ADL Enterprise Learner Record Repository System Architecture Report

21 October 2021

This work was supported by the U.S. Advanced Distributed Learning (ADL) Initiative (24361820D0001). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ADL Initiative or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
06-14-2021	Document	09/21/2020- 09/20/2021

4. TITLE AND SUBTITLE

ADL Enterprise Learner Record Repository System Architecture Report

5a. CONTRACT NUMBER

FA701420STS10

5b. GRANT NUMBER

N/A

5c. PROGRAM ELEMENT NUMBER

N/A

6. AUTHOR(S)

Flamer, Eric
Adusei-Poku,Adjoa
Franki Schafrik

5d. PROJECT NUMBER

0603769D8Z

5e. TASK NUMBER

N/A

5f. WORK UNIT NUMBER

N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Deloitte Consulting LLP
1919 North Lynn Street
Arlington, VA 22209

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Support Services Center
Advanced Distributed Learning Initiative
13501 Ingenuity Drive, Suite 248
Orlando, Florida 32826

10. SPONSOR/MONITOR'S ACRONYM(S)

OUSD/P&R/DSSC/ADLI

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution A

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The Enterprise Learner Record Repository (ELRR) System Architecture Report (SAR) outlines the ELRR solution architecture by describing the specification and deployment of hardware and software, interfaces, and port/protocol usage. This report focuses on the standards, specifications, services, and systems necessary to enable individual DoD organizations to contribute their local, authoritative learning records to an enterprise aggregation of these local, authoritative training records, using a growing number of supporting data services (e.g., course information, identity management, competency data) to federate a learner's portfolio of records into an ELRR capability. The intended audiences for this document are developers or engineers who are familiar with TLA standards and core cloud computing principles.

15. SUBJECT TERMS

Learner Records, System Architecture, Enterprise

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Joe Camacho
U	U	U	UU	36	19b. TELEPHONE NUMBER *(Include area code)* 757-572-6206

Reset

Deloitte.

System Architecture Report in Support of:

Advanced Distributed Learning (ADL) Initiative

Enterprise Learner Record Repository Development

Document Control Information

Document Name	System Architecture Report
Client	Advanced Distributed Learning Initiative
Task Order Name	Enterprise Leaner Record Repository Prototype Development
Document Owner	**Eric Flamer, Adjoa Adusei-Poku, Franki Schafrik (Deloitte Consulting LLP)**
Document Version	2.0
Document Status	Baseline

Document Edit History

Version	Date	Additions/Modifications	Prepared/Revised by
1.0	11/16/2020	Baseline submitted to ADL Initiative	Deloitte Consulting LLP
1.1	12/03/2020	ADL Initiative Revisions	ADL Initiative
2.0	6/14/2021	Final submitted to ADL Initiative	Deloitte Consulting LLP

Distribution of Final Document

The following people are designated recipients of the final version of this document:

Name	Organization/Title
Dr. Sae Schatz	Contracting Officer Representative
Ashley Reardon	Government Technical Point of Contact
Brent Smith	Government Technical Lead

Table of Contents

1.0 Solution Overview

In July 2018, the Department of Defense (DoD) Chief Management Officer (CMO) and the Reform Management Group formally initiated the Enterprise Digital Learning Modernization (EDLM) reform initiative. The goal of the EDLM initiative is to build an enterprise-wide integrated digital learning ecosystem that enables efficient acquisition and spending management for DoD education and training products and services. The Enterprise Learner Record Repository (ELRR) is one of the three EDLM lines of effort supported by the Advanced Distributed Learning (ADL) Initiative.

The end goal of this effort is to prototype a system that demonstrates a minimum viable product (MVP) that is 1) integrated into the ADL Initiative's Total Learning Architecture (TLA) DevSecOps sandbox environment, 2) to connect to three different DoD Learning Record Providers (LRPs) to capture various learner record use cases, evidentiary chains, and credentialing/competency frameworks, and 3) to use the ADL Initiative's DATASIM project to stream simulated learner data across the ELRR data fabric.

The prototype will demonstrate at scale (i.e., thousands of simulated) learner records aggregated from multiple and disparate (simulated) digital learning-delivery platforms in a sandbox environment replicating real DoD systems. The learning records data will also federate with semantic information (e.g., competencies, credentials) and simulated profile data (e.g., learner attributes, individual identities). The resulting records will maintain their data lineage at the enterprise level and be collectively interrogable from a single human-usable portal.

Figure 1: ADL Initiative ELRR System Architecture Overview

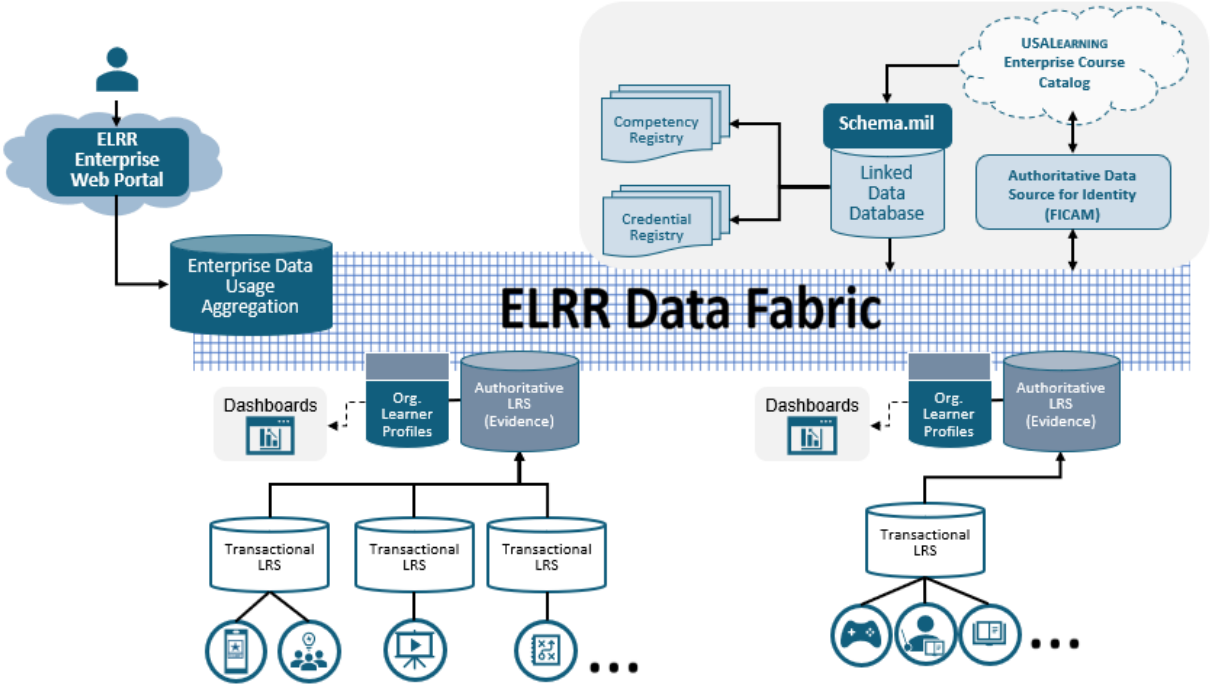

For the Prototype, ELRR will not pull data directly from Integrated Product Team (IPT) edge learning systems. Instead, the ADL Initiative's designated TLA Learning Record Store will serve as the consolidated learner record for the prototype using simulated data. A process will be needed to create a

backlog for intake of additional learner data when a DoD organization has a unique data requirement that differs from the xAPI standard and takes place outside of a Learner Record Store (LRS).

1.1 ADL Initiative Enterprise Architecture

The ADL Initiative's TLA data strategy includes three functional pillars—information about specific runtime performance outcomes (Learning Activity), a description of the domain in general (Experiences), and individuals' aggregate records and personal descriptors over time (Learner Profile) —connected via a fourth pillar (Competencies) that creates a "common currency" for semantic interoperability across the other pillars. In an operational TLA system, the creation of authoritative statements would be logic driven, including layers of business logic encoded in the Competency Framework. Consequently, ELRR's scope includes all four data pillars of the ADL Initiative's TLA data strategy, highlighted in Figure 2 below. The business objective is to develop a method – via a data fabric and microservices application – for connected and searchable record repositories of detailed learner performance data to be made available to any command, learning system, or activity across the DoD. This will support adaptive instruction, improved decision making, and analytical insights into learners and the systems they interact with.

