Monitoring Aquaponics Facilities Around the Globe

Hunter Mason | Seattle University

Ryan O'Connor | Splunk

splunk> .conf19

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results may differ materially. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

Hunter Mason
Seattle University
Electrical and Computer Engineering

Ryan O'Connor
Splunk & University of Connecticut

splunk> .conf19

“An aquaponics farm in a town called Manchay, near Lima, Peru is utilizing Raspberry Pi’s to measure Dissolved Oxygen (DO), pH, and water temperature. This farm is monitoring data on both the fish and plants in order to make necessary adjustments that will facilitate the maintenance of a healthy ecosystem”

Tara Watrous, University of Connecticut

Connecting Universities

Partnership Between UConn + SeattleU

splunk> .conf19

How Does Splunk Fit In

- A Global IoT Company
- Great at collecting both structured, and unstructured data
- Helps organizations derive value from their data in real-time
- Has a philanthropic program designed for just this sort of project

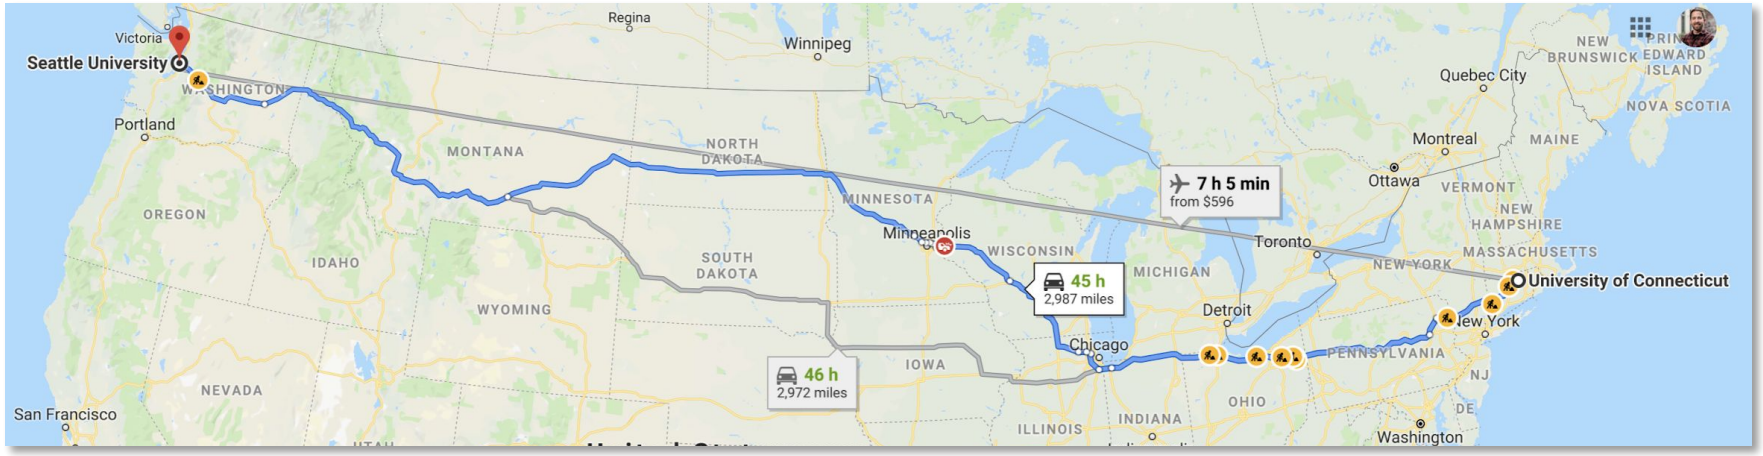

splunk> .conf19

UConn and Splunk for Good

splunk> for good

splunk> .conf19

How did the University of Connecticut Get Started

University of Connecticut

- Registered their University for Splunk Pledge

They had also navigated standing up Splunk at a University

- Planned a Splunk Deployment using various Best Practice Guides
  - Splunk Validated Architecture
  - Reference Hardware
  - Summary of Performance Recommendations
- Worked with Local IT to secure servers

How did the University of Connecticut Get Started

With all of that legwork, UConn was able to:

- Teach courses using Splunk
- Worked on a real aquaponics facility using Splunk.

Courses Included:

- Data Analytics
- Network Security
- Security, Governance, and Audit
- Industrial IoT

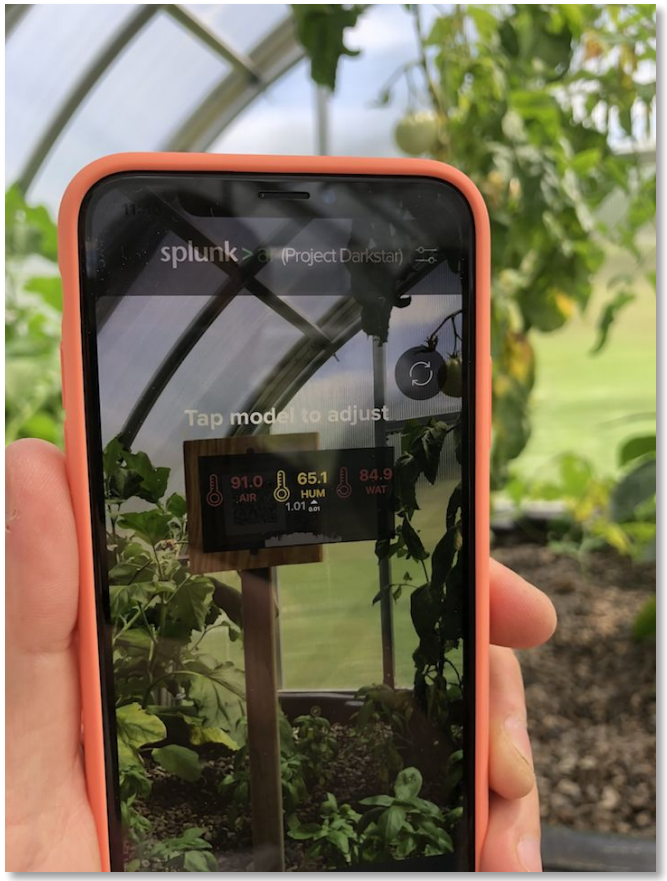

Powering Projects with Splunk for Good

Splunk for Good is a group within Splunk that represents our commitment to research, education and community service.

