

SOL ENGINEERING SERVICES

The Daniel Group LLC

The Organization as an Operational Variable:

An ALICE Workforce Intelligence Architecture for Defense Logistics Agency Mission Readiness

Prepared in support of SBIR Topic DLA26BZ03-NV011

Defense Logistics Agency | DoD SBIR 2026 Round 3

July 2026

Classification: Unclassified

Executive Summary

The Defense Logistics Agency faces a workforce management challenge that has no adequate solution within its current toolset. DLA employs approximately 26,000 personnel worldwide across nine distinct supply chain domains, and the agency's 2025-2030 Strategic Plan identifies workforce readiness as one of its four foundational pillars. Yet the tools available to DLA leadership for assessing workforce health, anticipating personnel risk, and planning for operational surges are largely retrospective and disconnected from mission performance data. Static headcount dashboards and annual workforce planning cycles cannot support the decision speed that great power competition demands.

SBIR Topic DLA26BZ03-NV011 names this gap precisely: DLA seeks to develop a dynamic workforce digital twin capable of generating synthetic data, enabling scenario planning, and accelerating the adoption of AI-enabled workforce management. The U.S. Department of Defense's January 2026 AI Acceleration Strategy adds strategic urgency, directing every defense agency to identify and execute pace-setting AI projects at wartime speed while eliminating bureaucratic barriers to capability deployment.

SOL proposes **ALICE** — the **Adaptive Labor Intelligence and Capacity Engine**: a workforce digital twin purpose-built for defense logistics enterprise management. ALICE maintains a persistent, continuously updated model of the DLA organization, generates privacy-preserving synthetic workforce data under differential privacy guarantees, and enables leaders to simulate the effects of workforce disruptions, staffing gaps, and operational surges before those events materialize. It is not a dashboard; it is a simulation environment that gives DLA the capacity to test futures before committing to them.

This white paper establishes the strategic rationale for the ALICE architecture, examines the technical foundations that make it achievable, addresses the privacy and data governance framework governing implementation, and defines the Phase I feasibility path through which SOL will demonstrate the concept's viability. The document draws on DLA's publicly stated strategic priorities, the U.S. Department of Defense's 2026 AI strategy, peer-reviewed research on organizational digital twin architectures, and the operational lessons of contested logistics in the era of great power competition.