Figure 2: ADL Initiative TLA Data Strategy

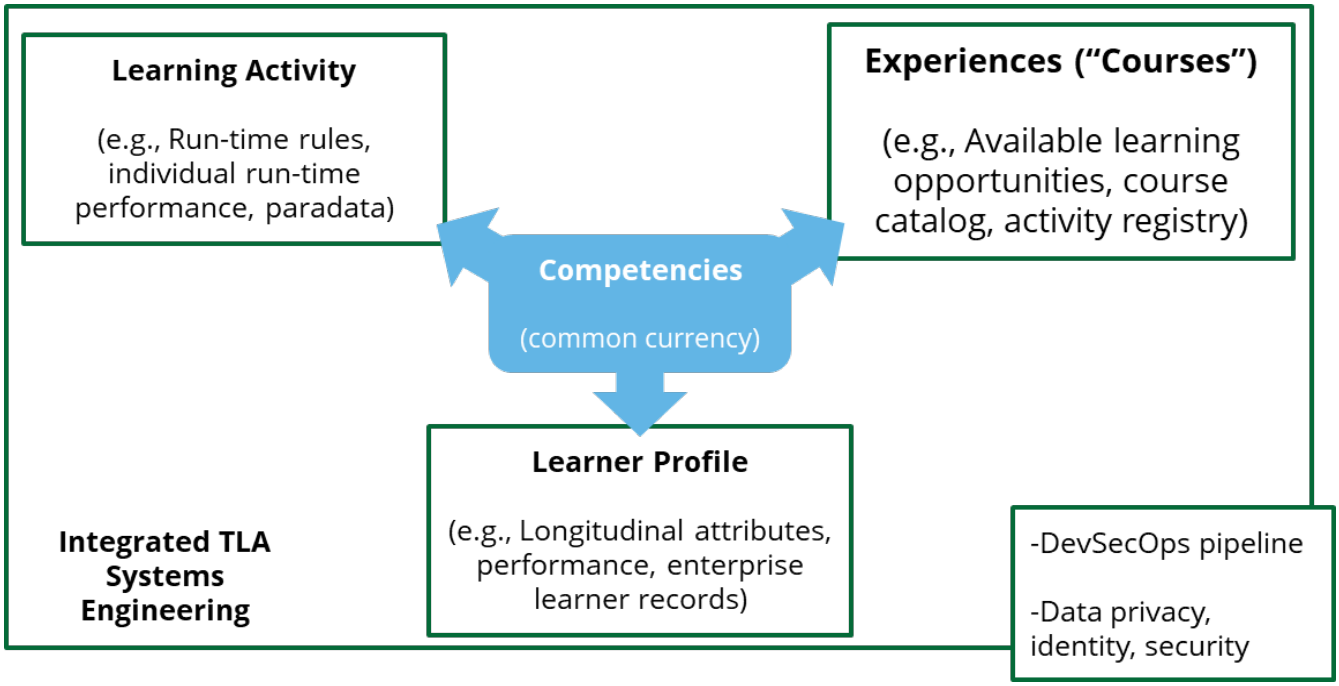

Over the last several years, the ADL Initiative has shepherded the greater adoption of capturing learning activities with the Experience Application Programming Interface (xAPI), contributing to a data set with formatted consolidated learning event and learner state records, adhering to the JavaScript Object Notation (JSON) open standard file format and data interchange format defined for the xAPI (IEEE P9274). Critical to ELRR, the ADL Initiative is focused on establishing tools, technologies, and policies that maximize the adoption of xAPI to enable the collection of learner-related data to support enterprise learner analytics. In its 2019 TLA report, the ADL Initiative noted that there were several key factors hindering the quality of an enterprise learner record data set that would allow TLA to meet its business objectives, including: inconsistency in statement naming patterns, missing statement fields, a loss of

traceability into competency management, inconsistency in labeling competency levels, and difficulty describing relationships between external competency documentation, internal competency levels, and learning resources.

Based on the xAPI standard, adjudicated performance records are routinely collected in a LRS, the system with specific data storage and retrieval capabilities for xAPI statements. However, these statements typically exist to capture the first and primary control loop for abstracting and organizing learner records throughout the continuum of learning, capturing the learner activity.

Assertions include other information such as physical/psychological/behavioral attributes, personal preferences, and competencies not associated with a credential. Assertion of competence – the progress made based on the mapping of the learning event towards a defined credential – is based on the combination of individual learner records, the environment in which the learning activity occurred, and represents the second control loop; multiple credentials together form competency-based assertions in relation to a learner's job, duty, gig, or assignment.

Taken together across multiple jobs, duties, gigs, or assignments in a learner's lifecycle of work and training performance, the sum total of learning activities, learning credentials, and the current job form the basis of a career trajectory, which may or may not include a separate career trajectory (e.g., Control Loop 5) should a learner entertain or embark on a new career trajectory unrelated to the expected one described above. The data exchange across this control loop hierarchy, demonstrated in Figure 3 below, forms the abstract basis of the evidentiary chain critical to the reference implementation of ELRR.

Figure 3: Learning Data Exchange Across Organizational Echelons

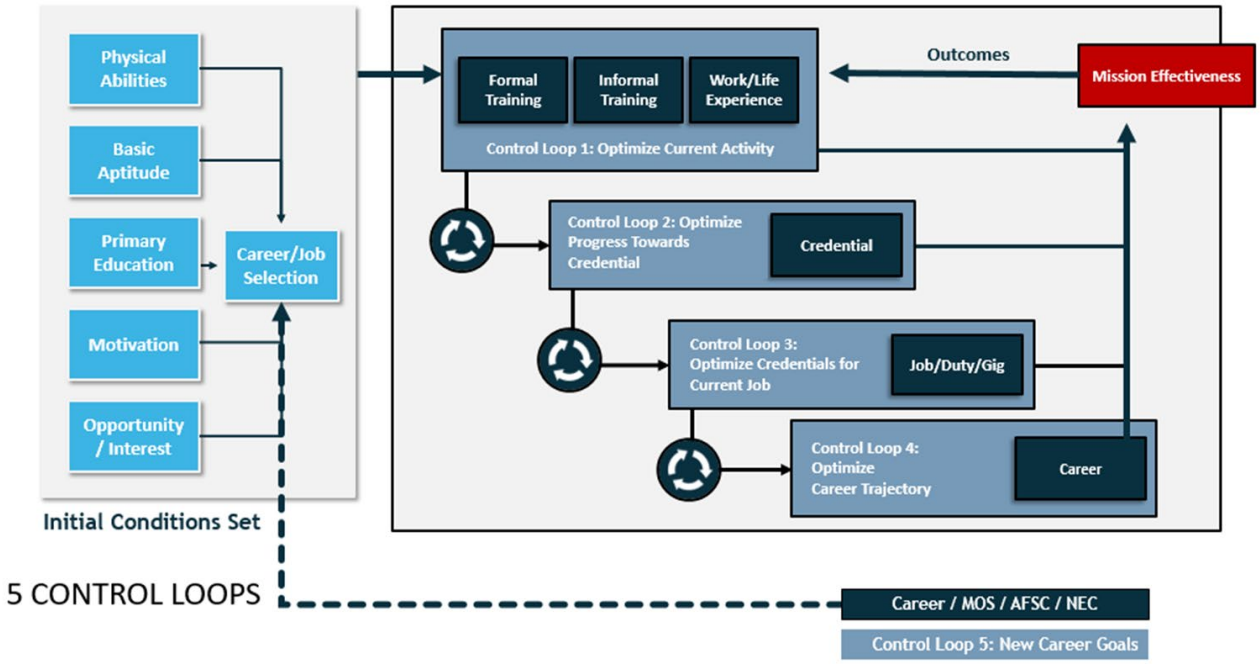

In addition to this persistent reconciliation of credential, job, and career information that "rolls up" at the enterprise level based on individual learner activity records, the evidentiary chain also includes the reconciliation and management of the individual local learner profiles that, taken together, contribute to the view of an enterprise learner record. Again, this is important given that a great many learners will have or maintain a local learner profile associated with one or more assignments taking place across

organizational boundaries and across edge systems. In the management of federated learner identities, one challenge has been that not all systems store user information the same way. Learners may also use different local account names for different purposes. There is no guarantee these accounts, or systems know about each other, so a federated approach to Identity Management is required to resolve this issue. Taken together, the identity and credentialing evidentiary chains follow the pattern from local learning record systems to the ELRR, as demonstrated in Figure 4.

Figure 4: Demonstration of the Evidentiary Chain Across Local and Enterprise Systems

