Systems Integration Plan

Advanced Distributed Learning
Enterprise Course Catalog Initial Operational Capability

8 March 2022

Distribution Statement A
Approved for public release: distribution unlimited.

REPORT DOCUMENTATION PAGE

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
03-08-2022	Report	03/08/2021 - 03/8/2022

4. TITLE AND SUBTITLE

Enterprise Course Catalog Systems Integration Plan

5a. CONTRACT NUMBER

FA701420D007

5b. GRANT NUMBER

N/A

5c. PROGRAM ELEMENT NUMBER

0603769D8Z

6. AUTHOR(S)

Eric Flamer, Sam Schenkman, Michael Son (Deloitte Consulting LLP)

5d. PROJECT NUMBER

N/A

5e. TASK NUMBER

FA701421F0184

5f. WORK UNIT NUMBER

N/A

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

Deloitte Consulting LLP

8. PERFORMING ORGANIZATION REPORT NUMBER

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

OUSD Personnel & Readiness
Advanced Distributed Learning Initiative
13501 Ingenuity Drive, Suite 248
Orlando, Florida 32826

10. SPONSOR/MONITOR'S ACRONYM(S)

OUSD/P&R/FE&T/ADLI

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

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution A

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The ADL ECC Systems Integration Plan provides a summary overview of the ECC Initial Operational Capability's implementation, guidance on the required interfaces between the ECC and the ADL TLA reference implementation in an AWS Platform One environment, a corresponding integration approach, and stakeholder engagement strategy to support a reference implementation of the ECC IOC.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	33	Sae Schatz
U	U	U			19b. TELEPHONE NUMBER (Include area code) 571-480-4640

Systems Integration Plan in Support of:

Advanced Distributed Learning

Enterprise Course Catalog Initial Operational Capability

Review periods. The Government will review and comment on each delivered data item within 10 business days of its receipt. The Contractor shall provide updates or resolutions to all comments within 10 business days of receiving the DIT to close out the deliverable.

Document Control Information

Document Name	Systems Integration Plan
Client	Advanced Distributed Learning (ADL) Initiative
Task Order Name	Enterprise Course Catalog Initial Operating Capability
Document Owner	Eric Flamer, Sam Schenkman, Michael Son (Deloitte Consulting LLP)
Document Version	1.1
Document Status	Final

Document Edit History

Version	Date	Additions/Modifications	Prepared/Revised by
1.0	12/22/2021	Draft submitted to ADL	Deloitte
1.1	1/18/2022	Final submitted to ADL after original feedback on draft	Deloitte

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
Steve Faber	Government Technical Point of Contact
Brent Smith	Government Technical Lead
Trey Hayden	Government Solution Engineer

Table of Contents

1.0 Introduction

1.1 Deliverable Purpose

The Enterprise Course Catalog (ECC) is one of the three Enterprise Digital Learning Modernization (EDLM) lines of effort supported by ADL. Today, course catalogs for the Department of Defense (DoD) organizations are stored in disparate locations with inconsistent data schemas which complicates the transport, management, consolidation, and governance of the course catalogs across and within DoD organizations.

The goal of the ECC is to aggregate course catalog across multiple systems and organizations to provide a centralized location for DoD personnel to view and interact with their learning and development data. The aim of this document is to provide a summary overview of the ECC Initial Operational Capability's implementation, guidance on the required interfaces between the ECC and the ADL Total Learning Architecture (TLA) reference implementation in an Amazon Web Services (AWS) Platform One environment, a corresponding integration approach, and stakeholder engagement strategy to support a reference implementation of the ECC Initial Operational Capability (IOC).

The Deloitte Team, comprised of Deloitte practitioners and our Lingatech sub-contractors, will perform system integration activities to successfully deploy ECC on Platform One in preparation for an IOC test and evaluation period, which is expected to occur during the months of June and July 2022. ECC IOC implementation efforts will begin within the DoD Platform One environment starting in December 2021 and concluding in August 2022 after the completed government acceptance testing activities. Beyond August 2022, ADL may continue to mature, test, and harden the implemented ECC IOC in pursuit of the appropriate accreditations required to support additional testing within stakeholder environments using live course data, which may include integration with other DoD systems in preparation for the Full Operational Capability (FOC).

1.2 Deliverable Scope

The scope of the ECC IOC Systems Integration Plan (SIP) is to provide technical guidance for integrating ECC IOC with DoD organizations to extract, transform and load course metadata within the Platform One Cloud environment. The intended audiences for this document are developers or engineers who are familiar with TLA standards and core cloud computing principles. Readers are not required to know all the details about each standard. However, some knowledge of basic containerization concepts is advantageous for reading this document. Readers should also have a fundamental understanding of how data is exchanged across the internet including how to represent and exchange data using JavaScript Object Notation (JSON), and how to access Representational State Transfer (REST) Application Programming Interfaces (API) on the web.

2.0 ECC IOC Integration Plan and Strategy

The following section outlines the core components of the ECC IOC, the ECC IOC deployment approach, and integration strategy with target ADL TLA systems within the Platform One Cloud. The ECC IOC consists of eight core components: Experience Discovery Service (Front and Back End), Experience Management Service (Front and Back End), Experience Search Engine, Experience Index Agent(s), Experience Index Service, Experience Schema Service. Components are divided into three tiers: ECC Client Tier, ECC Application Tier, and ECC Data Tier. Refer to **Figure 1** below for visual representation. Components are defined as independently deployable software services which, combined together, comprise the system. Components are made up of one or more software packages, libraries, or modules, depending on the technology used to construct the solution; component source code is typically modeled as a repository, although packages, libraries, and modules are also often modeled as repositories as well. A detailed description of the ECC IOC components is defined within **Section 2.1**.

The ECC IOC application deployment approach relies on repeatable, well-defined steps that involve compiling code, building the application, conducting automated unit tests, provisioning servers, and uploading code to Platform One GitLab. Our deployment approach is orchestrated by a continuous integration / continuous delivery (CI/CD) pipeline. Leveraging the CI/CD pipeline enables quicker releases and improves code quality resulting in less bugs following a release. Refer to **Section 2.2** for more detailed information on the ECC IOC deployment approach.

2.1 ECC Overview

The ECC is designed to offer flexibility in connecting to multiple source catalogs. Once invoked, the ECC Index Agents load the data into the Experience Index Service for a global search capability of ingested courses. Once the data has been stored within the XIS, data is pushed to Elasticsearch where high performance search engine is used to retrieve course data. Please refer to **Figure 2** below for a visual representation of the components and connections comprising the ECC. While the figure depicts an AWS-hosted environment, the ECC can likely run on other major cloud providers (e.g., Azure, Google Cloud Platform) or on-premise environments with minor modifications.