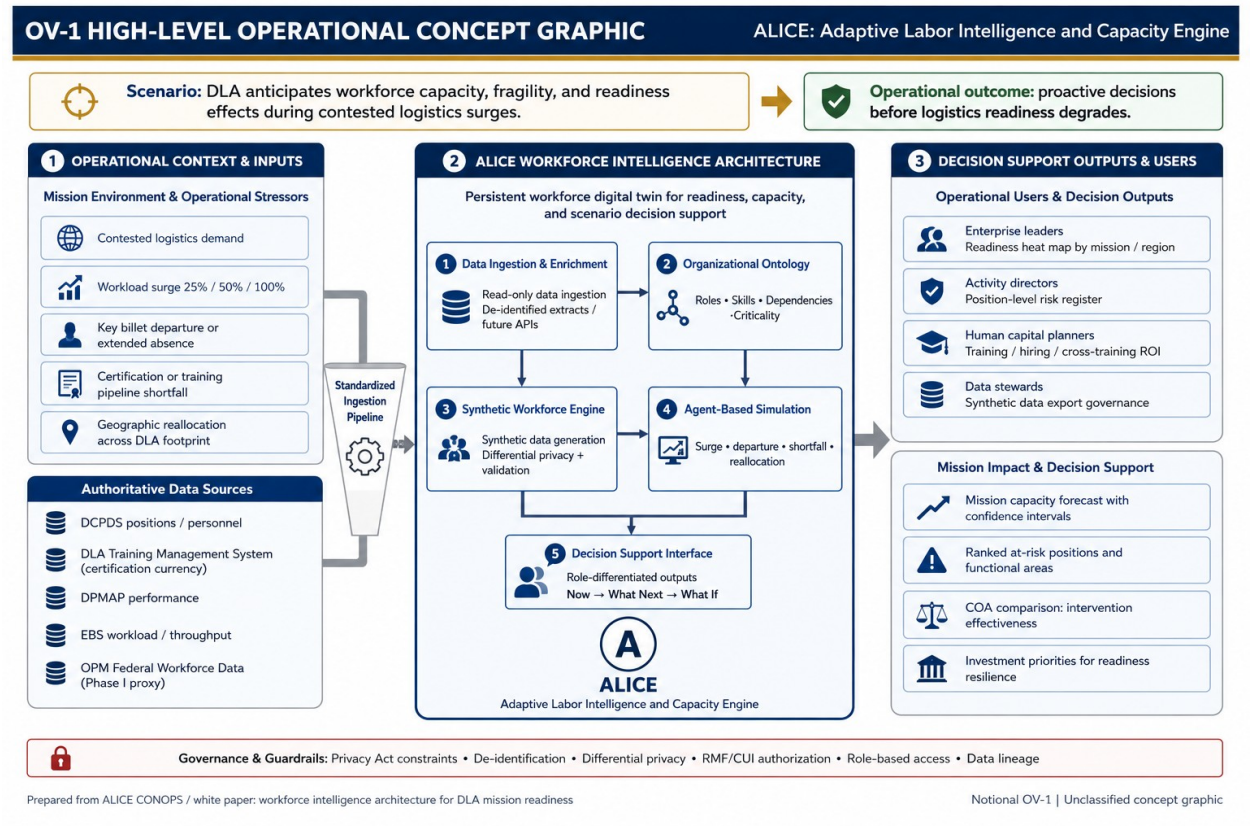

Figure 1. OV-1 High-Level Operational Concept Graphic — ALICE: Adaptive Labor Intelligence and Capacity Engine

I. The Strategic Context: Logistics as a Decisive Variable

The Contested Logistics Environment

The strategic literature is unambiguous on one point: in a high-intensity conflict with a near-peer adversary, logistics performance does not support the outcome; it determines it. The Defense Acquisition University's *Defense AT&L* journal observed in late 2024 that the character of 21st-century warfare under great power competition has changed every dimension of logistics, eliminating the freedom of maneuver that allowed U.S. forces to accumulate iron mountains of materiel in prior conflicts.[1]

The Center for Strategic and International Studies reinforced this assessment in September 2025, concluding that future conflicts with peer competitors will require states to sustain readiness through engagements that may extend for years, and that investments in logistics readiness and resilience will determine which combatant can continue fighting when the other cannot.[2] Against this backdrop, the workforce

that plans, coordinates, and executes DLA's nine supply chains is not a support function; it is an operational variable with direct bearing on mission success or failure.

DLA's Workforce Imperative

DLA Director Army Lieutenant General Mark Simerly has stated directly that the agency must think, act, and operate differently to support the joint warfighter in a contested logistics environment.[3] The agency's 2025-2030 Strategic Plan identifies four pillars: People, Posture, Precision, and Partnerships. The People pillar is not incidental; it is listed first, and DLA leadership has described building data competency across the entire workforce as one of its primary objectives under that pillar.[4]

DLA's Chief Information Officer, Adarryl Roberts, has framed the agency's AI strategy in explicitly human terms. In January 2026, Roberts described the goal as creating digital employees that can work at scale, enabling the human workforce to redirect its capacity toward higher-value mission functions. Roberts has also emphasized that DLA's AI approach differs from many agencies' generative AI deployments: DLA is focused on deterministic models with high-confidence outputs, because the mission requires certainty, not approximation.[5]

> *"We need to move at pace to keep up with our adversaries in China and Russia. We're really looking to put the AI and the automation technology in the hands of the operator, because they know the mission, they know the problem." — Adarryl Roberts, DLA Chief Information Officer, November 2025*

By March 2025, DLA had deployed more than 55 AI models across production, testing, and operational use, with active programs in demand planning, supply chain risk management, and process automation.[6] ALICE represents the logical next step: extending AI-enabled intelligence from supply chain objects to the human organization that manages them.

The U.S. Department of Defense's AI Acceleration Mandate

On January 12, 2026, Secretary of Defense Hegseth issued the U.S. Department of Defense AI Acceleration Strategy, establishing 2026 as the year the Department would decisively raise the bar for military AI dominance. The strategy directed every military department, combatant command, and defense agency to identify a minimum of three AI pace-setting projects for prioritization within 30 days and to demonstrate measurable progress to Department leadership on a monthly basis beginning in July 2026.[7]

The strategy's Enterprise Agents initiative is directly relevant to the DLA workforce digital twin concept. Enterprise Agents directs the development of playbooks for

rapid, secure AI agent deployment to transform enterprise workflows across the DoD. A workforce digital twin that models organizational behavior, generates synthetic populations, and executes scenario simulations is precisely the class of enterprise AI agent the strategy envisions: mission-grounded, data-driven, and capable of operating at the scale of a large defense enterprise.[8]

The strategy also mandates the enforcement of DoD Data Decrees, requiring military departments and defense agencies to expose their data assets and system interfaces to the Chief Digital and Artificial Intelligence Office within 30 days and to make those assets available to cleared users with valid purpose. For DLA, this mandate is an enabling condition for ALICE: it creates the data access authority that Phase I system integration will require.

II. The Capability Gap: What Current Tools Cannot Do

The Limits of Retrospective Workforce Management

DLA's current workforce management apparatus is built for steady-state administration, not dynamic operational risk management. Personnel data flows through the Defense Civilian Personnel Data System into reporting outputs that describe what the workforce looked like at a point in time. Vacancy rates are reported quarterly. Training completion is tracked against annual plans. Attrition is observed after it occurs.

