



## In-Memory Computing Patterns for High Volume, Real-Time Applications

Narendra Paruchuri American Airlines Murali Ande American Airlines



#### 20/10/18

#### Outline

- Who We are
- Our Use Cases
- Our Journey
- Evaluation Criteria
- Our Journey with In-Memory data models
- Architecture Patterns
- Advantages
- Being mindful



#### Who we are

500,000 daily customers 6700 flights 350 destinations 50 countries



0

Multiple Hubs CLT, ORD, DFW, LAX, MIA, JFK, LGA, PHL, PHX, DCA



Complex aircraft turn activities before on time departures

0



Oneworld alliance 14,250 flights, 1,000 destinations 150 countries



Safety and Regulations



C

 $(\mathbf{0})$ 

0

Complex weather situations, reroutes, off-scheduled operations, travel plan changes, last minute

 $(\mathbf{O})$ 

## **Our Use Cases**

- Customer connections and Service recovery
  - 12 Hubs.
  - Connections delays, tight and missed connections.
  - Rebooking.
  - Expedite Immigration.
  - Taking care of AAdvantage customers and their bags.

Aircraft Turn Management Aircraft turn events Customer boarding **Bags** loading Crew check-ins. **Customer check-in** Fueling **Cabin Cleaning** Pro actively finding flight delays **Alerts and notifications** 

## **Our Journey**



While systems are modified for key value pair keeping in view the trends in latest technologies, Our business operations still require us to perform joins to correlate and 5 coalesce data to facilitate business decisions.



#### **Evaluation Criteria**

- ACID compliance with native persistence and third party persistence integration
- Readability and Maintainability of System
- Support High Volume Transactions with simultaneous updates to multiple attributes (12K / Min – Writes , 25K / Min – Reads)
- Replication and Distribution support with Multi DC
- Streaming connectors like Kafka
- Generate events on data updates
- Docker and Kubernetes support



#### 

Pain Points

- Most of the solutions does not offer Joins
- Lot of code to correlate and Coalesce data
- Slow response times.



#### Pain Points:

- Pre-defined Object Structures
- Reduced Flexibility
- Longer Deployment Cycles







8





9



- Our business operations still require us to perform joins to correlate and coalesce data for the business. 10
- While there may be several options, we have chosen to implement Ignite for our use cases.



#### Apache Ignite (Open Source) – Active – Passive topology



In-Memory Computing

#### Apache Ignite (Open Source) – Active – Passive topology



#### Apache Ignite and Cassandra Active-Passive topology



# Advantages of using Ignite along with Cassandra

- Bandwidth and Response times improvement (order of magnitude improvement).
- Improved Availability than standalone Cassandra system, as Ignite offers sophisticated clustering support, such as detecting and remediating split brain conditions.
- Horizontal and Vertically scalable.
- More efficient, as Ignite can use all the memory available on a node, and not only JVM memory.
- ANSI-99 SQL and ACID Transaction Guarantees (Improved Consistency)
- You can run other analytics off of Cassandra
- Support for jdbc and odbc make it easier to integrate with existing tech, such as Hibernate and Spring Data.
- No data remodeling required for existing Cassandra deployment, as Apache Ignite can read from it as well as relational databases.



#### Apache Ignite (GridGain) – Active – Active topology



#### 16/10/18

#### Being mindful

- Not a magic bullet
- Evaluate use case
- Consider future data growth
- Take advantage of all available technologies
- Lack of standards
- Not plug and play
- Vendor challenges
- Governance



## Why not traditional RDBMS database

- Database will work but at what cost ?
- Complexity of Architecture with Multi DC and Multi Cloud providers.
- Tool consistency
- Scalability Vertical vs Horizontal
- No collocated data processing
- Leveraging auto scaling using Containers and Kubernetes is not available.

Cloud Scaling Breaks Domain Driven Design





#### Advantages

- •Extremely fast response times
- •Highly scalable
- •Improved complex processing on events.
- Improved real time analytics
- Reduced cost of operation
- Inexpensive with improvements in memory technology



# **QUESTIONS?**







## In-Memory Computing Patterns for High Volume, Real-Time Applications

Narendra Paruchuri American Airlines Murali Ande American Airlines

