraction), LC-04 (Organizational Ontology), LC-05 (Synthetic Workforce Engine), LC-06 (ABM Simulation Engine), LC-07 (Decision Support Interface), LC-08 (IAM), LC-09 (Immutable Audit Logging), LC-10 (Configuration Management), LC-11 (Observability), and LC-12 (Secrets Management).

Rationale

- Independent testing: Phase I requires evidence of functional capability for each OV-1 layer separately. A modular architecture enables each layer to be tested in isolation, producing discrete deliverable artifacts for Phase I acceptance.
- Independent scaling: The ABM simulation engine (LC-06) is computationally intensive and subject to concurrent user load. It must scale horizontally without requiring co-scaling of data ingestion or decision support components.
- Blast radius containment: A failure in one service component must not cascade to other components. Service boundaries enforce this by design.
- Phase-gated delivery: Phase I can deliver a subset of components (ontology, synthetic engine, simulation prototype, decision support prototype) while deferring enterprise-grade components (live integration, full observability) to Phase II.

- Federal IT program risk reduction: Monolithic architectures in federal programs have produced well-documented failure modes including inability to test incrementally, inability to scale components independently, and difficulty obtaining ATOs due to undefined system boundaries.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Monolithic architecture: all capabilities in a single deployable unit.	Simpler initial deployment; fewer interface contracts to maintain.	Single point of failure prevents independent OV-1 layer validation. Cannot scale computationally intensive simulation service without co-scaling all services. ATO boundary becomes the entire monolith, increasing assessment scope.
Two-tier architecture: presentation layer plus data/processing backend.	Familiar pattern; reduces number of service boundaries.	Insufficient separation of simulation concerns from data governance concerns. Synthetic data engine and simulation engine would share a service boundary, violating the architectural requirement that the simulation engine have no access to de-identified records (see ADR-008).
Microservices at maximum granularity (20+ services).	Maximum scalability and deployment flexibility.	Exceeds Phase I operational complexity budget. Configuration management, service discovery, and inter-service authentication overhead is disproportionate to Phase I scale. Twelve components provide sufficient decomposition without microservices operational burden.

Consequences

- Each of the twelve service components requires a documented interface contract. Interface contracts must be maintained as versioned artifacts throughout Phase I and Phase II.
- Service-to-service authentication is required. Every inter-service call must present a service identity credential. This drives the design of LC-12 (Secrets Management) as a cross-cutting service.
- Integration testing must validate interface contracts between adjacent service layers, not just end-to-end user workflows.
- Phase II can deploy additional service components (enterprise integration adapters, additional scenario types) without redesigning the core six-layer architecture.

Implementation Guidance

Interface contracts shall be defined before implementation of any component begins. Each contract shall specify: (1) the operations exposed; (2) the data exchanged per operation; (3) the authentication mechanism; (4) error response formats; and (5) the version identifier. Contracts are versioned; breaking changes require a new contract version.

The service decomposition does not prescribe a specific deployment topology. Components may be deployed as separate process boundaries, containers, or logical modules within a process, provided that the interface contract discipline is maintained regardless of deployment topology. Technology selection for service runtime environments is deferred to Phase I implementation planning.

Risks

- **RISK-005 (Simulation execution time):** Horizontal scaling of LC-06 is necessary to meet the 5-minute preset scenario threshold under 20 concurrent users. Modular architecture enables

this scaling; it does not guarantee it without appropriate design of the simulation execution environment.

- **Inferred risk:** Interface contract drift — if interface contracts are not maintained as living documents, implementations may diverge from contracts over time. Version control of interface contracts is a required process artifact, not an optional best practice.

Cross-References

Reference Type	IDs
Architecture Principles	AP-5 (Modular Service Decomposition); AP-8 (Traceability)
Logical Components	LC-01 through LC-12
Functional Domains	FD-1 through FD-8
PRD Requirements	REQ-NF-006 (scalability); REQ-NF-003 (concurrent users); REQ-NF-002 (scenario execution time)
Research Log	RL-11 (NIST 800-160 SSE — security embedded in component design); RL-28 (ISO 42010 architecture views)
Related Risks	RISK-005
Related ADRs	ADR-004 (service-to-service authentication requirement); ADR-005 (dedicated audit service as one of twelve components)
Open Questions	None specific to this ADR. Hosting model (OQ-004) affects deployment topology but not the logical decomposition.

ADR-002: Four-Zone Data Architecture

ADR-002	**Four-Zone Data Architecture**
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Privacy and Data Architecture
Decision Owner	SOL Lead Architect and Privacy Officer
Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0

Context

ALICE processes data from four authoritative DLA source systems. This data originates as Privacy Act-protected workforce records. It passes through two transformation stages before reaching any analytical function: de-identification (which removes or tokenizes individual identifiers) and synthetic generation (which produces a statistically faithful population with differential privacy guarantees). The resulting analytical outputs — readiness indicators, scenario results — contain no individual-identifiable data.

Four materially different sensitivity states exist within ALICE: (1) raw source data that may contain Privacy Act records; (2) de-identified records that are CUI but are not Privacy Act records; (3) synthetic population records that are not Privacy Act records but remain CUI pending a formal Privacy Impact Assessment; and (4) analytical outputs that are CUI and contain no individual data. A single-zone or two-zone data model cannot enforce the distinct access control rules required for each state. Application-layer access control — the common alternative — cannot be relied upon for Privacy Act compliance because it depends on application correctness rather than architectural enforcement.

Decision

Implement four logical data zones with enforced boundaries, distinct storage domains, distinct access credentials, distinct encryption keys, and distinct audit logging:

- **Zone 1 — Raw Zone:** Transient. Source system data in native format, potentially containing Privacy Act-protected records. Data exists in this zone only in-flight during ingestion processing. Raw records are never written to persistent storage within ALICE.
- **Zone 2 — Processed Zone:** Persistent. De-identified records with all individual-identifying fields removed or tokenized. CUI. Accessible by: LC-02 (Ingestion Service, write); LC-04 (Ontology Service, read); LC-05 (Synthetic Engine, read). Not accessible by LC-06 (Simulation Engine) or LC-07 (Decision Support Interface).
- **Zone 3 — Synthetic Zone:** Persistent. Validated synthetic population. Classification: CUI by default pending formal PIA (see ADR-007). Accessible by: LC-05 (Synthetic Engine, write during generation); LC-06 (Simulation Engine, read-only); Data Steward (read-only for export authorization). Not accessible by LC-04 or LC-07 directly.
- **Zone 4 — Analytical Zone:** Persistent. Readiness indicators, scenario run records, simulation outputs, historical state snapshots. CUI. Accessible by: LC-04 (Ontology Service, write); LC-06 (Simulation Engine, write); LC-07 (Decision Support Interface, read). Contains no individual-identifiable data.

No data may cross a zone boundary except through a defined, logged transformation pipeline. Each zone boundary crossing is an auditable event.