With the Splunk Pledge program, we offer:

- 10GB License
- Access to Free Training
- Splunk Fundamentals 1
- Splunk Fundamentals 2

Access to https://workplus.splunk.com

Summary: Why Splunk is Great Inside the Classroom

- Teaching Splunk yields marketable skills (source: past students)
- These skills are free
- Roughly 170 students have registered for Splunk Fundamentals 1 and 2 just at UConn alone
  - The interest is there
- Splunk Fundamentals courses can be easily used to augment and enhance pre-existing curriculum
- It can spark interdisciplinary courses and research

Seattle University's Use Case

splunk> .conf19

Why Seattle University needs Splunk

- The chapter of Engineers for a Sustainable World at Seattle University is working to harness data from Aquaponics facilities around the globe
  - http://www.seattleu-esw.org/
- Goal of the project was to collect data from facilities to monitor them from anywhere in the world
- By making sure facilities are kept up and running, organizations can grow food at a faster rate than traditional farming

Benefits of These Projects

- Food being used by orphanages in Thailand
- Food being sold to fund programs in Schools in Peru
- Learning opportunities in Middle Schools and a Boys and Girls Club in Seattle
- Learning opportunities in several programs and Universities around the United States

How Can We Make These Projects Even Better?

Being able to collect data in real-time sounds like such a nice benefit to a global aquaponics project

UConn was able to help disseminate information around Splunk

And it should be easy… right?

The Problem(s) so far:

- Collecting data in remote locations comes with several challenges
- Physical challenges (Weather, corrosion, rodents, loose wires, etc.)
- Technology challenges (Failed sensors, bad wiring, maintenance, etc.)