Figure 1 – ECC System Tiers
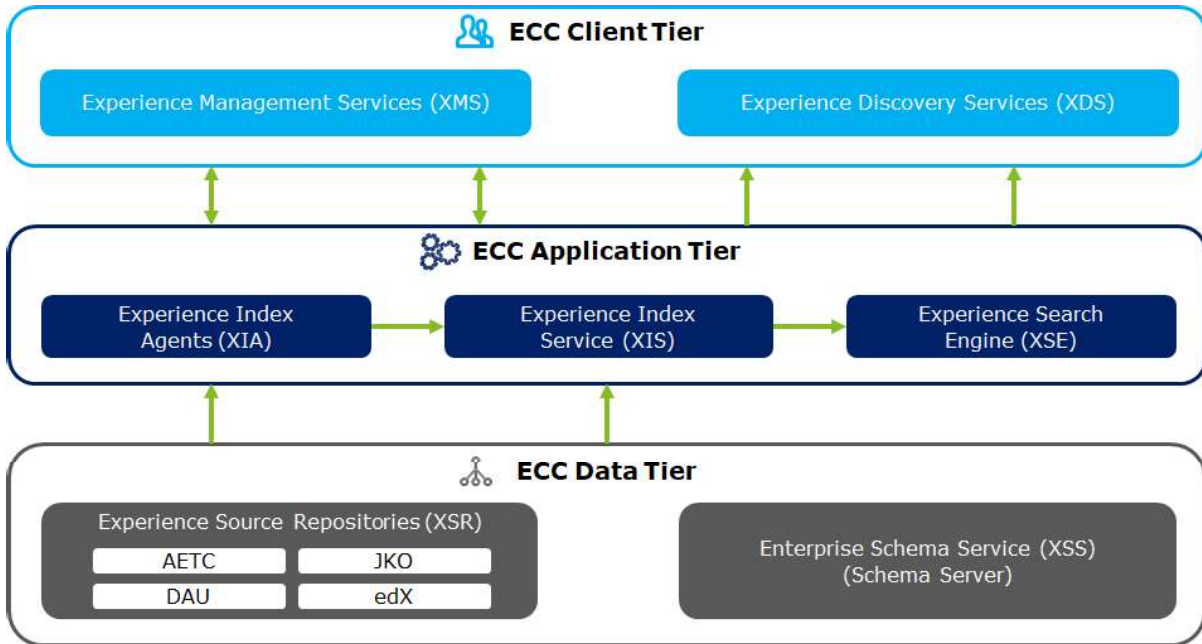
Figure 2 – ECC Infrastructure
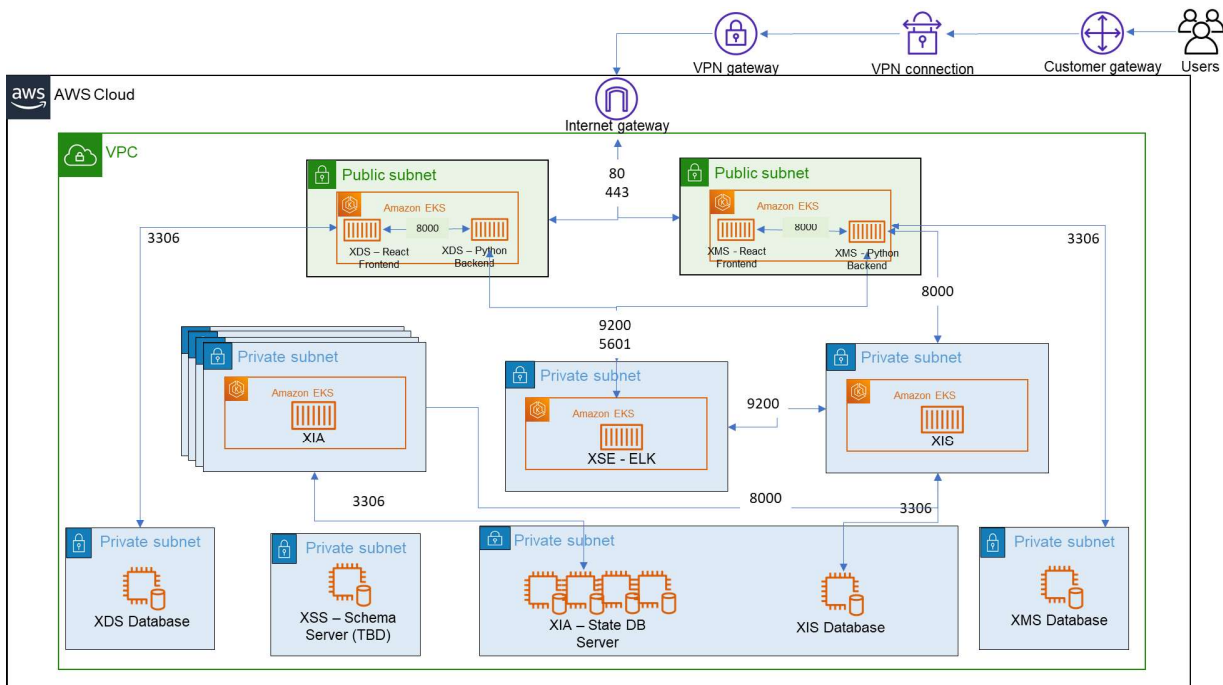
Tier 1 - ECC Presentation/Client

The **ECC Presentation/Client** layer represents the user-facing components of the ECC to enable users to interact with the aggregated course catalog. The front-end offers views for Experience Participant (Facilitator and Consumer), system operators, experience owners, and experience managers with tailored data and role-based access control (RBAC). The ECC Presentation/client layer consists of the following proposed components:

- **Experience Discovery Service** – The Experience Discovery Service (XDS) is the human-facing application enabling an Experience Consumer or Experience Facilitator to easily locate a pertinent learning experience metadata record that has been indexed by the Experience Index Service (XIS). Because the XDS is a separate application, it can be deployed in a separate environment from the XIS, and can even be configured to point to a different XIS as needed. In addition, multiple XDS applications can be deployed and point to the same XIS, allowing for a great amount of installation and configuration flexibility.

  The Experience Discovery Service (XDS) focuses on enabling two main use cases:

  o Direct Search: An Experience Consumer or Experience Facilitator interacts directly with the XDS in order to locate one or more pertinent learning experiences using search/browse features presented by the component.

  o Indirect Search: An Experience Consumer or Experience Facilitator receives communications asynchronously from the XDS in order to become aware of one or more pertinent learning experiences.

- **Experience Management Service -** The Experience Management Service is the human-facing application enabling an Experience Owner or Experience Manager to modify or augment a learning metadata record ingested by the Experience Index Service (XIS). Because the XMS is a separate application, it can be deployed in a separate environment from the XIS, and can even be configured to point to a different XIS as needed. In addition, multiple XMS applications can be deployed and point to the same XIS, allowing for a great amount of installation and configuration flexibility.

  In addition to managing the XIS, the XMS also enables an Experience Owner to monitor Experience Index Agent (XIA) operational health in order to identify problems with the aggregation workflow.