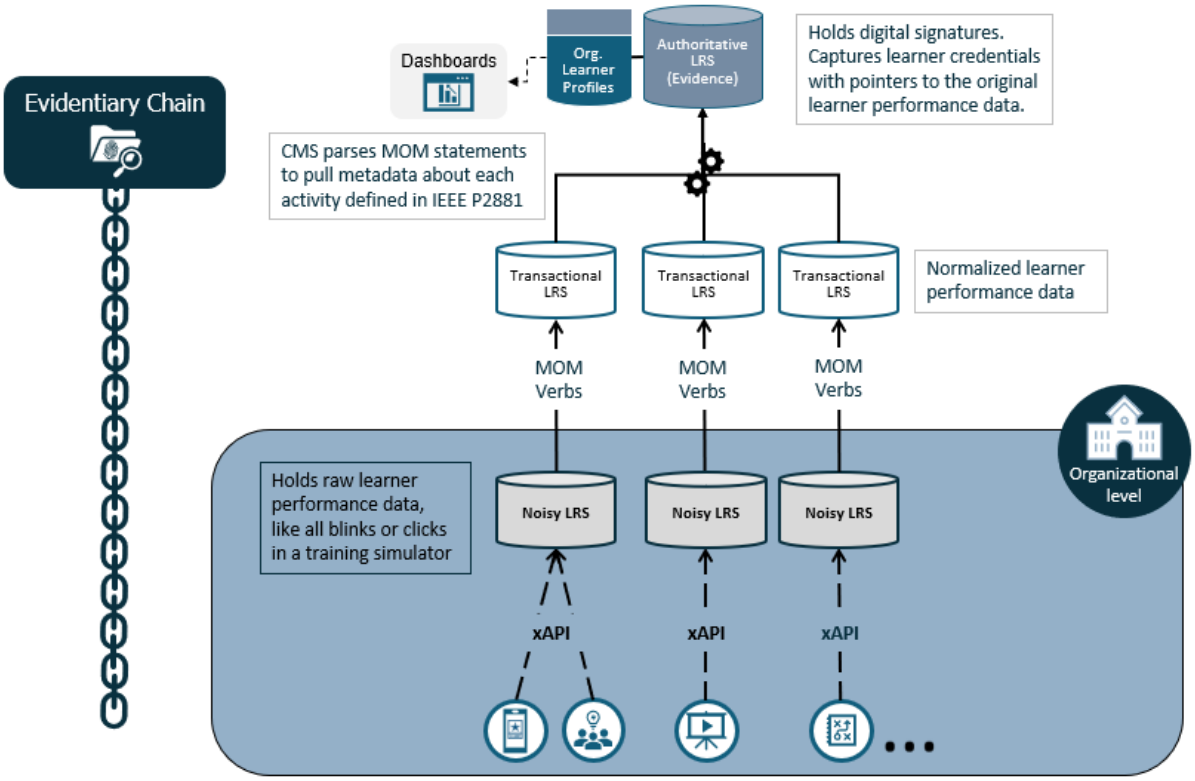

1.2 ELRR Solution Architecture Assumptions

Our efforts work will focus on the standards, specifications, services, and systems necessary to enable individual DoD organizations to contribute their local, authoritative learning records to an enterprise aggregation of these local, authoritative training records, using a growing number of supporting data services (e.g., course information, identity management, competency data) to federate a learner's portfolio of records into an ELRR capability. To meet these goals, the Deloitte team began development on the ELRR Prototype with the following assumptions related to the overall solution and resulting system architecture:

- Learner records within the DoD have varying levels of centralization and standardization
- The ADL Initiative's current authoritative learning record store will serve as the consolidated learner record for the prototype
- The prototype will use simulated learner records
- Learner records are formatted using the xAPI specification

- Courses (e.g., formal learning experience) are the only learning activities used in the prototype
- The prototype uses an approved sub-set of the Draft 1484-2 Integrated Learner Record v4_20201111
- The prototype will support the enterprise learner, training manager, career manager, and system administrator roles/personas, as defined in this document

1.3 ELRR Solution Guiding Principles

In developing the prototype, attention is paid to the standards, specifications and guiding principles as found in the table below.

Table 1: ELRR Solution Guiding Principles

Guiding Principles	Rationale
Use Open Source First, Where Possible	The team should avoid use of licensed software to meet the PWS requirements; the spirit of this development effort, like all ADL Initiative projects, is to avoid proprietary software dependencies while ultimately encouraging broad community participation and adoption.
Remain Cloud Agnostic, Where Possible	The technical/development team should avoid the use of licensed or proprietary cloud engineering software to meet the PWS requirements. Along with the costs associated with this software, the ADL Initiative anticipates a future migration to a different cloud environment from the one currently in use.
ELRR Solution is Evolutionary in Nature	The prototyping process is one used to accelerate capability development; our agile software development approach will provide ADL Initiative and the ELRR stakeholders with the ability to evaluate working code during this period of performance. As a result, the solution and the process behind its development should remain flexible as the overall ELRR team implements the initial set of requirements and learns lessons on the outcomes from this approach.
Use of Existing TLA Services, Where Possible	The technical/development team continues to review existing/developed core services that are described in the TLA reference implementation. For example, Competency and Skills System (CASS) services are critical to the ELRR effort to process and align evidence of learner performance against competency and credentialing frameworks. The team will avoid creating duplicative services to meet the requirements of ELRR if the service already exists.
Draft IEEE Standards Provide Guidance to ELRR Data Model	The draft Integrated Learner Record and Learner Activity Metadata data models are important data standards, currently undergoing review led in part by the ADL Initiative. The updated standards include required, recommended, and optional data fields that match some of the attributes expected in current and future learner records. The ELRR should leverage these data standards only to the extent that they are needed to meet the requirements so that ELRR works in accordance with these data models. The ELRR prototype will adhere to an agreed upon sub-set of data attributes as defined in the Draft 1484-2 Integrated Learner Record v4_20201111.
Future Design Considerations for Cybersecurity	The ELRR Prototype will conform to a certain degree of security requirements, but ultimately will use simulated learner record data because the key outcome of the prototype is the successful demonstration of the evidentiary chain, rather than full or partial attainment of future security level. The development team will work with the ADL Initiative to document security-related requirements as the prototype progresses. document these requirements, prioritize them as required.

1.4 ELRR Solution Architecture Attributes

The architecture is developed in adherence to enterprise architecture principles and best practices. The following non-functional requirements will be incorporated into the architecture.

Table 2: Solution Attributes for the ELRR Prototype

Architecture Attribute	Rationale
Common Components	As described above, we are planning for the use of open-source technologies to avoid custom-built implementations to the degree possible. We have included recommended development tooling and solution environment technologies in Appendixes A and B, respectively, that highlight this approach.
Portability	The development team will leverage a schema service to transform metadata from one metadata format to another, maximizing the portability of data across the ELRR solution. As described above, we plan to remain cloud agnostic and limit proprietary software use so as to not lock the ELRR solution into any single vendor.
Capacity	One of the hallmarks of the planned ELRR solution is the Elasticsearch search engine, which can demonstrably scale to billions of records. This search tool and associated capability can also be used for the ELRR. Our overall technical solution will use cloud solutions and infrastructure to provide additional avenues for solution capacity and scale.
Performance	The team intends to leverage known/recognized technologies (e.g., Kubernetes, Docker), which are best of breed solutions that are known to scale and perform well at a high-level.
Availability and Reliability	Beyond the portability, capacity, and performance attributes described above, the development team will work within the existing practices defined by the ADL Initiative and in use in the deploying the ADL Initiative's TLA Sandbox environment.
Scalability	The prototype needs the capacity to process increasing amounts of data, in the form of additional learners, additional learner profiles/records, additional learning activities, etc.
System Management, Monitoring, and Administration	The team has identified a number of PWS requirements that pertain specifically to the system management, monitoring, and administration activities. We have also specifically planned for several system workflows that support this attribute, including system installation, configuration, validation, monitoring, and interoperability.
Business Continuity/Disaster Recovery and Continuity of Operations	Beyond the capacity, performance, and availability/reliability attributes described above, the development team will work within the existing guidelines and practices defined by the ADL Initiative and used in maintaining the ADL Initiative's TLA Sandbox environment.

2.0 Solution Architecture

The ADL Initiative authoritative LRS, hosted within the ADL Initiative TLA sandbox environment will interface with the prototype to retrieve the consolidated learner records. The criterion for success of the ELRR Prototype will be a demonstration, at scale, using thousands of simulated learner records from the ADL Initiative authoritative LRS. The ELRR prototype architecture is a three-tiered architecture comprised of the presentation tier, the application tier, and the data tier. Within each tier, the components that comprise that tier are identified and described below. Please note that solidly-colored components listed below represent current/expected components as part of the MVP, while shade-colored components represent services that will be or may be added in development beyond the MVP.