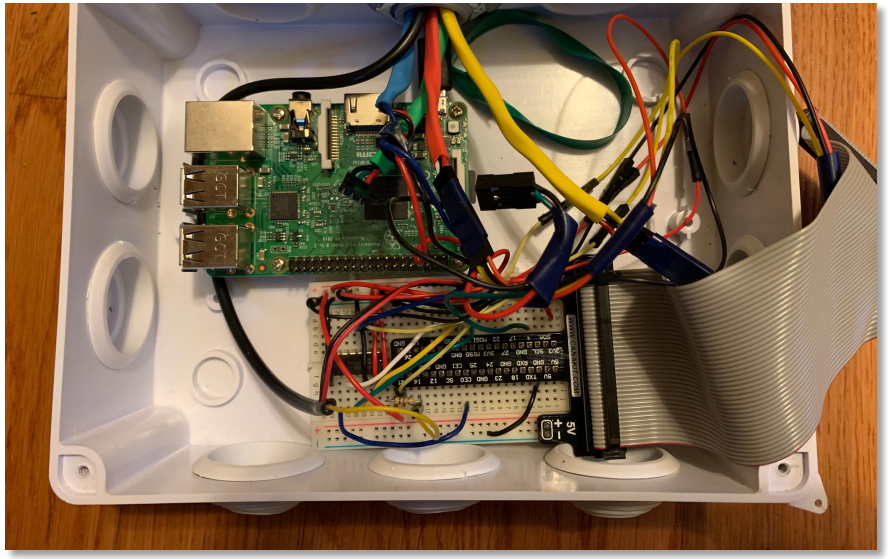

Hardware and Software Development

splunk> .conf19

Our Hardware Solutions

Custom PCB Design

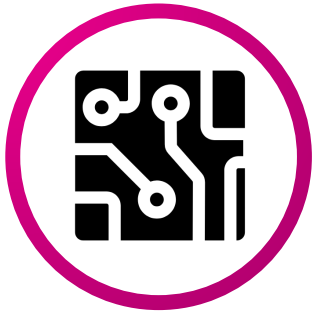

RasberryPi

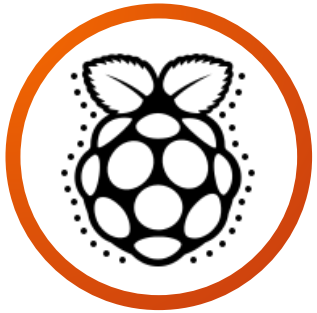

Atlas Scientific Sensors

KVM

Enclosure and Mounting Hardware

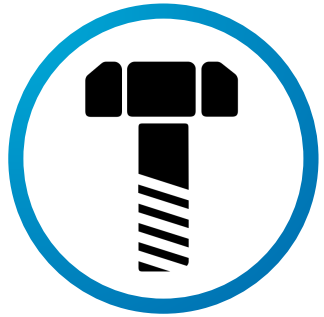

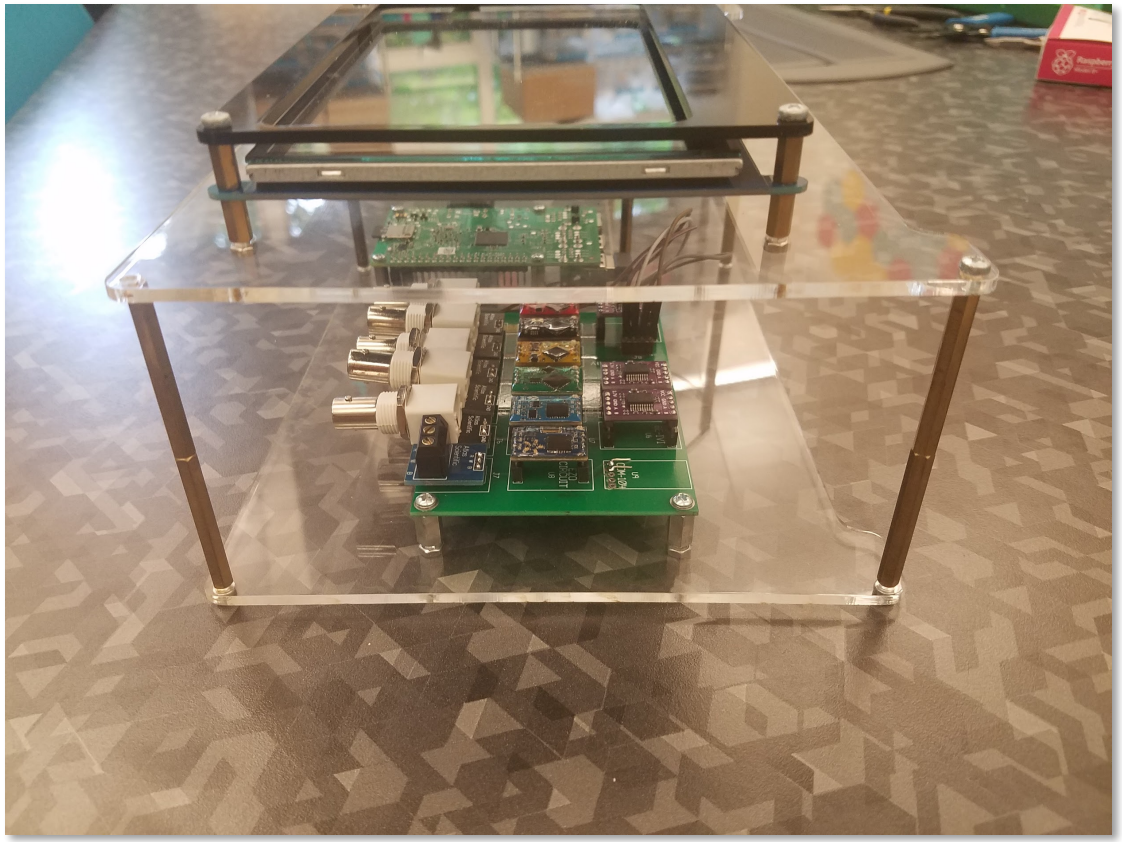

Printed Circuit Board (PCB) Design

The fastest and cheapest way to increase durability and reliability

- Sockets for each circuit
- Mechanical durability by soldering parts
- Electrical reliability
- Less that $2/Board means…
  - A low cost system
  - Rapid prototyping
  - Multiple iterations

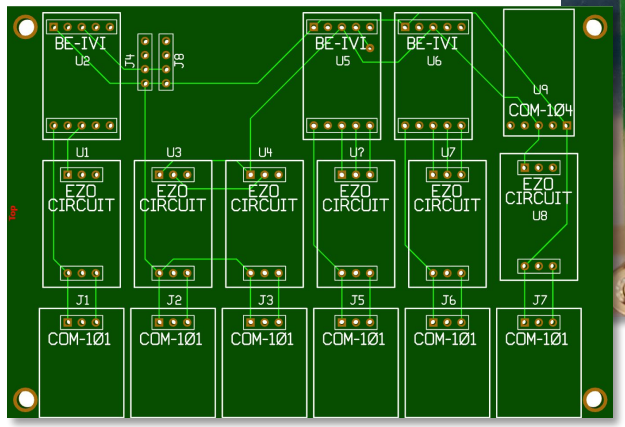

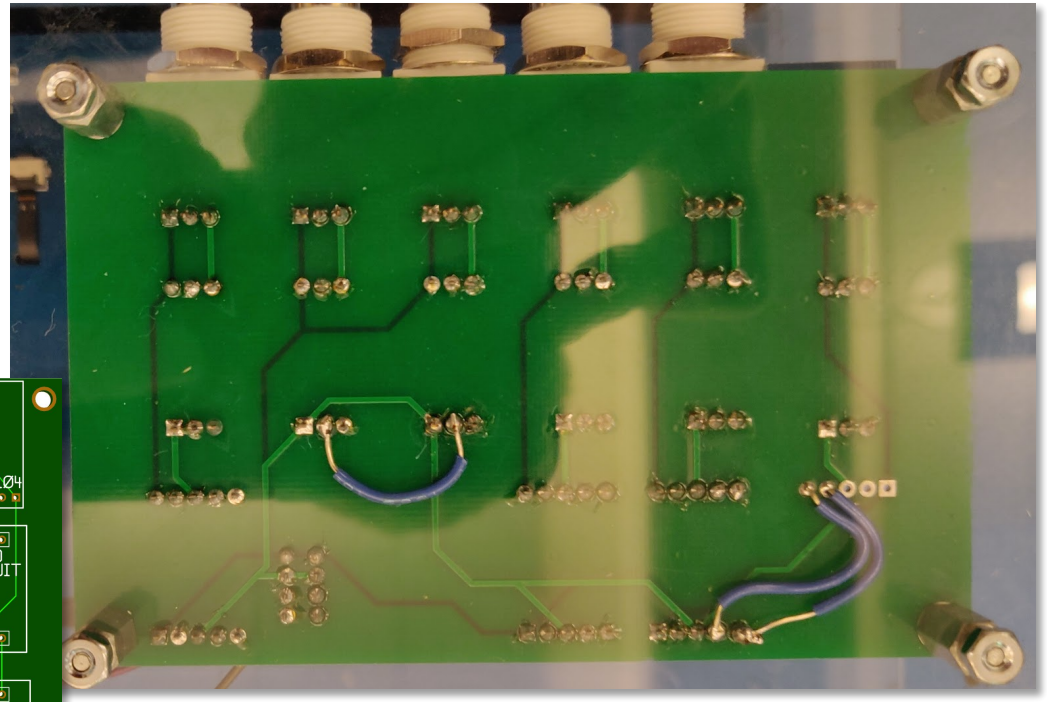

Raspberry Pi

RaspberryPi's are an easy choice for general IoT applications. They are…

- Cheap, Replaceable
- Internet enabled
- Built with SPI & I2C IO
- Widely supported by open source
- Compact

