Collecting IoT Data in InfluxDB

DAVID G. SIMMONS

SENIOR DEVELOPER EVANGELIST

@DAVIDGSIOT

influxdata®

The Modern Engine for Metrics and Events

Who, What and Why is InfluxData

Overview

- ✓ Founded in 2013
- ✓ Delivering a modern open-source platform for metrics and events
- ✓ Guiding principles:
  - Developer Happiness
  - Ease of Development + Scale Out
  - Time to Awesome
- ✓ Results
  - 70,000+ Active Servers
  - 300+ Customers

Specialized Platforms

Orders, Customers etc. *(Data with Relationships)* | **SQL**

Web pages, Documents, etc.(Text Data) | **Search**

Volume/Variety of Data. *(Big Data)* | **Big Data**

Events and Metrics (Time-Series Data) | **Time-Series**

What makes

time series

different?

1. **High volumes of streaming data**
2. **Support for Time-based Functions**
3. **Need for Scalability and Availability**

InfluxData Open Source Projects

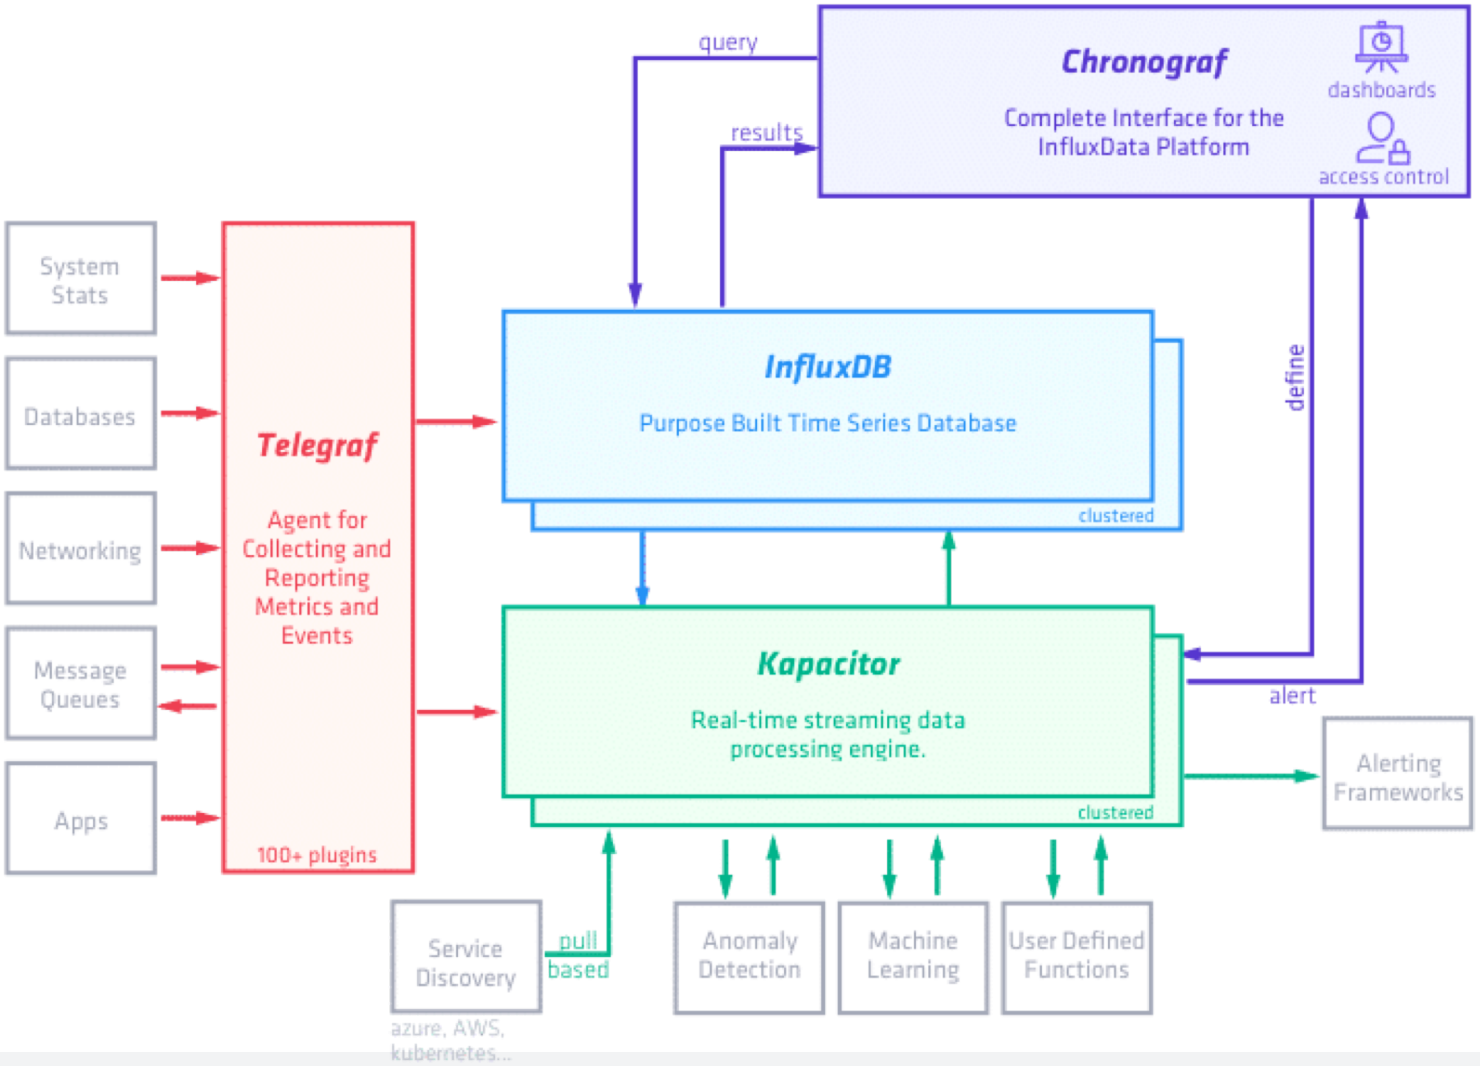

*Note: Clustering only in Commercial Offerings

IoT Data Layer

IoT Data Needs

- IoT data MUST be
  - Timely – ingestion rates and query efficiency is key
  - Accurate – data integrity and platform reliability is important
  - Actionable – data visualization, anomaly detection & alerting are essential
  - Deployable in the datacenter and at the edge
- IoT deployments are struggling to find efficient, scalable, data platform that meets all of these criteria

IoT Platform Needs

- Extensive protocol support for data ingestion
  - BLE, Zigbee, Zwave, CoAP, MQTT, MODBUS
- Robust offline data capability
  - Normalization of data
  - Storage for local data
  - Local event processing engine
- Hardware-agnostic, scalable architecture
- Comprehensive analytics and visualization tools

Source: https://www.networkworld.com/article/3247801/internet-of-things/the-top-5-user-requirements-of-iot-edge-platforms.html

Protocol Support

- Telegraf has over 150 plugins
  - More every release
- Many IoT-Specific plugins
  - MQTT
  - RabbitMQ
  - Particle.io
- Many more to come
  - CoAP
  - PPMP

Robust Offline Data Capability

- Normalization of data
  - Use Kapacitor for data normalization
- Storage for local data
  - Store locally on the edge
  - Forward upstream to cloud
  - Downsample before forwarding to reduce bandwidth usage
- Local Even Processing Engine
  - Use Kapacitor for event processing

Hardware Agnostic, Scalable

- Runs on x86 and ARM
- Runs on servers and embedded platforms
- Single code-base
- Simple deployment

Comprehensive Analytics and Visualization

- Use Chronograf for visualization
  - Also compatible with Grafana
- Analytics via visualization, Kapacitor or CQs
- Forward data to external analytics platforms

What Does an IoT Data Solution Need?

