

## DS120: Data Integration with MongoDB Atlas

OEM: MongoDB • Code: DS120

### COURSE DESCRIPTION

MongoDB Atlas provides a comprehensive suite of data services that enable developers to query, process, analyze, and visualize data across multiple sources without complex integrations. In this speciality training course, students will master key Atlas data services, including Atlas Data Federation, Atlas Stream Processing, and MongoDB Charts. Atlas Data Federation allows you to natively query and combine data across MongoDB Atlas databases and AWS S3, enabling seamless access to both operational and archival data. Atlas Stream Processing empowers you to process and transform data in real-time as it flows through your systems, enabling immediate insights and actions. Together, these services provide a complete data analysis pipeline — from real-time streaming data to historical data lakes — all accessible through familiar MongoDB tools and APIs. Throughout this course, you'll gain hands-on experience implementing these technologies and learn how to leverage the full power of Atlas for your data analysis needs.

### AUDIENCE PROFILE

- This training is intended for developers and data engineers who have a solid understanding of the MongoDB Query API, attained through the completion of MongoDB Developer Fundamental courses, and a basic familiarity with MongoDB Atlas and aggregations.

### PREREQUISITES

- OFA400: Atlas Admin
- DEV400: MongoDB Developer Extension

### COURSE OBJECTIVES

**By the end of this course, participants will be able to:**

- This section doesn't contain any content.

### COURSE OUTLINE

#### Module 1: Course Outline

- Understanding what Atlas Data Federation and Atlas Stream Processing are and what problems they address
- Deploying a Federated Database Instance and configuring data sources

- Processing real-time data streams using Atlas Stream Processing
- Querying and manipulating federated data using the MongoDB Query API, including federated queries across AWS S3 and Atlas databases
- Building stream processing pipelines to validate, transform, and route streaming data
- Transferring data between Atlas databases and S3 using Atlas Data Federation aggregation operators
- Pre-processing and transforming data using scheduled triggers and aggregation pipelines
- Using MongoDB Charts to visualize both real-time and historical data
- Deploying Atlas Online Archive to manage data lifecycle and enable cost-efficient long-term data retention