Rationale

- Privacy Act compliance: 5 U.S.C. § 552a prohibits unauthorized processing of records about individuals. Architectural enforcement at the zone boundary is the only mechanism that provides a technically verifiable guarantee. Application-layer filtering can be bypassed by code defects; zone boundaries enforced by distinct access credentials cannot.
- NIST SP 800-160 SSE: security must be embedded in architecture from inception, not applied as a post-design overlay. Zone separation is the primary architectural mechanism for embedding Privacy Act compliance.
- Differential evidence: zone boundaries create discrete, verifiable evidence points for RMF assessment, CMMC audit, and Privacy Act compliance review. An assessor can verify that the simulation engine's service account does not have credentials to access the Processed Zone.
- Synthetic data governance: the four-zone model provides the structural basis for the export governance workflow — exports can only be authorized from the Synthetic Zone, and the Steward authorization gate is enforced at the zone delivery boundary.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Single database with row-level security.	Simpler operational model; well-understood pattern; reduces storage provisioning complexity.	Application-layer control (row-level security) is insufficient for Privacy Act compliance. A code defect or privilege escalation attack could expose Privacy Act records to the simulation or analytical

Alternative	Why Considered	Why Rejected
		environment. Auditors cannot independently verify protection without reviewing application code.
Two-zone model: Identifiable vs. De-identified.	Simpler than four zones; adequate for de-identification boundary enforcement.	Does not address the distinct governance requirements of the synthetic population (CUI pending PIA, distinct export governance) or the distinct read access patterns of the analytical zone. The synthetic population requires its own zone to enforce Data Steward governance of exports.
Three-zone model: Raw, Processed, Analytical (no separate Synthetic Zone).	Avoids the overhead of managing a fourth zone.	The synthetic population has distinct access controls (simulation engine reads it; decision support interface does not), a distinct generation governance workflow, distinct export governance, and a distinct classification status (CUI pending PIA vs. confirmed CUI for other zones). Merging synthetic and analytical zones would conflate these governance requirements.

Consequences

- Four distinct storage provisioning domains are required, each with separate access credentials and encryption keys managed by LC-12 (Secrets Management).
- Every zone boundary crossing is an auditable event. LC-09 (Audit Logging) must record all zone boundary transitions with: source zone, destination zone, service identity, data batch identifier, timestamp, and validation results.
- Phase I must implement the full four-zone model even if only a subset of zones contains data. The zone architecture is a foundational design element, not a Phase II addition.
- Data retention policies differ by zone: Raw Zone records are never persisted; Processed, Synthetic, and Analytical Zone records have distinct retention schedules documented in TNA Section 12.3.

Implementation Guidance

Zone boundaries shall be implemented as logically separate storage domains with distinct access credential sets. The implementation technology for each zone (relational store, object storage, graph store, or other capability class) is deferred to Phase I technical design and is not prescribed by this decision. What is prescribed is that each zone uses distinct credentials, distinct encryption keys, and contributes distinct audit events.

The PII detection capability integrated into LC-02 (Data Ingestion Service) is the enforcement mechanism for the Raw-to-Processed zone boundary. If PII detection detects a residual identifying field, ingestion is aborted and the record is not written to the Processed Zone. This detection capability must be tested as part of Phase I acceptance validation.

Risks

- **RISK-007 (Membership inference attack):** Even with zone separation, if the synthetic population in Zone 3 has insufficient differential privacy noise (epsilon > 1.0), zone separation alone does not protect against re-identification. ADR-007 and the synthetic validation protocol (LC-05) are the complementary mitigations.
- **Inferred risk — credential management complexity:** Managing four distinct access credential sets increases the operational complexity of LC-12 (Secrets Management).

Credential rotation must occur across all four zones on schedule. This risk is mitigated by the dedicated secrets management service design.

Cross-References

Reference Type	IDs
Architecture Principles	AP-1 (Privacy as Architectural Invariant); AP-4 (Data Zone Separation); AP-6 (Immutable Audit by Design)
Logical Components	LC-02 (ingestion boundary enforcement); LC-04 (Processed Zone reader); LC-05 (Synthetic Zone writer/reader); LC-06 (Synthetic Zone reader); LC-09 (zone boundary audit); LC-12 (per-zone credentials)
Data Entities	All 12 data entities in TNA Section 12.2 — each mapped to a specific zone
Information Flows	IF-06 (Raw to Processed); IF-07 (Processed to Ontology); IF-08 (Processed to Synthetic Engine); IF-09 (Synthetic to Synthetic Zone); IF-10 (Synthetic Zone to Simulation)
PRD Requirements	REQ-NF-008 (no Privacy Act record in analytical/simulation layers); REQ-COMP-001 (Privacy Act compliance); REQ-F-006 (PII detection at ingestion); CON-003 (Privacy Act constraint)
Research Log	RL-20 (Privacy Act statute); RL-21 (5 C.F.R. Part 293); RL-11 (NIST 800-160 SSE)
Related ADRs	ADR-007 (CUI classification of Synthetic Zone); ADR-008 (Candidate — simulation engine restricted to Synthetic Zone); ADR-005 (zone boundary crossing events logged to dedicated audit service)
Open Questions	OQ-006 (NIST SP 800-188 finalization — may affect Synthetic Zone classification after PIA)

ADR-003: Integration Abstraction Layer for DCPDS/DCHRMS Migration

ADR-003	**Integration Abstraction Layer for DCPDS/DCHRMS Migration**
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Integration Architecture
Decision Owner	SOL Lead Architect
Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0

Context

DCPDS — the Defense Civilian Personnel Data System and authoritative source for DLA position and person-position data — is in phased migration to DCHRMS (Defense Civilian Human Resources Management System), an Oracle HCM-based successor. The migration is a DoD-wide program

proceeding on a rolling schedule by component. The migration timeline for any specific DLA activity is not fixed in publicly available documentation.

If ALICE integrates directly with the DCPDS schema, and the pilot sub-component's DCHRMS migration occurs during Phase I or Phase II, the entire data ingestion service (LC-02) would require redesign to accommodate the new schema. This would represent a major unplanned scope change with budget and schedule consequences. The risk is rated HIGH in RISK-002 because the impact of mid-program schema change is severe even if the likelihood is currently assessed as low.

Decision

Implement LC-03 (Schema Abstraction and Integration Adapter Service) as an independent, dedicated service component that sits between source systems and the ALICE data ingestion pipeline. LC-03 maintains versioned schema mapping definitions for both DCPDS and DCHRMS data structures and translates all source system records into the ALICE canonical data model before handing them to LC-02 for de-identification processing.

The ALICE canonical data model is the stable internal schema that all downstream ALICE components depend upon. LC-03 is the sole component responsible for translating from external source schemas to this canonical model. All other ALICE components depend on the canonical model, not on any external source schema.