Figure 5: ELRR Prototype Architecture

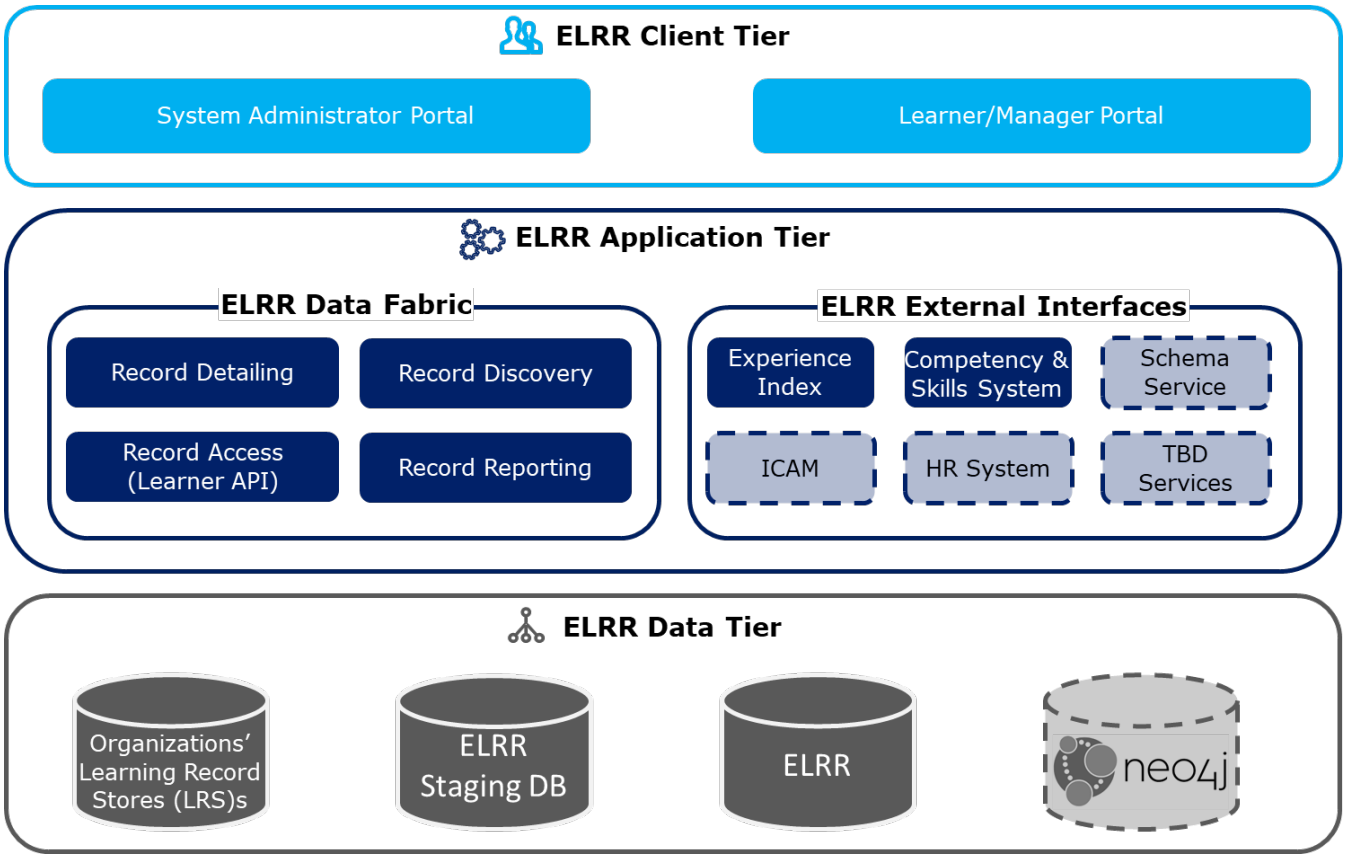

2.1 ELRR Solution Personas

The Deloitte development team, consulting with the ADL Initiative's provided technical guidance, has identified the following solution personas based on an analysis of the ADL Initiative Performance Work Statement (PWS), as well as additional stakeholder interviews and design sessions conducted at the beginning of the project. It is expected that the personas for the system will evolve as additional stakeholder informational and working sessions are held over the course of the project.

- **Enterprise Learner:** all personnel with a learning record acquired by either formal or informal means. The learner may view their training/learning from various institutions, in aggregate form, within the ELRR.

- **Training Manager:** responsible for ensuring that Learners meet their local and general training requirements. The Training Manager also monitors completion of courses and expiring courses, while reviewing competency or credential status of assigned learners.
- **Career Manager:** subject-matter expert who identifies the training requirements of a career field. The Career Manager is a more strategic role that will plan for personnel readiness within a career field. This may include participating in progression and succession planning to ensure that there are sufficient personnel with the necessary skills and training to accomplish the mission. Unlike the Training Manage, the focus of this persona is to survey an assigned organization for competency or credential information, rather than drill down at the learner level for this type of information.
- **System Administrator**: responsible for validating and monitoring the ELRR system. The system administrator will view configuration data stored in the ELRR application, as well as monitor imports from the authoritative LRS, and from other external services such as, ELRR, CaSS, etc.

2.2 ELRR Solution Components

The ELRR **Data Tier** contains the persistent storage (relational database management system) and has a data access layer which enables the data to be independent of the application logic. In the future, it may also contain an open-source graphical database, (e.g., Neo4j), as the prototype addresses data analytics and reporting. The following components comprise the data tier:

- **Organizations' Learning Record Store (LRS)** - This component is not part of the ELRR however, it is included since there are essential interactions with the LRS. The ADL Initiative LRS stores the consolidated xAPI statements, from learning record providers and listening services within the TLA core, (simulated data for prototype), according to the TLA Master Object Model (MOM) (IEEE P9274.3.1) standard. Learning event and learner state records are in the JSON format defined for the xAPI (IEEE P9274). For the purposes of the ELRR Prototype, this will be the ADL Initiative's Authoritative LRS, hosted in the TLA environment. Although the LRS is not contained within the ELRR prototype's data tier, the ELRR prototype will develop data services to retrieve the data and store within the ELRR Staging Database. In the future, this component will accommodate other DoD organizations' learner record stores (e.g., Navy's MyNavy Learning and Army's ATIS LRS).
- **ELRR Staging Database** – a Postgres database used to store the learner information retrieved from the learning record store(s).
- **ELRR Database** – a Postgres database used to store the aggregated learner records.

The ELRR **Application Tier** contains the business or functional logic for the ELRR. At the TLA's enterprise level, application services will be advertised through a service registry. The application tier is comprised of the following components:

- **ELRR Data Fabric** includes microservices designed to support the IEEE ELR standard, which enables interoperability with numerous other learner record standards in use today. This will include the Learner API, published to allow 3rd party software systems, tools, or platforms to access and use ELRR data. The ELRR Learners API is discussed further in Appendix C as part of the Interface Control Document. Specific microservices for consideration include:

- **Record Detail -** required to establish the evidentiary chain linking individual local learner profiles and individual local learner records to the enterprise learner profile and learning record.
- **Record Discovery -** required to efficiently locate learning records with the highest possible level of accuracy and efficacy.
- **Record Access (Learner API) -** required to direct a third-party system, tool, or platform to access ELRR data. The ELRR Learners API, as described in Appendix C Interface Control Document (ICD), is used to synchronize the data in the ELRR with the data in the local learning systems.
- **Record Reporting –** planned for ongoing/future enhancement and used to facilitate various Enterprise Learner Record queries at scale. This may be used to support specific dashboard views in the Client Tier portal services, or for specialized Learner API queries requiring numerous, detailed steps.

- **ELRR External Services –** non-ELRR-based APIs containing information that is relevant to the enterprise learner record and contributes to the evidentiary chain, the enterprise learner profile, and the overall data services provided by ELRR. Identity Credential and Access Management (ICAM), HR Systems, Schema Service, and any additional data services are all out of scope for the MVP.

Finally, the ELRR **Presentation Tier** represents the user-facing aspect of the system designed to enable human or external system/machine interactions. The ELRR portal provides role-based access for learners, training managers, career managers and system administrators. It provides a common access point for accessing all information about learners, career fields, occupational specialties, and analytics derived from the different groups of learners. In the future, a series of dashboards will be incorporated into the ELRR that will demonstrate how learner data may be used at different stages to support different stakeholders. The presentation tier consists of the following components:

- **Learner/Manager Portal** – user interface for enterprise learners, training managers and career managers to view their aggregated learning records as well as limited learning record data for their reports.
- **System Administrator Portal** – user interface for system administrators to monitor imports and configurations; allows the system administrator to monitor overall flow of enterprise learner records through the data fabric to observe potential issues arising when adding additional data services for detail of source data records.

2.3 ELRR Component Architecture

Figure 6: “ELRR Prototype Components” illustrates the components in both the extraction of the learning records from the LRS, as well as the aggregation of those learner records. Below are descriptions of the components.

Figure 6: ELRR Prototype Components

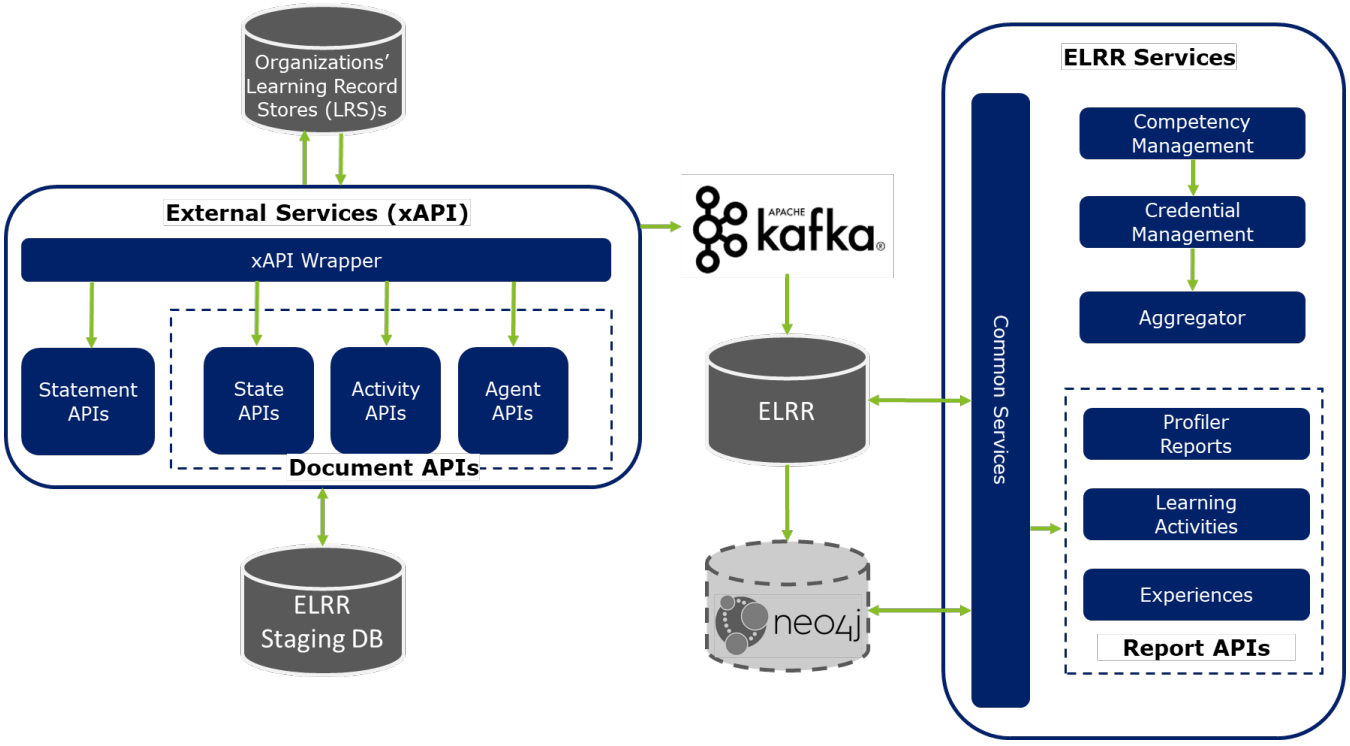

External Services (LRS Services)

