Replay Service: Powering the Future of Sports Betting with On-Demand Event Reproduction

A Case Study of Huddle

by Zvonimir Bedi, Senior Backend Engineer

INTRODUCTION

Overview of Replay Events Functionality

In the fast-paced and highly competitive world of sports betting, the ability to reliably reproduce past event data is an invaluable asset. Huddle Tech's Replay service is a powerful **Kotlin-based** backend system designed to enable exactly this functionality. It was built using Spring WebFlux, Kafka, S3 bucket storage, and Kubernetes.

Replay allows for the seamless capture, storage, and re-publication of raw sports event data.

This system enables internal development teams, external operators, and potential customers to **access and replay historical events as if they were occurring in real time**. With robust support for multiple data feeds, replay ensures maximum data fidelity while providing the ability to adapt the payload dynamically for testing, integration, and production simulations.

Importance in the Sports Betting Industry

Sports betting is increasingly driven by real-time data and rapid odds generation. The ability to simulate live betting environments using historical data is essential for maintaining competitive advantages, improving user experience, and ensuring operational stability.

Replay serves as a crucial tool for:

- **Feature development and QA testing** – Allowing product teams to refine and validate new functionalities before deployment.
- **Operator onboarding and system integration** – Helping partners seamlessly integrate Huddle Tech's solutions without relying on live event data.
- **Demo environments for prospective clients** – Showcasing Huddle's cutting-edge technology in real-time pricing scenarios.
- **Performance** – [illegible] models generate precise and reliable odds under various market conditions.
- **Load and** [illegible] **ture to guarantee** scalability and performance.

By addressing these key operational needs, Replay becomes a cornerstone service in Huddle Tech's product suite, enabling both internal teams and external partners to operate with confidence in a fast-paced, data-driven industry.

UNDERSTANDING REPLAY EVENTS

How It Works: Capturing and Storing Events

At its core, Replay functions as a service that consumes raw data messages from multiple Kafka topics. These topics correspond to data feeds from which are then mapped to Huddle Tech's internal IDs.

Once ingested, this data is bundled into compact files (approx. 5MB each), timestamped, and stored in an organized hierarchy within an S3 bucket. This setup ensures:

- Efficient long-term storage
- Easy retrieval by event or date
- Data integrity for later replay

The modular structure also supports future scalability, enabling additional feeds or formats to be incorporated with minimal refactoring. Additionally, Redis is leveraged as a lightweight, high-performance state store to maintain the state of ongoing replay sessions. This enables fault tolerance by allowing the system to recover gracefully from failures and continue replaying from the last known state.

Tracking Market Movements & Odds Changes

Replay doesn't merely store data—it preserves market context and odds changes as they occurred.

This capability is essential for:

Analyzing algorithmic improvements over time
Validating price generation in response to game incidents
Demonstrating real-world performance against competitors

Each snapshot captures live market fluctuations, which are vital for ensuring that replays mimic real-life dynamics.

Real-Time Replay: Reproducing Events On-Demand

Upon receiving a REST API call, Replay initiates a replay session:

Download Phase: Locates and retrieves relevant files from S3.
Caching Phase: Stores the data locally for repeated or future use.
Transformation Phase: Adapts payloads by updating internal IDs and timestamps, ensuring compatibility with current systems.
Publishing Phase: Pushes messages back to the Kafka topics, creating a real-time illusion of a live event.

Replay also supports advanced functionality such as:

Adjustable playback speed: Replay can simulate events at different speeds to suit various testing or demo requirements.
Phase-specific replay: Events can be initiated from the start of a particular period—such as a quarter, inning, or other significant phase of the match—allowing for precise control over the replay timeline.

This methodology allows developers and operators to experience events exactly as they unfolded, with the added benefit of seeing how today's system would react.

Multi-Feed Integration: Combining External Data Sources

Replay was built with the understanding that modern sports data ecosystems rely on multiple sources. It gracefully handles:

Timestamp Harmonization: Ensuring chronological consistency across feeds.

Event Synchronization: Aligning the timing of events from various data sources to ensure accurate reproduction.

Sequence Integrity: Maintaining the correct order of incidents as they originally occurred, regardless of source feed latency.

The integration layer ensures Replay's output is coherent, accurate, and production-ready.

KEY BENEFITS OF REPLAY EVENTS

Faster and More Efficient Operator Integration

Operators integrating with Huddle Tech's platform face fewer roadblocks thanks to Replay. Rather than waiting for live events, they can replay actual historical games and incidents:

- Speed Up QA Cycles
- Identify Bugs in Controlled Environments
- Enable Parallel Development & Testing Teams
- Perform Load Testing and Smoke Tests on Infrastructure

This improves go-to-market timelines and overall integration quality.

Customization & On-Demand Replay

Replay is not a static tool—it is adaptable to the needs of different teams. Users can:

- Select specific matches, leagues, or dates
- Modify playback speed
- Inject test data to simulate edge cases
- Start replays from any point in the event
- Replays can be scheduled in advance for specific times, enabling automated testing or demo sessions using pregame data.
- Replay in loop can be configured to run repeatedly—e.g., every 3 hours—ideal for continuous integration pipelines, long-term testing, or always-available demo environments.

This high degree of customization empowers teams to test more thoroughly and confidently.

USE CASES & REAL-WORLD APPLICATIONS

Operator Benefits: Seamless Integration & Optimization

Replay is essential for **operators** at every stage of their platform's lifecycle. During initial integration, it provides real-game data for testing **UIs, APIs, and internal logic** without relying on live events. Post-launch, it enables ongoing optimization, allowing operators to refine models, adjust trading strategies, and experiment with new market types.

Replay also supports system training and simulation, offering a risk-free environment for onboarding traders and testing operational workflows. By leveraging Replay, operators gain deeper insights into Huddle Tech's odds engine.

Product Development: Improving Market Offerings

Internally, Replay has transformed how Huddle Tech builds and validates new features. It's used for **regression testing**, ensuring updates don't introduce errors, and for benchmarking pricing models, measuring improvements in odds accuracy.

Replay also enhances **incident handling**, allowing teams to refine responses to unpredictable events like lineup change [illegible] periods to ensure platform stability.

This continuous feedback loop ensures each new release is more robust, efficient, and aligned with market demands, benefiting both Huddle Tech and its partners.

Image 1: Huddle Staging Environment, Replay Events

Continue Reading The Article With Huddle Newsletter Subscription

Image 2: Huddle Staging Environment, Replay Events

Continue Reading The Article With Huddle Newsletter Subscription

Image 3: Huddle Staging Environment, Replay Events

CONCLUSION

Summary of Benefits & Competitive Differentiation

Replay is more than a testing utility—it is a strategic advantage for Huddle Tech. By offering real-time replays of historical sports events, it supports better testing, faster integration, and more compelling demonstrations. It:

- Eliminates the dependence on live data for testing
- Delivers accurate, incident-rich simulations
- Enhances pricing algorithm development and validation
- Provides operators with realistic integration environments
- Enables load testing and phase-specific event simulation

This capability sets Huddle Tech apart in the sports betting industry.

How Operators & Clients Can Leverage Replay Events

Operators are encouraged to [illegible] whether onboarding to Huddle Tech for the first time or optimizing post-launch, Replay provides the flexibility, realism, and speed required to [illegible]

By combining innovative engineering with a deep understanding of industry needs, Replay positions Huddle Tech and its partners for long-term success.