Rationale

- Risk isolation: the schema migration impact is isolated to LC-03 alone. When DCHRMS migration occurs, only the LC-03 schema mapping configuration for DCHRMS requires updating. LC-02, LC-04, LC-05, and all downstream components continue to operate on the unchanged canonical data model.
- Dual-schema support: the DCHRMS migration is phased. During the transition, some DLA components may provide DCPDS-format data while others provide DCHRMS-format data. LC-03 supports both schemas simultaneously without requiring different ingestion pipelines for each format.
- Schema change alerting: LC-03 can detect schema deviations in source data and alert administrators before attempting to process records against a stale mapping. This provides early warning of schema changes rather than silent data corruption.
- Technology neutrality: the abstraction layer allows ALICE to support future source system changes (additional HR systems, third-party integrations) without redesigning ingestion or downstream components.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Direct DCPDS integration without abstraction, with planned redesign for DCHRMS.	Simpler initial implementation; one fewer service component.	Migration timing is outside SOL's control. If DCHRMS rollout reaches the pilot sub-component during Phase I or Phase II, the redesign would be an unplanned scope change within a fixed-price contract. The risk-to-mitigation cost ratio strongly favors the abstraction layer.
Parallel ingestion pipelines: separate pipeline for DCPDS and separate pipeline for DCHRMS.	Each pipeline can be optimized for its source schema independently.	Duplicates all ingestion logic. Doubles the maintenance burden. Requires two sets of tests, two sets of credentials, and two de-identification implementations. When new source systems are added in Phase II or III, the parallel pipeline count grows unboundedly.
Schema auto-	Eliminates the need for pre-	Requires machine-readable schema

Alternative	Why Considered	Why Rejected
detection with dynamic mapping.	defined schema mappings; adapts to any source schema automatically.	discovery capability that may not be available from DoD systems. Risk of silent schema mismatches producing incorrectly mapped records without detection. Pre-defined and tested mappings provide higher reliability for a compliance-sensitive system.

Consequences

- LC-03 is an additional service component beyond the five OV-1 layers. It is a cross-cutting integration capability, consistent with the modular decomposition in ADR-001.
- Schema mapping configuration in LC-03 must be version-controlled. Changes to schema mappings must be treated as system configuration changes, logged per REQ-ADM-006, and tested before deployment.
- The ALICE canonical data model must be formally specified and documented as an independent artifact. This specification is a Phase I deliverable that must be completed before LC-02 implementation begins.
- If the DCHRMS migration reaches the pilot sub-component before Phase I is complete, LC-03 DCHRMS schema mappings must be developed and tested within Phase I. This is a schedule risk (RISK-002) mitigated by the abstraction design.

Implementation Guidance

The ALICE canonical data model specification must be produced as a Phase I Month 1 deliverable. The specification must include: entity names, field names and types, nullability rules, and the mapping from each DCPDS and DCHRMS source field to the corresponding canonical field. The specification is the authoritative reference for both LC-03 (schema translation) and LC-02 (de-identification field definitions).

LC-03 schema mappings must be stored as configuration data in LC-10 (Configuration Management), not hardcoded in LC-03 application logic. This enables schema mapping updates without code deployment.

Risks

- **RISK-002 (DCHRMS migration timing):** The primary risk this ADR mitigates. Residual risk: if DCHRMS migration for the pilot sub-component begins within Phase I Month 1 before LC-03 DCHRMS mappings are ready, a brief integration outage may occur. Mitigation: confirm DCHRMS migration schedule at Phase I kickoff (OQ-004).
- **OQ-008 (DLA TMS extract format):** LC-03 abstraction must also accommodate DLA TMS certification data schema. If the DLA TMS extract format is non-standard or requires negotiation, LC-03 mapping development scope increases.

Cross-References

Reference Type	IDs
Architecture Principles	AP-9 (Open Standards for Data Exchange); AP-7 (Configuration Over Customization); AP-5 (Modular Service Decomposition)
Logical Components	LC-03 (Schema Abstraction Service); LC-02 (Data Ingestion — consumer of LC-03 output); LC-10 (schema mapping configuration storage)
Integrations	INT-01 (DCPDS/DCHRMS); INT-02 (DLA TMS)
PRD Requirements	REQ-INT-001 (integration abstraction for DCPDS/DCHRMS transition); REQ-NF-014 (standard data exchange formats)

Reference Type	IDs
Research Log	RL-05 (DCHRMS migration finding — drives this decision); RL-04 (DCPDS latency — context for integration design)
Related ADRs	ADR-001 (LC-03 is one of the twelve modular components)
Related Risks	RISK-002 (DCHRMS migration timing)
Open Questions	OQ-004 (DCHRMS migration schedule for pilot sub-component); OQ-008 (DLA TMS extract format)
Assumptions	ASM-004 (DCHRMS migration does not reach pilot sub-component during Phase I)

ADR-004: Identity-Centric Security with ZTA Principles

ADR-004	**Identity-Centric Security with ZTA Principles**
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Security Architecture
Decision Owner	SOL Security Lead
Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0

Context

ALICE handles Controlled Unclassified Information. Its user population spans enterprise-level DLA leadership and activity directors who may access the system from multiple locations and device types across DLA's global footprint. Traditional perimeter-based security — trusting any request that originates from an approved network segment — is insufficient for a CUI system accessed across distributed environments, particularly in the context of OMB M-22-09, which directs federal agencies to move toward Zero Trust Architecture.

NIST SP 800-207 defines ZTA through three logical components: the Policy Engine (evaluates access requests against policy), the Policy Administrator (manages policy and communicates decisions to enforcement points), and the Policy Enforcement Point (allows or denies access at resource boundaries). The fundamental principle is that no implicit trust is granted based on network location: every access request is authenticated, authorized, and validated before the resource is accessed.

Decision

Apply Zero Trust Architecture principles from NIST SP 800-207 throughout the ALICE security architecture. The implementation applies these principles through the following mechanisms:

- **Every user session:** authenticated at LC-01 (Interface Control and API Management) via DoD CAC, producing a role-scoped session token. The token explicitly encodes the user's role and activity scope. No session may persist without active token validation.

- **Every service request:** service-to-service calls use authenticated service identity credentials managed by LC-12 (Secrets Management). No service grants access to another service based solely on network origin.
- **Every data zone access:** zone-specific access credentials are required. A service's credentials for the Synthetic Zone are distinct from its credentials for the Processed Zone. Credentials are scoped to the minimum access required per service.
- **Every administrative action:** System Administrator actions require authenticated identity and are logged to LC-09 (Audit Logging) with individual attribution. No shared administrative credentials.

Full ZTA deployment (device trust, continuous behavioral validation, micro-segmentation) is targeted for Phase II. Phase I applies ZTA principles through the mechanisms above within the available implementation budget.

Rationale

- OMB M-22-09 (January 2022) directs federal agencies to adopt ZTA. Designing ALICE consistent with ZTA from inception is both a current policy expectation and the direction of DoD security architecture.
- NIST SP 800-207 is the authoritative reference for ZTA implementation. Its three-component model (Policy Engine, Policy Administrator, Policy Enforcement Point) maps directly to the LC-01/LC-08 identity control architecture.
- ALICE's user population is distributed across DLA's global footprint. A perimeter trust model would require network-level access control to the ALICE system boundary for every authorized user location, which is operationally complex and inconsistent with DLA's cloud strategy.
- CUI handling requires defense-in-depth. ZTA principles add an identity layer of defense beyond the network boundary, consistent with NIST SP 800-160 SSE requirements.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Network-perimeter-only security: trust all requests from approved IP ranges or network segments.	Simpler to implement; familiar operational model.	Insufficient for CUI under DoD IL4. Inconsistent with OMB M-22-09 and DoD ZTA direction. Provides no protection against compromised internal hosts or lateral movement after a perimeter breach. Does not support DLA's cloud-based access model.
Full ZTA with device trust and continuous behavioral validation in Phase I.	Provides strongest available security posture.	Exceeds Phase I budget and schedule constraints. Full ZTA deployment requires device management infrastructure that is outside the Phase I scope boundary. Principles are applied in Phase I; full deployment is a Phase II target.
Username and password authentication (no CAC).	Simpler user provisioning; no dependency on DoD PKI infrastructure.	Explicitly prohibited by DoDI 8520.02 for DoD systems. Inconsistent with DoD IL4 requirements. Eliminates the non-repudiation property required for RMF compliance.