- In the Cloud
  - High-capacity data ingestion
  - Flexible Data retention policies
    - Not all data needs to be around forever!
  - Dashboarding and visualization
  - Alerting
    - Response to data events

What Does an IoT Data Solution Need?

- At the edge
  - Same needs as the cloud +
  - Small, lightweight, yet powerful
  - Data downsampling
  - Multi-protocol ingestion

An IoT Architecture

- Run the same stack on the edge as the cloud
  - Lower cost of development
  - Faster Time to Awesome™
  - Flexible deployment
- Location-specific analytics
  - Shop-floor Dashboard
  - Backend Dashboard

IoT Edge Monitoring

- Monitor the sensors
- Monitor the network
  - Did the sensor die, or was it the wireless connection?
- Monitor the platform
  - Battery?
  - CPU?
  - Storage?

Edge Dashboard

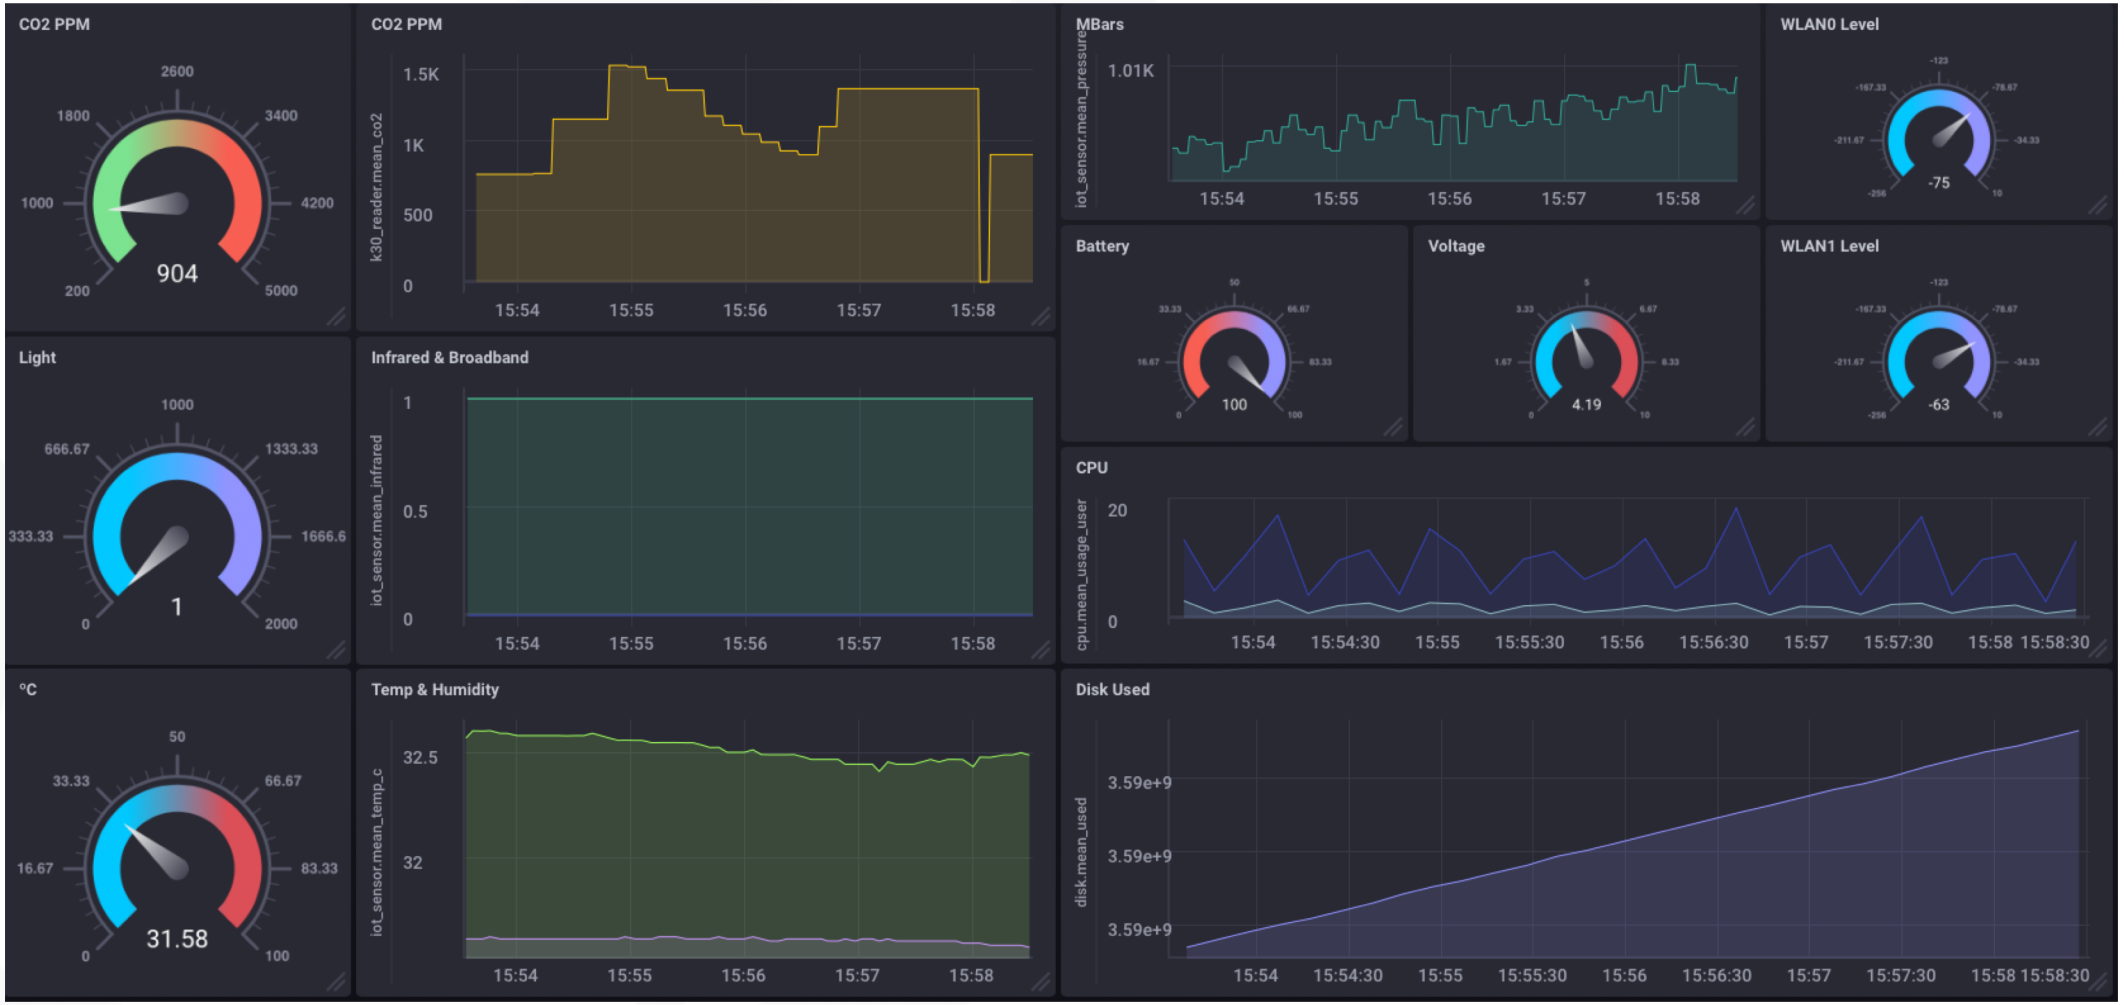

What is Time Series Data?

