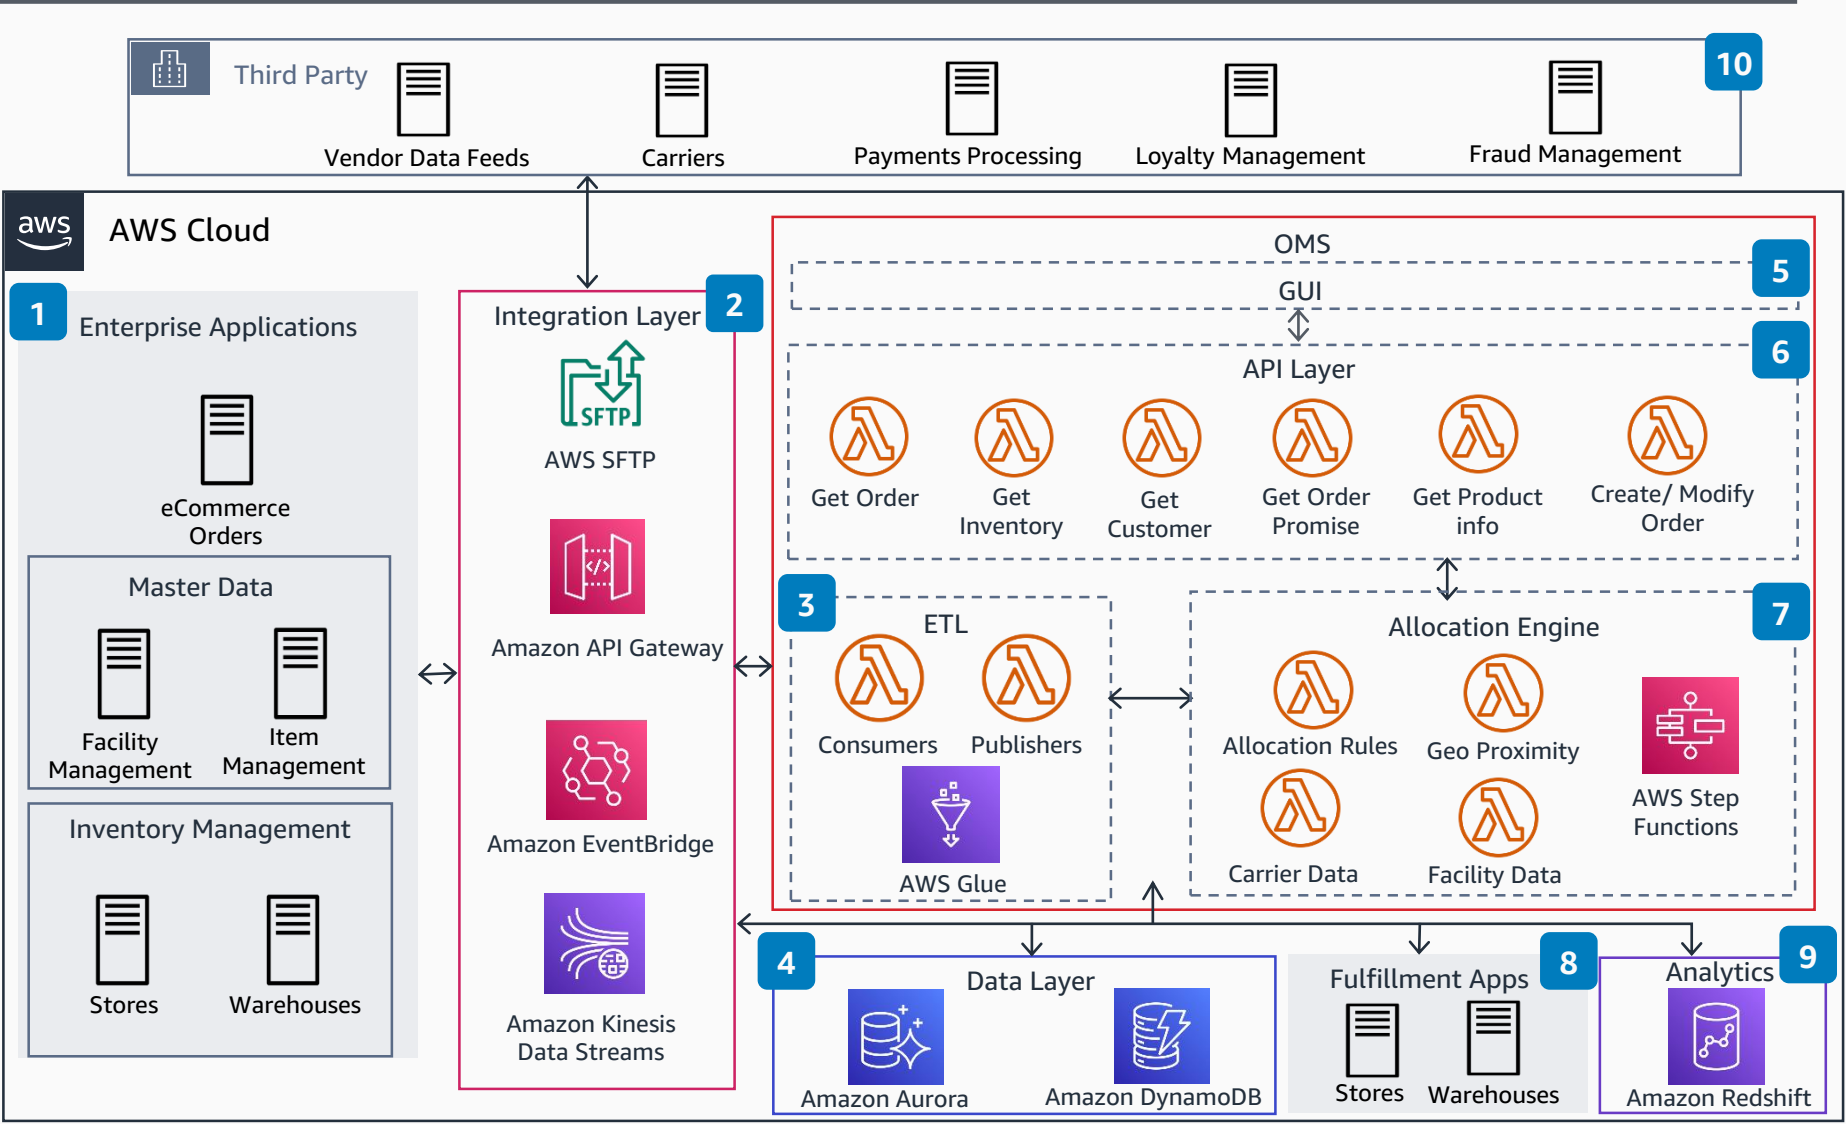


Guidance for Distributed Order Management on AWS

This architecture shows how to set up an OMS and related integrations on AWS.



- 1 Enterprise applications feed data into OMS. This includes facility attribute data (such as store and warehouse data), product data, and inventory data. eCommerce order data is also sent to OMS for allocation and release.
- 2 The integration layer consists of multiple AWS services that support file transfer for external file feeds, APIs, event-driven patterns, and streaming for inventory and master data.
- 3 The extract, transform, load (ETL) layer consists of **AWS Lambda** functions that consume and publish data to **Amazon Kinesis Data Streams** and **Amazon EventBridge**. **AWS Glue** loads and transforms data for batch transactions.
- 4 The data layer consists of **Amazon Aurora** for transactional data and **Amazon DynamoDB**, which serves requests at low latency.
- 5 The OMS exposes a graphical user interface (GUI) that associates will use to create and modify orders, which in turn calls the necessary APIs from the API layer.
- 6 The API layer consists of **Lambda** functions. The presentation layer of OMS and other applications, such as eCommerce, front-end, and customer care, invoke these functions.
- 7 The allocation engine consists of **Lambda** functions and **AWS Step Functions**. These services execute the optimal allocation logic, publish eCommerce orders to **EventBridge**, and identify the appropriate facility to fulfill the order.
- 8 **EventBridge** sends the orders to fulfillment applications. Associates pick and pack the items and send shipment confirmations to OMS.
- 9 Data moved by the integration layer is sent to the analytics layer. **Amazon Redshift** generates insights on order processing efficiency.
- 10 Third-party applications provide functionality based on specific tasks and interact with OMS through the integration layer.