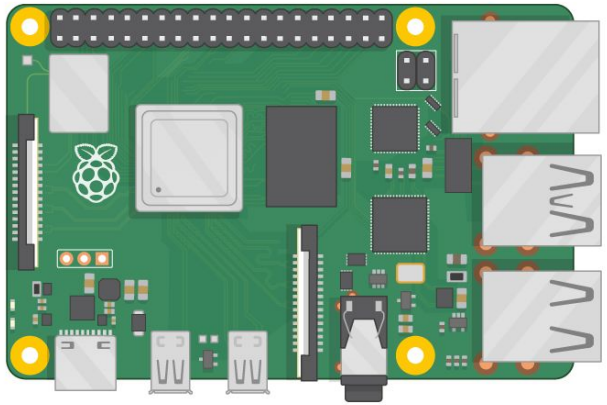

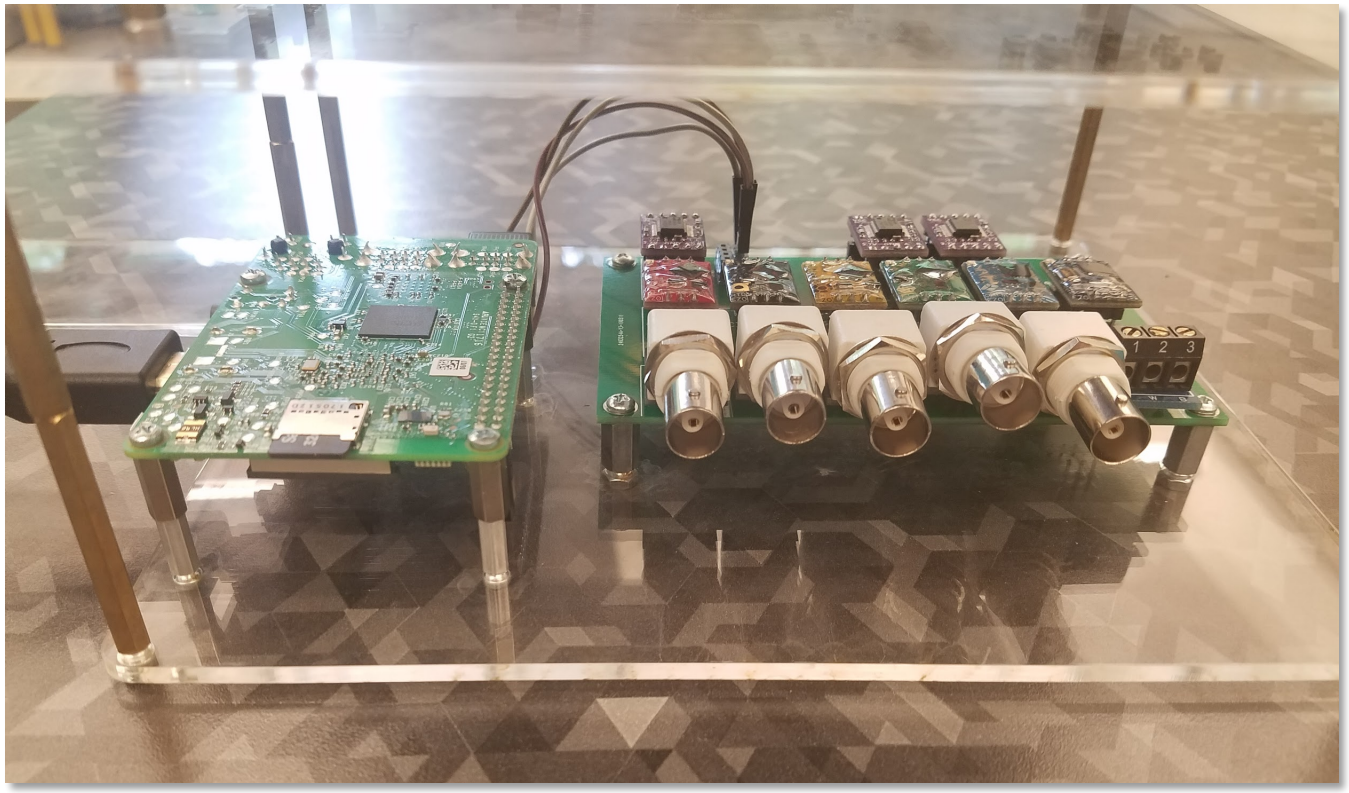

Full Keyboard/Video/Mouse Solution

- Allows maintenance and calibration in the field
- Users won't necessarily have skills or resources to use SSH or remote desktop
- Features:
  - 7in touch screen
  - Bluetooth/USB wireless keyboard and trackpad

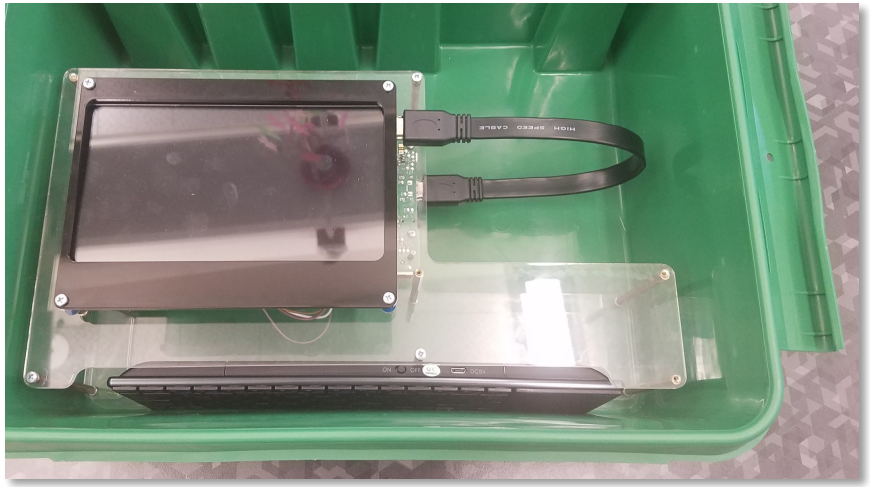

Enclosure

- Water ingress to old enclosures
- Cable glands incompatible with sensors

- Sizing up for new design
- Ease of access
- Similar cost

Mounting Plates

- Increases mechanical durability
- Laser cut acrylic for cheap and rapid prototyping
- Many iterations as the project has evolved

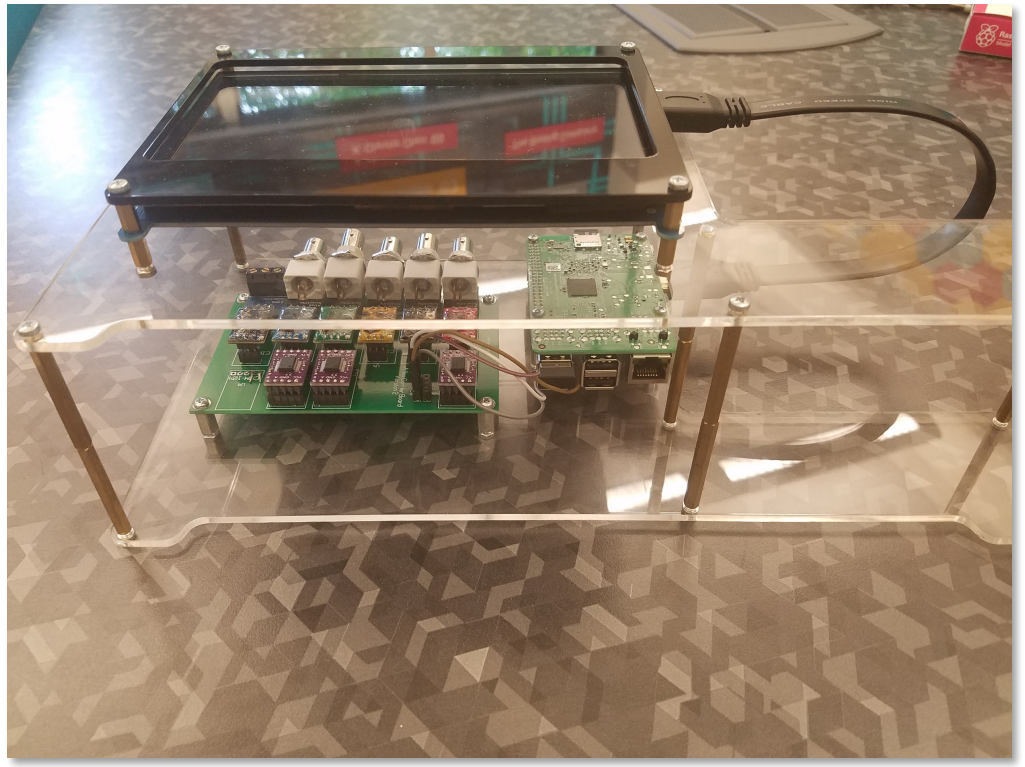

Why custom hardware?