Consequences

- LC-01 (Interface Control) must integrate with the DoD PKI infrastructure for CAC certificate validation. This integration must be designed and tested before Phase I demonstration.

- Role-scoped access tokens must encode: user identity, assigned role, assigned activity scope (for Activity Directors), and session expiration time. Token format must be defined as part of the interface contract for LC-01 and LC-08.
- Service identity credentials (for service-to-service calls) must be provisioned by LC-12 and rotated on a defined schedule. Rotation schedule must be defined in the System Security Plan.
- Phase II must complete the ZTA implementation with device trust evaluation, continuous session validation, and micro-segmentation as resources permit.

Risks

- **ASM-008 / RISK-004 (CAC integration in Phase I):** CAC integration in the Phase I demonstration environment is an assumption. If the DoD PKI trust anchor is not accessible in SOL's Phase I environment, an alternative MFA mechanism must be implemented as an interim measure, with CAC planned for Phase II. This must be confirmed at Phase I kickoff.
- **Inferred risk — token scope creep:** If activity scope is not precisely encoded in access tokens, Activity Directors may be able to query data for activities outside their authorization. Scope encoding must be defined in the token specification and tested with negative test cases.

Cross-References

Reference Type	IDs
Architecture Principles	AP-3 (Defense-in-Depth with Identity-Centric Controls); AP-6 (Immutable Audit by Design)
Logical Components	LC-01 (Interface Control — authentication enforcement point); LC-08 (IAM — identity and token management); LC-12 (service credential management)
PRD Requirements	REQ-SEC-001 (MFA/CAC); REQ-SEC-002 (RBAC); REQ-SEC-003 (least privilege); REQ-SEC-004 (activity-scoped data access); REQ-SEC-005 (admin separation of duties)
Research Log	RL-10 (NIST 800-207 ZTA); RL-11 (NIST 800-160 SSE); RL-14 (DoDI 8520.02 CAC); RL-17 (OMB M-22-09 ZTA mandate)
NIST Controls	IA-2 (Identification and Authentication); AC-2 (Account Management); AC-3 (Access Enforcement); AC-6 (Least Privilege); AC-12 (Session Termination)
Related ADRs	ADR-001 (service boundaries enable per-service enforcement points); ADR-005 (audit logging captures all authentication and authorization events)
Related Risks	RISK-004 (CAC integration feasibility in Phase I)
Assumptions	ASM-008 (DoD PKI accessible in Phase I environment)
Open Questions	OQ-003 (Phase II ATO level — IL4 vs IL5 — may affect identity requirements); OQ-004 (hosting model affects PKI integration approach)

ADR-005: Immutable Audit Logging as a Dedicated Service

ADR-005	**Immutable Audit Logging as a Dedicated Service**

Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Security and Compliance Architecture
Decision Owner	SOL Security Lead
Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0

Context

ALICE must maintain immutable audit logs for all user access events, data refresh events, scenario runs, and synthetic data exports for a minimum of seven years (REQ-NF-012; REQ-SEC-010). NIST SP 800-53 AU-9 requires that audit logs be protected against unauthorized access, modification, and deletion. The RMF authorization process requires audit evidence that is independently verifiable.

If audit logging is implemented within each service component (each service writes its own logs to its own storage), two failure modes are introduced: (1) a misconfigured or compromised service could modify or delete its own logs, eliminating the tamper-evidence property required by AU-9; and (2) log coverage becomes dependent on the correctness of each service's logging implementation, which is an application-layer dependency rather than an architectural guarantee.

Decision

Implement LC-09 (Immutable Audit Logging Service) as a dedicated, independent service. All ALICE service components write audit events to LC-09 via a write-only interface. LC-09 is the sole service with write access to the audit log store. The audit log store is append-only: no record can be modified or deleted after it is written. No role — including System Administrator — has delete capability on audit records. The integrity of the audit log is verifiable through a hash-chain or equivalent mechanism that detects any tampering attempt.

Rationale

- NIST SP 800-53 AU-9 (Protection of Audit Information) requires that audit records be protected against unauthorized access and modification. Architectural separation of the audit service from service operation is the technically enforceable mechanism for this requirement.
- An independent audit service is technically verifiable: an assessor can confirm that no service other than LC-09 has write credentials to the audit log store, without reviewing application code. In-service logging provides no equivalent verification mechanism.
- Separation of the audit service from service operation means that a compromised service cannot selectively suppress its own audit trail. The audit trail remains intact even if a service is compromised.
- Seven-year retention at enterprise scale requires dedicated storage planning and automated retention policy enforcement. A dedicated service can manage retention lifecycle for all log records centrally rather than requiring each service to manage its own log lifecycle.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Each service maintains its own audit log in its own storage.	Simpler architecture; each service is self-contained.	A compromised or misconfigured service can modify or delete its own logs. Log coverage depends on application correctness. Cannot provide AU-9-compliant tamper evidence without reviewing every service's logging implementation. Fragmented log storage

Alternative	Why Considered	Why Rejected
		complicates audit queries and compliance reporting.
Centralized log aggregation without immutability (standard logging infrastructure).	Standard operational logging pattern; widely available capability.	Aggregation without write-once protection does not satisfy NIST AU-9. An administrator with write access to the aggregated log store could modify records. Does not provide independently verifiable tamper evidence.
Blockchain-based immutable log.	Mathematically strong immutability guarantee.	Operationally complex for a system of ALICE's scale. Adds a novel technology dependency that may create supply chain risk and maintenance burden. A hash-chain mechanism over an append-only store provides equivalent functional immutability with significantly lower operational complexity.

Consequences

- Every ALICE service component must implement log event generation and delivery to LC-09. Event generation must be designed into each component, not added after implementation.
- The write-only interface from each service to LC-09 uses service identity credentials (managed by LC-12). These credentials must be distinct from read credentials issued to authorized log reviewers.
- System Administrators receive read-only access to audit logs for compliance review. They do not have delete capability. This separation must be enforced by the LC-09 access control design.
- Log integrity verification (hash-chain or equivalent) must be designed into LC-09 from the outset. Retrofitting integrity verification to an existing log store is operationally disruptive.
- Storage provisioning for seven-year retention at enterprise scale (30,000 positions, 20+ concurrent users, daily refresh cycles) must be estimated during Phase I technical design. This is an implementation planning task, not a design constraint on the architectural decision.