- **xAPI Wrapper**: Rest-based xAPI clients used to simplify the connecting to and communicating to an LRS. It encapsulates the xAPI Verbs project, allowing a single object to contain both the ADL Initiative verbs and the ADL Initiative xAPI Wrapper.
- **Statement APIs**: xAPI statements used to retrieve records from the LRS. It also defines requirements for the storage and retrieval of xAPI statements. The xAPI statement structure can be expressed in the form of "actor verb object", which conveys learning activity in a simple JSON format.
- **State API**: Document API to provide internal state of learner Activity that do not have their own internal storage or need to persist state across devices. Stores, fetches, or deletes the document specified by the given State ID that exists in the context of the specified Activity, Agent, and registration
- **Activity (Profile) API**: Document API, similar to the State API, it allows arbitrary key/document pairs, related to a learning activity, to be saved. It also includes a method to retrieve a full description of an activity from the LRS.
- **Agent (Profile) API**: Document API, similar to the State API, it allows arbitrary key/document pairs, related to an agent, to be saved. It also includes a method to retrieve a special Object with combined information about an Agent derived from an outside service, such as a directory service.

ELRR Services

- **Competency Management (TLA Core Service)** – ELRR will use the ADL Initiative hosted competency management services which track overall competency state for selected goals and

makes assertions based on evidence provided via xAPI. An instance of the Competency and Skills System (CaSS) hosted the competency framework definitions in the 2019 implementation.

- **Credential Management (TLA Core Service)** – ELRR will use an ADL Initiative-determined credential management service which track overall credential state for selected goals and makes assertions based on evidence provided via xAPI.
- **Aggregator** – consolidates a learners' enterprise learner records to include learning activities undertaken, competencies attained, and credentials conferred throughout a learner's career and from differing DoD organizations.
- **Report APIs** – Although this is not in scope for the MVP, the Report APIs are used to support dashboard reporting and data analytics.

2.4 ELRR Solution Workflows

While many consider workflows to only represent those operations embodied by a system, it is important to also consider and include all pertinent manual operations needed to complete the required work. The following high-level workflows are identified based on analysis of the ADL Initiative's system requirements.

Table 3: Proposed ELRR Solution Operational Workflows

Workflow	Description
System Installation	Automated operations required to download, extract, and deploy the system components comprising the ELRR solution.
System Configuration	Automated operations required to enable the system to interact with both internal components and external systems.
System Monitoring	Automated operations required to enable observation of system activity and resource consumption.
Record Detailing	Automated operations required to establish the evidentiary chain linking individual local learner profiles and individual local learner records to the enterprise learner profile and elearning record.
Record Discovery	Automated operations required to efficiently locate learning experience records with the highest possible level of accuracy and efficacy.
Record Access	Automated operations required to direct a third-party system, tool, or platform to access ELRR data.
Record Reporting	Automated operations required to gain insights into learners' learning activities and over-all learning/training trends.

3.0 Data Architecture

The Data Architecture will be following applicable ADL Initiative and DoD data standards and will be defined in compliance with the data architecture standards for the ELRR system. The Data Architecture will adhere to the below listed standards:

- IEEE P9274.1 Experience API (xAPI) 2.0 – Learning activity tracking uses the xAPI to capture learning activity streams. The xAPI standard also includes xAPI profiles such as cmi5 and the TLA's Master Object Model. xAPI 2.0 is targeted for approval in 2020.
- IEEE P2997 Enterprise Learner Records (ELR) or IMS Global Comprehensive Learner Record - Learner profile standards do not currently meet all TLA requirements. These new standards are actively being developed and modified based on input from numerous industry groups and associations.
- IEEE 1484.20.1 Reusable Competency Definitions - The definition of a competency, the relationship to other competencies, and the alignment of evidence to help measure proficiency of the competency, are included in this standard. This standard is expected for approval in 2020.
- IEEE P2881 Learning Activity Metadata – Descriptions of learning activities and their associated content are stored in the TLA's Experience Index and use a modified version of the Learning Resource Metadata Initiative standard. A draft standard is being submitted for finalization in early 2020.

In the ELRR prototype database design description, the development team will develop the logical and physical data models which will have progressively detailed entities, attributes and relationships defined. Key functional requirements such as the minimum set of data elements to support a learner record, the roles involved, the auditing, the monitoring and record update requirements will influence the data architecture. The ELRR conceptual data model (DIV-1) shows the data concepts that will be addressed by the MVP as well as how those concepts relate to each other. The conceptual data model is based on the above learning standards.

Information for the ELRR can come from multiple external sources and will be combined within the system to generate a learner record. Information about courses will come from the ECC. The Course Catalog will identify the Learning Source of the course. The Learning Source will also provide Evidence and Assertions for the course to confirm the learner completed the course and that the completion has be validated by an evidentiary chain. The Learning Source has an associated Organization which is responsible for accrediting the Learning Source.

An Accreditation will be provided to a Course to indicate the intended use of the learning material and the validity of the learning material. A Course will confer a Credential to indicate the outcome of the completion of the course. A Course can also indicate a Competency, which will have a measured Proficiency (e.g., a member has a Competency of data modeling with a proficiency rating such as "experienced"). Competencies are contained within a Framework, with the Framework external to the ELRR.

A Person will have a Profile, which indicates personal, employment, and learning information. The Profile will have a Status to show system related characteristics such as last update, active/inactive, etc. When the Person takes a Course, the Profile will be updated. Employment history is made up of Assignments which have been given to a Person. This information is represented as the conceptual prototype in the figure below, with supporting details on these entities in the corresponding table.