Relational Data

- Relationships between different kinds of data
- Transactional
- Frequently updated

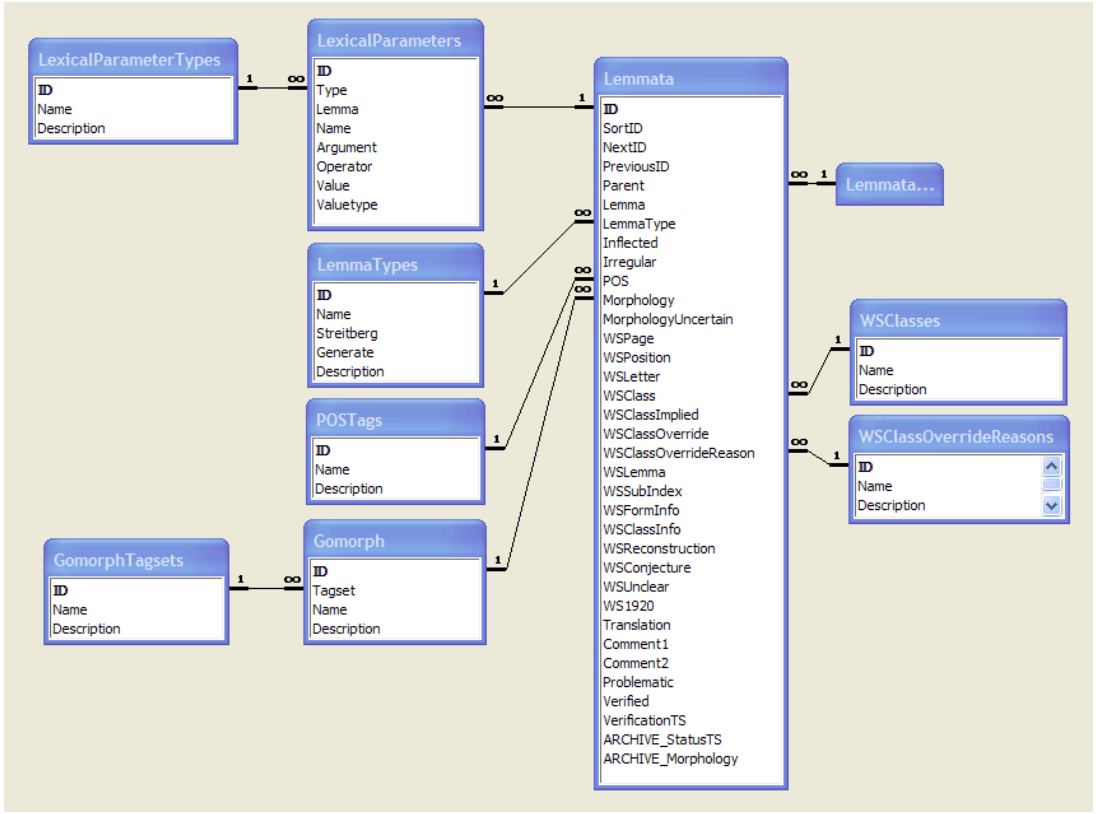

Document Data

- Document storage
- JSON Objects
- NoSQL Databases

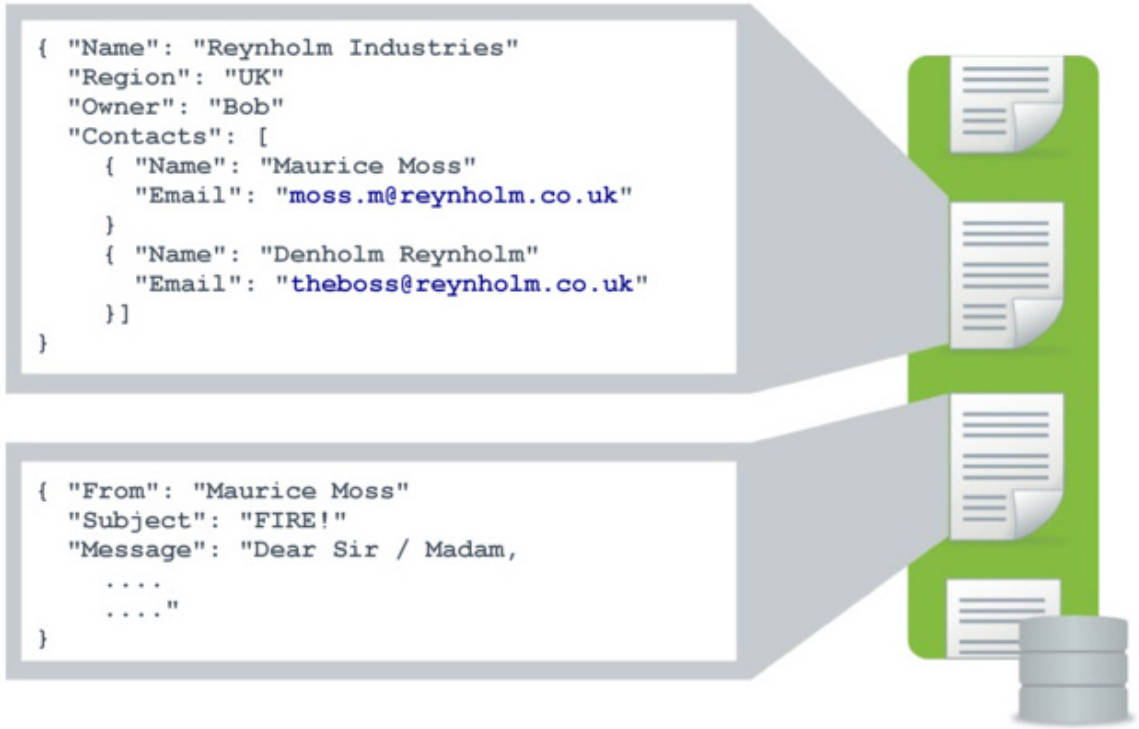

Time Series Data

- Data over time
  - Stock prices
- Event Data
  - event@time
- Sensor Readings
  - temperature@time

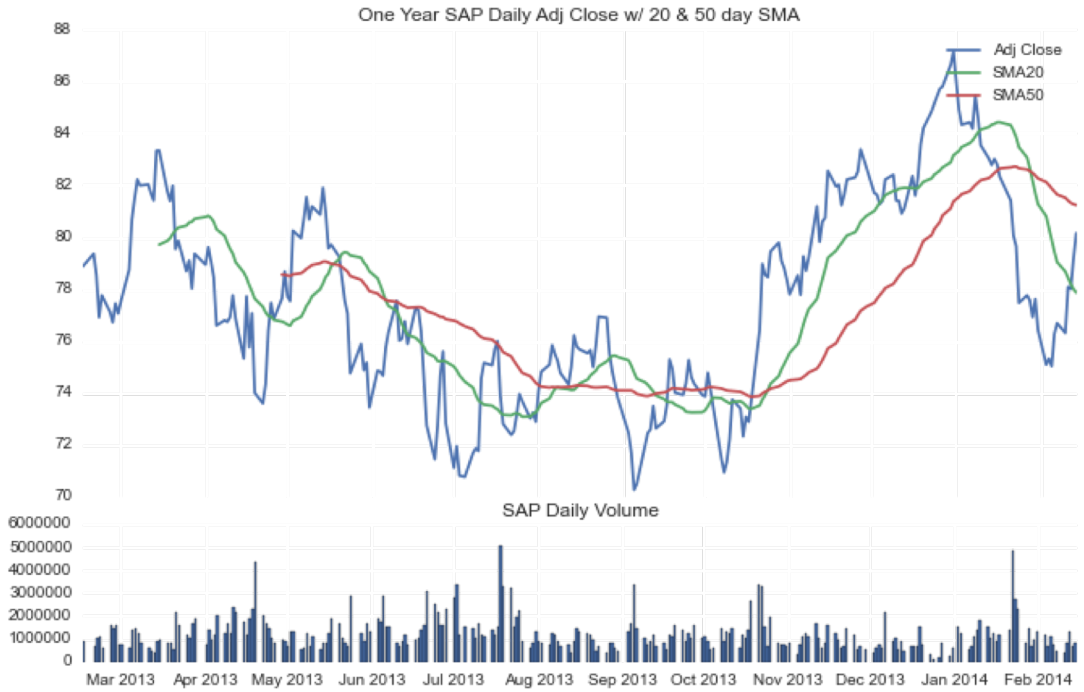

Why is IoT Data Time Series Data?