Minimum Audit Event Catalog

The following event types must be captured by LC-09 regardless of which service generates the event:

Event Category	Events to Log	Minimum Fields
User Access	Login, logout, failed login attempt, session expiration	Entry ID; timestamp (UTC); user identity (CAC-derived); event type; source IP; outcome
Data Access	Read access to readiness indicators, position risk register, scenario results, historical views	Entry ID; timestamp; user identity; data scope accessed; service origin
Scenario Execution	Scenario initiation, completion, cancellation, timeout	Entry ID; timestamp; user identity; scenario type; parameters (excluding simulation model internals); duration; output reference
Synthetic Data	Generation initiation,	Entry ID; timestamp; actor identity; event

Event Category	Events to Log	Minimum Fields
Events	validation pass/fail, population promotion, export request, export authorization, export delivery	type; synthetic population version; epsilon value (for generation events); export file hash (for export events)
Administrative Actions	User provisioning/de-provisioning, role assignment change, threshold configuration change, data source connection change, export workflow configuration change	Entry ID; timestamp; administrator identity; action type; prior value; new value
System Events	Ingestion job start/completion/failure, schema validation failure, PII detection event, service health alert, authentication failure threshold breach	Entry ID; timestamp; service identity; event type; record counts (for ingestion events); error details
Data Zone Transitions	Every crossing of a zone boundary (Raw-to-Processed, Processed-to-Synthetic, Synthetic Zone export)	Entry ID; timestamp; service identity; source zone; destination zone; record count or file identifier; validation outcome

Risks

- **Inferred risk — event volume at enterprise scale:** At full enterprise scale (30,000 positions, daily refresh, 20+ concurrent users), audit event volume may be substantial. Storage provisioning and LC-09 throughput capacity must be evaluated during Phase I technical design.
- **Inferred risk — log delivery failure:** If a service cannot reach LC-09, events may be lost. LC-09 must provide a durable event queue that accepts events asynchronously and provides delivery confirmation. Services must log a delivery failure locally if LC-09 is unreachable, and deliver buffered events when connectivity is restored.

Cross-References

Reference Type	IDs
Architecture Principles	AP-6 (Immutable Audit by Design); AP-8 (Traceability)
Logical Components	LC-09 (Immutable Audit Logging Service); all other components (event sources)
PRD Requirements	REQ-NF-012 (7-year retention); REQ-SEC-009 (immutable audit log); REQ-SEC-010 (audit log retention); REQ-F-018 (synthetic export audit log); REQ-ADM-006 (configuration change log)
NIST Controls	AU-2 (Event Logging); AU-3 (Content of Audit Records); AU-9 (Protection of Audit Information); AU-11 (Audit Record Retention)
Research Log	RL-08 (NIST 800-53 AU control family); RL-11 (NIST 800-160 SSE embedded security)
Related ADRs	ADR-001 (dedicated audit service is LC-09, one of twelve modular components); ADR-002 (zone boundary crossings are audit events); ADR-004 (all authentication and authorization events captured)
Functional Domains	FD-8 (Governance, Audit, and Compliance)

Reference Type	IDs
Open Questions	None specific to this ADR. Log storage provisioning is a Phase I implementation planning task.

ADR-006: Phase-Gated Security Posture

ADR-006	**Phase-Gated Security Posture**
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Security and Compliance Architecture
Decision Owner	SOL Program Manager and Security Lead; requires DLA GTPOC and DLA Legal confirmation
Stakeholder Acceptance	Pending explicit Gate-2 confirmation AND confirmation from DLA GTPOC and DLA Legal (ASM-007)

Context

Phase I is bounded to $100,000 and six months. Obtaining a DoD RMF Authority to Operate is a multi-step process (Prepare, Categorize, Select, Implement, Assess, Authorize) that typically requires more than six months and significantly more than $100,000 in assessment costs for a new system. A Phase I ATO requirement would make the SBIR program technically undeliverable within its stated constraints.

However, ALICE is designed to handle Controlled Unclassified Information. The question is whether Phase I — operating exclusively on de-identified DCPDS extracts or publicly available OPM FWD proxy data — triggers the same ATO requirement that applies to a system processing live CUI in production.

Decision

Phase I operates on de-identified data extracts or OPM FWD proxy data with strong access controls but without a formal RMF Authority to Operate. Phase I design, implementation, and documentation must build actively toward ATO: the System Security Plan must be initiated in Phase I, core security controls must be implemented (authentication, encryption, audit logging), and all RMF evidence must be collected. The system is not operationally gated by ATO during Phase I. An ATO at DoD Impact Level 4 is a mandatory gate condition before Phase II processes any live DLA workforce data.

The boundary condition is clear: the moment live Privacy Act-protected DLA workforce data enters the ALICE environment, ATO is required. Phase I data use agreements must ensure that any data delivered by DLA has been de-identified before delivery.

Rationale

- PRD REQ-NF-009 and CON-003 explicitly state that ATO is required before Phase II operational data processing. Neither the PRD nor the CONOPS states that ATO is required for Phase I.

- Operating on de-identified or publicly available proxy data within a controlled development and demonstration environment is consistent with federal guidance for research and feasibility work, provided that appropriate security controls are in place.
- Requiring full ATO in Phase I would eliminate ALICE's eligibility as a Phase I SBIR program. The SBIR program structure is expressly designed for feasibility demonstration, not operational deployment.
- The phase-gated posture also serves a practical security benefit: the security documentation produced during Phase I (SSP, control implementation evidence, architectural security documentation) is directly reusable in Phase II RMF assessment, reducing Phase II ATO timeline.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Require full Phase II-equivalent ATO before Phase I begins.	Maximum security posture; no ambiguity in data handling requirements.	Makes Phase I undeliverable within $100,000 and six months. RMF assessment for a new system typically requires months of assessment time and costs that exceed the Phase I funding ceiling. This alternative is inconsistent with the SBIR program structure.
No security controls in Phase I; defer all security implementation to Phase II.	Maximizes Phase I development velocity within budget.	Unacceptable risk. Even de-identified data contains sensitive workforce population information. CUI controls are appropriate even for de-identified extracts at DoD IL4. More critically, security controls implemented in Phase I create the evidence base for Phase II ATO; deferring implementation entirely would delay Phase II ATO by the time required to implement all 287 Moderate baseline controls from scratch.
Operate Phase I on purely synthetic or fictitious data only, with no real workforce records at all.	Eliminates any Privacy Act or CUI handling concern in Phase I.	Prevents Phase I from demonstrating ontology fidelity against real organizational structure. The Phase I deliverable requires demonstrating that the system can accurately represent a real DLA organizational structure. A purely fictitious dataset cannot support this demonstration.

Consequences

- Every Phase I data use agreement with DLA must explicitly confirm that data delivered is de-identified before delivery. SOL must not receive raw DCPDS extracts containing Privacy Act records.
- The System Security Plan must be initiated in Phase I Month 1. The SSP is a Phase I deliverable as well as a Phase II ATO prerequisite.
- Phase I must implement at minimum the following security controls, regardless of ATO requirement: MFA (with CAC or interim equivalent); role-based access control; encryption in transit and at rest; audit logging per LC-09; and secrets management per LC-12. These controls are architectural requirements, not ATO-gated requirements.
- Phase II must not ingest live workforce data before ATO issuance by the designated DLA Authorizing Official.

Risks

- **ASM-007 / Critical risk:** DLA legal counsel must confirm that Phase I operation on de-identified data in SOL's environment does not trigger CUI ATO requirements under DoD policy. This confirmation is required before Phase I begins. If DLA's legal interpretation is more restrictive, the Phase I plan requires revision.
- **Inferred risk — Phase I data access scope creep:** If DLA provides broader data access than de-identified extracts (e.g., raw DCPDS exports), Phase I would inadvertently receive Privacy Act-protected data without ATO authority. Data use agreements must explicitly define the scope and format of data delivery.

