

Overview

Energy Data Insights (EDI) on AWS is now delivered through AWS Managed Services (AMS), which provides managed support for your subsurface data on AWS. Now, you can automatically deploy EDI on AWS and accelerate your data ingestion from weeks to hours and intelligently process and organize your subsurface data with minimal manual effort. Through AMS, the operational capabilities of your team are now extended, allowing you to focus on innovation and accelerate time to value with your subsurface data.

EDI delivers its capabilities through three main components: the Data Platform Portal for managing the underlying EDI Data Platform services, EDI IQ for at-scale data transformation and ingestion, and the Operations Portal for centralized management of multiple EDI instances. With AMS support, you receive 24/7 infrastructure monitoring, streamlined incident response, and enhanced operational reliability for your EDI environment. To learn more about EDI, visit the [product detail page](#).

M24 Data Platform Updates

Note: Refer to the OSDU® Forum [M24 Release Notes](#) for more details and the [appendix](#) for the Project Capability Maturity table.

New Features

The following new forum features have been adopted by AWS for the EDI v1.1.0-M24 release.

Project and Workflow service

The Project and Workflow Service enables separation of data in the Storage service through dedicated namespaces within the OSDU® Data Platform. These namespaces create isolated environments that function as systems of engagement, allowing applications to interact with and modify data without affecting the primary system of record (record of truth). This new capability enables teams to collaborate on data entities in a controlled workspace before committing changes back to the master record, significantly improving data governance and team collaboration workflows. In M24, this service has been implemented following AWS best practices, utilizing Kubernetes Deployments and Istio Service mesh for robust traffic management. The service integrates with existing Storage, Search, and Schema Services, requiring minimal AWS-specific modifications to maintain compatibility with the OSDU® Forum specifications. Refer [to Project and Workflow project](#) for more details. Key Benefits include:

- Enables creation and management of records of engagement separate from records of truth
- Facilitates team collaboration in isolated workspaces without impacting production data
- Provides controlled pathway to promote changes from workspace to master record
- Maintains data integrity while enabling parallel work

Existing feature enhancement/Improvement

AWS contributed towards the community effort to update existing Sandbox/Incubator features, with [Schema Upgrade](#) service being promoted from Incubator to Graduated status. Additionally, we implemented the following enhancements to EDI.

1. Geospatial Consumption Zone (GCZ) service:
 - Now support Authorization Token and data-partition-id via GET query parameters instead of HTTP Headers, improving integration capabilities with third-party mapping services like ArcGIS Online.
2. Seismic DDMS:

- Extended credential model in Seismic DDMS service (core code) with optional field 'region'.
 - Added AWS_REGION information to API responses for getToken APIs.
3. RE DDMS Improvement:
 - Addressed conflict between PostgreSQL automatic updates and Terraform upgrades by upgrading the PostgreSQL server version, ensuring consistent database versioning, and preventing unexpected downgrades during deployments.
 4. **Ingestion Performance improvements**
 - Improved ingestion performance by using the thread pool. Enable scaling of the index-queue by the queue size.

Bug fixes

In addition to the general Forum bug fixes, AWS has implemented the following additional bug fixes for EDI.

1. Fix issues with entitlement service, so that listAllPartitionGroups return the right data.
2. Configured Service Partition Host variable and updated osdu-airflow lib version from 0.27.0 to 0.27.1.
3. Pinned specific compatible versions of packages in requirements.txt to prevent issues with Airflow-sqs-buffer stopping message pulls from SQS.
4. Corrected Elasticsearch namespace issue in the upgrade script, replacing hardcoded 'shared' with the appropriate environment variable.
5. Resolved Indexer service memory leak issue causing Out of Memory (OOM) kills.
6. Corrected Karpenter scaling issue caused by misconfigured podDisruptionBudget in Istio Discovery deployment, resolving scenarios where node scaling was prevented due to stuck Istio Discovery pods.

Technology Updates

1. **Amazon EKS and Amazon EKS add-ons version upgrade**
 - Upgraded the Amazon EKS version to 1.32 in the latest M24 release (from 1.29 in M23) to keep up with the Amazon EKS version's support plan.
 - Amazon EKS add-ons provide installation and management of a curated set of add-ons for Amazon EKS clusters. All Amazon EKS add-ons include the latest security patches and bug fixes, and they are validated by AWS to work with Amazon EKS. Amazon EKS add-ons help you to consistently ensure that your Amazon EKS clusters are secure and stable, and that they reduce the amount of work that you need to do to install, configure, and update add-ons. With release this, we also upgraded the versions of Amazon EKS add-ons to the compatible versions.
2. **Spring framework upgrade**

Shared responsibility with other Cloud Service Providers (CSPs) to upgrade select services from Spring v5 to Spring v6. This upgrade reduces Common Vulnerabilities and Exposures (CVEs) related to Spring v5, enhancing overall security. For details, refer to the forum [M24 Release Notes](#) and the [Spring framework forum epic](#).
3. **Other technology updates**

We made other technology updates to keep the M24 up to date with security enhancements while maintaining EKS compatibility. These updates include eks-blueprints-addons upgraded from 1.16.3 to 1.19.0, Amazon MWAA version upgrade, MongoDB updates, Istio version upgrade, Prometheus update, Karpenter upgrade.
4. **Security vulnerability fixes**

AWS continues to contribute to the ongoing resolution of OSDU® core code security vulnerabilities. For the latest security vulnerability status, visit the OSDU® Forum [vulnerability dashboard](#).

