NSTREAM.

Going Customer 360 with Kafka, Flink, & SwimOS

Real-time data streaming for entities & events

03/17/2023

© 2023 Nstream Inc. Reproduction and distribution of this presentation without written permission of Nstream is prohibited. All rights reserved.

Fred Patton

- > Nstream Developer Evangelist
- Long-time backend engineer, generalist, and polyglot
- Serverless and fullstack by night
- > Avid technologist
- Lover of the arts and various spoken languages





Customer 360 Use Case for

large online fashion retailer

- > Purchase history aware
- **Browse/Search history aware**
- Social Media action aware
- > Geo-location aware
- > Inventory aware
- > Weather aware
- > Physical event aware
- > Third-party data aware

Relevant Business Entities



relates to relates to CustomerDetails +id: integer ThirdPartyPurchases name: string +id: integer email: string customer id: integer phone: string purchase_history: string gender: string birthdate: date address: string relates to has OrderHistory +id: integer Reviews Returns customer_id: integer Wishlist order date: date +id: integer +id: integer total amount: float +id: integer customer_id: integer customer_id: integer order_status: string customer_id: integer product_id: integer order_id: integer product_id: integer product_id: integer review_text: string return_date: date date added; date rating: integer tier: string quantity; integer return reason: string product_price: float date_added: date return_status: string product_discount: float lovalty points earned; integer contains Products PaymentDetails WeatherEvents +id: integer product_name: string +id: integer +id: integer category: string customer_id: integer customer id: integer subcategory: string order_id: integer event_type: string price: float payment_date: date event_date: date discount: float payment method: string temperature: float brand; string payment amount: float weather condition: string gender: string age_group: string generates extends generates generates 0..* 0...* LocationEvents SearchEvents **SocialMediaEvents** +id: integer +id: integer +id: integer customer_id: integer customer id: integer customer_id: integer event_type: string event_type: string event_type: string event_date: date event_date: date event date: date latitude: float search_query: string event_details: string longitude: float search_results: string

ThirdPartyPsychographics

customer_id: integer

interests: string

values: string

personality_traits: string

+id: integer

ThirdPartyDemographics

customer_id: integer

+id: integer

age: integer

income: float education_level: string



marketing_opt_in: boolean

© 2023 Nstream Inc. Reproduction and distribution of this presentation without written permission of Nstream is prohibited. All rights reserved.

Relevant Business Entities



N

without written permission of Nstream is prohibited. All rights reserved.

Focal streams

- > Customer details
- Order history
- > Wishlist
- Loyalty program stats
- Products
- > Weather events
- Social media events
- Location events
- Search (and browse) events
- Third-party purchases
- Third-party demographics
- Third-party psychographics





Taking Customer 360 to Its Logical Conclusion

CRMs with up to the moment information, and decision automation in lockstep with analytics

66

Righteous Analytics: Right Data. Right Time. Right Way — David W. Johnson"

The right time is...

- » 🛛 Within days 🔽
- 📎 🛛 Within hours 🔽
- 📎 Within minutes 🔽
- 📎 🛛 Within seconds 😲
- Within milliseconds 😱

Kafka + Flink + SwimOS

ETL + Analytics + Business Logic



High-Level Architecture



We begin with an event-driven application where microservices store events in Kafka. Flink will be used to provide enrichments for analytics, such as windowing, summations, and aggregations. SwimOS will ingest application events from Kafka along with analytics results from Flink and apply business logic in the same real-time. Where a moment later, a real-time UI will consume streaming entity state directly from SwimOS, despite being outside the network.



Entities

- Simple object state transmitted as a binary diff to its consumers
- Multiplexed with other entity state needed by respective consumers
- Consumer is notified via an update handler at which point the binary diff has already been applied
- Consumer simply executes its specific logic in the same way as it would without streaming

Events

- Event is replicated to its consumers
- The event captures the minimal fields needed to make record of an occurrence with fidelity
- Each consumer processes the given event and takes appropriate action
- Any information not contained in the event has to be retrieved somehow adding latency

Real-Time Event Data

≫ The application and analytics data streams are made up of events partitioned by event type



- This is standard operating \gg procedure to maximize throughput for ETL and analytics
- ≫ It is also critical to optimize for data-loss prevention





Real-Time Entity Data

- Streams coming out of SwimOS contain entity state partitioned by entity ID so there is no need for stream joins
- Authorized consumers outside the network may sync entity data in real-time even at the field level
- Consumers may also subscribe to lifecycle methods to get context beyond state changes
- Consumers may also invoke commands that stream back the other way



Entity Partitioning



Real-Time Data Observability

- Data does not accumulate unbounded, rather it moves through a window of relevance for business logic processing
- Entity-specific business logic dictates the point at which information decays, either immediately or scaled within the step function of roll-ups
- Data that will never be stored can still be processed, in which case the extracted value can be conditionally stored



Exhibit A: An Intractable Many Stream Join

To add insult to injury, this 501-line massive stream join in Flink can't even lift itself off the tarmac let alone satisfy any real-time SLAs. It is not Flink's fault. It is the wrong algorithm for the problem.



Exhibit B: Entity Partitioning w/ Differential State Sync

});

Here, we find simple application-style logic, where a function called runBusinessLogic is an application defined method that re-runs its business logic when associated state changes. In this case, the logic does not have to be expressed in terms of mathematical operators, but is a simple user defined function. Outside of the user defined method, the code is less than 200 lines.

@Override

```
public void didStart() {
super.didStart();
orderHistoryJoin = this.<Integer, OrderHistory>joinValueLane()
.didUpdate((key, newValue, oldValue) -> {
    runBusinessLogic(ORDER_HISTORY_CHANGED);
    });
```

paymentDetailsJoin = this.<Integer, PaymentDetails>joinValueLane()
.didUpdate((key, newValue, oldValue) -> {
 runBusinessLogic(PAYMENT_DETAILS_CHANGED);
});

```
wishlistJoin = this.<Integer, Wishlist>joinValueLane()
.didUpdate((key, newValue, oldValue) -> {
    runBusinessLogic(WISHLIST_CHANGED);
});
```

reviewsJoin = this.<Integer, Reviews>joinValueLane()
.didUpdate((key, newValue, oldValue) -> {
 runBusinessLogic(REVIEWS_CHANGED);



About Nstream...

Nstream has been pursuing a vision of a real-time streaming Internet powered by Web Agents — web-addressable as the name suggests — and delivering on much of the promise of Digital Twins and Distributed Actors along the way.

https://www.linkedin.com/pulse/st ate-streaming-data-christophersachs





Get started!

- READMEs: <u>https://github.com/swimos/swim</u>
- Developer site: <u>https://www.swimos.org/</u>
- Sample Applications: <u>https://github.com/swimos</u>
- Cookbook: <u>https://github.com/swimos/cookbook</u>
- > Tutorials: <u>https://github.com/swimos/tutorial</u>
- Yours truly: <u>https://www.linkedin.com/in/thoughtpoet/</u>



Thank You!



© 2023 Nstream Inc. Reproduction and distribution of this presentation without written permission of Nstream is prohibited. All rights reserved.