- <sensor>@<time>
  - The @time component!
  - Temperature reading, valve flow-rate, etc.
- IoT Data **is** time series data
  - See above.

  <sensor>@<time>

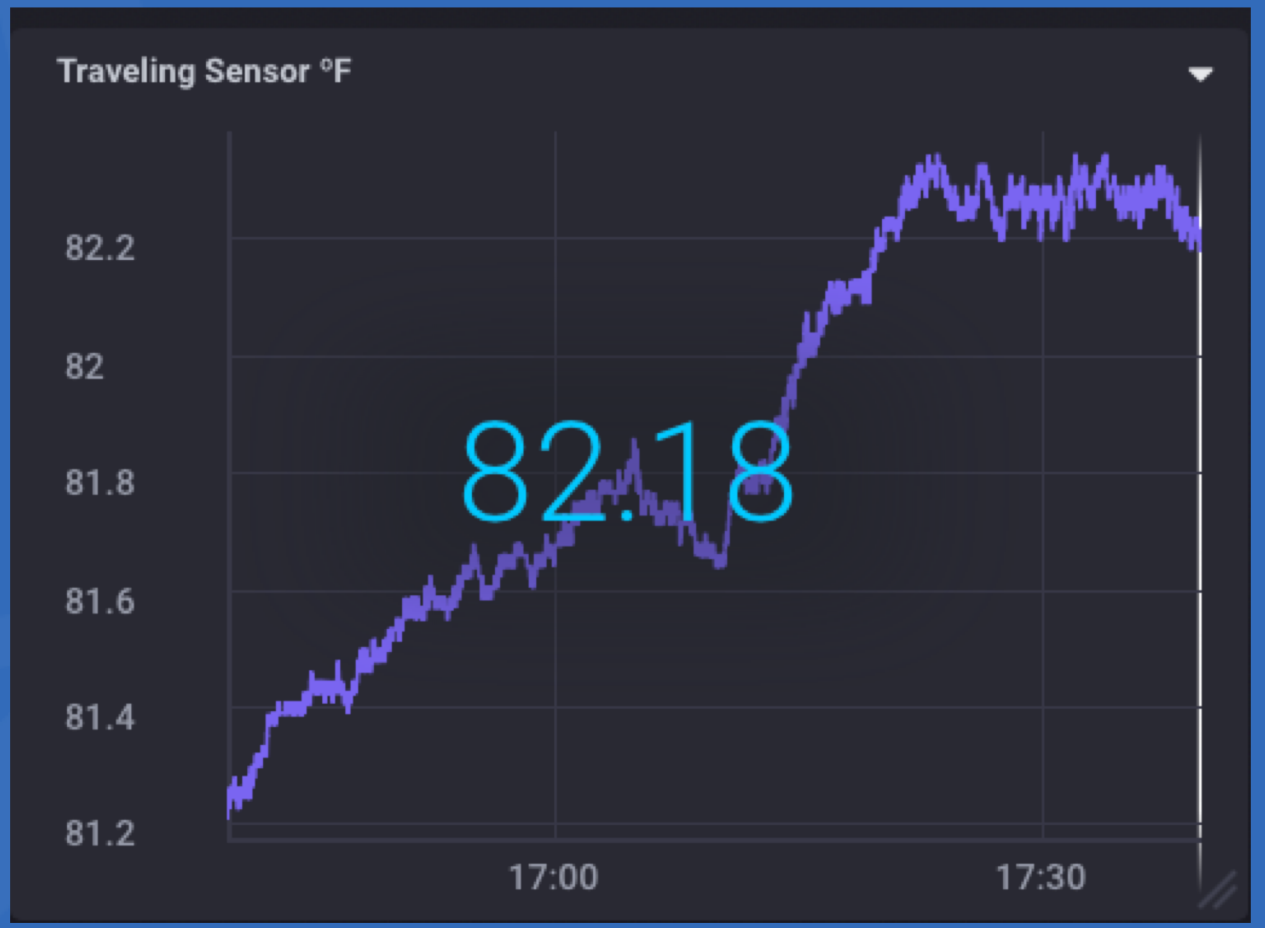

Installing InfluxDB

Mac OS X

- Install Homebrew
  - `/usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/master/install)"`
- Install InfluxDB
  - `brew install influxdb chronograf telegraf kapacitor`
- Enable InfluxDB
  - `brew services start influxdb`
  - …

Linux

- https://docs.influxdata.com/influxdb/v1.5/introduction/installation/
  - curl -sL https://repos.influxdata.com/influxdb.key | sudo apt-key add –
  - source /etc/lsb-release
  - echo "deb https://repos.influxdata.com/${DISTRIB_ID,,} ${DISTRIB_CODENAME} stable" | sudo tee /etc/apt/sources.list.d/influxdb.list
  - sudo apt-get install influxdb chronograf telegraf kapacitor
  - systemctl enable influxdb
  - …

Windows

- https://portal.influxdata.com/downloads
- Not well supported and not recommended for production

Running InfluxDB

- Go to http://localhost:8888/
  - Graphical front-end to InfluxDB, Chronograf, and Kapacitor

Getting Data In

How does InfluxData help?

- Extremely efficient data collection
- High-volume data collection
  - IoT generates huge volumes of data very quickly
  - Being able to ingest, analyze and query that data is key to IoT success
- Ease of Deployment
  - Easy to deploy InfluxDB and the TICK stack for data collection, analysis and action
  - Very low time to value – Time To Awesome™
- Dashboards and visualization
  - Easy to build useful, easy to read dashboards.

Getting Data from IoT Device to InfluxDB

Line Protocol

<measurement>[,<tag>=<value>[,<tag>=<value>]] <field>=<value>[,<field>=<value>] [<timestamp>]

Element	Optional/Required	Description	Type
Measurement	Required	The measurement name. InfluxDB accepts one measurement per point.	String
Tag Set	Optional	All tag key-value pairs for the point.	Tag keys and tag values are both strings.
Field Set	Required. Points must have at least 1 field.	All field key-value pairs for the point.	Field keys are strings. Field values can be floats, integers, strings, or booleans.
Timestamp	Optional. InfluxDB uses the server's local nanosecond timestamp in UTC if the timestamp is not included with the point.	The timestamp for the data point. InfluxDB accepts one timestamp per point.	Unix nanosecond timestamp. Specify alternative precisions with the HTTP API.

On success, the `response.status` should be 204.

Getting Data from IoT Device to InfluxDB

- Using Telegraf Plugins
- Enormous list of available plugins
  - https://github.com/influxdata/telegraf/tree/master/plugins/inputs
- IoT Plugins
  - AMQP (RabbitMQ) Plugin
  - http_listener – Telegraf plugin for line-protocol
  - Mqtt_consumer – subscribe to mqtt topics and add messages to InfluxDB
  - **New plugin for Particle.io webhooks available!**

Analyzing Your Data

Analyzing Your Data

"The more simple your data, the more complex the analysis you can run."