Well it always comes down to cost doesn't it…

When you remove the shared price of sensing equipment…

Our PCB and IoT Integration
for up to 6 sensors cost:

$92

Comparable Off-the-shelf system
for one sensor costs:

~ $188

Our Software Solutions (so far)

splunk> .conf19

How does our data get to Splunk?

Splunk HTTP Event Collector (HEC)

- A versatile approach to getting data in (GDI)
- With a few lines of Python, a user can send an event (or multiple events) to a Splunk installation

Splunk Metrics Store

- Metrics is a feature introduced in 7.0
- Focuses on collecting, investigating, monitoring, and sharing metrics from your technology infrastructure

Sample scripts for Raspberry Pi can be found on Github

HTTP Event Collector Architecture

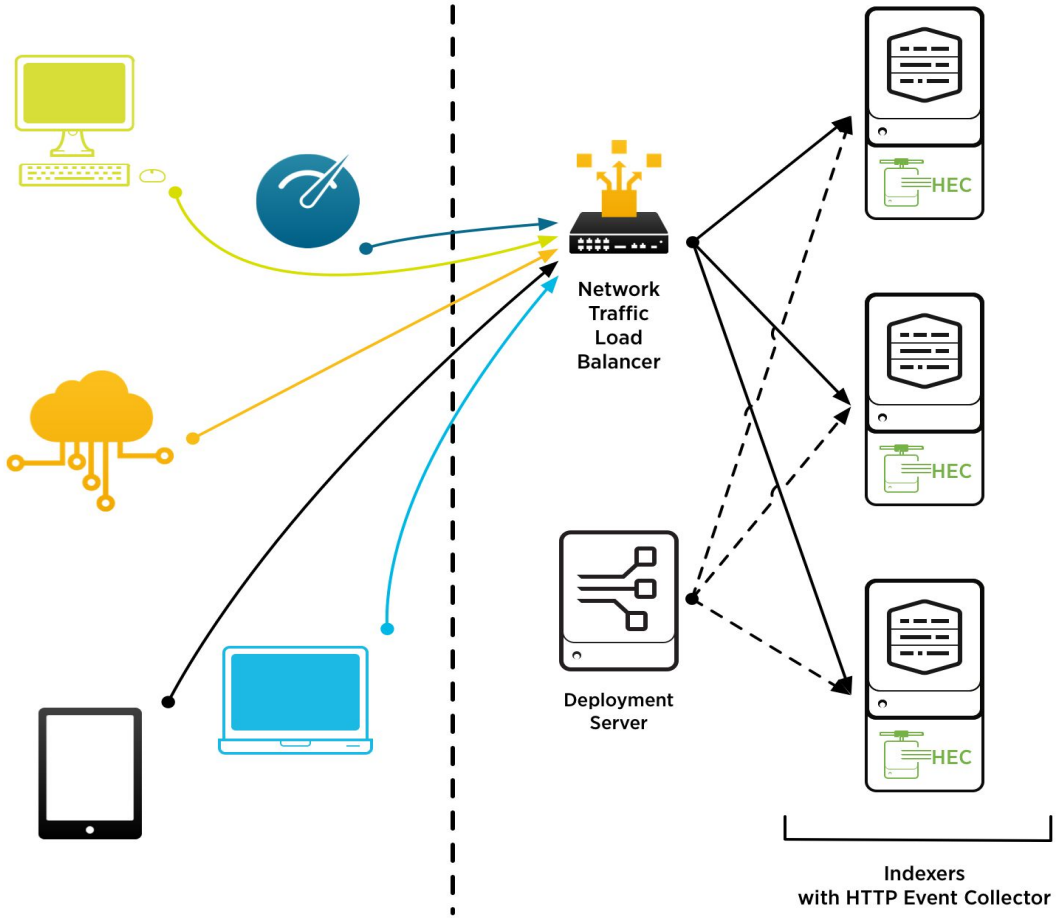

Simple Python Example - Metrics

```
# Import Requests Library
import requests
# Create a JSON Dictionary
jsonDict = {'index':'main_metrics', 'host': 'localhost', 'sourcetype': 'sensor_data', 'event':
'metric', 'fields':{'temp_f':67,'_value': 67,'metric_name':'temp_f', 'affiliation':'Splunk'}}
# Post the Dictionary to Splunk
requests.post('http://localhost:8088/services/collector' ,headers={'Authorization': 'Splunk
7493b017-0f10-44c0-a72c-9cf20afeb888' },json=jsonDict,verify= False)
```

Simple Python Example - Batch Metrics

```
# Import Requests Library
import requests
# Create a JSON Dictionary
jsonDict = [{'index':'main_metrics', 'host': 'localhost', 'sourcetype': 'sensor_data',
'event': 'metric', 'fields':{'temp_f':62,'_value': 62,'metric_name':'temp_f',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'pH':7,'_value': 7,'metric_name':'pH',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'DO':7,'_value': 6.5,'metric_name':'DO',
'affiliation':'Seattle U'}}]
# Post the Dictionary to Splunk
requests.post('http://localhost:8088/services/collector' ,headers={'Authorization': 'Splunk
7493b017-0f10-44c0-a72c-9cf20afeb888' },json=jsonDict,verify=False)
```

Simple Python Example - Batch Metrics

```
# Import Requests Library
import requests
# Create a JSON Dictionary
jsonDict = [{'index':'main metrics', 'host': 'localhost', 'sourcetype': 'sensor_data',
'event': 'metric', 'fields':{'temp_f':62,'_value': 62,'metric_name':'temp_f',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'pH':7,'_value': 7,'metric_name':'pH',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'DO':7,'_value': 6.5,'metric_name':'DO',
'affiliation':'Seattle U'}}]
# Post the Dictionary to Splunk
requests.post('http://localhost:8088/services/collector' ,headers={'Authorization': 'Splunk
7493b017-0f10-44c0-a72c-9cf20afeb888' },json=jsonDict,verify=False)
```