This architecture is adequate for a personnel department managing routine HR functions. It is inadequate for a combat support agency whose workforce performance directly affects the joint force's capacity to sustain operations under contested conditions. The temporal lag embedded in retrospective reporting means that by the time a workforce risk is visible in official data, the window for proactive intervention has already closed.

Three Questions Leadership Cannot Currently Answer

SOL's analysis, informed by operational experience across logistics readiness squadron command and defense enterprise management, identifies three categories of decision that DLA leadership cannot currently support with adequate analytical confidence.

Surge capacity: Can DLA absorb a 30 percent increase in workload across Class V ammunition and Class IX repair parts within 72 hours? Which functional nodes will degrade first, and in what sequence? Current tools offer headcount data; they do not offer throughput modeling tied to specific workforce configurations.

Fragility mapping: Which positions represent single points of organizational failure, where one departure or extended absence triggers cascading downstream effects? In a contracting shop where a single warranted contracting officer processes the majority of complex modifications, the departure of that individual has consequences far beyond the staffing gap it creates. Identifying these nodes before they fail requires a relational model of the organization that current HR systems do not maintain.

Investment optimization: Given constrained human capital budgets and competition for specialized talent, which training investments, cross-training programs, or targeted recruitment actions yield the highest return in mission readiness terms? Without a model that connects workforce configuration to operational output, this question is answered by intuition rather than evidence.

The Synthetic Data Problem

A second and equally important limitation is the privacy constraint on existing workforce data. DLA's personnel records are protected under 5 U.S.C. § 552a, the Privacy Act of 1974. The richest datasets in DLA's possession cannot be freely shared across organizational boundaries, used in simulation environments without appropriate protections, or exposed in analytical tools accessed by users without specific need-to-know authorization.

Synthetic data generation offers a resolution to this constraint. A synthetic workforce population derived from real data under formal differential privacy guarantees preserves the statistical structure of the underlying organization — including grade distributions, occupational series concentrations, tenure profiles, and skill densities — without containing records of actual individuals. Research published in *Decision Support Systems* in early 2026 confirmed that digital twin architectures employing synthetic workforce data can achieve sufficient fidelity to support consequential human capital decisions while maintaining realistic privacy protections.[9]

III. ALICE: Concept and Architecture

System Vision

Figure 1 presents the ALICE OV-1 High-Level Operational Concept Graphic, which frames the complete system across three columns: the operational context and authoritative inputs that feed ALICE; the five-layer workforce intelligence architecture that processes and models those inputs; and the decision support outputs delivered to operational users. The scenario driving the concept is concrete: DLA must be able to anticipate workforce capacity, fragility, and readiness effects

during contested logistics surges — and must do so proactively, before logistics readiness degrades. The following sections describe each architectural layer in detail.

ALICE — the Adaptive Labor Intelligence and Capacity Engine — is a persistent, AI-powered simulation of the DLA organization. It does not replicate what HR dashboards already provide. Instead, it models the organization as a dynamic system of interdependent roles, skills, relationships, and workflows, and uses that model to generate three classes of operational output: continuous readiness assessments, forward-looking scenario simulations, and privacy-preserving synthetic data products.

The distinction that separates a digital twin from a conventional dashboard is architectural. A dashboard reports current state; a digital twin simulates future states. A dashboard displays headcount; a digital twin models how the organization would perform if that headcount changed. A dashboard shows vacancy rates; a digital twin shows which vacancies matter and what happens downstream when they persist. ALICE is built on this distinction as a design principle.

> *ALICE enables three capabilities that current tools cannot: seeing the organization as it is now, projecting the organization as it will be under defined conditions, and asking what the organization would look like if leadership acted differently. These are the three temporal dimensions of operational decision support.*

Data Ingestion Architecture

ALICE connects to four authoritative DLA data sources through read-only integration layers. The Defense Civilian Personnel Data System (DCPDS) provides the authoritative record of positions and personnel: grade, series, appointment type, duty station, supervisory status, and organizational assignment. The DLA Training Management System provides certification currency, course completion histories, and training pipeline data. The Defense Performance Management and Appraisal Program (DPMAP) contributes performance appraisal records used to calibrate individual and organizational productivity parameters. DLA's Enterprise Business System (EBS) contributes workload metrics and supply chain throughput data that allow the twin to correlate workforce configuration with mission output.