– Sarah Cooper, AWS IoT Solutions

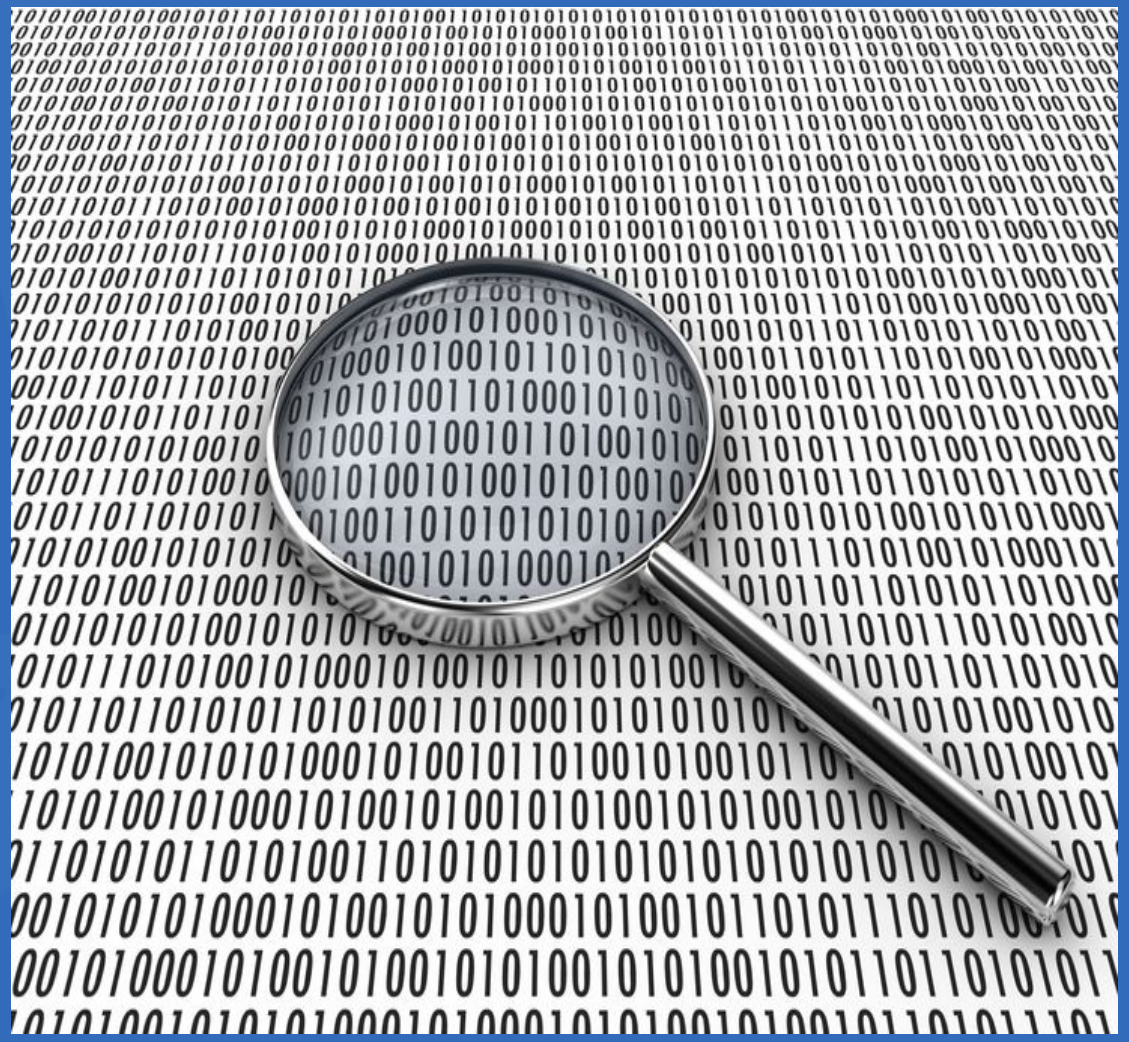

Visualize Your Data on Dashboards

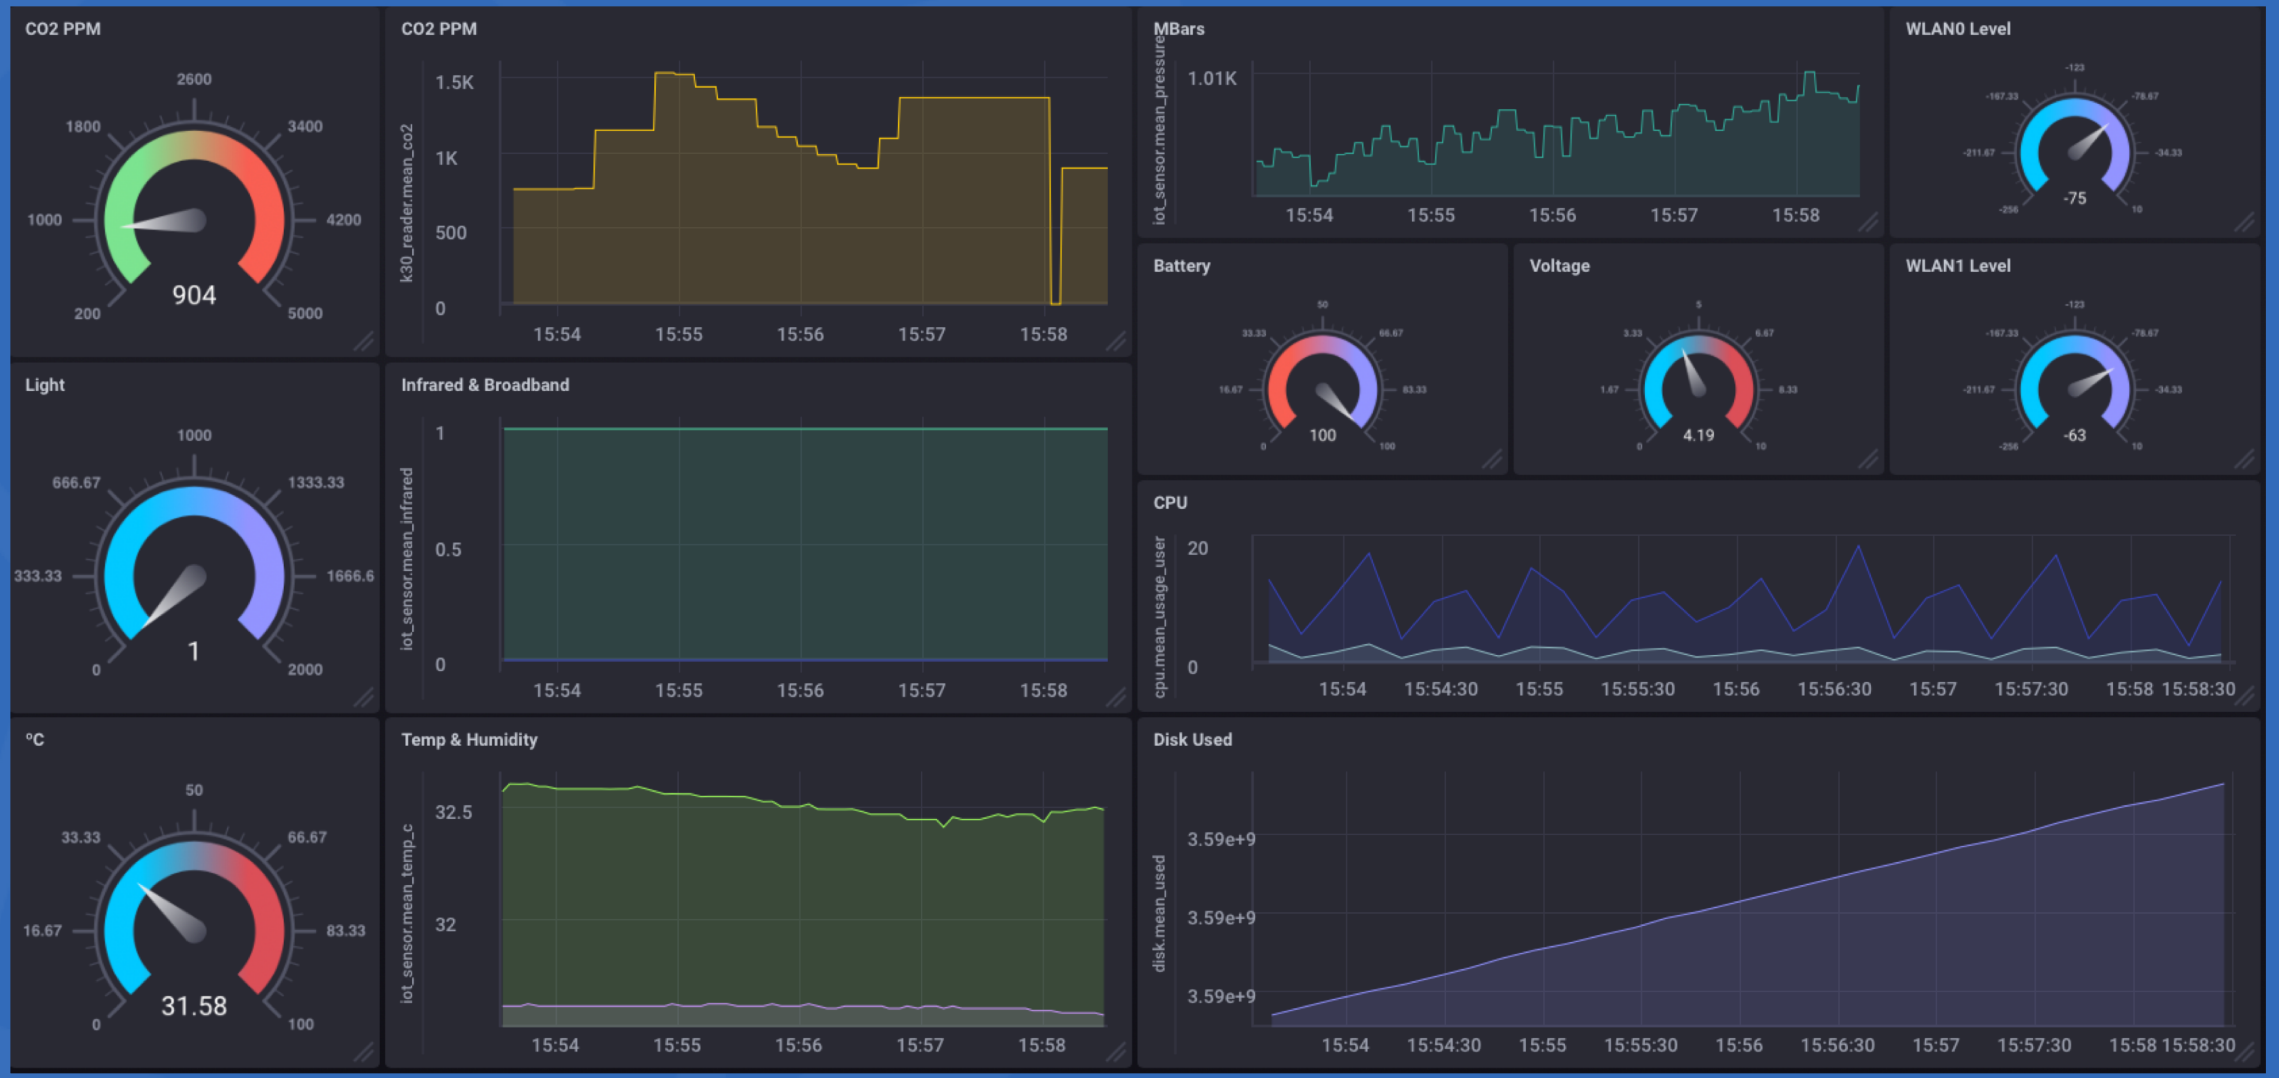