Tier 2 - ECC Application Tier

The **ECC Application Tier** contains the functional process logic for the ECC IOC. This process logic is critical in extracting, validating, transforming, validating, and loading the course metadata from each source. Below is a description of the expected common services and how they will be used for the ECC IOC:

- **Experience Index Agent** – Experience Index Agents (XIA) are the lifeblood of the ECC system. XIAs are purpose-built pieces of technology – network-based, independently-

deployed software services – configured to work with a specific Experience Source Repository (XSR).

Experience Index Agents are deployed within a containerized operating environment such as Docker. XIA containers can be deployed in proximity to their designated Experience Source Repository (XSR), for security and/or management reasons, or they can be deployed within the ECC environment – or another environment entirely – and connect to the configured target XSR over a network protocol. XIAs operate on a defined/configured interval utilizing an internal scheduler. When the XIA is able to connect to its corresponding XSR, it then executes the following workflow operations:

- o **Extract**: XIAs pull pertinent learning experience metadata records from their corresponding XSR, based on configuration.

- o **Validate**: XIAs compare extracted learning experience metadata against the configured standard XSR reference schema stored in the Experience Schema Service (XSS)

- o **Transform**: XIAs transform extracted learning experience metadata to the configured target schema using the XSR-to-Target transformation stored in the Experience Schema Service (XSS)

- o **Validate**: XIAs compare transformed learning experience metadata against the configured target XSR reference schema stored in the Experience Schema Service (XSS)

- o **Load**: XIAs push validated and transformed learning experience metadata to the destination Experience Index Service (XIS) for downstream processing

- o **Log**: XIAs log error, warning, informational, and debug events to a data store which can be monitored.

- **Experience Index Service** – The Experience Index Service (XIS) is the core of the Enterprise Course Catalog solution. It is the central sink for learning experience metadata collected by the Experience Index Agent (XIA) components. In addition, the XIS can receive supplemental learning experience metadata – field name/value overrides and augmentations – from the Experience Management Service (XMS).

Learning experience metadata is received from XIAs is stored in the Metadata Loading Area and processed asynchronously in order to enhance overall system performance and scalability. Processed metadata is combined with supplemental metadata provided by an Experience Owner or Experience Manager and the "composite record" is stored in the Metadata Repository. Metadata Repository record addition/modification events are logged to a job queue and the metadata is then sent to the Experience Search Engine (XSE) for indexing and high-performance location/retrieval.

It is possible for an XIS to syndicate its composite records to another XIS. The record set can be filtered by one or more facets/dimensions in order to transmit a subset of the overall composite record repository. In addition, the transmitted field set can be configured to contain redacted values for specified fields when information is considered too sensitive for syndication.

- **Experience Search Engine** – The Experience Search Engine (XSE) system component is responsible for enabling fast identification and retrieval of learning experience metadata records aggregated by the Experience Index Service (XIS). After being loaded into one or more XSE indices, learning experience metadata records can be identified specifically by identifier, or more broadly by a keyword and/or facet match.

  Because search technology is generally a solved problem, the ECC project team will not be implementing a custom service. Instead, the team will be utilizing Elasticsearch, which is well-known for performance and scalability. Elasticsearch supports a variety of features which should prove to be useful in the ECC solution, including search, faceting, and percolation.

Tier 3 - ECC Data Tier

ECC data tier leverages Amazon Simple Storage Service (S3) and Amazon Relational Database Service (RDS) in Platform One. S3 is a highly reliable, scalable, fault-tolerant, and secure object storage system. Fault tolerance and availability of S3 buckets are managed by AWS. ECC will leverage S3 buckets to store schema files for ECC components to extract, validate and transform source metadata. ECC will leverage Amazon RDS MySQL for persistent data store and index service data. RDS MySQL offers a highly scalable, reliable, fault-tolerant, and secure database. RDS will be designed to fail over between availability zones that eliminates interruptions to the database availability.

The **ECC Data Tier** contains the following components:

- **Experience Source Repositories (XSR) -** Experience Source Repositories are the primary data resources supporting the entire Enterprise Course Catalog solution. Without the learning experience metadata extracted from the XSRs, the entire system would depend on Experience Owners and Experience Managers to enter all of the metadata manually. This would introduce an additional workload on top of the work they are already performing to populate their own experience delivery systems and related catalog services. ECC can extract data via scheduled file drops, direct connection to the source database, upload data and connection to API endpoints. Source data can be in XML, JSON, and CSV file formats when being extracted into ECC.

- **Experience Schema Service (XSS) -** The Experience Schema Service maintains referential representations of domain entities, as well as transformational mappings that describe how to convert an entity from one particular schema representation to another. The XSS is also responsible for storing schema mappings used by the XIA for source metadata transformation and validation. Both metadata and supplemental metadata ledgers get extracted from respective XSR repositories and transformed using the P2881 standards defined in the XSS. **Figure 3** below represents ETL pipeline performed by XIA using the XSS to validate the schema standards.

Figure 3 – ECC ETL Pipeline

Schema Files on Schema Server and ETL Pipeline on XIA :
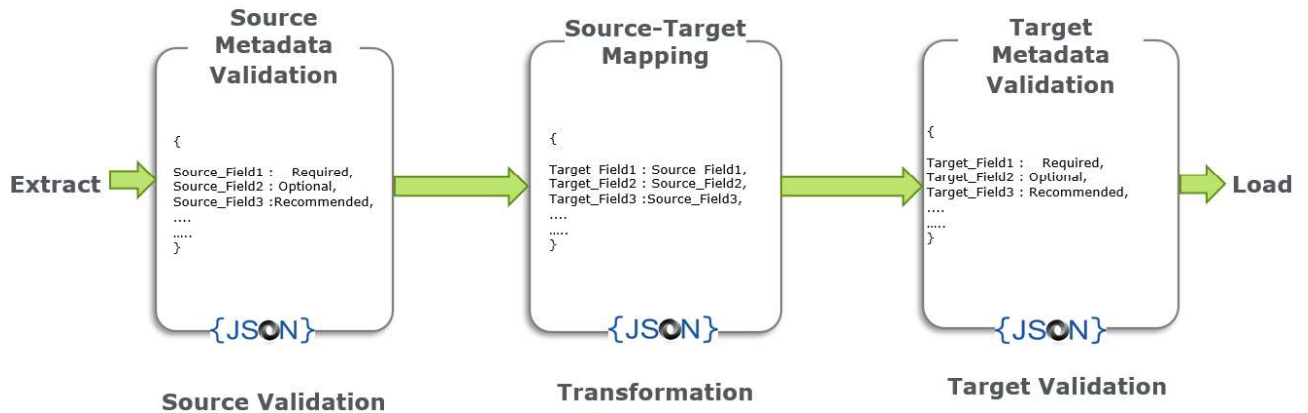
2.2 ECC IOC Deployment Strategy

The internal ECC IOC application leverages a robust CI/CD pipeline that enables quality and secure application builds of the ECC. The CI/CD pipeline offers more consistency that the latest versions of the ECC will operate as intended and to meet Platform guidelines within the US Air Force's Platform One Cloud environment. ECC IOC will be integrated with the Platform One Gitlab CI/CD pipeline to help orchestrate the build, testing, and deploying applications to Kubernetes. Code quality is achieved through SonarQube, performing static code analysis for the ECC to catch errors within code prior to a deployment. Platform One CI/CD pipeline scans each repository for a Dockerfile that meets compliance, unit tests using lint, static code analysis using Fortify, Sonarqube and Trufflehog. Figure 4 provides a summary representation of each stage in Platform One's CI/CD process based on the standardized Platform One pipeline.