Simple Python Example - Batch Metrics

```
# Import Requests Library
import requests
# Create a JSON Dictionary
jsonDict = [{'index':'main_metrics', 'host': 'localhost', 'sourcetype': 'sensor_data',
'event': 'metric', 'fields':{'temp_f':62,'_value': 62,'metric_name':'temp_f',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'pH':7,'_value': 7,'metric_name':'pH',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'DO':7,'_value': 6.5,'metric_name':'DO',
'affiliation':'Seattle U'}}]
# Post the Dictionary to Splunk
requests.post('http://localhost:8088/services/collector' ,headers={'Authorization': 'Splunk
7493b017-0f10-44c0-a72c-9cf20afeb888' },json=jsonDict,verify=False)
```

Simple Python Example - Batch Metrics

```
# Import Requests Library
import requests
# Create a JSON Dictionary
jsonDict = [{'index':'main_metrics', 'host': 'localhost', 'sourcetype': 'sensor_data',
'event': 'metric', 'fields':{'temp_f':62,'_value': 62,'metric_name':'temp_f',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'pH':7,'_value': 7,'metric_name':'pH',
'affiliation':'Seattle U'}},{'index':'main_metrics', 'host': 'localhost', 'sourcetype':
'sensor_data', 'event': 'metric', 'fields':{'DO':7,'_value': 6.5,'metric_name':'DO',
'affiliation':'Seattle U'}}]
# Post the Dictionary to Splunk
requests.post('http://localhost:8088/services/collector' ,headers={'Authorization': 'Splunk
7493b017-0f10-44c0-a72c-9cf20afeb888' },json=jsonDict,verify=False)
```

Searching Metrics Store

Fast approach to accessing data

Can be accessed using the mstats command

Example:

```
| mstats max(pH) as pH max(DO) as DO max(temp_f) as temp_f where
index=main_metrics span=1s
```

How do I find out about my metrics data?

- mcatalog to the rescue!
- Command designed to display schema of data stored in the metrics data store
- Example finding metrics names:
  - `| mcatalog values(metric_name) where index=main_metrics`
- Example finding dimensions:
  - `| mcatalog values(affiliation) where index=main_metric`

Combining Dimensions and Metrics

Example utilizing the Metadata we posted:

```
| mstats max(pH) as pH max(DO) as DO max(temp_f) as temp_f where
index=main_metrics span=1s by affiliation
| timechart avg(temp_f) span=1s by affiliation
```

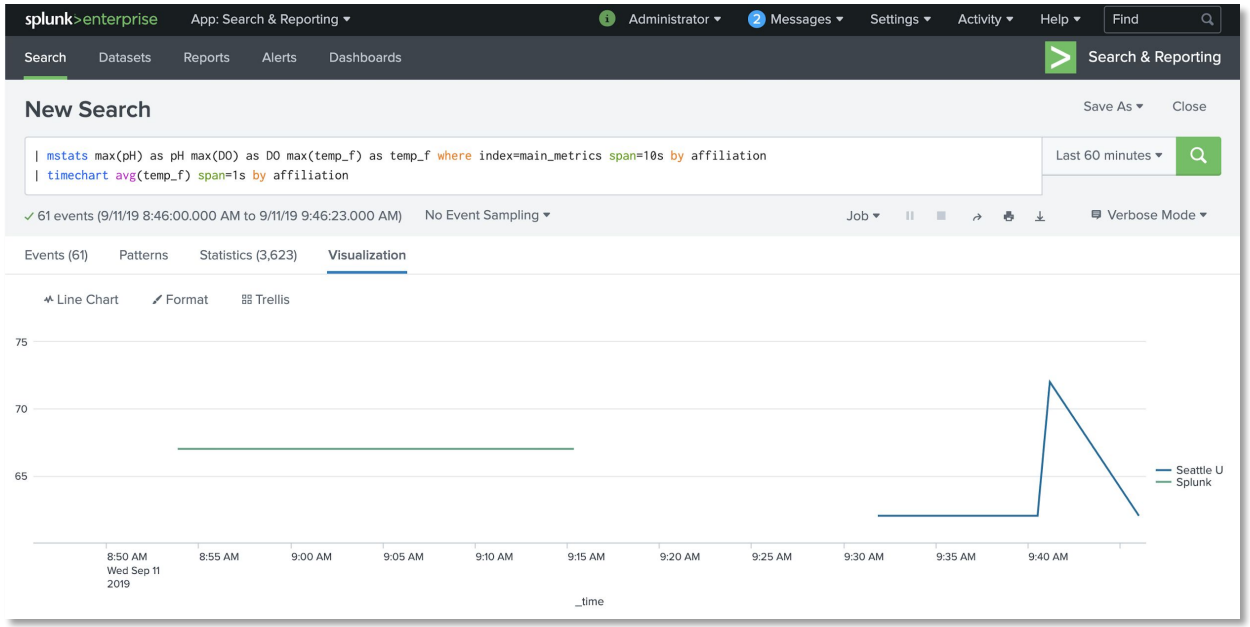

Exactly how fast is Metrics?

