

Ingestion Pipeline

Design & Operations

AUTHOR

K. Adeyemi

DATE

1 April 2026

VERSION

1.0.0

STATUS

Approved

CLASSIFICATION

Internal

OWNER

Platform Eng

1 Overview

This document describes the design and operation of the ingestion pipeline that backs Folio Analytics. It covers the intake path, the processing stages, and the operational guarantees the service is expected to meet under load.

1.1 Scope

The scope is limited to the server-side pipeline from the public intake endpoint to the durable event store.

NOTE

Throughput figures assume the default batch size of 500 events. Tuning guidance is given in Section 2.1 below.

2 Architecture

The pipeline is a sequence of four stages, each independently scalable.

Stage	Responsibility	p99
Gateway	Auth, rate-limit, enqueue	20 ms
Validate	Schema + dedupe	40 ms
Enrich	Geo, session stitch	60 ms
Commit	Durable write + ack	30 ms

2.1 Configuration

```
ingest:  
  batch_size: 500      # events per flush  
  flush_interval: 200ms # max wait before commit
```