Figure 4 – ECC-Adopted Platform One CI/CD Pipeline for Party Bus Deployment
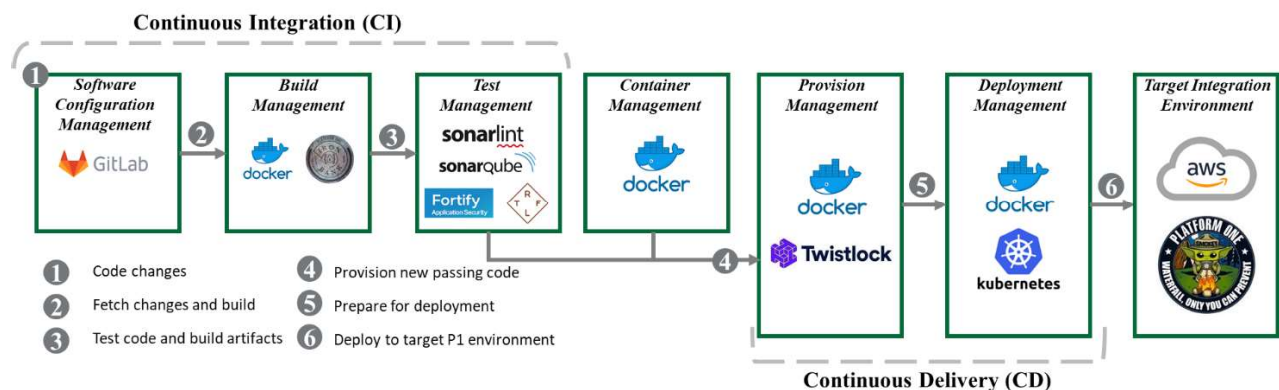
2.3 System Integration Roles and Responsibilities

To support the deployment of the ECC IOC into the Platform One Cloud environment, the following staffing plan, roles, and quantity are accounted for in **Table 1** below.

Table 1 – ECC Integration Roles and Responsibilities

ECC IOC Role	ECC IOC Responsibilities
Deloitte Team Development Engineers (x6)	▪ Develop ECC code ▪ Verify ECC code meets Platform One standards ▪ Verify ECC code passed all analysis and quality checks to ensure pipeline is green ▪ Update README files accordingly for deployment steps
Deloitte Team Testers (x2)	▪ Conducts ECC tests in the QA environment on Deloitte AWS sandbox and in Platform One
Deloitte Team Project Manager (x1)	▪ Manages scope and schedule as defined in the ECC technical management work plan ▪ Supports overall project functional activities, including stakeholder management, risk and issue resolution, and stakeholder communications
ADL Engineers	▪ Verify ADL EDLM code meets Platform One standards ▪ Verify ADL EDLM code passed all analysis and quality checks to ensure pipeline is green ▪ Work with Deloitte Team Development Engineers to verify ECC IOC deployment for ECC IOC Test and Evaluation activities.
Platform One DevOps Engineers	▪ Develop ECC pipelines matching ECC components for deployment to test, pre-production, and production Platform One environments. ▪ Verify ECC code passed all analysis and quality checks to ensure the pipeline is green. ▪ Work with Deloitte Team and ADL Engineers to troubleshoot deployment and integration error within the Platform One environments.

2.4 Assumptions

In relation to the guidance provided within the ECC IOC SIP, the following assumptions have been made within **Table 2** below.

Table 2 – ECC IOC SIP Assumptions

Assumption	Rationale
Development will rely on Platform One's DevOps processes to successfully deploy ECC	ECC Development team will work closely with the Platform One DevOps team to integrate ECC with Platform One's Cloud environment. Development team will be responsible of keeping the CI/CD pipeline green through various quality & security checks. To the extent needed, Deloitte will update this SIP deliverable to match the expected integration and deployment practices established by Platform One for the ECC IOC.

Assumption	Rationale
Development Team will have significantly limited access to the Platform One cloud environment after code is pushed to Gitlab	Development team will have limited access/control of the deployment into AWS. All deployment activities will be handled by Platform One's DevOps team. To the extent needed, Deloitte will update this SIP deliverable to match Platform One's deployment processes, workflows, and practices once our ECC pipelines are operational.
Development Team will rely on Platform One to confirm system integration with ADL's newly deployed EDLM systems/enclave	In the previous Operational Prototype development cycle, Deloitte relied on the ADL's TLA Sandbox environment for integration with other TLA components, allowing for a greater degree of access/control via ADL. With Platform One, the assumption above relating to access impacts ADL's own control of its other planned pipeline deployments for services that may impact the ECC IOC (specifically the xAPI Profile Server). Deloitte will coordinate with ADL on the deployment and integration practices for these other systems and update this SIP deliverable to match the expected integration steps with these other EDLM systems.

3.0 ECC Integration Approach with Platform One Party Bus

The ECC IOC relies on connections to source repositories to extract, transform and load course metadata. The four source repositories used during the prototype phase were Defense Acquisition University (DAU), Air and Education Training Command (AETC), Joint Knowledge Online (JKO) and an external source edX. The ECC communicates with the source repositories via API endpoints or from an exported CSV/Excel file from the source provider. The ECC development team will follow Platform One's Dockerfile guidelines to deliver secure and approved containers. Dockerfiles will be limited to non-root users in the scripts. The base image in each Dockerfile will be pointing to a hardened container in Registry One's repository. Even in multi-stage Dockerfiles, the team will ensure hardened containers are referenced.

3.1 Experience Index Agent (XIA)

The XIA is the data processing engine for the ECC. XIA agents are configured to connect to source repositories, extract the data via an API endpoint or CSV. Once extracted, XIA agents will automatically validate the extracted data, transform the data to the mapped data schema, validate the transformed data, then load the data into XIS.

The identified system interfaces between an Authoritative XSR and XIA are depicted within **Table 3** below.