Web Application Updates

Data Platform Portal Enhancements

The Data Platform Portal provides a graphical user interface for managing your OSDU data platform. It streamlines common administrative tasks and data management operations that traditionally required multiple tools or complex command-line interactions. The following minor updates were made to the Data Platform Portal in this release:

User Interface Enhancements:

- Enabled right-click functionality on menu items
- Improved experience for users lacking permissions
- Enhanced user experience for timeout logout scenarios
- Improved User Management list view

Functionality Improvements:

- Enhanced Search Service view service API integration
- Improved FileCollection Load example
- Upgraded platform status check functionality
- Added View Storage Records page in the Data Portal. This allows for viewing a record directly from the OSDU Storage Service, viewing and comparing previous versions of a record, and creating a new version of a record.

Bug Fixes:

- Resolved issue in Manage Schemas view query request
- Fixed login redirect issue

Operations Portal Enhancements

The Operations Portal is an administrative interface for managing multiple EDI instances across your organization. It provides IT Administrators with a centralized view over your EDI deployments. The following minor updates were made to the Operations Portal in this release:

User Interface Enhancements:

- Enabled right-click functionality on menu items
- Added console URL display in Product Fulfillment Information
- Enhanced error message clarity and visibility

Deployment Management Improvements:

- Strengthened development deployment capabilities
- Enhanced deployment parameter management across components

Technical Updates:

- Updated dependency versions
- Enhanced logging and debugging capabilities for improved troubleshooting

EDI IQ Updates

EDI IQ streamlines the process of ingesting data from OSDU-aware and unaware data sources, into your EDI data platform. The following updates were made to IQ in this release:

Existing feature enhancement/Improvement

Seismic Data Ingestion Workflow

The Seismic Data Management capability in EDI IQ has been enhanced to utilize Seismic-DMS APIs, replacing the standard Normal Dataset ingestion process previously used. This new implementation introduces a structured storage approach where .VDS files are organized in a hierarchical path format (TenantID/ProjectID/DataSetID), enabling better data organization and management. The enhancement includes configurable parallel processing for SEG-Y-to-VDS conversions, allowing organizations to optimize their ingestion performance based on available infrastructure resources. Key Benefits include:

- Project-specific file organization: A more structured ingestion approach enables better data organization and scalability.
- Faster ingestion performance: Users can significantly improve ingestion speed, especially for large seismic datasets, by configuring parallel processing.

Job Control Enhancement

Job cancellation functionality is now introduced, enabling users to halt job execution across Source Data Scan, Manifest Generation, Validation, and Ingestion workflow stages. This feature addresses the limitation in v1.0.0 where initiated jobs could not be terminated. Key Benefits include:

- User-driven correction mechanism: Reduces wasted compute and time by allowing users to stop erroneous jobs.
- Operational efficiency: Minimizes delays caused by invalid job executions and improves overall data workflow management.

Technology Updates

1. AWS SDK for Java upgrade

- Upgraded AWS SDK for Java from version 1.x to 2.27.21.
- This update improves performance and maintainability while ensuring compatibility with AWS services.