Cross-References

Reference Type	IDs
Architecture Principles	AP-10 (Phase-Gated Security Posture)
PRD Requirements	REQ-NF-009 (ATO before Phase II data processing); CON-003 (CUI impact level); CON-005 (CMMC Level 2)
Research Log	RL-12 (NIST 800-37 RMF process); RL-13 (DoDI 8510.01 ATO requirement)
NIST Controls	All 287 Moderate baseline controls — implementation phased across Phase I (core) and Phase II (full baseline)
Related ADRs	ADR-007 (synthetic population CUI handling — also applies a conservative security posture); ADR-004 (core security controls implemented regardless of ATO status)
Assumptions	ASM-007 (Phase I on de-identified data does not trigger ATO requirement — must be confirmed by DLA Legal)
Related Risks	RISK-001 (data sharing agreement delays); RISK-004 (CAC availability in Phase I)
Open Questions	OQ-003 (Phase II ATO level — IL4 vs IL5); OQ-004 (Phase II hosting model affects ATO pathway); OQ-005 (DLA AO designation and control inheritance availability)
Recommended Next Steps	N-01 (distribute TNA for stakeholder review); N-09 (develop SSP draft in Phase I); N-18 (archive Phase I RMF evidence for Phase II ATO)

ADR-007: Synthetic Population Treated as CUI Pending Formal PIA

ADR-007	**Synthetic Population Treated as CUI Pending Formal PIA**
Status	Accepted
Version	1.0 — July 2026 (initial)
Category	Privacy and Compliance Architecture
Decision Owner	SOL Privacy Officer; DLA Chief Data Officer (PIA authority); DLA Privacy Officer

Stakeholder Acceptance	Pending explicit Gate-2 confirmation; accepted implicitly through TNA v1.0

Context

The ALICE synthetic population is generated by a privacy-preserving generative model trained on de-identified DLA workforce records, with differential privacy noise applied during training (target epsilon ≤ 1.0). The synthetic population is not a record of any real individual. However, it is derived from — and statistically mirrors — a dataset of Privacy Act-protected personnel records.

A legitimate question arises: does differential privacy at epsilon ≤ 1.0 constitute sufficient protection to treat the synthetic population as non-CUI, enabling broader sharing and use? The answer in current federal guidance is unresolved. The IAPP (November 2024) documented that federal guidance does not formally authorize differential privacy as an automatic mechanism for reducing data sensitivity classification below CUI. NIST SP 800-188, which may provide formal guidance on this question, is in draft status as of July 2026.

Decision

The synthetic population shall be treated as CUI by default, with all CUI handling controls applied, until a formal Privacy Impact Assessment (PIA) conducted by DLA and the DLA Chief Data Officer determines otherwise. The PIA must explicitly address: (1) whether differential privacy at the applied epsilon value is sufficient to prevent individual identification; (2) whether the synthetic population constitutes a 'record about an individual' under 5 U.S.C. § 552a; and (3) whether a different classification applies.

The architecture does not prevent a future PIA determination that the synthetic layer may be reclassified. The four-zone data architecture (ADR-002) supports the synthetic zone having a different classification from other zones. This decision establishes the conservative default; a PIA-based determination can update the classification without requiring an architectural change.

Rationale

- Conservative default protects both DLA and SOL from regulatory risk. If the synthetic population were treated as non-CUI and a future policy determination or audit finding concluded that CUI handling was required, the remediation cost would be substantially higher than the cost of maintaining CUI controls from the outset.
- A formal PIA is the correct governance mechanism for a classification determination of this type. Classification decisions about novel data products derived from sensitive records require formal agency-level governance, not unilateral contractor determination.
- Epsilon ≤ 1.0 provides strong individual privacy guarantees as established by Dwork and Roth (2014), but mathematical privacy guarantees and regulatory classification are distinct questions. A mathematical guarantee does not automatically translate into a regulatory classification under federal CUI policy.
- Treating the synthetic population as CUI enables it to be used in the analytical and simulation environments immediately, without waiting for a PIA determination. The PIA determination may later enable broader use; it does not restrict current use.

Alternatives Considered and Rejected

Alternative	Why Considered	Why Rejected
Treat synthetic population as non-CUI immediately, based on epsilon ≤ 1.0 differential	Enables broader sharing; simplifies Phase III commercialization; consistent with the mathematical privacy guarantee.	Federal guidance does not currently authorize this determination. Without a formal PIA, treating synthetic data derived from Privacy Act records as non-CUI creates unacceptable regulatory risk

Alternative	Why Considered	Why Rejected
privacy guarantee.		for both DLA and SOL. (RL-23 — IAPP November 2024.)
Treat synthetic population as Privacy Act-protected records equivalent to source data.	Maximum conservative posture.	Overstates the restriction. The synthetic population contains no individual records; differential privacy guarantees that no individual's data can be inferred. Treating synthetic data as equivalent to source Privacy Act records would prevent the primary analytical use case the system is designed to enable.
Defer the classification determination until Phase II without any default classification.	Avoids a premature decision.	No default classification creates operational ambiguity during Phase I. Any synthetic data generated in Phase I would have no defined handling controls. CUI-by-default eliminates this ambiguity and is the operationally safe choice.

Consequences

- The Synthetic Zone (ADR-002, Zone 3) is subject to full CUI handling controls: encryption at rest, access credential control, audit logging of all access events, and export governance.
- All synthetic data export files carry CUI markings pending the PIA determination. Export recipients must be cleared to receive CUI.
- A Privacy Impact Assessment must be initiated in Phase II as a prerequisite to any reclassification of the synthetic layer. PIA initiation should be recommended as a Phase II Month 1 task (consistent with TNA N-13).
- Phase III commercialization scenarios that envision synthetic data as a freely shareable public dataset are contingent on the PIA determination. This is a known risk to Phase III business planning (documented in RISK-010).
- If NIST SP 800-188 is finalized before Phase II design, the PIA process may be able to reference SP 800-188 as the formal federal framework for the classification determination. This could accelerate the PIA timeline.

Risks

- **RISK-010 (Classification ambiguity delay):** PIA initiation may be delayed by DLA governance process. Synthetic data export governance will remain under CUI controls until the PIA is complete. This may limit the analytical use cases that can be served by synthetic data exports in Phase II.
- **Inferred risk — Phase III commercial impact:** If the PIA determines that the synthetic population remains CUI indefinitely, Phase III commercialization as a public synthetic dataset is not viable. Business planning should account for this contingency.

Cross-References

Reference Type	IDs
Architecture Principles	AP-1 (Privacy as Architectural Invariant); AP-4 (Data Zone Separation)
Logical Components	LC-05 (Synthetic Workforce Engine — generates the population); LC-09 (Audit Logging — tracks all synthetic zone access)
Data Entities	Synthetic Person; Synthetic Data Export; Generation Metadata (TNA Section 12.2)

Reference Type	IDs
Information Flows	IF-09 (Synthetic Population to Synthetic Zone); IF-15 (Synthetic Data Export Delivery)
PRD Requirements	REQ-COMP-001 (Privacy Act compliance); REQ-COMP-004 (DP documentation); REQ-F-015 (epsilon ≤ 1.0); REQ-F-018 (auditable export log)
Research Log	RL-20 (Privacy Act statute); RL-22 (Dwork & Roth DP foundations); RL-23 (IAPP — DP not automatic classification reduction); RL-24 (NIST 800-188 draft status)
Related ADRs	ADR-002 (Synthetic Zone is one of four data zones with distinct CUI controls); ADR-006 (conservative security posture applied consistently across both ADRs)
Related Risks	RISK-010 (classification ambiguity delay); RISK-007 (membership inference attack — privacy validation is the technical complement to this governance decision)
Open Questions	OQ-006 (NIST SP 800-188 finalization — may affect PIA timeline and framework)
Recommended Next Steps	N-13 (engage DLA Privacy Officer and CDO to initiate PIA in Phase II Month 1)

4. Candidate ADRs for Stakeholder Determination

The following three candidate ADRs were identified during cross-artifact consistency analysis as implicit architectural decisions documented in the TNA that were not formally captured as numbered ADRs. Each is raised to the level of a formal ADR because the decision involves a significant architectural commitment with specific consequences and risks that warrant explicit stakeholder awareness. Gate-2 approval of these candidates as Accepted ADRs would formally close three implicit decisions that the implementation team will otherwise resolve independently.