Table 3 – XIA System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XSR	API endpoint for XIA to connect and run ETL on the data	Ephemeral	On-demand
ECC XIA	Connects to Experience Source Repositories to run ETL and map course data to the schema	Ephemeral	On-demand

3.2 Experience Index Service (XIS)

The XIS is the component that holds the transformed data from the XIAs. When data is loaded from respective XIA agents, XIS validates the data against the XSS schema standards. Once data is stored in the XIS, data is pushed to XSE (Elasticsearch). The XDS component retrieves user queries directly from XSE and XMS component connect directly to XIS to retrieve detailed course data and allow experience managers to make modifications to the metadata

The identified system interfaces between XIA, XIS, XSE and XMS are depicted within **Table 4** below.

Table 4 – XIS System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XIS	API Endpoint	Ephemeral	On-demand
ECC XIA	API Endpoint	Ephemeral	On-demand
ECC XSE	API Endpoint	Ephemeral	On-demand
ECC XMS	API Endpoint	Ephemeral	On-demand

3.3 Experience Search Engine (XSE)

The XSE is the core search engine for ECC. Elasticsearch will handle indices, queries, and highly available search capabilities for the users. Data is pushed from XIS which is then used to retrieve queries from XDS.

The identified system interfaces between XSE, XDS, XIS are depicted within **Table 5** below.

Table 5 – XSE System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XSE	API Endpoint	Ephemeral	On-demand
ECC XIS	API Endpoint	Ephemeral	On-demand
ECC XDS	API Endpoint	Ephemeral	On-demand

3.4 Experience Discovery Service (XDS)

The XDS is the main interface for the users to interact with. This is where users will get the full experience as experience participants (facilitator and consumer). XDS will interact with the XSE component to retrieve queries, highlight course, save courses, and be taken directly to the course provider.

The identified system interfaces between XDS, Users, XSE are depicted within **Table 6** below.

Table 6 – XDS System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XDS	API Endpoint	Ephemeral	On-demand
Internet/Users	HTTPS endpoint	Ephemeral	On-demand
ECC XSE	API Endpoint	Ephemeral	On-demand

3.5 Experience Management Service (XMS)

The XMS is the management system where updates to the metadata can be made. Each experience facilitator can be granted roles to the XMS that matches their line of organization to modify the metadata for the organization's preferences. XMS will connect directly with the XIS to make these modifications.

Table 7 – XMS System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XDS	API Endpoint	Ephemeral	On-demand
Internet/Users	HTTPS endpoint	Ephemeral	On-demand
ECC XSE	API Endpoint	Ephemeral	On-demand

3.6 Experience Schema Service (XSS)

The Experience Schema Service is meant to serve as the TLA component responsible for managing pertinent object/record metadata schemas and the mappings for transforming records from a source metadata schema to a target metadata schema. The ECC Development team has developed a reference implementation of the schema service to support this linked data service for the ECC solution. Global metadata registry along with interfaces for indexing services and agents. This component will also be used to store and link vocabularies from stored schema. The schema service is used to validate metadata against the defined metadata standards. XSS can be used by other services as the global linked vocabulary service.

Table 8 – XSS System Interfaces

ECC Component	Interfaces	Connection Duration	Connection Frequency
ECC XIA	API Endpoint	Ephemeral	On-demand
Internet/Users	HTTPS endpoint	Ephemeral	On-demand
ECC XIS	API Endpoint	Ephemeral	On-demand

3.7 ECC Conformance Alerting (Notifications)

The ECC conformance alerting/notification component feature is designed to send email notifications to the system managers, system owners, and participants. Notifications will consist of warnings/error trap reports or record updates automatically sent to subscribers. ECC Conformance alerting is part of the automated scheduler celery tasks. The system operators can add/remove subscribers from the Django admin page. System operators can also customize the content of the email notifications. This component is an independently deployable PyPi package.

ECC conformance alerting consists of the configuration types below, which help System Operators quickly configure and set up conformance alerting.

- **Sender Email Configurations:** ECC OpenLXP-Notification package has the ability to configure sender email. System Operator can choose an email address to send conformance emails from.

- **Receiver Email Configurations:** ECC OpenLXP-Notification package has the ability to configure receiver email addresses. System Operators can add the list of email addresses to send conformance alerts.

- **Notification Email Content:** Conformance alert's email body contents is easily editable. System Operators can choose their own email body content. To learn more information on the editing of the email body, please review the ECC team's [GitHub documentation found here](#).

3.8 ECC Automated Scheduling & Tasks

The ECC ETL (Extract, Transform, Load) process is fully automated using Celery tasks in both XIA and XIS components. In the XIA, Celery is used to extract, validate, transform, validate, load, and validate the metadata before loading transformed data into the XIS. When the ETL process is completed, celery runs a task to load data into the XIS and send notifications to subscribed users. After the data is loaded into the XIS, celery tasks kick off in the XIS to consolidate metadata ledgers and load the metadata into the XSE.

Automating these tasks eliminates manual work performed by the system operator. Hence to gain better performance, ECC tasks run on a customizable schedule to minimize disruption. System Operators can schedule daily, weekly, or monthly tasks in ECC. In the case of emergency planning or meeting specific scheduling needs, system operators can choose to run tasks manually by running an API. The parameters for scheduling the ECC automated tasks using Celery tasks includes the following:

- The System Operator can specify the Name for the task.

- The System Operator can select the applicable tasks from the registered ECC tasks.

- The System Operator can enable or disable the task scheduling by selecting or unselecting the 'Enabled' option for each task.

- The System Operator can specify interval schedule, including the interval period and the number of periods.

- The System Operator can specify the Start Date Time for the scheduled task

- The System Operator can enable or disable the task as a one-time action by selecting or unselecting the 'One-Off' option for each task

4.0 Appendix

Appendix A - Hardware / Software

ECC IOC Hardware

Table 9 – *ECC Hardware*

Component	Service	Operating System	Count
XIA	Kubernetes	Docker	1
XIS	Kubernetes	Docker	1
XDS	Kubernetes	Docker	1
XDS-UI	Kubernetes	Docker	1
XMS	Kubernetes	Docker	1
XMS-UI	Kubernetes	Docker	1
XSE	Kubernetes	Docker	1
XSS	Kubernetes	Docker	1

ECC IOC Software

Table 10 – *ECC Software*

Configuration Item (CI)	Name	Version
Security / Authentication	KeyCloak	7.0.x
	Django Administration	3.1.x
Application	Nginx	1.14.0
	Elasticsearch	7.16
	Django	3.1.x
	React	
Data / Storage	RDS MySQL	13.1
Container	Docker	20.10.x
Container Orchestration	Kubernetes	1.19.x
Operating System	Linux/Ubuntu	18.04
Administration	Gitlab CICD	

Appendix B - Deployment Scripts

Appendix B-1 Security, Networking, and Identity and Access Management

The necessary firewall ports for ECC, based on the specified component/infrastructure, is outlined within **Table 11** below.