Figure 7: ELRR Prototype Conceptual Data Model

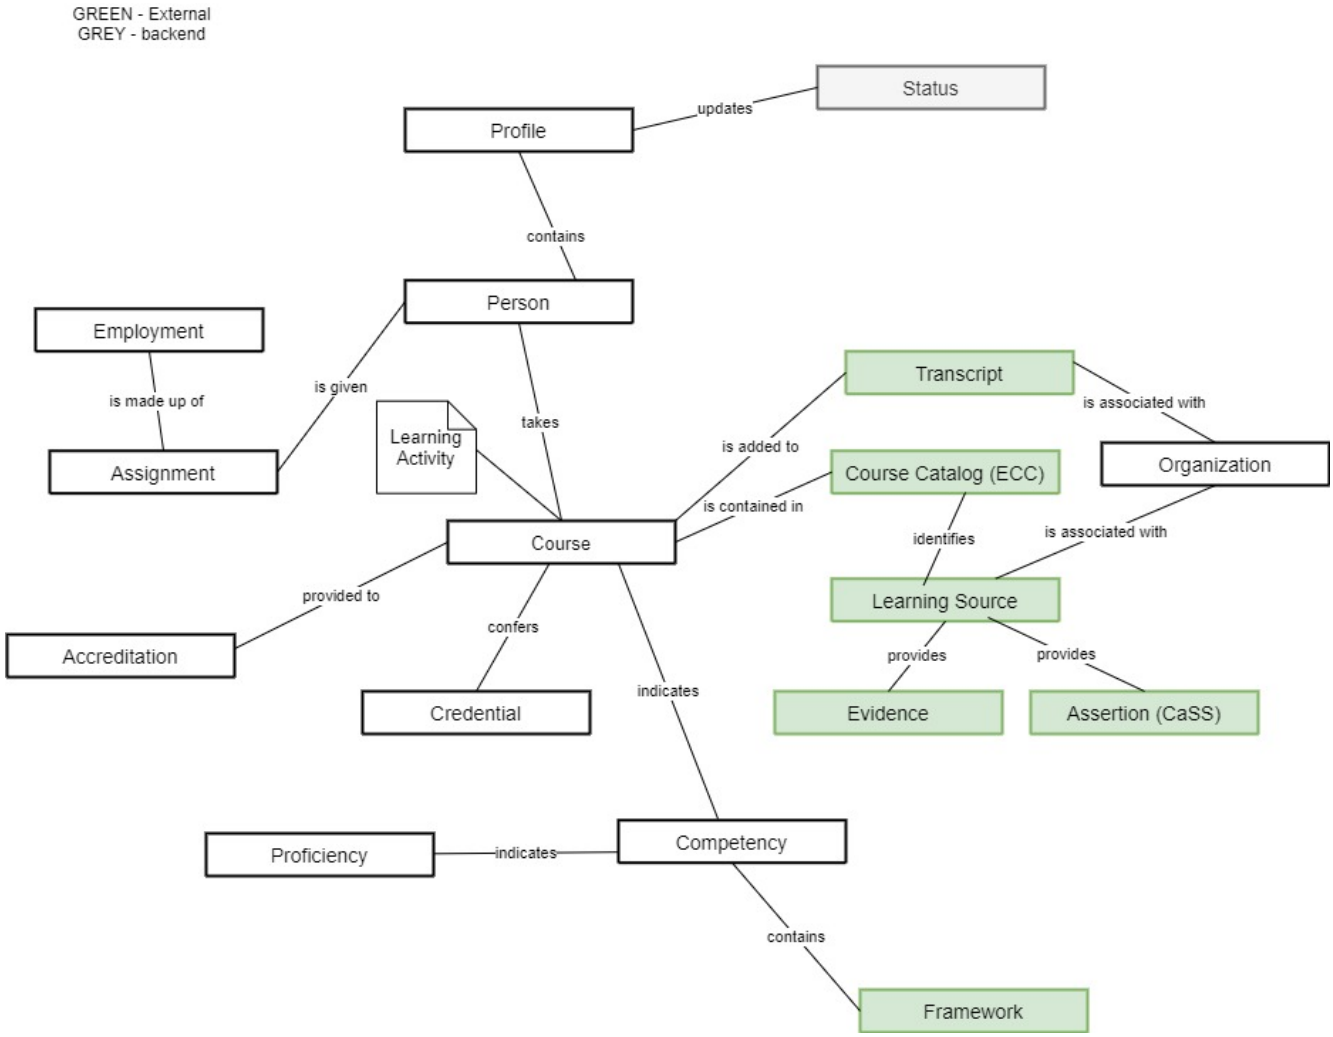

Table 4: Conceptual Data Model Entity Descriptions

Entinty	Description	Source
Accreditation	The official certification that a model, simulation, or federation of models and simulations and its associated data is acceptable for use for a specific purpose.	TLA
Assertion	In competency-based learning, it is a claim about a learner's mastery over a specific competency, or the elements that a competency is comprised of behaviors (e.g., knowledge, skills, abilities, aptitudes).	TLA
Assignment	The permanent change of learner's duty station from one location to another. The service will assign a learner with the necessary skills to valid manpower requirements in order to meet the mission objective.	AF Information Asset definitions
Competency	The set of demonstrable characteristics and skills that are required by an individual or team for performance of a job. Characteristics include knowledge, skills, abilities, and other (KSAO) such as attitudes, aptitudes, motivations, and social elements.	TLA

Entinty	Description	Source
Course	A series of learning activities about a specific subject that are sequenced together into a program of instruction (POI) or (Navy-Course of Instruction - COI). Each activity may include different types of instructional content.	TLA
Course Catalog	A catalog of learning activities.	ELRR Development Team
Credential	An attestation of qualification, competence, or authority issued to individual by a third to a party with a relevant or de facto authority or assumed competence to do so.	TLA
Employment	Service information and/or updates learner's duty status, specialties, experience identifier, and service dates.	AF Information Asset definitions
Evidence	Anything presented in support of an assertion of competency.	TLA
Framework	A model that broadly defines the blueprint for 'excellent' performance within an organization or sector. Generally, the framework will consist of numerous competencies, which can be applied to a broad number of roles within the organization or sector.	TLA
Learner	The generic term for a user in a distributed learning environment that generates data while interacting with elements within that environment.	TLA
Learning Source	Any source that provides a learning activity.	ELRR Development Team
Organization	Identifies various types of unit designators and organizational hierarchy recognized by personnel accounting symbol code.	AF Information Asset definitions
Person	Distinct descriptors of a person, to include name, address, email address, and demographic data.	AF Information Asset definitions
Proficiency	Army: Proficiency is the completeness of achievement of lower-level standards to an overarching standard. 2. Navy: Proficiency is equivalent to level of mastery. 3. Joint: Proficiency is meeting standard.	TLA
Profile	A set of rules and human- and/or machine-readable documentation of application-specific vocabulary concepts, statement patterns, extensions, and statement templates used when implementing xAPI in a specific context.	TLA
Status	The status of a learner record within ELRR. The status is success, partial, failure.	ADL Initiative
Transcript	Records learners' permanent academic records. Typically, a transcript only includes the most basic of information such as courses taken, dates, grades received, and degrees conferred from a formal academic institution.	TLA

4.0 Security Architecture

The security system will follow DoD principles, policies, and regulations, as determined by the ADL Initiative, and ultimately aligned to a timeline or roadmap for security objectives and resulting milestones. There will be different levels of security depending on external vendors and internal DoD personnel having access to the system. Depending on the different sources (governmental and non-governmental) a high-level security solution with detailed description will be a part of future versions of this initial system architecture. Such a security architecture model may include the following depending on the application's detailed security requirements:

- "Least Privilege" principle
- Role based access control (RBAC), with expected future change to policy-based access controls as part of ADL Initiative efforts to implement privacy controls for underlying ELRR data attributes.
- Access and Authorization (including user/identity access management)
- Audit and Accountability
- Cloud Access Point (CAP) Configuration
- Network and Server Configuration (i.e., Ports and Protocols, Security Controls, DoD STIG, and Security Scan requirements)
- Active Directory/Single Sign-On

Interactions with common enterprise security components (e.g., Active Directory/Single Sign-On) will be defined in the future depending on the detailed security requirements. When appropriate, the development team will work with the government cybersecurity point-of-contact (POC) to ensure application development is compliant with DoD cybersecurity principles, policies, and regulations. The level of security will depend on factors such as the type of data utilized in development (i.e., PII, sensitive data, etc.), access required by users, testers and developers, and the security posture of the development environment (Enterprise System ATO, AWS GovCloud, and DoD Impact Level).

We will work with our C5ISR project stakeholders to elaborate on the security architecture requirements, design, and features throughout the ELRR software development activities. Determining the level of security required for development will enable the generation of detailed security plan that will be a part of the overall system architecture. Assumptions include the development team will have root access within the development environment and in coordination with the government cybersecurity POC will be able to design and configure network and servers that will ultimately expedite DevOps security requirements and develop a backlog to drive further cybersecurity integration alongside application development.

5.0 Infrastructure and Deployment

Once the solution components and resulting technologies have been defined, the application architecture deployment models will be mapped to hardware and software infrastructure specifications (e.g., memory and CPU specifications). This information will include software, hardware, network, middleware, and cloud specifications, as applicable. The application architecture deployment models will also include architecture guidance and specifications for all environments required for developing, testing, deploying, and operating the solution. Once the technology and infrastructure has been designed, the deployment models will be aligned to the technology and infrastructure requirements and will be documented in future versions of this document.

ELRR will be deploying components using Docker containers to manage the development of the microservices architecture. Docker containers are portable, scalable, and reliable. ELRR Docker images are stored in the public DockerHub registry to allow anyone to pull and deploy the image. Docker images are also cloud agnostic which allows for deployment on any cloud provider. Images can be deployed as a single Docker image, a cloud provided container service, and Kubernetes for orchestration. ELRR will deploy component images on Kubernetes to orchestrate containers for scalability, reliability, and high availability. Figure 8 below refers to the proposed diagram of the ELRR Amazon Web Services (AWS) Architecture.