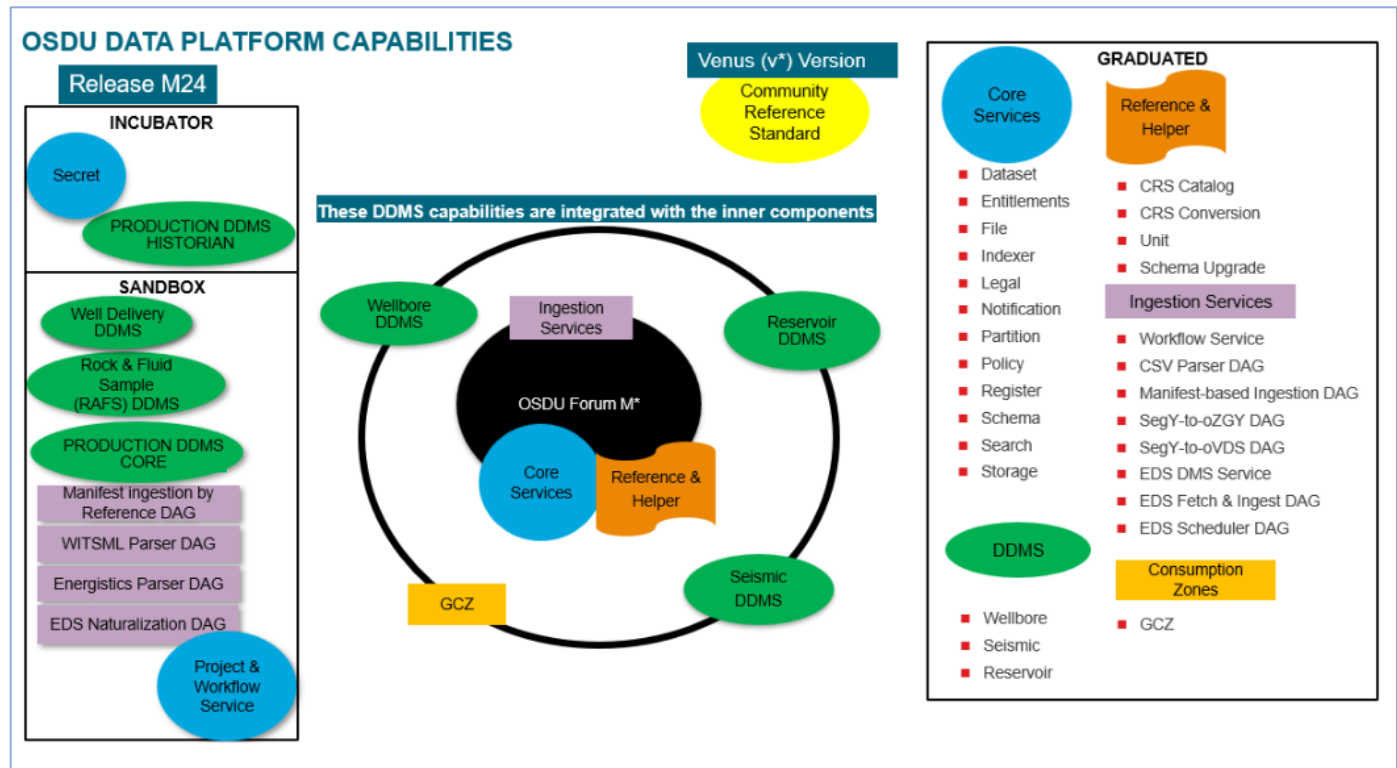
2. Poetry Package Management Implementation

- Implemented Poetry for 'API' and 'Transformer' Python repositories.
- Provides reproducible builds via lockfile and standardized pyproject.toml configuration.
- Implementation for ETL repository in progress.

Appendix

OSDU Community software maturity – M24: What is in AWS

Project Capability Maturity - M24 - Last updated October 21, 2024. Note: OSDU forum authentication is required to access Capability Maturity link.



The table below depicts OSDU® platform features that have been adopted by AWS, and the releases in which they were adopted. See the OSDU® Data Platform Status [here](#).

Maturity	Capability / service group	Projects	Project Management Committee (PMC) notes	AWS notes
Graduated	Core services	<ul style="list-style-type: none"> Dataset Entitlements File Indexer Legal Notification Partition Policy Register Schema Search Storage 		M16 M16 M16 M16 M16 M16 M16 M16 M16 M16 M16 M16
	References and helper	<ul style="list-style-type: none"> CRS catalog CRS conversion Unit Schema Upgrade 		M22 M22 M22 M22
	Domain Data Management Service (DDMS/DMS)	<ul style="list-style-type: none"> Seismic DDMS v4 Wellbore DDMS Reservoir DDMS 		M16 M16 M20
	Ingestion services	<ul style="list-style-type: none"> Workflow service CSV Parser DAG Manifest-based Ingestion DAG SegY-to-oZGY DAG SegY-to-oVDS DAG External Data Services (EDS) DMS service EDS Fetch and Ingest DAG EDS Scheduler DAG 		M22 M18 M16 M16 M16 M16 M18 M18
	Consumption services	<ul style="list-style-type: none"> Geospatial Consumption Zone (GCZ) 		M22
	Core services	<ul style="list-style-type: none"> Secret service 	<ul style="list-style-type: none"> Secret - Pending implementation of https://community.opengroup.org/osdu/platform/security-and-compliance/secret/-/issues/8 	M16
Sandbox	Core Services	<ul style="list-style-type: none"> Project & Workflow Service 		M24
	DDMS	<ul style="list-style-type: none"> Well Delivery DDMS Rock and Fluid Sample DDMS 		M16 M22
	Ingestion services	<ul style="list-style-type: none"> Manifest Ingestion by Reference DAG WITSML Parser DAG Energetics Parser DAG (WITSML Parser v2, RESQML Parser, ProdML Parser) EDS Naturalization DAG 		M23 M16 - M20