Table 11 – ECC Firewall Ports

Component	Ports (Ingress)	Ports (Egress)
Experience Index Agent (XIA)	:22 (per IP) :443 Public :8000 (per IP)	All Traffic
Experience Index Service (XIS)	:22 (per IP) :443 Public :8100 (per IP)	All Traffic
Experience Search Engine (XSE)	:22 (per IP) :443 Public :9200 (per IP)	All Traffic
Experience Schema Service (XSS)	:22 (per IP) :443 Public :8080 (per IP)	All Traffic
Experience Discovery Service (XDS)	:22 (per IP) :443 (per IP) :8000 (per IP) :3000 (per IP)	All Traffic
Experience Management Service (XMS)	:22 (per IP) :443 Public :8000 (per IP) :3000 (per IP)	All Traffic

Appendix B-2 Deployment Scripts

ECC is deployed and integrated using microservices approach for each component. Although some components rely on others for the data to be transformed and pushed into the global catalog system, each component can be deployed independently. This allows for flexible configuration modifications based on each organization's requirements. The Experience Index Agent and Notification components are python packages available on PyPi. Eventually, XIA and XIS syndication will be implemented. Organizations will have the options to syndicate data into the deployed ECC application or from ECC into a local index service. ECC allows any component to be deployed individually and connect to/from the ADL deployed ECC.

Integration Use Case 1: Deploying an Experience Index Agent to Contribute to the ADL ECC IOC

Experience Index Agents are implemented using docker-compose. Experience Index Agents is configurable to allow any organizations to connect to their source repository. This is done by configurable API endpoints for the metadata within the XIA Django administration page. Whether the new organization exports CSV files or has API endpoints integrated with their

source repository, new source repositories can be added to OpenLXP. Docker-compose.yaml file presented below is a generic representation of the ECC docker-compose file that can be customized where <XSR> can be updated based on source repository.

```yaml
1  version: "3"
2
3  services:
4    db_xia_<XSR>:
5      image: mysql:5.7
6      ports:
7        - '3306:3306'
8      environment:
9        MYSQL_DATABASE: "${DB_NAME}"
10       MYSQL_PASSWORD: "${DB_PASSWORD}"
11       MYSQL_ROOT_PASSWORD: "${DB_ROOT_PASSWORD}"
12       MYSQL_HOST: ''
13     networks:
14       - openlxp
15
16   app_xia_<XSR>:
17     build:
18       context: .
19     ports:
20       - "8000:8020"
21     command: >
22       sh -c ". /opt/app/start-app.sh"
23     environment:
24       DB_NAME: "${DB_NAME}"
25       DB_USER: "${DB_USER}"
26       DB_PASSWORD: "${DB_PASSWORD}"
27       DB_HOST: "${DB_HOST}"
28       DJANGO_SUPERUSER_USERNAME: "${DJANGO_SUPERUSER_USERNAME}"
29       DJANGO_SUPERUSER_PASSWORD: "${DJANGO_SUPERUSER_PASSWORD}"
30       DJANGO_SUPERUSER_EMAIL: "${DJANGO_SUPERUSER_EMAIL}"
31       BUCKET_NAME: "${BUCKET_NAME}"
32       AWS_ACCESS_KEY_ID: "${AWS_ACCESS_KEY_ID}"
33       AWS_SECRET_ACCESS_KEY: "${AWS_SECRET_ACCESS_KEY}"
34       AWS_DEFAULT_REGION: "${AWS_DEFAULT_REGION}"
35       REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
36       AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
37       SECRET_KEY_VAL: "${SECRET_KEY_VAL}"
38       LOG_PATH: "${LOG_PATH}"
39       CELERY_BROKER_URL: "${CELERY_BROKER_URL}"
```

```
40        CELERY_RESULT_BACKEND: "${CELERY_RESULT_BACKEND}"
41      volumes:
42        - ./app:/opt/app/openlxp-xia-<XSR>
43      depends_on:
44        - db_xia_<XSR>
45      networks:
46        - openlxp
47
48    redis:
49      image: redis:alpine
50      networks:
51        - openlxp
52
53    celery:
54      build:
55        context: .
56      command: celery -A openlxp_xia_edx_project worker -l info --pool=solo
57      volumes:
58        - ./app:/opt/app/openlxp-xia-<XSR>
59      environment:
60        REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
61        AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
62      env_file:
63        - ./.env
64      depends_on:
65        - db_xia_<XSR>
66        - redis
67        - app_xia_<XSR>
68      networks:
69        - openlxp
70      restart: on-failure
71
72    celery-beat:
73      build:
74        context: .
75      command: celery -A openlxp_xia_<XSR>_project beat --scheduler
76  django_celery_beat.schedulers:DatabaseScheduler --loglevel=info --
77  pidfile=/tmp/celerybeat.pid
78      volumes:
79        - ./app:/opt/app/openlxp-xia-<XSR>
80      environment:
81        REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
```

```
82        AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
83    env_file:
84      - ./.env
85    depends_on:
86      - db_xia_<XSR>
87      - redis
88      - app_xia_<XSR>
89    networks:
90      - openlxp
91    restart: on-failure
92
93  flower:
94    image: mher/flower:0.9.7
95    command: [ "flower", "--broker=redis://redis:6379/0", "--port=8888" ]
96    ports:
97      - 8888:8888
98    networks:
99      - openlxp
100
101 networks:
    openlxp:
      external: true
```

System Operators can use the current implementations of edX, JKO, DAU, and AETC Experience Index Agents by pulling the code from below repositories. Each of these implementations demonstrates a catalog contribution from a distributed system to the ADL's EDLM ECC deployment.

- **edX** (https://github.com/OpenLXP/openlxp-xia-edx)
- **JKO** (https://github.com/OpenLXP/openlxp-xia-jko)
- **DAU** (https://github.com/OpenLXP/openlxp-xia-dau)
- **AETC** (https://github.com/OpenLXP/openlxp-xia-aetc)

Integration Use Case 2: Deploying an Experience Index Service for Local Use of the Full ECC Solution:

The Experience Index Service is implemented using docker-compose. Clone the XIS Github repository and run 'docker-compose up -d'. XIS component is developed to allow syndication with other XIS components. System Operators can run their own XIS if they choose to. Whether it is external to Platform One or internal, any organization can pull or push data into/from other XIS components.