Phase I will employ de-identified data extracts rather than live system connections, pending the execution of data use agreements and Risk Management Framework authorization. For any gap period during Phase I where DLA data access is pending, OPM's Federal Workforce Data platform (data.opm.gov), launched January 2026, provides a functionally adequate proxy for developing and validating ALICE's statistical architecture before DLA-specific data is available.[10]

Organizational Ontology and Agent Model

The ontology is ALICE's internal language for representing the organization. Each position in DLA is modeled as a node with a defined attribute set: occupational series, grade, required certifications and their currency status, mission function assignment, workflow dependencies, knowledge criticality rating, current fill status, and attrition risk score derived from tenure, time-in-grade, and career stage data.

The relationships between nodes are what distinguish the ALICE ontology from a conventional position management database. The system maps which roles are prerequisites to which workflows, where single-point-of-failure conditions exist, how knowledge flows through informal professional networks, and how workload distributes across functional teams under varying staffing configurations.

Agent-based modeling (ABM) provides the simulation paradigm. Each role is instantiated as a computational agent carrying its attribute set and operating according to behavioral rules derived from organizational process documentation, subject matter expert input, and historical workflow data. ABM is the appropriate choice for organizational simulation because organizations are not governed by differential equations; they are governed by the behavior of individuals interacting within structures. ABM captures this directly.[11]

Synthetic Data Generation Engine

The synthetic data generation engine produces a privacy-preserving replica of the DLA workforce population. The architecture employs a variational autoencoder trained on de-identified workforce records to learn the joint distribution of workforce attributes, combined with differential privacy mechanisms applied during training to provide formal mathematical privacy guarantees. The resulting synthetic population preserves the marginal and joint distributions of grade, series, tenure, certification coverage, geographic distribution, and skill density while ensuring that no individual record can be reconstructed from the synthetic outputs.

Differential privacy provides a formal guarantee: the probability that any specific individual's data contributed to the synthetic population is bounded by a mathematically defined privacy parameter, epsilon. SOL's implementation targets an epsilon value not exceeding 1.0, the threshold that provides strong individual privacy protection while maintaining high population-level fidelity.[12]

Validation of the synthetic population follows a two-stage protocol. Statistical distributional testing using Kolmogorov-Smirnov tests for continuous variables and chi-squared tests for categorical variables confirms that the synthetic population matches the source population within defined tolerance bands. Membership inference attack simulation confirms that an adversary who attempts to determine whether a specific individual's record was used to train the model cannot do so at a rate better than random chance.

Scenario Planning and Simulation Module

The scenario planning module allows DLA leaders and planners to define, execute, and analyze simulations of workforce disruption and response. Scenarios are configured along three axes: the initiating event, the duration of the disruption, and the management response option being evaluated.

ALICE ships with a preset library of operational scenario types calibrated to the disruptions DLA most frequently confronts. Key departure scenarios model the effect of a single high-impact individual leaving a critical role and simulate the readiness trajectory under different replacement timelines. Workload surge scenarios model the supply chain throughput implications of a 25, 50, or 100 percent increase in demand signal across defined commodity classes. Pipeline shortfall scenarios model the effect of training program delays or recruiting shortfalls on the future certification profile of the workforce. Geographic reallocation scenarios model the readiness effects of redistributing personnel across DLA's global footprint in response to shifting operational requirements.

Simulation outputs provide a time-series readiness forecast showing projected mission capacity as a percentage of the current baseline, a ranked list of at-risk positions and functional areas, and a comparison of the modeled effectiveness of alternative management responses across the scenario horizon. All outputs carry confidence intervals derived from Monte Carlo sampling across the synthetic population.

Decision Support Interface

The decision support interface delivers simulation outputs through a role-differentiated dashboard environment designed for operational use without analyst intermediation. Enterprise-level leaders receive a readiness heat map by mission function and geographic region. Activity directors receive a position-level risk register with configurable alert thresholds. Human capital planners receive an investment prioritization panel that ranks training, hiring, and cross-training interventions by their modeled impact on mission readiness.

Interface design follows one governing principle: the tool must save time, not consume it. If field-level supervisors are required to enter data manually, adoption will not occur. ALICE pulls from existing authoritative systems automatically. Scenario planning executes through preset workflows accessible to non-technical users. Report outputs integrate with DLA's existing briefing and planning cycle formats. The system exists to put analytical leverage in the hands of the operators who understand the mission; it cannot require those operators to become data scientists to use it.

IV. Privacy, Governance, and Data Stewardship

The Legal Framework

Federal workforce data is among the most legally constrained data categories in the government's possession. The Privacy Act of 1974 (5 U.S.C. § 552a) governs the collection, maintenance, use, and dissemination of records about individuals maintained by federal agencies. DLA's workforce data is maintained under OPM's system of records notices and subject to OPM's data release policy, which governs what data can be released at the individual record level and under what conditions.[13]

ALICE's synthetic data architecture is designed to operate fully within these constraints. The system does not store or process Privacy Act records in its analytical or simulation environments. It processes de-identified data extracts or, in its simulation and scenario planning functions, synthetic populations derived from those extracts. The transformation from Privacy Act-protected records to de-identified extracts to synthetic populations is documented with full data lineage, and the synthetic population itself is treated as a DLA analytical asset rather than a personnel record.

