## THE IOT APPLICATION CHALLENGE HANDLING MASSIVE STREAMING DATA

In Nemound

COLIN MACNAUGTHON NEEVE RESEACH

## WHO IS NEEVE RESEARCH?

- Headquartered in Silicon Valley
- Creators of the X Platform<sup>TM</sup> Memory Oriented Application Platform.
- Passionate about high performance computing.
- Running in production at Fortune 100-300





# What is IoT ... What are the Challenges?

# How The X Platform tackles Streaming

# Streaming Usecase: IoT Fleet Tracking



## WHAT IS IOT

The "Internet of Things": "real world" stuff (often augmented with sensors) streaming data to a network



### WHAT WE ARE REALLY TALKING ABOUT IS: LARGE SCALE STREAMING



## WHAT IS NEEDED FOR IOT

	>	
--	---	--

**EVENT-DRIVEN** Its all about *streaming* lots of events



#### SCALABILITY

Lots of things  $\longrightarrow$  LOTS of events



#### SPEED

100s of thousands to millions of events/sec, response latency in microseconds or low millis.



#### RELIABILITY

CANNOT lose mission critical events No Dups / No Loss (Exactly Once)

#### AVAILABILITY

Always On, Always available in the face of network/process/machine/data center failure

#### AGILITY/EASE

Applications are infinite need to be able to evolve organically



## STREAMING APP CHARACTERISTICS

#### What do they do?

- I. Consume Inbound Messages
- 2. Read / Update State
- 3. ... and Produce Outbound Messages





## MICROSECONDS MATTER

A processing time of Ims limits your throughput to 1000 messages / sec.

Same applies to any synchronous callouts in the stream.

To achieve >10k Transactions/Second you must leverage In Memory technologies



## MICROSECONDS MATTER

Memory	Latency
LI Cache	~lns
L2 Cache	~3ns
L3 Cache	~I2ns
Remote NUMA Node	~40ns
Main Memory	~100ns
Network Read	100μs
Random SSD Read 4K	l 50μs
Data Center Read	500µs*
Mechanical Disk Seek	10ms

#### MEMORY ORIENTED COMPUTING!



All State in Memory All The Time!



Non Starters For Performance We're Talking About!

Sources: https://gist.github.com/jboner/2841832 http://mechanical-sympathy.blogspot.com/2013/02/cpu-cache-flushing-fallacy.html



## THE CHALLENGES

- Exactly Once Semantics
  - Messaging No Loss / No Dups / Atomic
  - Storage and Access to State No Loss / No Dups
- Atomicity between Message Streams and State Updates Receive-Process-Send atomic





## TRADITIONAL TP APPLICATION ARCHITECTURE





## LAUNCH DATA INTO MEMORY





## DATA GRAVITY (DATA STRIPING + SMART ROUTING)





## WHY STILL SLOW AND COMPLEX

#### How Slow?

#### Latency

- I0s to 100s of milliseconds
- Throughput
  - Very low with single pipe
  - Few 1000s per second with high concurrency
- Why Still Slow?
  - Remoting out of process
  - Synchronous data management and stabilization
  - Concurrent transactions are not cheap!

#### • Why Complex?

- Transaction Management still in business logic
- Thread management for concurrency (only way to scale)
- Data transformations due to lack of structured data models

## THE X PLATFORM APPROACH



In-Memory Computing 2017

## X PLATFORM TRANSACTION PIPELINING (HA)





## NOW WHAT IS THE PERFORMANCE?

#### How Fast?

- Latency
  - I0s of microseconds to low milliseconds
- Throughput
  - I00s of thousands of transactions per second
- How Easy?
  - Model Objects and State in XML, generated into Java objects and collections.
  - Annotate methods as event handlers for message types.
  - Single threaded processing
  - Work with state objects treating memory as durable.
  - Send outbound messages as "Fire And Forget"
  - Shard applications by state, messages routed to right app.



## **RELIABILITY – EXTERNAL DATA STORES**



## STREAMING APPS ON THE X PLATFORM



State as Java State in Local Memory Ultra Performance Zero Garbage Fully Fault Tolerant Zero Loss Horizontally Scalable

✓ Message Driven
 ✓ Stateful
 ✓ Multi-Agent

Totally Available
 Horizontally Scalable
 Ultra Performant





# Building a Fleet Tracking System with The X Platform



## IMPLEMENTING GEOFENCING

- We have a fleet of vehicles.
  - (cars, trucks, whatever)
- Each vehicle Should be following a route defined by Administrators
- Our Fleet Management System needs to:
  - Track location of vehicles to ensure routes are being followed.
  - If a vehicle leaves its route, trigger alerts.



## FLEET GEOFENCING





## THE CODE

<u>Messaging</u> Annotation based handler discovery,> Single Threaded	<pre>@EventHandler final public void onLocationEvent     try {</pre>	<u>Message</u> Pkain Old Java Object Generated from XML Model	State Management Plain Old Java Objects Generated from XML Mo	odel ) throws Exception {
State Management Plain Old Java objects and Java Collections	<pre>validate(message); } catch (Exception e) {     invalidEventCount.increme     return; } // look up vehicle's route: final VehicleRoute route = _v&gt; if (route == null) {     droppedCount.increment();     return; }</pre>	nt(); ehicleRoutes.get(message.getVehicl	eID());	
<u>State Management</u> Object Pooling and Preallocation for Zero Garbage	<pre>// update state Vehicle vehicle = state.getVe if (vehicle == null) {     vehicle = Vehicle.create(     vehicle.setVehicleID(mess     state.getVehicles().put(m }</pre>	hicles().get(message.getVehicleID( ); age.getVehicleID()); essage.getVehicleID(), vehicle);	)); St Re BC	<u>ate Management</u> ate Changes transparently eplicated to Hot ackup and/or Disk Based Journal
<u>Messaging</u> Create and populate –––– "Fire and Forget"	<pre>// check if vehicle is within if (!updateVehicleLocationIfI</pre>	<pre>its route bounds, and alert if no nRouteBoundary(vehicle, route, mes = populateRoutViolationEvent(vehi e("alerts", alert);</pre>	t sage.getLocation(), messag cle, message.getLocation()	ge.getSpeed())) { ), message.getSpeed());
	<pre>// update stats processedCount.increment(); }</pre>			
In-Memory Pure Busine Computing	ess Logic – E	xactly Once	Processin	g

## **IOT FLEET GEOFENCING**



## Location Updates Events/sec: > 30k

Single Shard, I Processor Core, Replicated.



Ims Response Time. Full HA (Replicated), Exactly Once

## WHY X?

- Easy to Build
  - Focus on domain
    - Pure Java
- Easy to Maintain
  - Pristine domain
    - No infrastructure bleed
- Easy to Support
  - Stock hardware
  - Small Footprint
  - Simple abstractions
  - Easy tools
- Very, very fast

Memory mputing

## ✓ No Compromise

Agility, Availability, Scalability, Performance

## GETTING STARTED WITH X PLATFORM™

**Getting Started Guide** 

https://docs.neeveresearch.com

Get the Demo Source

https://github.com/neeveresearch/nvx-apps

We're Listening

contact@neeveresearch.com



## QUESTIONS