Location of XIS GitHub Repository: https://github.com/OpenLXP/openlxp-xis

```yaml
1 version: "3"
2
3 services:
4   db_xis:
5     image: mysql:5.7
6     ports:
7       - '3310:3306'
8     environment:
9       MYSQL_DATABASE: "${DB_NAME}"
10#      MYSQL_USER: 'root'
11      MYSQL_PASSWORD: "${DB_PASSWORD}"
12      MYSQL_ROOT_PASSWORD: "${DB_ROOT_PASSWORD}"
13      MYSQL_HOST: ''
14    networks:
15      - openlxp
16   app_xis:
17    container_name: openlxp-xis
18    build:
19      context: .
20    ports:
21      - "8080:8020"
22    command: >
23      sh -c ". /opt/app/start-app.sh"
24    environment:
25      DB_NAME: "${DB_NAME}"
26      DB_USER: "${DB_USER}"
27      DB_PASSWORD: "${DB_PASSWORD}"
28      DB_HOST: "${DB_HOST}"
29      DJANGO_SUPERUSER_USERNAME: "${DJANGO_SUPERUSER_USERNAME}"
30      DJANGO_SUPERUSER_PASSWORD: "${DJANGO_SUPERUSER_PASSWORD}"
31      DJANGO_SUPERUSER_EMAIL: "${DJANGO_SUPERUSER_EMAIL}"
32      BUCKET_NAME: "${BUCKET_NAME}"
33      AWS_ACCESS_KEY_ID: "${AWS_ACCESS_KEY_ID}"
34      AWS_SECRET_ACCESS_KEY: "${AWS_SECRET_ACCESS_KEY}"
35      AWS_DEFAULT_REGION: "${AWS_DEFAULT_REGION}"
36      REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
37      AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
38      SECRET_KEY_VAL: "${SECRET_KEY_VAL}"
39      LOG_PATH: "${LOG_PATH}"
40      CELERY_BROKER_URL: "${CELERY_BROKER_URL}"
41      CELERY_RESULT_BACKEND: "${CELERY_RESULT_BACKEND}"
42    volumes:
```

```
43        - ./app:/opt/app/openlxp-xis
44      depends_on:
45        - db_xis
46      networks:
47        - openlxp
48
49    es01:
50      image: docker.elastic.co/elasticsearch/elasticsearch:7.11.1
51      container_name: es01
52      environment:
53        - discovery.type=single-node
54      ulimits:
55        memlock:
56          soft: -1
57          hard: -1
58      volumes:
59        - data01:/usr/share/elasticsearch/data
60      ports:
61        - 9200:9200
62        - 9300:9300
63      networks:
64        - openlxp
65
66    redis:
67      image: redis:alpine
68      networks:
69        - openlxp
70
71    celery:
72      build:
73        context: .
74      command: celery -A openlxp_xis_project worker -l info --pool=solo
75      volumes:
76        - ./app:/opt/app/openlxp-xis
77      environment:
78        REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
79        AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
80      env_file:
81        - ./.env
82      depends_on:
83        - db_xis
84        - redis
```

```
 85        - app_xis
 86      networks:
 87        - openlxp
 88      restart: on-failure
 89
 90    celery-beat:
 91      build:
 92        context: .
 93      command: celery -A openlxp_xis_project beat --scheduler
 94 django_celery_beat.schedulers:DatabaseScheduler --loglevel=info --
 95 pidfile=/tmp/celerybeat.pid
 96      volumes:
 97        - ./app:/opt/app/openlxp-xis
 98      environment:
 99        REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
100        AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
101      env_file:
102        - ./.env
103      depends_on:
104        - db_xis
105        - redis
106        - app_xis
107      networks:
108        - openlxp
109      restart: on-failure
110
111    flower:
112      image: mher/flower:0.9.7
113      command: [ "flower", "--broker=redis://redis:6379/0", "--port=8888" ]
114      ports:
115        - 5555:5555
116      networks:
117        - openlxp
118
119 volumes:
120    data01:
121      driver: local
122 networks:
     openlxp:
       external: true
```

The Experience Search Engine is implemented using docker-compose. Clone the Github repository and run 'docker-compose up -d'. There are no additional configurations needed to set up XSE for OpenLXP. Once the Elasticsearch instances are running, indices are created based on configuration from the XIA component.

 Location of XSE GitHub Repository: https://github.com/OpenLXP/openlxp-xse

```
1 version: '2.2'
2 services:
3   es01:
4     image: docker.elastic.co/elasticsearch/elasticsearch:7.16.2
5     container_name: es01
6     environment:
7       - node.name=es01
8       - cluster.name=es-docker-cluster
9       - discovery.seed_hosts=es02,es03
10       - cluster.initial_master_nodes=es01,es02,es03
11       - bootstrap.memory_lock=true
12       - "ES_JAVA_OPTS=-Xms512m -Xmx512m"
13     ulimits:
14       memlock:
15         soft: -1
16         hard: -1
17     volumes:
18       - data01:/usr/share/elasticsearch/data
19     ports:
20       - 9200:9200
21     networks:
22       - elastic
23
24   es02:
25     image: docker.elastic.co/elasticsearch/elasticsearch:7.16.2
26     container_name: es02
27     environment:
28       - node.name=es02
29       - cluster.name=es-docker-cluster
30       - discovery.seed_hosts=es01,es03
31       - cluster.initial_master_nodes=es01,es02,es03
32       - bootstrap.memory_lock=true
33       - "ES_JAVA_OPTS=-Xms512m -Xmx512m"
34     ulimits:
35       memlock:
36         soft: -1
37         hard: -1
```

```
38    volumes:
39      - data02:/usr/share/elasticsearch/data
40    networks:
41      - elastic
42
43  es03:
44    image: docker.elastic.co/elasticsearch/elasticsearch:7.16.2
45    container_name: es03
46    environment:
47      - node.name=es03
48      - cluster.name=es-docker-cluster
49      - discovery.seed_hosts=es01,es02
50      - cluster.initial_master_nodes=es01,es02,es03
51      - bootstrap.memory_lock=true
52      - "ES_JAVA_OPTS=-Xms512m -Xmx512m"
53    ulimits:
54      memlock:
55        soft: -1
56        hard: -1
57    volumes:
58      - data03:/usr/share/elasticsearch/data
59    networks:
60      - elastic
61
62  kib01:
63    image: docker.elastic.co/kibana/kibana:7.16.2
64    container_name: kib01
65    ports:
66      - 5601:5601
67    environment:
68      ELASTICSEARCH_URL: http://es01:9200
69      ELASTICSEARCH_HOSTS: '["http://es01:9200","http://es02:9200","http://e
70 s03:9200"]'
71    networks:
72      - elastic
73
74 volumes:
75  data01:
76    driver: local
77  data02:
78    driver: local
79  data03:
```

```
80      driver: local
81
82  networks:
83    elastic:
        driver: bridge
```

The Experience Discovery Service is implemented using docker-compose. Clone the Github repositories and run 'docker-compose up -d' to launch both XDS and XDS-UI. System Operators will configure the XDS component if there is a need to stand up their own discovery service UI.