ADR-008 (Candidate): Simulation Engine Access Restricted to Synthetic Zone Only

ADR-008	Simulation Engine Access Restricted to Synthetic Zone Only
Status	Proposed — Awaiting Gate-2 Stakeholder Confirmation
Version	1.0 — July 2026 (candidate, identified during SDP cross-artifact analysis)
Category	Privacy and Security Architecture
TNA Basis	TNA Section 10.2 (LC-06 component description): 'The simulation engine has no access to the raw or processed zones. It reads exclusively from the synthetic zone.' Also ADR-002 consequence: data zone access controls enforced at service account level.
Decision Owner	SOL Lead Architect and Privacy Officer

Context

The ABM simulation engine (LC-06) requires workforce population data to instantiate agents for scenario execution. Two population sources exist within ALICE: the de-identified records in the Processed Zone, and the synthetic population in the Synthetic Zone. Using de-identified records directly in simulation would expose population-level workforce structure in simulation outputs, even if individual identifiers are absent. It would also create a data flow path between the simulation environment and the Processed Zone — a zone subject to Privacy Act-adjacent controls — that is not required for simulation function.

The synthetic population is designed precisely to enable simulation without these constraints. Its use as the sole simulation input is both architecturally cleaner and provides stronger privacy protection.

Decision

The ABM simulation engine (LC-06) shall have read access to the Synthetic Zone only. LC-06 service account credentials shall be provisioned to the Synthetic Zone exclusively. LC-06 shall have no access credentials to the Processed Zone, the Raw Zone, or the Analytical Zone (except for writing scenario results).

Rationale

- The Processed Zone contains de-identified records that, while not Privacy Act records, retain the population structure of real DLA workforce data. Exposing this structure in the simulation environment — even without individual identifiers — creates re-identification risk through population inference.
- The synthetic population exists specifically to provide simulation-quality data without these risks. Using it as the sole simulation input fulfills the purpose for which it was designed.
- Restricting LC-06 credentials to the Synthetic Zone provides verifiable, architecturally enforceable protection. An assessor can confirm the restriction by reviewing LC-06 service account permissions, without reviewing simulation engine code.

Alternatives Considered

Alternative	Why Rejected
Allow simulation engine to read from Processed Zone with application-layer filtering.	Application-layer filtering cannot be independently verified without code review. Creates a dependency on simulation engine correctness for Privacy Act protection. Inconsistent with the architectural invariant in AP-1.
Allow simulation engine to read from both Processed and Synthetic zones, using Processed data for validation only.	Validation functions can be executed in the Synthetic Engine (LC-05) or against the Ontology (LC-04), not in the simulation environment. Cross-zone read access for validation purposes is architecturally unnecessary.

Cross-References

Reference Type	IDs
Architecture Principles	AP-1 (Privacy as Architectural Invariant); AP-4 (Data Zone Separation); AP-3 (Least Privilege — ZTA)
Logical Components	LC-06 (Simulation Engine — affected); LC-05 (Synthetic Engine — provides Synthetic Zone data); LC-12 (zone-specific credential management)
Information Flows	IF-10 (Synthetic Zone to Simulation Engine — the only inbound data flow

Reference Type	IDs
	to LC-06)
PRD Requirements	REQ-NF-008 (no Privacy Act record in simulation layer); REQ-SEC-003 (least privilege); REQ-COMP-001 (Privacy Act compliance)
Research Log	RL-20 (Privacy Act statute); RL-22 (DP — synthetic data as the privacy mechanism)
Related ADRs	ADR-002 (zone access controls); ADR-007 (synthetic zone CUI classification)

ADR-009 (Candidate): Read-Only External Integration Enforced at Service Account Level

ADR-009	**Read-Only External Integration Enforced at Service Account Level**
Status	Proposed — Awaiting Gate-2 Stakeholder Confirmation
Version	1.0 — July 2026 (candidate, identified during SDP cross-artifact analysis)
Category	Integration and Security Architecture
TNA Basis	AP-2 (Read-Only External Integration): 'All external integration interfaces shall be implemented as read-only or pull mechanisms. Service accounts used for data ingestion shall be granted read-only access with the minimum data scope required.' Also CON-002 (system is read-only relative to all source systems) and CON-004 (no write-back to DCPDS, DPMAP, EBS, or DLA TMS).
Decision Owner	SOL Lead Architect; DLA source system owners (service account provisioning)

Context

ALICE is defined by both the PRD and CONOPS as a read-only analytical overlay. It consumes data from four DLA source systems but does not initiate personnel actions, update position records, or write data back to any source system. This boundary is fundamental to the ALICE operational concept: ALICE provides decision support, but execution of personnel decisions occurs through existing HR systems and human decision-makers.

The read-only boundary is stated as a design principle (AP-2) and a constraint (CON-002, CON-004). However, the TNA does not formally document the specific enforcement mechanism — namely, that read-only access must be enforced at the service account credential level, not through application-layer logic. This distinction matters for security assessment: an application-layer read-only implementation can be overridden by a code change or privilege escalation; service account-level enforcement cannot.

Decision

All ALICE service accounts provisioned for external source system integration (INT-01 through INT-04) shall be granted read-only database or API access by the source system administrators. The service accounts shall not have insert, update, delete, or execute permissions on any source system object. SOL shall provide DLA source system administrators with a written service account

permission specification before any service account is provisioned. The permission specification shall be retained as a system security artifact.

Rationale

- Service account-level enforcement provides a technically verifiable read-only guarantee. A source system administrator, DLA ISSO, or independent assessor can confirm read-only permissions by reviewing source system access control lists, without requiring a code review of ALICE ingestion services.
- Application-layer read-only logic can be overridden by a defect, update, or compromise. Service account-level restrictions cannot be overridden from within ALICE; they are enforced by the source system's access control mechanism.
- Source system data integrity is a mission-critical concern. An unintended write operation to DCPDS, DPMAP, or EBS could corrupt authoritative personnel records, with serious operational and legal consequences. The service account restriction is an irreversible architectural safeguard against this outcome.

Alternatives Considered

Alternative	Why Rejected
Application-layer read-only enforcement (no external write calls in ALICE code).	Cannot be independently verified without code review. Can be overridden by a code change or injection attack. Does not provide the RMF-quality evidence that source system administrators and assessors require.
No formal service account permission specification (rely on DLA source system defaults).	Source system defaults vary; default permissions may include write access to some objects. Without a formal specification, there is no documented evidence of the read-only enforcement for RMF purposes.