Figure 8: ELRR Prototype Infrastructure

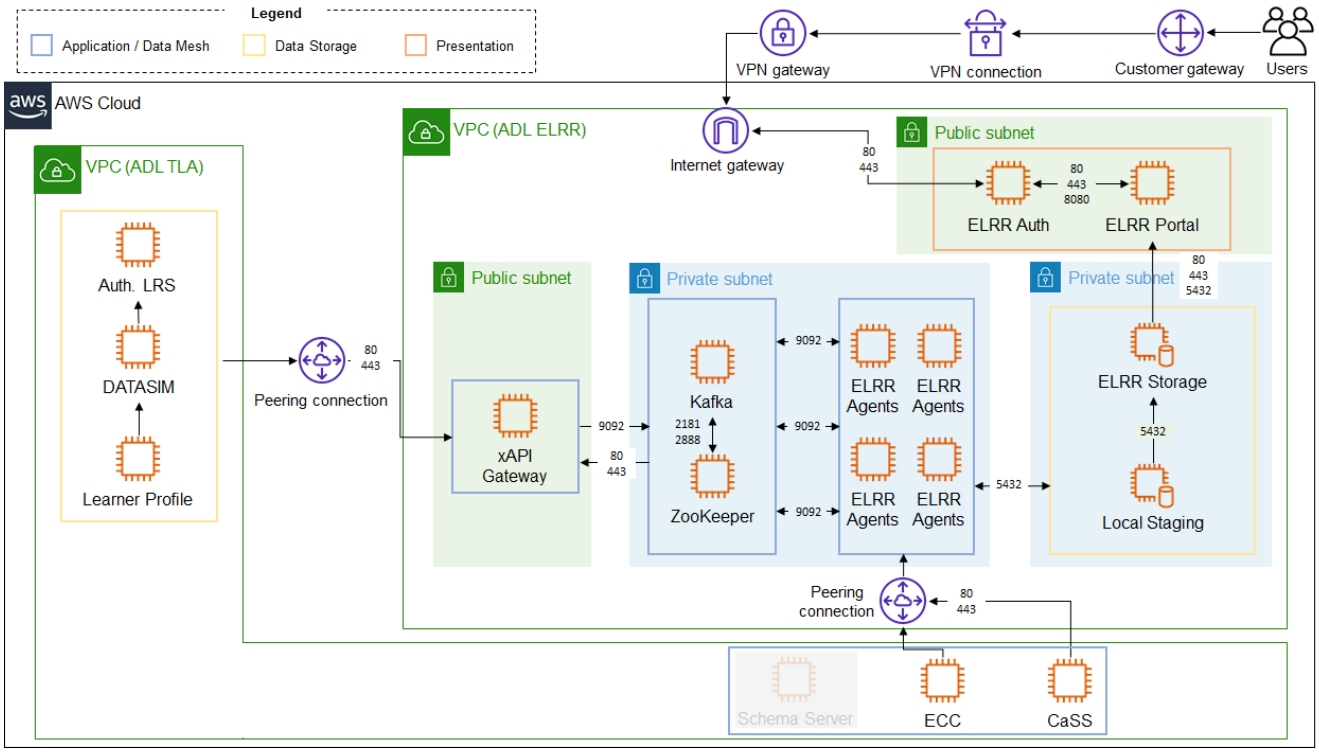

A more complete view of the infrastructure decisions, the hardware and software determination, is provided in the ELRR Prototype System Integration Plan (SIP), System Integration Document. The above illustration of the ELRR prototype infrastructure shows the hardware needed, the different availability zones, and the internal and external application interactions.

The ELRR development team is currently testing and deploying the ELRR Prototype in a Deloitte-managed AWS sandbox environment, while code coverage testing is performed/automated in Github using Github Workflows. The table below provides a summary view of the development environments we are currently using to test the ELRR Prototype.

Table 5: ELRR Environments for Software Testing

Instance/Environment	Description
GitHub	GitHub allows the development team to test code on each commit through GitHub Actions workflow. GitHub Actions builds a stand-alone virtual container that runs the code and coverage test from the repository.
Deloitte-Managed AWS	AWS' commercial cloud environment managed by Deloitte; Deloitte provides a platform for accessing and managing AWS' various services, including additional security and role-based access controls.
ADL TLA Sandbox	AWS' government cloud environment managed by the ADL Initiative; the ADL Initiative will perform prototype testing in this environment with various ADL Initiative stakeholders, both on the ELRR Prototype IPT and outside of the core team involved with the ELRR Prototype.

APPENDIX A: Technology Recommendations for Project Tooling

Identifying the technology and tools to effectively plan, deliver, and support project management and software development activities is critical to the ability to install, configure, load, and test the identified tools planned for the ELRR. Deloitte's tooling technology decisions rely on plans for meeting the guiding attributes and solution architecture requirements.

Technology Name	Technology Function	Expected Use
SharePoint	Project Documentation/Artifact Management	Provides reference to GFI and delivery of project artifacts (ADL Initiative SharePoint site)
Confluence	Product Management	Captures ELRR-related development context, notes, and documents
Jira	Requirements Management	Used for managing ELRR-related solution deliverables, features, and fixes.
GitHub	Source Code Management	Used for collaborating on all implementation--oriented aspects of ELRR solution.
Jenkins	Build Automation	Used for testing and compiling component code packages.
SonarQube	Code Analysis	Containerize SonarQube Instance for code analysis.
Selenium	Test Automation	Lightweight continuous testing solution.
Idera	Data Model Development and Management	Licensed tool for developing robust data models (conceptual, logical, physical), creating data definition language, and exporting various data files into commonly accessible formats.

APPENDIX B: Technology Recommendations for Development Tooling

Identifying the methods and tools to effectively plan, deliver, and support project management and software development activities is critical to the ability to install, configure, develop, deploy, and test the prototype. Deloitte's technology decisions and recommendations are informed by the guiding principles and the requirements.

Technology Name	Technology Function	Expected Use
Amazon Web Services (AWS)	Cloud Infrastructure	Building and deploying services in AWS. No native services -- prepare for migration to Azure in the future.
Terraform	Component Deployment	Used to deploy cloud infrastructure securely and efficiently.
KeyCloak (ADL Initiative TLA)	Details on TLA reference implementation of KeyCloak for identity and access management	Used for security/authentication management.
Kubernetes	Component Hosting/Management	Used for managing deployed components/clusters.
Docker	Component/Application Foundation	Used for specifying containers deployed to cloud infrastructure.
Postgres	Relational Database	Open-source relational database management system (RDBMS).
Apache Kafka	Asynchronous Job Processing	Used for distributed publish/subscribe messaging topology built for event streaming of different data topics
Java	Object-oriented programming language	Development of prototype

Technology Name	Technology Function	Expected Use
Nginx	Web Server	Hosting of web applications.
HTML	Client: Page Structure	Specifies web browser page/screen structure.
CSS	Client: Presentation	Used for decorating web browser page/screen.
React	Client: Application Framework	JavaScript framework used for development of client-side applications

APPENDIX C: Learner API Interface Control Document (ICD)

This ICD will provide the details for the interfaces between the ELRR and external systems. For the prototype, the ICD documents the ELRR Learners interface.

The ELRR Learner application programming interface (API) provides a mechanism to synchronize the data in the ELRR with the data in the local learning systems. This API allows the local learning systems to request data from the ELRR by date range, or by starting date to ending date, and by organization (e.g., DAU, NAVY, etc.). Organization (orgName) and start date (fromDate) are mandatory, whereas, end date (toDate) is optional. If the end date is not provided by the requester, then the current date will be used as the end date.

Below are the details for the request and the response:

Http Method: GET

Request:

fromDate:

Description: The date from which learner records needs to be fetched from ELRR.

Format: YYMMDD

Mandatory: Yes

toDate:

Description: The date until learner records needs to be fetched from ELRR.

Format: YYMMDD

Mandatory: No

orgName:

Description: Organization, (“NAVY” or “DAU”), for which records are requested from ELRR.

Format: String

Mandatory: Yes

responseEntities:

Description: List of entities to be returned in the response. Currently, only “Course” and/or “Competency” may be used.

Format: String

Mandatory: No

startRecord:

Description: optional element which determines with which record the response begins. If nothing is provided, then the response starts from the beginning, (default is 1). "startRecord" must be > 0.

Format: int > 0

Mandatory: No

URI: https://<elrr_hostname>:<elrr_hostport>/learners

Sample Request:

https://<elrr_hostname>:<elrr_hostport>/learners?fromDate=060121&orgName=NAVY

Response:

The response is either failure, (failure message returned), or success, (request fulfilled).

Failure:

Validation failure: one or more of the parameters provided are invalid, e.g., incorrect date format

Error type returned: Validation Error

Error message returned: message includes parameters that failed

System failure: any error that occurs during processing of a valid

request, e.g., database time out

Error type returned: System Error

Error message returned: includes action for user, e.g. "Please try

again later" or "Contact SysAdmin"

Success:

If the "startRecord" is provided, then only records from the "startRecord" until "maxSize", (internally), defined will be returned. If, however, the "endRecord", in the response, is larger than the "totalRecords" retrieved, then another request must be made, providing a "startRecord" beginning after the "endRecord" count.

Header:

The header contains the "fromDate", "toDate", "orgName" and "totalCount" (totalRecords=50,000, startRecord=1; endRecord=10000). The header is pre-

populated with the request parameters for "fromDate", "toDate", “orgName”, and “responseEntities”.

Payload:

The payload contains a series of learner records comprised of personnel information, courses and competencies. If the “responseEntities” is null, then all the information is returned, otherwise, the person category and only the categories selected by the “responseEntities” are returned.