OMB Data Governance Alignment

OMB Memorandum M-19-15 and the Federal Data Strategy's ten principles of data stewardship require agencies to implement appropriate safeguards when deriving new data products from existing records and to document the data lineage of those products. ALICE's data governance layer fulfills both requirements: it maintains an auditable record of every data transformation from source system to synthetic population, documents the privacy parameters applied at each stage, and records the intended use of each synthetic data export with requestor identity and authorization basis.

The DoD Data Decrees issued alongside the January 2026 AI Acceleration Strategy create a complementary authorization framework. By directing the CDAO to facilitate data sharing with cleared users with valid purpose, and by requiring defense agencies to expose their data assets through federated catalogs, the Decrees establish the governance basis for the data use agreements ALICE requires. SOL will pursue these agreements beginning at contract award.[14]

Security Architecture

ALICE's baseline operating environment targets the Controlled Unclassified Information impact level, consistent with the sensitivity of de-identified workforce data and the unclassified nature of the organizational models it maintains. The system architecture is designed for Risk Management Framework authorization at the CUI level, with a Phase II extension path to IL4/IL5 for classified scenario planning applications.

V. Technical Foundations and State of the Art

Digital Twins in Defense Applications

Digital twin technology has migrated from its origins in manufacturing and aerospace into defense applications across a widening range of domains. Research published in *Sensors* in December 2025 documented the progression of AI-driven digital twins from passive monitoring systems into prescriptive, self-optimizing simulation environments that enable real-time visibility across organizational levels.[15]

Federal application is accelerating. The Atlantic Council called for a dedicated federal digital twins strategy in 2024, noting that digital twins, when appropriately applied, can contribute to national security and a more effective, productive, and resilient workforce.[16] By March 2026, Maximus had formalized a partnership with Skyral specifically to deploy digital twin technology in federal defense, civilian, and healthcare mission environments, enabling agencies to simulate workforce coverage, resource allocation, and service delivery scenarios before changes are implemented in live programs.[17]

Organizational Digital Twins: Research Foundation

A peer-reviewed study published in *Decision Support Systems* in February 2026 demonstrated that a digital twin architecture employing real-time workforce data, synthetic population modeling, and scenario simulation can function as a living, responsive decision-support system that accurately reflects the current state of the workforce while enabling the simulation of future scenarios and the evaluation of strategic decisions before implementation.[18]

That same study identified five leading workforce indicators with the strongest predictive relationship to mission performance degradation: position vacancy rate by critical occupational series, average days-to-fill for key billets, certification currency ratio, supervisor-to-subordinate span of control relative to task complexity, and informal knowledge network density. ALICE's ontology is built to track all five indicators continuously, providing DLA with earlier warning of workforce risk than any current reporting cycle allows.

Agent-Based Modeling for Organizational Simulation

Agent-based modeling has emerged as the dominant simulation paradigm for complex organizational systems. ABM is appropriate for organizations because it captures the emergent properties of collective behavior — the patterns that arise from the interactions of individual agents operating under defined rules — rather

than assuming that organizational behavior can be predicted from aggregate statistics alone.

The National Academy of Sciences has documented ABM's capacity to simulate human systems of substantial complexity.[19] In the organizational context, each agent represents a role with defined attributes and behavioral rules governing task assignment, knowledge transfer, collaboration, and workload response. The simulation engine executes scenario runs by modifying agent attributes or relationships and observing the resulting system-level effects on throughput capacity, readiness metrics, and organizational resilience.

Synthetic Data Generation: Technical Approach

Variational autoencoders, trained on de-identified records with differential privacy noise injection during training, have demonstrated the ability to preserve complex joint distributions across multiple workforce attributes while providing formal mathematical privacy guarantees.

The SYN-DIGITS calibrated digital twin simulation framework, published in early 2026, documents methodologies for using synthetic control approaches to align synthetic populations with real-world base rates, extending the technical foundation directly applicable to ALICE's data generation architecture.[20]

Machine Learning Integration Architecture

ALICE is built on a layered intelligence model in which machine learning functions are introduced where they produce verifiable analytical gains and withheld where rule-based or simulation-based approaches are more auditable, more explainable to DLA leadership, or more defensible before an RMF assessor. This layered approach reflects a deliberate design principle: ML is a capability amplifier within the ALICE architecture, not a replacement for the transparent computational logic that operational decision-makers must be able to interrogate.