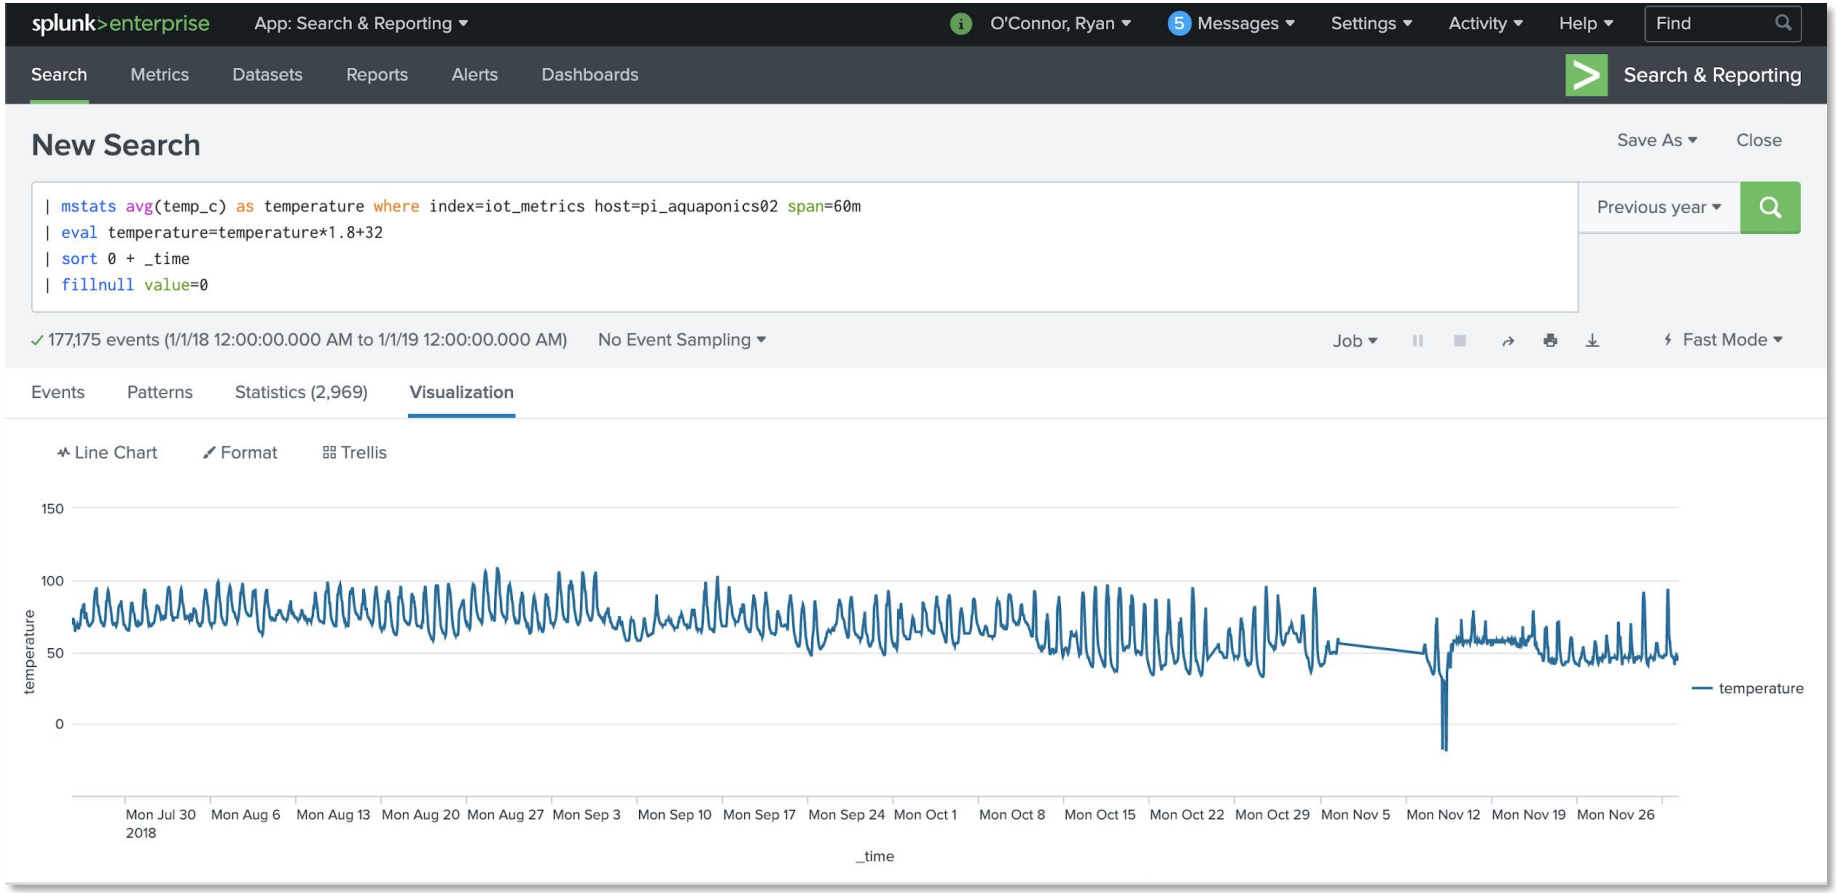

Exactly how fast is Metrics?