Cross-References

Reference Type	IDs
Architecture Principles	AP-2 (Read-Only External Integration); AP-3 (Least Privilege)
Integrations	INT-01 (DCPDS/DCHRMS); INT-02 (DLA TMS); INT-03 (DPMAP); INT-04 (EBS)
Constraints	CON-002; CON-004; CON-010
PRD Requirements	PRD Section 2.6 (System Boundary — read-only); PRD Section 11.2 (Out of Scope — write-back)
Related ADRs	ADR-001 (LC-02 and LC-03 are the only components with source system access); ADR-004 (least-privilege service account access is consistent with ZTA principles)

ADR-010 (Candidate): Configuration-Driven Operational Parameters

ADR-010	**Configuration-Driven Operational Parameters**

Status	Proposed — Awaiting Gate-2 Stakeholder Confirmation
Version	1.0 — July 2026 (candidate, identified during SDP cross-artifact analysis)
Category	Maintainability and Operations Architecture
TNA Basis	AP-7 (Configuration Over Customization); LC-10 (Configuration Management Service); REQ-ADM-002 (threshold configuration without code modification); REQ-ADM-004 (export authorization workflow configuration); REQ-ADM-006 (configuration change log)
Decision Owner	SOL Lead Architect

Context

DLA operational requirements will change over time. Readiness thresholds that define when an alert is generated will need adjustment as DLA leadership calibrates the tool to their operational standards. Scenario parameter ranges may need expansion as field users identify additional disruption types. Export authorization workflow configuration — required approver roles, required documentation fields — may evolve as data governance practices mature. If these parameters are hardcoded in application logic, every operational adjustment requires a code change, a test cycle, and a deployment event, creating friction that will inhibit adoption and increase maintenance cost.

Decision

All operational parameters that must be adjustable without system reinstallation or code modification shall be stored as configuration data in LC-10 (Configuration Management Service) and exposed through the System Administrator interface. This includes: readiness indicator threshold values by mission function area; alert delivery conditions and recipient lists; scenario library parameter ranges; export authorization workflow settings (required approver roles, required documentation fields); data source ingestion schedules; and synthetic population generation schedule and parameters (excluding model architecture).

Configuration data is version-controlled, changes are logged per REQ-ADM-006, and rollback to prior configuration versions is supported.

Rationale

- Operational flexibility: DLA's use of ALICE will evolve over Phase II. Threshold values calibrated in Phase I may require adjustment based on actual field experience. Configuration-driven parameters allow this adjustment without developer involvement.
- Adoption risk reduction: if operational users cannot adjust ALICE to their specific context without requesting a code change, adoption friction increases. Configuration-driven parameters give System Administrators the autonomy to calibrate the system to operational requirements.
- Audit defensibility: changes to operational parameters are logged by LC-09 with administrator identity, prior value, new value, and timestamp. This provides a complete, auditable record of how the system's behavior has evolved over time — an RMF evidence requirement.

Alternatives Considered

Alternative	Why Rejected
Hardcode all operational parameters in application logic.	Every operational adjustment requires a code change, testing, and deployment. Unacceptable maintenance burden for an operational system whose threshold values are expected to evolve as DLA calibrates its use.
Allow direct database modification for parameter changes (no administrative	Uncontrolled database modifications bypass the configuration change log requirement (REQ-ADM-006) and create audit gaps. Cannot be independently tracked or rolled back.

Alternative	Why Rejected
interface).	

Risks

- **Inferred risk — configuration as a security surface:** Configuration data that governs system behavior is itself an attack surface. If an unauthorized party can modify configuration data, they could disable alert thresholds or alter export authorization requirements. LC-10 configuration data must be access-controlled and all changes must be logged. Read access to configuration data must be authorized; write access must be restricted to System Administrator role and logged.

Cross-References

Reference Type	IDs
Architecture Principles	AP-7 (Configuration Over Customization)
Logical Components	LC-10 (Configuration Management Service)
Functional Domains	FD-7 (Administration and Configuration)
PRD Requirements	REQ-ADM-002 (threshold configuration without code change); REQ-ADM-004 (export authorization workflow configuration); REQ-ADM-006 (configuration change log); REQ-ADM-007 (configuration backup)
Related ADRs	ADR-001 (LC-10 is one of twelve modular components); ADR-005 (all configuration changes are logged to LC-09)

5. Gate-2 Summary and Decision Required

What Was Created

Artifact 2 (ADR Register) contains ten fully specified Architecture Decision Records: seven formally accepted ADRs from TNA Section 21 (ADR-001 through ADR-007), each expanded to include decision version history, implementation guidance, alternative rejection rationale, risk cross-references, and an expanded cross-reference table; and three candidate ADRs (ADR-008 through ADR-010) identified during SDP cross-artifact consistency analysis as implicit decisions in the TNA that warrant explicit stakeholder confirmation.

Every ADR preserves the ID, status, context, decision, rationale, alternatives considered, consequences, and related requirements from TNA v1.0. No IDs have been changed. The three candidate ADRs are labeled Proposed and do not require implementation action until Gate-2 acceptance.

Major Assumptions in This Artifact

- The seven TNA ADRs (ADR-001 through ADR-007) are treated as accepted decisions pending explicit Gate-2 confirmation, consistent with their Accepted status in TNA v1.0. If the stakeholder group wishes to revise any of these decisions, the revision must be documented in this register with a version increment and a record of the prior decision.
- ADR-008, ADR-009, and ADR-010 reflect implicit decisions documented elsewhere in the TNA. They are raised to formal ADR status because the implementation implications of these decisions warrant explicit stakeholder awareness and record.

- ADR-006 carries a critical assumption (ASM-007): DLA legal counsel must confirm that Phase I operation on de-identified data does not trigger a CUI ATO requirement. This confirmation is independent of Gate-2 and must be obtained before Phase I execution begins.

Open Issues

- **[Critical — ADR-006]** ASM-007: DLA legal counsel confirmation of Phase I data handling permissibility. This is a prerequisite for Phase I execution, not just for Gate-2 approval.
- **[High — ADR-004]** ASM-008: CAC integration feasibility in the Phase I demonstration environment must be confirmed at Phase I kickoff.
- **[High — ADR-007]** OQ-006: NIST SP 800-188 finalization status should be monitored. If finalized before Phase II design, it must be incorporated into the PIA framework for the synthetic layer classification determination.
- **[Medium — ADR-003]** OQ-004: DCHRMS migration schedule for the pilot sub-component must be confirmed at Phase I kickoff to determine whether LC-03 DCHRMS schema mappings are in scope for Phase I.
- **[For Stakeholder Determination]** ADR-008, ADR-009, and ADR-010 require explicit acceptance or rejection at Gate-2. The architecture currently reflects these decisions implicitly; formal acceptance converts them to binding implementation guidance.

Decision Required

Gate-2 Decision: Please confirm one of the following: (A) Approve Artifact 2 as documented, including explicit acceptance of ADR-001 through ADR-007 and candidate ADR-008 through ADR-010, and proceed to Artifact 3 (Component Catalog). (B) Approve with modifications — specify which ADR(s) require revision. (C) Revise the status of candidate ADR-008, ADR-009, or ADR-010 to Deferred or Rejected, with documented rationale. Once Gate-2 is approved, Artifact 3 will be produced using the same artifact IDs and consistent cross-references to this ADR Register.