Phase I: Generative Modeling as the Core ML Function

The synthetic workforce engine is the primary machine learning component in Phase I. A privacy-preserving generative model — trained on de-identified workforce records — learns the joint statistical distribution of grade, series, tenure, certification profile, performance tier, and geographic distribution across the pilot sub-component. Differential privacy noise injection during training provides a formal mathematical privacy guarantee, with an epsilon target not exceeding 1.0. The resulting synthetic population is the substrate for all simulation and analysis: it enables agent instantiation, scenario execution, and analytical export without Privacy Act exposure. This is the function for which ML is uniquely suited — learning complex high-dimensional distributions from constrained

training corpora while preserving privacy guarantees that rule-based data transformation cannot provide.

Readiness indicator computation in Phase I is intentionally rule-based. The five leading workforce degradation indicators identified in the research literature — position vacancy rate by critical occupational series, average days-to-fill for key billets, certification currency ratio, supervisor-to-subordinate span of control relative to task complexity, and informal knowledge network density — are computed from a formally specified algorithm validated by DLA subject matter experts. This choice is deliberate. An ML-learned readiness model requires labeled historical outcome data that Phase I has not yet generated. The rule-based computation is transparent, explainable, and SME-validatable, and it produces the historical snapshot corpus that will serve as the training dataset for a learned successor in Phase II.

Phase II: Four ML Maturation Paths

As Phase I generates operational data and Phase II scales to enterprise scope, four ML capability paths become viable. Each is architecturally compatible with the modular design: the configuration-driven parameter approach means each enhancement is a component-level upgrade, not a system redesign.

Learned readiness prediction. After 12 to 18 months of Phase II operational data, the historical snapshot corpus is large enough to train a supervised prediction model on actual workforce outcome events. This model identifies non-obvious leading indicator combinations that a manually specified formula would miss and replaces or augments the rule-based computation. The transition requires no architectural change: the computation algorithm is a configuration parameter, switchable without code modification. The Phase II evaluation methodology must characterize the learned model's performance against the rule-based baseline before any transition to operational use.

Natural language scenario initiation. A natural language interface layer within the decision support interface would allow an Activity Director to describe a workforce disruption in plain language — the interface translates the description into a structured scenario parameter object conforming to the preset library schema. The translation model processes only the user's natural language query and produces a validated parameter set before any request reaches the simulation engine. It does not touch workforce data. This capability addresses user adoption friction directly: it eliminates the requirement for field-level supervisors to understand preset scenario structure in order to initiate a simulation.

ML-augmented PII anomaly detection. The de-identification specification that governs data ingestion defines a static field removal list as the primary Privacy Act compliance control. In Phase II, an ML-trained PII classifier adds a defense-in-depth layer capable of detecting quasi-identifier combinations that the static list

does not cover — unusual co-occurrences of tenure, grade, series, and duty station that could allow re-identification even after direct identifiers are removed. Deployment at the Privacy Act compliance boundary requires formal validation of the classifier's false negative rate and DLA Privacy Officer acceptance before operational activation.

Advanced generative model architectures. If Phase I validation reveals that the initial generative model underperforms on rare subpopulation categories — small occupational series, low-density duty stations, or thin certification specialties critical to specific mission functions — Phase II can introduce a more capable architecture. Conditional tabular GAN variants or Gaussian copula models with minority-class strategies can improve tail fidelity while preserving the epsilon constraint. This is a component-level improvement with no interface contract changes required.

The Boundary Between ML and Simulation

Two design boundaries are held against ML replacement throughout the ALICE development path. The agent-based simulation engine remains the simulation paradigm because organizational decision-makers require outputs they can interrogate at the agent level: the causal chain from a specific role departure to a specific mission throughput reduction must be traceable and explainable without a machine learning interpretability framework. An opaque ML predictor that outputs a readiness percentage without an auditable causal chain is not appropriate for the personnel and operational decisions ALICE supports and would face resistance in the RMF authorization process.

Phase I readiness indicator computation similarly remains rule-based. The SME validation protocol — a minimum of two DLA subject matter experts reviewing and approving the ontology and its computational outputs — is the Phase I deliverable acceptance mechanism. The rule-based foundation is not a limitation of the ALICE architecture; it is the correct entry point for a system whose operational credibility depends on its ability to explain every readiness score to the operational leader who acts on it. The Phase I architecture is designed to accumulate the historical data that makes a learned successor both possible and well-founded.

VI. Phase I Feasibility Demonstration

Scope and Objectives

SOL's Phase I program is bounded by a six-month performance period at a funding ceiling of $100,000, consistent with the DLA26BZ03-NV011 award structure. The program demonstrates the technical feasibility of the ALICE architecture across four

discrete objectives, each producing a tangible deliverable and providing the evidence base for Phase II development funding.

Objective 1: Organizational Ontology Construction

SOL will develop and validate a DLA-representative organizational ontology for a defined pilot sub-component, selected in coordination with the DLA Government Technical Point of Contact at program kickoff. Validation will be conducted through structured subject matter expert review sessions with a minimum of two DLA subject matter experts.

