

THE LEADING IN-MEMORY COMPUTING PLATFORM

NODE 01

NODE 04

NODE 05

# Low-Latency Machine Learning Applications

John DesJardins – CTO N. America



NODE 06 NODE 06 ob.select ob.selected

# > Al Techniques Continue to Expand & Evolve



### Online Machine Learning within an In-Memory Platform

Move from Reactive to Pro-Active

Take Action <u>before</u> negative impact or <u>ahead of</u> opportunity

Hazelcast In-Memory Platform – Fast Data



#### **Online Inference & Nearline Machine Learning**



# The Hazelcast Difference

#### > Example: Credit Card Processing





#### **Hazelcast - High Performance Platform**

		IMDG	In-Memory Data Grid		
	-	Integrate APIs, Microservices, Notifications	<b>Communicate</b> Serialization, Protocols	I.	
Mobile		<b>Store/Update</b> Caching, CRUD Persistence	<b>Compute</b> Query, Process, Execute	4	
Apps		Scale Clustering & Cloud, High Density	<b>Replicate</b> WAN Replication, Partitioning	((	
Social		<b>Secure</b> Privacy, Authentication, Authorization	<b>Available</b> Rolling Upgrades, Hot Restart		
ooda		Jet In-Memory Streams			
Commerce	₩	Ingest & Transform Events, Connectors, Filtering	<b>Combine</b> Join, Enrich, Group, Aggregate		
	•	<b>Stream</b> Windowing, Event-Time Processing	Compute & Act Distributed & Parallel Computations		
Communities	$\star$	<b>Secure</b> Privacy, Authentication, Authorization	<b>Available</b> Job Elasticity, Graceful shutdown	Ľ	
		Managemen	t Center		
		Secure   Mana Embeddable   Scala Secure   Resilie	age   Operate able   Low-Latency ent   Distributed		





Live Streams Kafka, JMS, Sensors, Feeds

Databases JDBC, Relational, NoSQL, Change Events

Files HDFS, Flat Files, Logs, File watcher

> Applications Sockets



# Evolution of Stream Processing

1 <sup>st</sup> Gen (2000s) Hadoop(batch) <u>or</u> Apama(CEP) <i>hard choices</i>	Distributed Batch Compute – MapReduce – scaled, parallelized, distributed, resilient, - not real- time or   Siloed, Real-time – Complex Event Processing – specialized languages, not resilient, not distributed(single instance), hard to scale, fast, but brittle, proprietary
2 <sup>nd</sup> Gen (2014) Spark <i>hard to manage</i>	<b>Micro-batch distributed</b> – heavy weight, <u>complex to manage</u> , not elastic, require large dedicated environments with many moving parts, not Cloud-friendly, <u>not low-latency</u>
3 <sup>rd</sup> Gen (2017 Jet & Flink) flexible & scalable True "Fast Data"	Distributed, real-time streaming – highly parallel, true streams, advanced techniques (Directed Acyclic Graph) enabling reliable distributed job execution <u>Flexible deployment</u> - Cloud-native, elastic, embeddable, light-weight, supports serverless, fog & edge. <u>Low-latency</u> Streaming, ETL, and fast-batch processing, built on proven data grid



## Streaming Performance

#### Streaming Word Count - Average Latency (lower is better)

1 sec Tumbling Windows



Messages / sec

\* Spark had all performance options, including Tungsten, turned on

Hazelcast.org – Data Grid

https://github.com/hazelcast

Jet.hazelcast.org – Streaming Analytics

Demos: <a href="https://jet.hazelcast.org/demos/">https://jet.hazelcast.org/demos/</a>

My info:

•John DesJardins

• john.desjardins@hazelcast.com - email

•@johnmdesjardins - Twitter



# Thank You

**hazelcast** 

519648891000x58094565