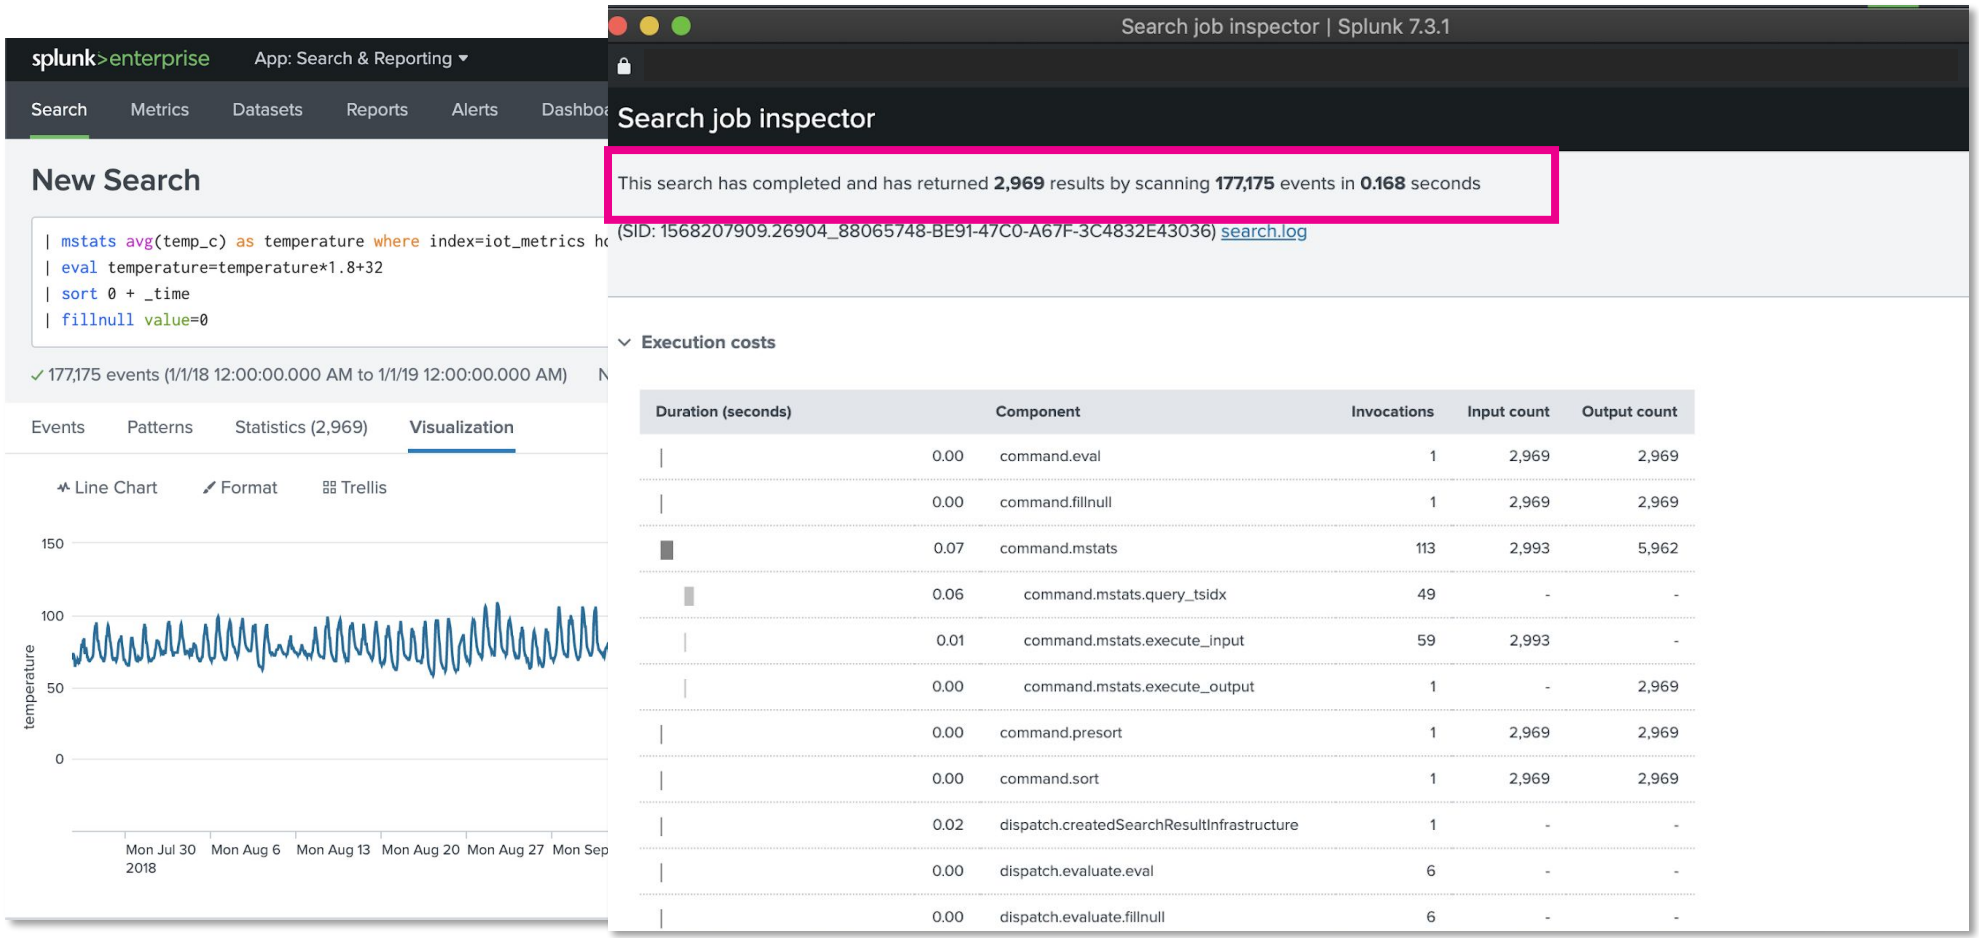

splunk> .conf19

Tips for Speeding up HEC

HTTP vs. HTTPS

- Performance improvement when sending data over HTTP versus HTTPS

Batching

- If you batch multiple events into single requests, it can speed up data transmission.

HTTP Keep-alive

- Setting keepalive on your connection can increase performance. As long as the client sending the data supports HTTP 1.1 and is set up to support a persistent connection, you're taking advantage of keep-alive.

Splunk AR and Mobile

Two solutions that we've utilized to make data more accessible are:

- Splunk AR - Augmented Reality App for Splunk (iOS Only)
- Splunk Mobile - App for Viewing Splunk Dashboards on a Mobile Device

Note: Splunk AR and Splunk Mobile are Generally Available on the App Store

Live Demo of Splunk AR and Mobile

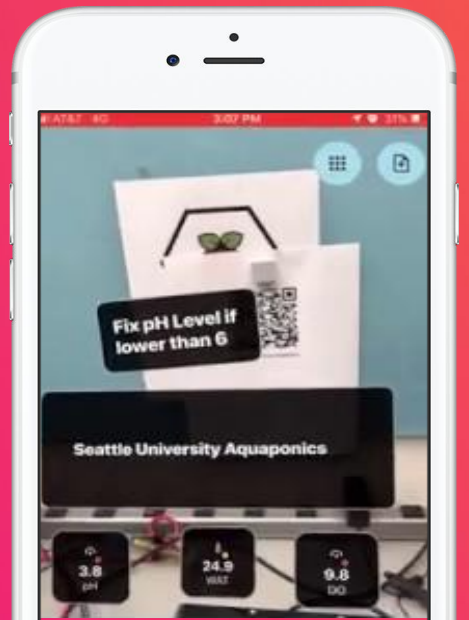

splunk> .conf19

Special Thanks

Special Thanks to…

- Seattle University Faculty
  - Dr. Shiny Abraham
  - Dr. Phillip Thompson
  - Richard Bankhead
- W. M. Keck Foundation
- Splunk for Good Team

UCONN
UNIVERSITY OF CONNECTICUT

W.M. KECK FOUNDATION

SEATTLE UNIVERSITY 1891
COLLEGE OF SCIENCE AND ENGINEERING

Links

- Github Repository
  - https://github.com/ryanwoconnor/iot_class_scripts
- Splunk AR
  - iOS Download Link: https://apps.apple.com/us/app/splunk-ar/id1420233757
- Splunk Mobile
  - iOS Download Link: https://apps.apple.com/us/app/splunk-mobile/id1420299852
- Hackaday.io Project Page
  - https://hackaday.io/project/166680-seattle-univeristy-aquaponics-sensing-system
- Splunk Pledge
  - https://www.splunk.com/en_us/about-us/splunk-pledge.html
- Seattle Engineers for a Sustainable World
  - http://www.seattleu-esw.org/

QA

splunk> .conf19

.conf19

splunk>

Thank You !

Go to the .conf19 mobile app to

RATE THIS SESSION