Objective 2: Synthetic Data Generation Prototype

SOL will train and validate a synthetic workforce population generator. If DLA data sharing agreements are executable within the Phase I period, the generator will be trained on de-identified DLA workforce extracts. If data access is pending, OPM's Federal Workforce Data platform will serve as a functional proxy. The synthetic population will be validated against statistical distributional targets and differential privacy thresholds defined in the ALICE architecture.

Objective 3: Agent-Based Simulation Prototype

SOL will implement a functional agent-based simulation of the pilot sub-component capable of executing a minimum of three preset scenario types. The simulation prototype will execute historical backtesting against a minimum of 12 months of workforce state data, with a Phase I accuracy target of 80 percent within the defined tolerance band.

Objective 4: Decision Support Interface Prototype

SOL will develop and demonstrate a working prototype of the ALICE decision support interface for a minimum of two user role types. The prototype will be demonstrated in structured usability evaluation sessions with a minimum of three potential DLA end users.

Phase I Timeline

Months 1-2: Program kickoff, pilot sub-component selection, data use agreement initiation, ontology construction and subject matter expert review. Months 3-4: Synthetic data generator development and validation, agent model construction and calibration, simulation prototype development. Month 5: Historical backtesting, interface prototype development, user evaluation sessions. Month 6: Report preparation, Phase II plan development, final deliverable submission.

VII. Phase II Development Path

Phase II expands the validated Phase I prototype into a deployable enterprise capability across an estimated 24-month performance period. Development is organized around five parallel work streams.

Enterprise Data Integration: Execute data use agreements and implement live API connections to DCPDS, DLA Training Management System, DPMAP, and EBS, transitioning from de-identified exports to authorized near-real-time data feeds.

Enterprise Scale: Extend the organizational ontology from the Phase I pilot sub-component to DLA enterprise scope, encompassing all major activities and functional domains across DLA's nine supply chains and global distribution network.

RMF Authorization: Conduct Risk Management Framework assessment and pursue Authority to Operate at the CUI impact level. Establish the security architecture required for operational deployment within the DLA IT environment.

Operational Pilot: Deploy ALICE to a minimum of three DLA activities for a minimum six-month operational pilot. Conduct longitudinal evaluation against the key performance parameters defined in the ALICE architecture.

Machine Learning Maturation: Transition readiness indicator computation from rule-based to ML-learned prediction using the Phase II historical snapshot corpus, pending model validation against the rule-based baseline. Deploy a natural language scenario initiation layer to reduce adoption friction for field-level supervisors. Introduce ML-augmented PII anomaly detection at the ingestion boundary as a defense-in-depth layer, subject to DLA Privacy Officer acceptance of the classifier's validated false negative rate. Evaluate advanced generative model architectures for tail-fidelity improvement on rare workforce subpopulations.

Phase III Commercialization Path: The ALICE architecture developed for DLA has direct application across the defense enterprise and the broader federal civilian workforce. Phase III commercialization targets DLA enterprise-wide deployment, extension to other U.S. Department of Defense combat support agencies, and adaptation for federal civilian agencies with large geographically distributed workforces where workforce readiness directly affects mission performance.

VIII. Why SOL

SOL brings a combination of credentials that is rare in the defense technology space: deep operational experience in defense logistics at the command level, sustained expertise in defense policy research and enterprise analytics, and active engagement in federal proposal development and technology assessment for DoD clients.

The firm's founder commanded a Logistics Readiness Squadron at Ramstein Air Base, Germany, the largest logistics readiness organization in U.S. Air Forces in Europe. That experience provides direct insight into the workforce dynamics that the ALICE system is designed to address: the brittleness of specialized certification structures, the compounding effect of simultaneous personnel turbulence across a multi-functional organization, and the gap between what workforce dashboards show and what operational leaders actually need to know. SBIR topic DLA26BZ03-NV011 is not an abstract research challenge for SOL; it describes a problem the firm's principals have managed firsthand.

SOL's research practice, operating under the STORM multi-perspective methodology for defense policy analysis, has produced white paper series on contested logistics, multi-domain operational effects, and AI-enabled acquisition for defense audiences. The firm's analytical register and commitment to evidence-grounded research translate directly into the ontology construction, model calibration, and validation methodology that Phase I requires.

The firm's active work in federal proposal development and GovCon market intelligence provides continuous visibility into the U.S. Department of Defense's stated priorities, funding mechanisms, and transition pathways. SOL understands not only how to build the ALICE system but how to position it for the transition to operational deployment and enterprise-scale adoption that Phase III requires.

Conclusion

The Defense Logistics Agency's workforce is an operational variable, not an administrative function. In the strategic environment that great power competition has created, the capacity to manage that variable with the speed, foresight, and analytical precision that the mission demands is a readiness requirement, not an administrative preference. Current tools do not meet that requirement.

