


Sensor Data Analytics Acceleration with Apache Edgent



Brandon Swink– swink@us.ibm.com
IT Architect, IBM

Apache Edgent is currently undergoing Incubation at the Apache Software Foundation.

Edge Analytics with Edgent



- What is the edge?
- Edge Examples
- The value of analytics at the edge
- Apache Edgent (formerly Quarks): Open source analytics acceleration
- Example and Demonstration
- Broader architecture
- Getting Started



What is the Edge?

- Constrained compute platform
 - Due to cost, weight restrictions, space constraints, borrowing resources, ...
- Limited connectivity to central systems
 - Limited by expense, bandwidth
 - Periods of being disconnected
- Access to sensors for system being analyzed
 - Directly or through a bus (e.g. CAN bus)
 - Potential to control system based on analysis
- Can be mobile or static
- Expected to be thousands to millions of devices
 - Internet of Things




Edge Examples

- Vehicle
 - Car, truck, race car, bike, train, bus, boat, drone, plane, ...
 - Analyze engine sensors to predict/reduce chance of failure
 - Mobile, may lose connectivity
- Building/Store/Warehouse
 - HVAC, climate, energy use, motion sensors, ...
- Server in machine room
 - Analyze load, cpu temps, rack temps
- Raspberry Pi with a couple of sensors
 - Cheap \$5+
- Smartphone
- Industrial machine

The value of analytics at the edge?

- Reduce communication cost
 - Send relevant data when an event of interest occurs
 - Heartbeats alone may not contain enough data or be too late to take action
- React locally to events
 - More intelligent decision making on the device
 - Execute analytics while disconnected
- Collaborate with related devices
 - Learn from devices with similar characteristics
 - or location





Apache Edgent

- **A community** for accelerating Edge Analytics
 - Open Source, incubating at Apache Software Foundation
 - <http://edgent.incubator.apache.org/>
 - <http://wiki.apache.org/incubator/QuarksProposal>
 - Extensible SDK with functional flow API for streaming analytics
 - Initial support for Java 8,7 & Android,
 - Goal is to support multiple languages with priorities driven by the community
- A modular, lightweight and embeddable runtime



Current Features

- Functional Flow API for **streaming analytics**
 - Per-event and windowed processing with basic analytics
 - [Map](#), [FlatMap](#), [Filter](#), [Aggregates](#), [Split](#), [Union](#), [Join](#), [DeadbandFilter](#)
- Connectors
 - Messaging systems & data stores
 - [MQTT](#), [HTTP](#), [Web Sockets](#), [JDBC](#), [File](#), [Apache Kafka](#) and [IBM Watson IoT Platform](#)
- Micro-kernel style runtime with multi-platform support
 - Small-footprint edge devices or sensors
 - [Including Raspberry Pis or smart phones](#)
- Development mode
 - Web-console for viewing application graph and metrics
- Testing mechanism
 - Junit integration



Streaming Analytics Paradigm

- A stream is a infinite sequence of tuples
 - Events, sensor readers, location updates, ...
- Everything is a stream ...
 - Source streams bring the raw data to be analyzed
 - Functions **are applied to each tuple** on a stream to produce new streams
 - Filters – Only temperatures greater than 100°C
 - Map – Convert a position to a distance from another position
 - Sink streams send data to external systems (e.g. messages to a back-end)
- Or a window
 - A window is an ever changing subset of a stream
 - Last ten minutes
 - Last 30 tuples (data values)





Example Demonstration



Water conservation

- Using Smarter sprinklers
 - LOCAL, Continuous, real time analysis to determine the need to water
 - Connects to back end systems for in-depth analysis
 - Checks soil moisture
 - Takes weather into account
 - Honors government water restrictions






Smart sprinkler – Raspberry Pi
Local analytics running on Apache Edgent





Smart sprinkler device

The hub that connects the actual sprinkler to back end systems



Analytics to determine whether the sprinkler can be turned on

- Weather forecast
- Governance (water ban)



Rapid web UI development



Visualization server

Connects the browser UI to other components



The Weather Company API
Provides weather data




Why Use Edgent?