Sample Response:

{
"header": {
"fromDate": "xs:date",
"toDate": "xs:date",
"orgName": "xs:string",
"responseEntities": "xs:string",
"startRecord": xs:int,
"endRecord": xs:int,
"totalCount": xs:int
},
"data": [{
"learner": {
"personnel": {
"person": {
"name": "xs:string",
"firstName": "xs:string",
"middleName": "xs:string",
"lastName": "xs:string",
"namePrefix": "xs:string.",
"titleAffixcode": "xs:string”,
"nameSuffix": "xs:string”,
"qualificationAffixcode": "xs:string”,
"maidenName": "xs:string",
"preferredName": "xs:string",

```
    "humanResourceIdentifier": "xs:string",
    "personnelIdentificationSystem": "xs:string",
    "birthdate": xs:date,
    "sex": "xs:string",
    "primaryLanguage": "xs:string",
    "militaryVeteranindicator": "xs:string",
    "recordStatus": "xs:string"
},
"organization": {
    "organizationname": "xs:string",
    "organizationidentifier": "xs:string",
    "organizationidentificationcode": "xs:string",
    "organizationidentificationsystem": "xs:string",
    "industrytypeidentifier": "xs:string",
    "organizationfein": "xs:string",
    "organizationdescription": "xs:string",
    "parentorganization": "xs:string",
    "recordstatus": "xs:string"
},
"contactInformation": {
    "contactinformation": "xs:string",
    "telephonenumber": "xs:string",
    "isprimaryindicator": "xs:string",
    "telephonetype": "xs:string",
    "electronicmailaddress": "xs:string",
    "electronicmailaddresstype": "xs:string",
    "emergencycontact": "xs:string",
    "recordstatus": "xs:string"
},
"employment": [
    {
        "employerName": "xs:string",
```

```
            "employerdepartment": "xs:string",
            "hiredate": xs:date,
            "employmentstartdate": xs:date,
            "employmentenddate": xs:date,
            "joblevel": "xs:string",
            "occupation": "xs:string",
            "employed": "xs:string",
            "primarycarrercategory": "xs:string",
            "secondcarrercategory": "xs:string",
            "recordstatus": "xs:string"
        }
    ]
},
"courses": [
    {
                "name": "xs:string",
        "coursesubjectmatter": "xs:string",
        "coursesubjectabbreviation": "xs:string",
        "courseidentifier": "xs:string",
        "courselevel": xs:int,
        "coursenumber": "xs:string",
        "courseinstructionmethod": "xs:string",
        "coursestartdate": xs:date,
        "courseenddate": xs:date ,
        "courseenrollmentdate": xs:date ,
        "courseacademicgrade": xs:int,
        "courseprovidername": xs:string,
        "departmentname": xs:string,
        "coursegradescalecode": xs:int,
        "coursemetadatarepository": xs:int,
        "courselrsendpoint":  "xs:string",
        "coursedescription": "xs:string",
```

```
            "recordstatus": "xs:string"
        }
    ],
    "competencies": [
        {
            "competencyframeworktitle": "xs:string",
            "competencyframeworkversion": "xs:string",
            "competencyframeworkidentifier": "xs:string",
            "competencyframeworkdescription": "xs:string",
            "competencyframeworksubject": "xs:string",
            "competencyframeworkvalidstartdate": xs:date,
            "competencyframeworkvalidenddate": xs:date,
            "competencydefinitionidentifier": "xs:string",
            "competencydefinitionidentifierurl": "xs:string",
            "competencytaxonomyid": "xs:string",
            "competencydefinitionvalidstartdate": xs:date,
            "competencydefinitionvalideenddate": xs:date,
            "competencydefinitionparentidentifier": "xs:string",
            "competencydefinitionparenturl": "xs:string",
            "competencydescriptionparentcode": "xs:string",
            "competencydefinitioncode": "xs:string",
            "competencydefinitionstatement": "xs:string",
            "competencydefinitiontypeurl": "xs:string",
            "competencydefinitiontype": "xs:string",
            "recordstatus": "xs:string"
        }]
}]
}
```


Sample Response:

```
{
"header": {
        "fromDate": "210601",
        "toDate": "210610",
        "orgName": "NAVY",
        "startRecord": 1,
        "endRecord": 40000,
        "totalCount": 100000
},
 "data": [ {
   "learner": {
     "personnel": {
        "person": {
           "name": "Alexandrina Annabelle Arredondo-Arteaga",
           "firstName": "Alexandrina",
           "middleName": "Annabelle",
           "lastName": "Arredondo-Arteaga",
           "namePrefix": "Ms.",
           "titleAffixcode": null,
           "nameSuffix": null,
           "qualificationAffixcode": null,
           "maidenName": "Atkinson-Abbas",
           "preferredName": null,
           "humanResourceIdentifier": "3599900000",
           "personnelIdentificationSystem": null,
           "birthdate": "12/31/2000",
           "sex": "F",
           "primaryLanguage": "Spanish",
           "militaryVeteranindicator": "Y",
           "recordStatus": null
        },
        "organization": {
```

```
    "organizationname": "AETC",
    "organizationidentifier": "D0DAF",
    "organizationidentificationcode": "G0V4",
    "organizationidentificationsystem": null,
    "industrytypeidentifier": null,
    "organizationfein": "1234573",
    "organizationdescription": "DoD AIR FORCE",
    "parentorganization": "AIR FORCE",
    "recordstatus": null
},
"contactInformation": {
    "contactinformation": "Email",
    "telephonenumber": "1-935-456-4578",
    "isprimaryindicator": "Y",
    "telephonetype": "Private",
    "electronicmailaddress": "alex91@gmail.com",
    "electronicmailaddresstype": "Personal",
    "emergencycontact": "Email",
    "recordstatus": null
},
"employment": [
    {
        "employerName": "NAVY",
        "employerdepartment": "",
        "hiredate": null,
        "employmentstartdate": null,
        "employmentenddate": null,
        "joblevel": "Petty Officer Third Class",
        "occupation": "Mass Communications Specialist",
        "employed": "Y",
        "primarycarrercategory": null,
        "secondcarrercategory": null,
```

```
            "recordstatus": null
        },
        {
            "employerName": "NAVY",
            "employerdepartment": "",
            "hiredate": null,
            "employmentstartdate": null,
            "employmentenddate": null,
            "joblevel": "Petty Officer Third Class",
            "occupation": "Mass Communications Specialist",
            "employed": "Y",
            "primarycarrercategory": null,
            "secondcarrercategory": null,
            "recordstatus": null
        }
    ]
},
"courses": [
    {
        "name": "Fundamentals of Systems Acquisition Management",
        "coursesubjectmatter": null,
        "coursesubjectabbreviation": null,
        "courseidentifier": "ACQ 101",
        "courselevel": null,
        "coursenumber": "101",
        "courseinstructionmethod": "Web",
        "coursestartdate": "12/21/2020",
        "courseenddate": "01/15/2021",
        "courseenrollmentdate": "12/01/2020",
        "courseacademicgrade": null,
        "courseprovidername": "DAU",
        "departmentname": "Defense Acquisition University",
```

"coursegradescalecode": null,
"coursemetadatarepository": null,
"courselrsendpoint": null,
"coursedescription": null,
"recordstatus": null
},
{
"name": "Mentoring the Acquisition Workforce",
"coursesubjectmatter": null,
"coursesubjectabbreviation": null,
"courseidentifier": "CLC 067",
"courselevel": null,
"coursenumber": "67",
"courseinstructionmethod": "Web",
"coursestartdate": "01/15/2021",
"courseenddate": "01/25/2021",
"courseenrollmentdate": "01/03/2021",
"courseacademicgrade": null,
"courseprovidername": "DAU",
"departmentname": "Defense Acquisition University",
"coursegradescalecode": null,
"coursemetadatarepository": null,
"courselrsendpoint": null,
"coursedescription": null,
"recordstatus": null
},
{
"name": "Facilities Capital Cost of Money",
"coursesubjectmatter": null,
"coursesubjectabbreviation": null,
"courseidentifier": "CLC 103",
"courselevel": null,

"coursenumber": "103",
"courseinstructionmethod": "Web",
"coursestartdate": "01/15/2021",
"courseenddate": "01/25/2021",
"courseenrollmentdate": "01/03/2021",
"courseacademicgrade": null,
"courseprovidername": "DAU",
"departmentname": "Defense Acquisition University",
"coursegradescalecode": null,
"coursemetadatarepository": null,
"courselrsendpoint": null,
"coursedescription": null,
"recordstatus": null
}
],
"competencies": [
{
"competencyframeworktitle": "Contract Principles: General Contracting Concepts",
"competencyframeworkversion": null,
"competencyframeworkidentifier": null,
"competencyframeworkdescription": null,
"competencyframeworksubject": null,
"competencyframeworkvalidstartdate": null,
"competencyframeworkvalidenddate": null,
"competencydefinitionidentifier": null,
"competencydefinitionidentifierurl": null,
"competencytaxonomyid": null,
"competencydefinitionvalidstartdate": null,
"competencydefinitionvalideenddate": null,
"competencydefinitionparentidentifier": null,
"competencydefinitionparenturl": null,
"competencydescriptionparentcode": null,

"competencydefinitioncode": null,
"competencydefinitionstatement": null,
"competencydefinitiontypeurl": null,
"competencydefinitiontype": null,
"recordstatus": null
},
{
"competencyframeworktitle": "Skill and Roles: Business Skills and Acumen",
"competencyframeworkversion": null,
"competencyframeworkidentifier": null,
"competencyframeworkdescription": null,
"competencyframeworksubject": null,
"competencyframeworkvalidstartdate": null,
"competencyframeworkvalidenddate": null,
"competencydefinitionidentifier": null,
"competencydefinitionidentifierurl": null,
"competencytaxonomyid": null,
"competencydefinitionvalidstartdate": null,
"competencydefinitionvalideenddate": null,
"competencydefinitionparentidentifier": null,
"competencydefinitionparenturl": null,
"competencydescriptionparentcode": null,
"competencydefinitioncode": null,
"competencydefinitionstatement": null,
"competencydefinitiontypeurl": null,
"competencydefinitiontype": null,
"recordstatus": null
}
]}
]
}