ALICE is SOL's answer to that requirement. It is a purpose-built Adaptive Labor Intelligence and Capacity Engine that gives DLA leadership a persistent, simulated representation of the organization; a privacy-preserving synthetic data environment that eliminates the Privacy Act constraint on sophisticated workforce analysis; and a scenario planning capability that allows leaders to test futures before committing to them. It is grounded in the DLA strategic plan, aligned with the U.S. Department of Defense's 2026 AI Acceleration Strategy, and built on a mature and validated technical foundation.

Phase I will demonstrate that the architecture is feasible, the data can be generated with sufficient fidelity and privacy protection, the simulation produces accurate readiness forecasts, and field-level leaders can use the tool without analyst

intermediation. Phase II will deploy it. Phase III will extend it across the defense enterprise.

The workforce that sustains the joint force in a high-intensity conflict deserves a management system as capable as the mission it supports. ALICE is that system.

Footnotes

1. Defense Logistics Agency, 'The Defense Logistics Agency's Role in Overcoming the Challenges of Contested Logistics,' Defense AT&L (November-December 2024), Defense Acquisition University, https://www.dau.edu/library/damag/nov-dec2024/thedefenselogisticsagencysrole.

2. Cynthia Cook, 'Power Projection and the Logistics of Modern War,' Center for Strategic and International Studies (September 26, 2025), https://www.csis.org/analysis/chapter-15-power-projection-and-logistics-modern-war.

3. Defense Logistics Agency, 'The Defense Logistics Agency's Role in Overcoming the Challenges of Contested Logistics,' Defense AT&L (November-December 2024).

4. Brad Bunn, as quoted in Jason Miller, 'How Data, AI Are Cornerstones of DLA's Digital Strategy,' Federal News Network (September 12, 2025).

5. Adarryl Roberts, as quoted in Jason Miller, 'DLA's Foundation to Use AI Is Built on Training, Platforms,' Federal News Network (January 9, 2026).

6. 'DLA Troop Support Teams Pitch AI, Digital Workforce Concepts,' Executive Government (January 12, 2026).

7. U.S. Department of Defense, 'Artificial Intelligence Strategy' (January 9, 2026), https://media.defense.gov/2026/Jan/12/2003855671/-1/-1/0/ARTIFICIAL-INTELLIGENCE-STRATEGY-FOR-THE-DEPARTMENT-OF-WAR.PDF.

8. Ibid.

9. Mariam Jabri et al., 'Digital Twin-Driven Decision Support for Human Capital Management Under Uncertainty,' Decision Support Systems (February 2026), https://www.sciencedirect.com/science/article/pii/S2949736126000199.

10. Data Foundation, 'OPM's Modernized Federal Workforce Data Platform' (January 8, 2026), https://datafoundation.org/news/press-releases-and-statements/796/.

11. Charles Macal and Michael North, 'Introductory Tutorial: Agent-Based Modeling and Simulation,' Proceedings of the Winter Simulation Conference 2014, IEEE (2014), pp. 6-20.

12. Cynthia Dwork et al., 'The Algorithmic Foundations of Differential Privacy,' Foundations and Trends in Theoretical Computer Science 9, no. 3-4 (2014): 211-407.

13. U.S. Office of Personnel Management, 'Privacy Act Issuance for the Office of Personnel Management, 2025,' Federal Register (January 1, 2025), https://www.govinfo.gov/app/details/PAI-2025-OPM.

14. U.S. Department of Defense, 'Artificial Intelligence Strategy' (January 9, 2026).

15. Jiaxin Chen et al., 'AI-Driven Digital Twins for Manufacturing: A Review Across Hierarchical Manufacturing System Levels,' Sensors 26, no. 1 (December 24, 2025): 124, https://doi.org/10.3390/s26010124.

16. Joseph T. Bonivel Jr. and Solomon Wise, 'Call for a Federal Digital Twins Strategy,' Atlantic Council (February 2024), https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/call-for-a-federal-digital-twins-strategy-unlocking-the-potential-for-digital-twins-in-the-federal-enterprise/.

17. 'Maximus Integrates Skyral's Digital Twin Technology into Federal Operations,' Maximus (March 31, 2026), https://maximus.com/insights/maximus-skyral-digital-twin-federal-operations.

18. Jabri et al., 'Digital Twin-Driven Decision Support for Human Capital Management Under Uncertainty,' Decision Support Systems (February 2026).

19. Joshua Epstein and Robert Axtell, Growing Artificial Societies: Social Science from the Bottom Up (Washington: Brookings Institution Press, 1996); Charles Macal and Michael North, 'Introductory Tutorial: Agent-Based Modeling and Simulation,' 2014.

20. SYN-DIGITS: A Synthetic Control Framework for Calibrated Digital Twin Simulation, arXiv:2604.07513 (2026).