Look For Hidden Gems

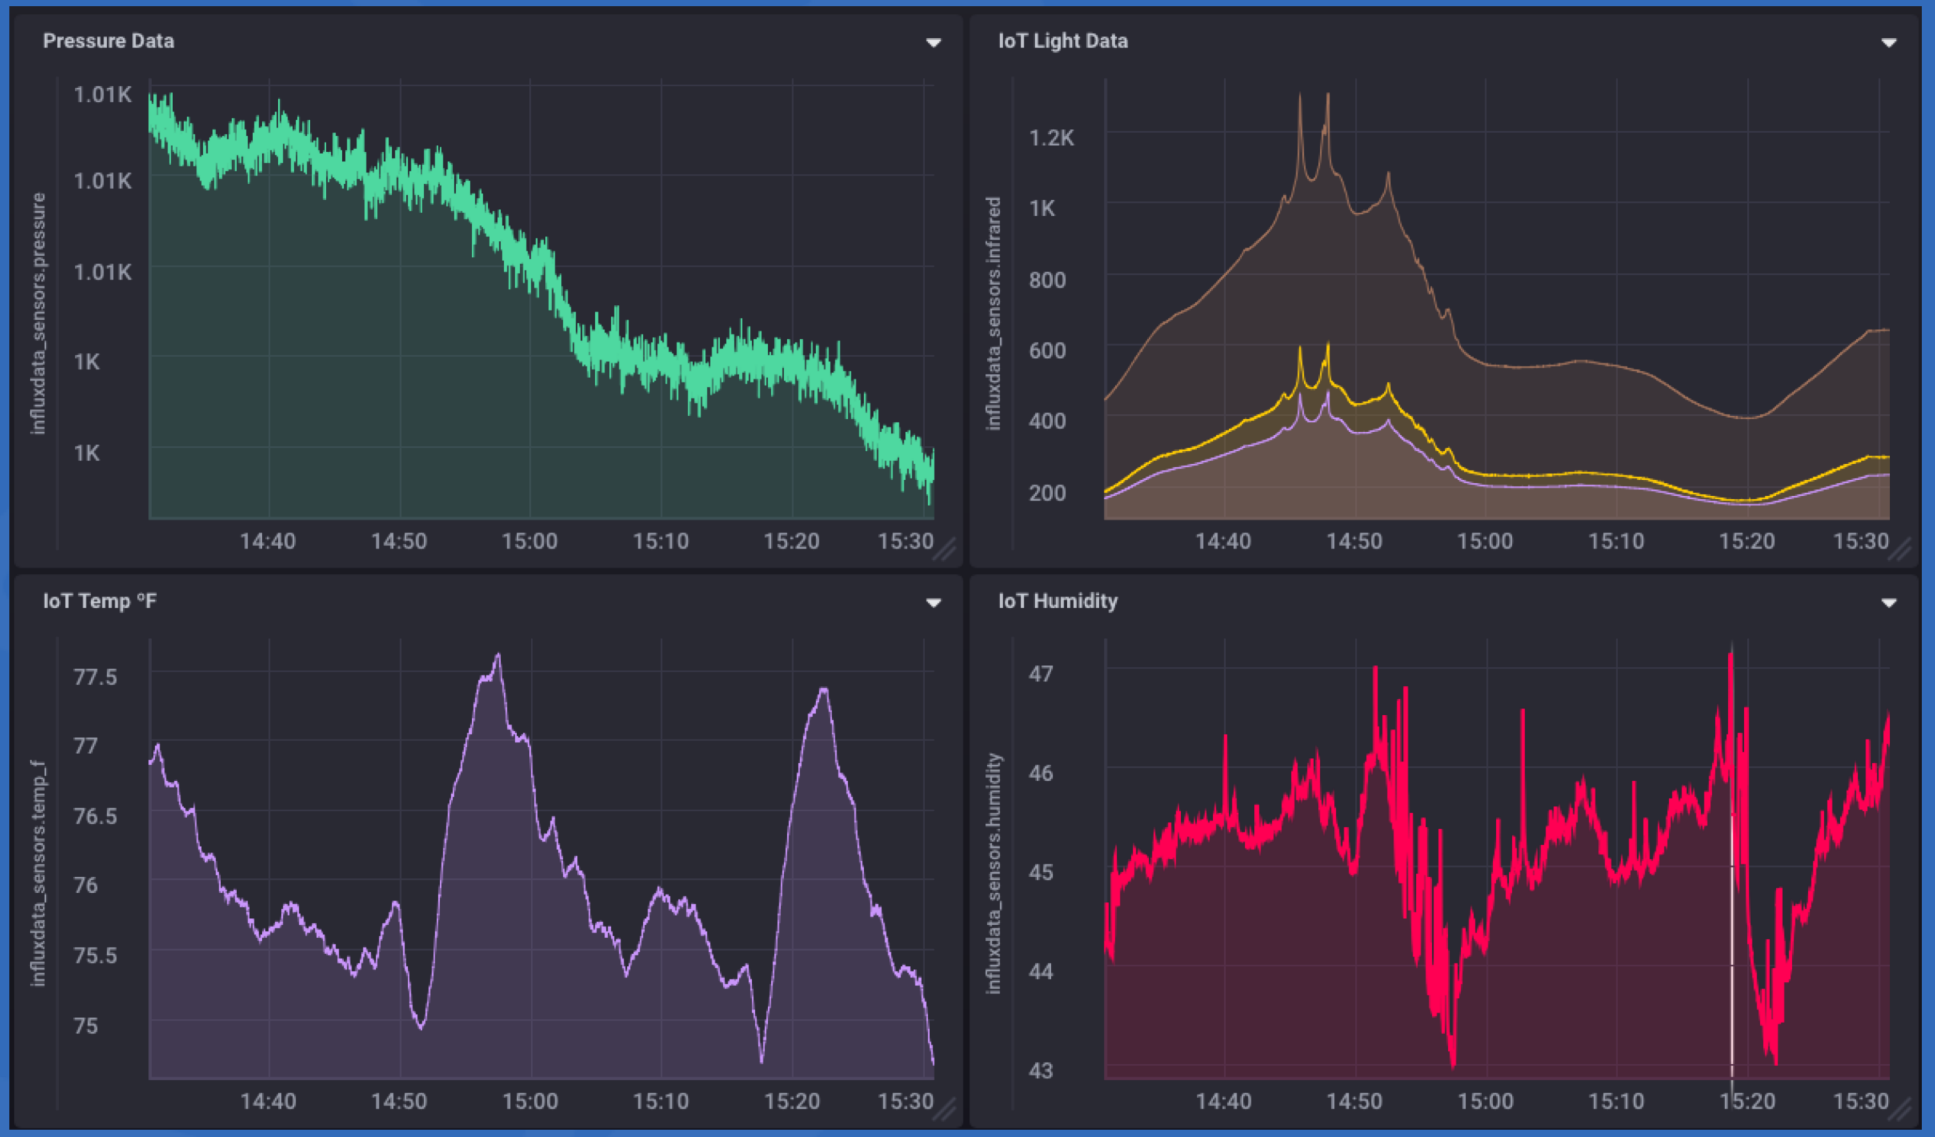

influxdata

“It’s amazing what you can discover when you can actually **see** your data!”