Location of XDS GitHub Repository: (https://github.com/OpenLXP/openlxp-xds):

```
1  version: "3"
2
3  services:
4    db:
5      image: mysql:5.7
6      ports:
7        - '3306:3306'
8      environment:
9        MYSQL_DATABASE: "${DB_NAME}"
10#       MYSQL_USER: 'root'
11        MYSQL_PASSWORD: "${DB_PASSWORD}"
12        MYSQL_ROOT_PASSWORD: "${DB_ROOT_PASSWORD}"
13        MYSQL_HOST: ''
14      networks:
15        - openlxp
16
17    app:
18      container_name: openlxp-xds
19      build:
20        context: .
21      ports:
22        - "8100:8020"
23      command: >
24        sh -c ". /opt/app/start-app.sh"
25      environment:
26        DB_NAME: "${DB_NAME}"
27        DB_USER: "${DB_USER}"
28        DB_PASSWORD: "${DB_PASSWORD}"
29        DB_HOST: "${DB_HOST}"
30        DJANGO_SUPERUSER_USERNAME: "${DJANGO_SUPERUSER_USERNAME}"
31        DJANGO_SUPERUSER_PASSWORD: "${DJANGO_SUPERUSER_PASSWORD}"
32        DJANGO_SUPERUSER_EMAIL: "${DJANGO_SUPERUSER_EMAIL}"
```

```
33        AWS_ACCESS_KEY_ID: "${AWS_ACCESS_KEY_ID}"
34        AWS_SECRET_ACCESS_KEY: "${AWS_SECRET_ACCESS_KEY}"
35        AWS_DEFAULT_REGION: "${AWS_DEFAULT_REGION}"
36        REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
37        AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
38        SECRET_KEY_VAL: "${SECRET_KEY_VAL}"
39        ES_HOST: "${ES_HOST}"
40        ES_INDEX: "${ES_INDEX}"
41        LOG_PATH: "${LOG_PATH}"
42        ENTITY_ID: "${ENTITY_ID}"
43        SP_PUBLIC_CERT: "${SP_PUBLIC_CERT}"
44        SP_PRIVATE_KEY: "${SP_PRIVATE_KEY}"
45      volumes:
46        - ./app:/opt/app/openlxp-xds
47      depends_on:
48        - db
49      networks:
50        - openlxp
51
52 networks:
53   openlxp:
54     external: true
```

Location of XDS-UI GitHub Repository: https://github.com/OpenLXP/openlxp-xds-ui):

```
1 version: '3'
2 services:
3   # list of containers to run
4   nextjs:
5     # pointing to the dockerfile locations
6     build: ./
7   nginx:
8     # pointing to the dockerfile location in nginx/
9     build: ./nginx
10    ports:
11      - 80:80
12      - 443:443
```

The Experience Management Service is implemented using docker-compose. Use the docker-compose file below to launch both XMS and XMS-UI. System Operators will configure the XMS component if there is a need to stand up their own management service UI.

Location of XMS GitHub Repository: https://github.com/OpenLXP/openlxp-xms

```
1 version: "3"
2
3 services:
4   db:
5     image: mysql:5.7
6     ports:
7       - '3306:3306'
8     environment:
9       MYSQL_DATABASE: "${DB_NAME}"
10 #     MYSQL_USER: 'root'
11      MYSQL_PASSWORD: "${DB_PASSWORD}"
12      MYSQL_ROOT_PASSWORD: "${DB_ROOT_PASSWORD}"
13      MYSQL_HOST: ''
14    networks:
15        - openlxp
16
17  app:
18    container_name: openlxp-xms
19    build:
20      context: .
21    ports:
22      - "8000:8020"
23    command: >
24      sh -c ". /opt/app/start-app.sh"
25    environment:
26      DB_NAME: "${DB_NAME}"
27      DB_USER: "${DB_USER}"
28      DB_PASSWORD: "${DB_PASSWORD}"
29      DB_HOST: "${DB_HOST}"
30      DJANGO_SUPERUSER_USERNAME: "${DJANGO_SUPERUSER_USERNAME}"
31      DJANGO_SUPERUSER_PASSWORD: "${DJANGO_SUPERUSER_PASSWORD}"
32      DJANGO_SUPERUSER_EMAIL: "${DJANGO_SUPERUSER_EMAIL}"
33      DATA_FILE_NAME: "${DATA_FILE_NAME}"
34      SCHEMA_FILE: "${SCHEMA_FILE}"
35      BUCKET_NAME: "${BUCKET_NAME}"
36      UPLOAD_BUCKET_NAME: "${UPLOAD_BUCKET_NAME}"
37      AWS_ACCESS_KEY_ID: "${AWS_ACCESS_KEY_ID}"
38      AWS_SECRET_ACCESS_KEY: "${AWS_SECRET_ACCESS_KEY}"
39      AWS_DEFAULT_REGION: "${AWS_DEFAULT_REGION}"
40      REQUESTS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
41      AWS_CA_BUNDLE: '/etc/ssl/certs/ca-certificates.pem'
42      SECRET_KEY_VAL: "${SECRET_KEY_VAL}"
```

```
43        LOG_PATH: "${LOG_PATH}"
44      volumes:
45        - ./app:/opt/app/openlxp-xms
46      depends_on:
47        - db
48      networks:
49        - openlxp
50
51 networks:
52   openlxp:
53     external: true
```

Location of XMS-UI GitHub Repository: https://github.com/OpenLXP/openlxp-xms-ui

```
1 version: '3'
2 services:
3   # list of containers to run
4   nextjs:
5     # pointing to the dockerfile locations
6     build: ./
7   nginx:
8     # pointing to the dockerfile location in nginx/
9     build: ./nginx
10    ports:
11      - 80:80
12      - 443:443
```

Finally, the notifications implementation of the OpenLXP enables ECC users to receive notifications from the application. This component is also configurable through the Django administration page. Sender/receiver emails, notification body message, and attachment options can be configured. This allows ECC users to be notified of any errors that occurred during the metadata ETL process. OpenLXP notifications is packaged in PyPi to allow installation during the container build of the XIA component. Below is the Github code repository for the notifications.

- Notifications: https://github.com/OpenLXP/openlxp-notifications

Appendix C: Key Terms

Table 12 below summarizes the acronyms referenced in this document.

Table 7 – Acronyms

Acronym	Term
ADL	Advanced Distributed Learning
API	Application Programming Interface
AWS	Amazon Web Services
CI/CD	Continuous Integration/Continuous Deployment
DoD	Department of Defense
EC2	Elastic Compute Cloud
ECC	Enterprise Course Catalog
EDLM	Enterprise Digital Learning Modernization
EKS	Elastic Kubernetes Service
IAC	Infrastructure as Code
IAM	Identity Access Management
IOC	Initial Operating Capability
JSON	JavaScript Object Notation
OAS	OpenAPI Specification
SIP	Systems Integration Plan
TCP	Transmission Control Protocol
TLA	Total Learner Architecture
UI	User Interface
VM	Virtual Machine
XDS	Experience Discovery Service
XIA	Experience Index Agent
XIS	Experience Index Service
XMS	Experience Management Service
XSE	Experience Search Engine

Acronym	Term
XSS	Experience Schema Service
YAML	YAMLAin't Markup Language