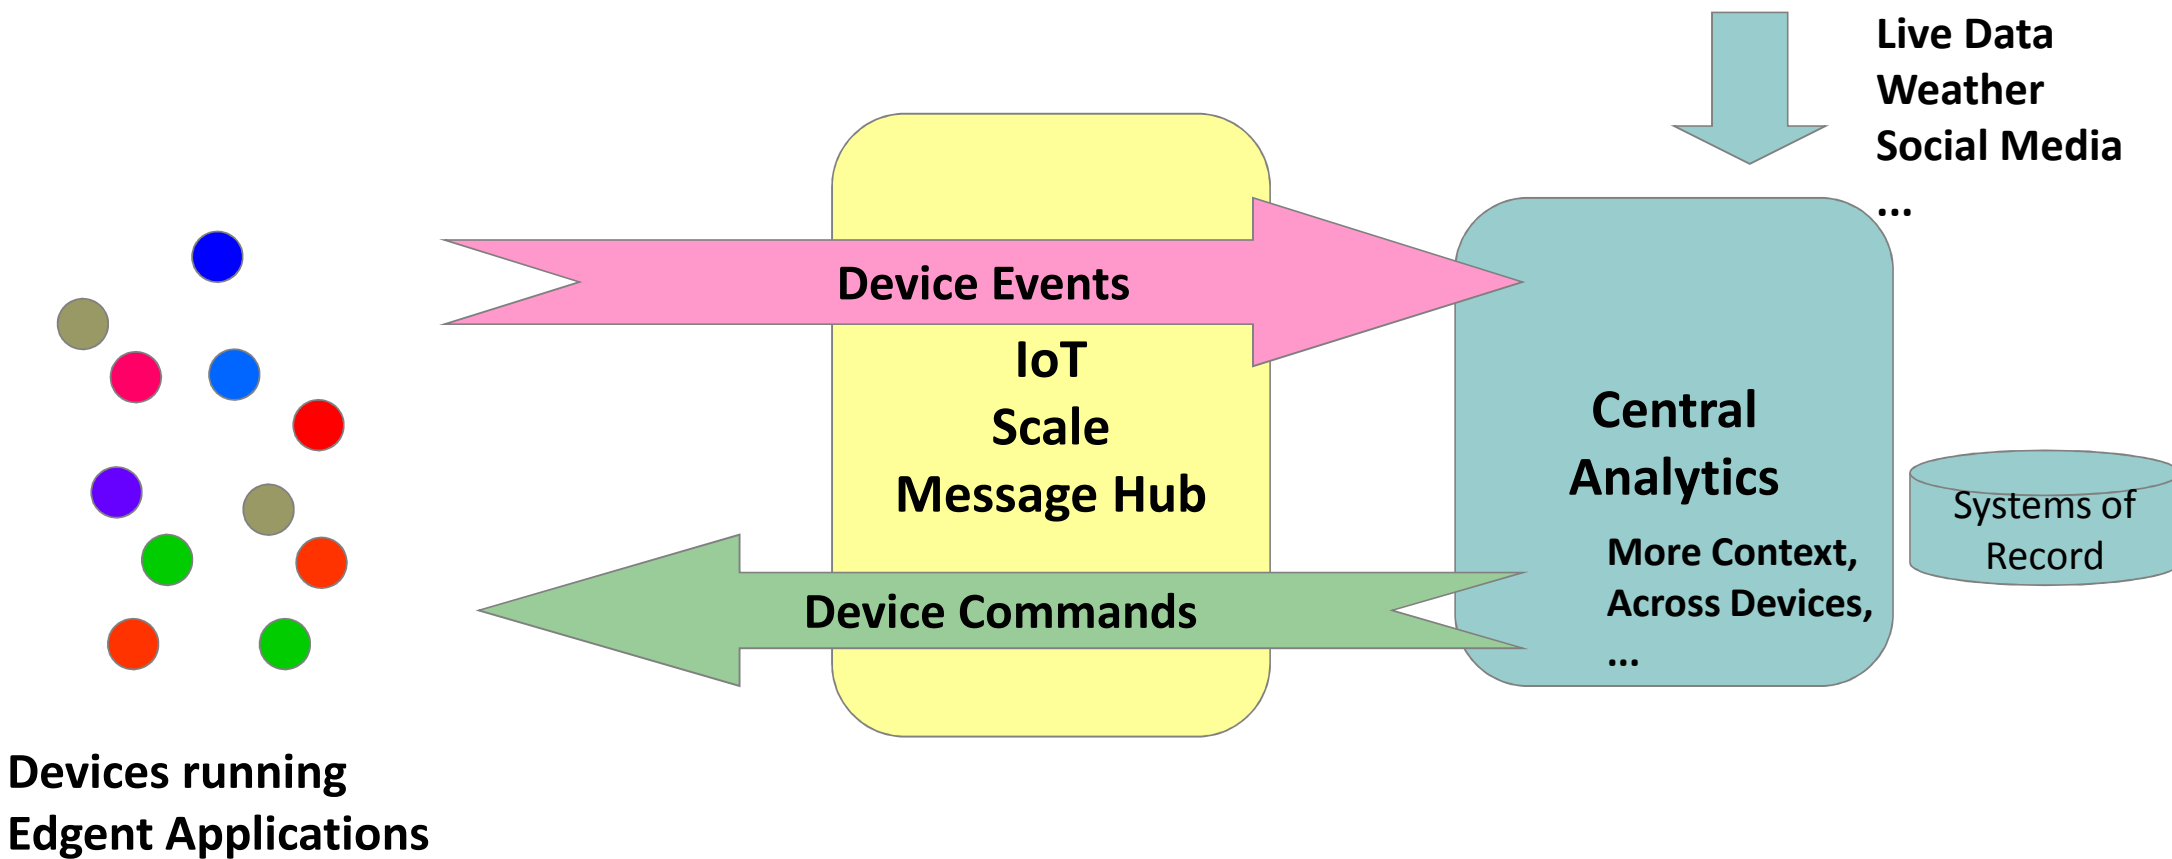
Broader Architecture: Integration with Centralized Deeper Analytical Platforms



- Integrates with centralized analytics systems through IoT scale message hub
 - Any hub
 - Any central system



Broader Architecture: Control loop through Central Analytics





Getting started with Apache Edgent

- Build from source to get the latest version
 - Fork/Clone/download source from github.com
 - Apache/incubator-edgent
 - <https://github.com/apache/incubator-edgent>
 - Download Java 8, Apache Ant, Junit, Jacoco
 - Details see: DEVELOPMENT.md
 - <https://github.com/apache/incubator-edgent/blob/master/DEVELOPMENT.md>
- Getting started guide
 - <http://edgent.incubator.apache.org/docs/edgent-getting-started>
- Other Meetups and Hangouts
 - https://www.youtube.com/channel/UC_uXzbJmQzPkODsEE2E7PgA