Get InfluxDB
https://influxdata.com

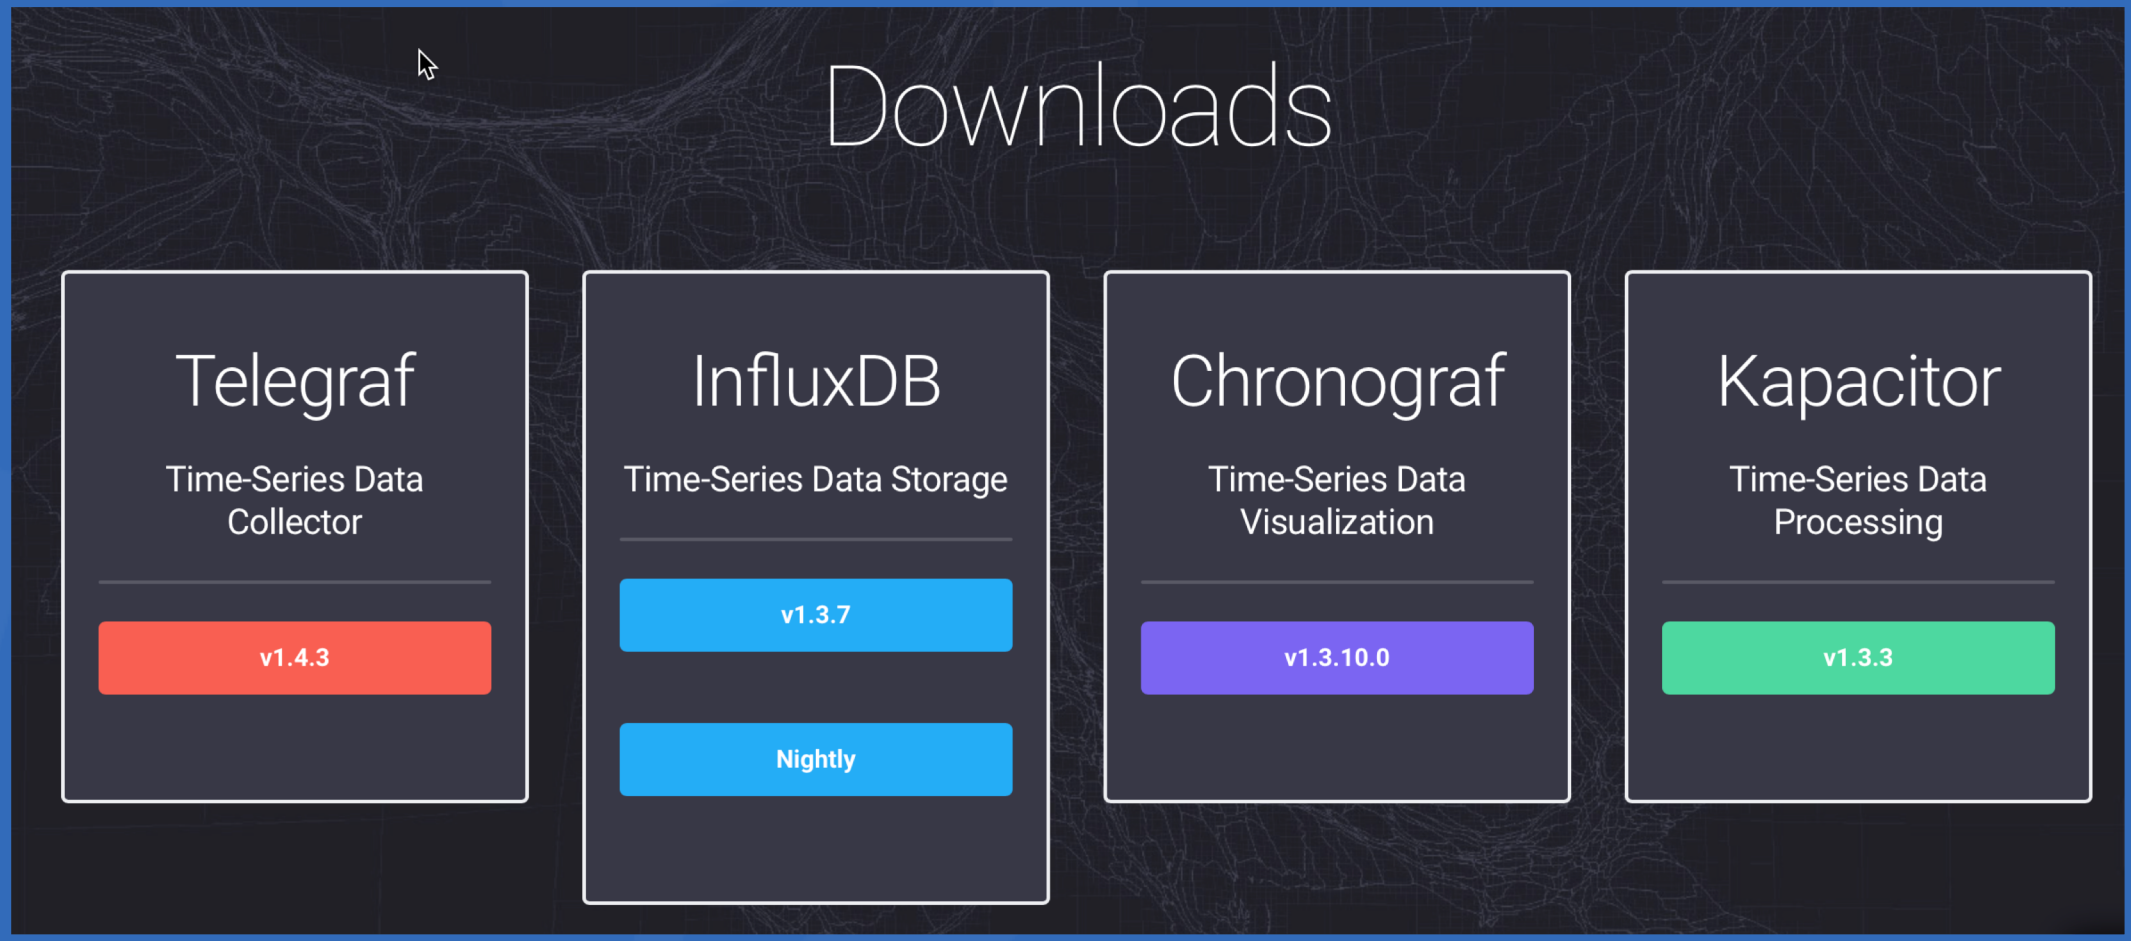

influxdata

How Can InfluxData Fit Into My IoT Architecture?

An IoT Architecture

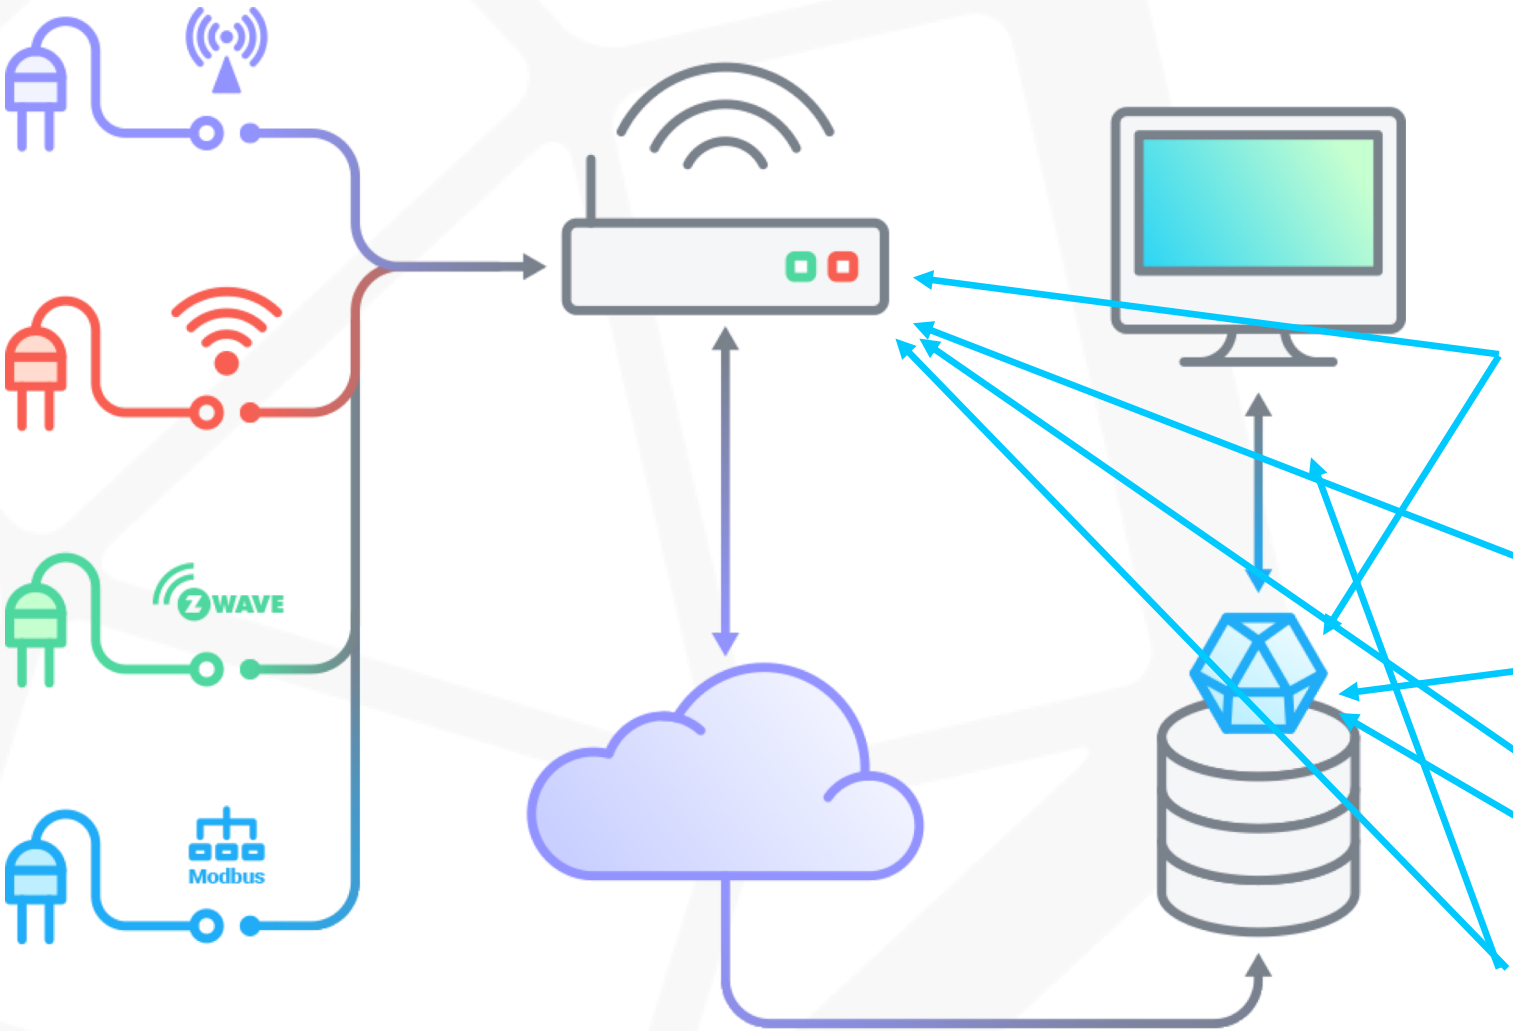

InfluxData

- Telegraf
  - Data Collection
- InfluxDB
  - Short-term storage
  - Long-term storage
- Kapacitor
  - Local Alerts
  - System-wide alerts
- Chronograf
  - Dashboards

Getting Data Out

Event Processing Engine

- Kapacitor
  - Data co-processing engine
  - Downsampling, etc.
  - Anomaly detection
  - Alerting

Kapacitor Alerting

- Alert Types
  - Threshold Alert
  - Deadman Alert
  - Alerts based on Calculated Value
  - …
- Alert Mechanisms
  - TCP/HTTP
  - Pager Duty
  - MQTT
  - …

Kapacitor Alert - Calculated Value

```
var parseTemp =
'''  {{ if (gt ( index .Fields "max.value" ) 90.0 ) }} #ff0000
{{ else if (gt ( index .Fields "max.value" ) 88.0 ) }} #ff4000
{{ else if (gt ( index .Fields "max.value" ) 86.0 ) }} #ff8000
{{ else if (gt ( index .Fields "max.value" ) 84.0 ) }} #ffbf00
{{ else if (gt ( index .Fields "max.value" ) 82.0 ) }} #ffff00
{{ else if (gt ( index .Fields "max.value" ) 80.0 ) }} #bfff00
…
 {{ else }} #bf00ff {{ end }}'''
```

Kapacitor Alert - MQTT

```
trigger
    |alert()
        .mqtt('colorChange')
        .brokerName('127.0.0.1')
        .topic(message)
```

Join Our Community

- On the Web
  https://community.influxdata.com
- On GitHub
  https://github.com/influxdata
- On Slack
  https://gophersinvite.herokuapp.com
  #influxdb
- On Twitter:
  @InfluxDB and @davidgsIoT

Questions?


Thank You
influxdata®
davidgs@influxdb.com
@davidgsIoT