What's Cooking?

Recipes for bringing data to everything

splunk>
turn data into doing™

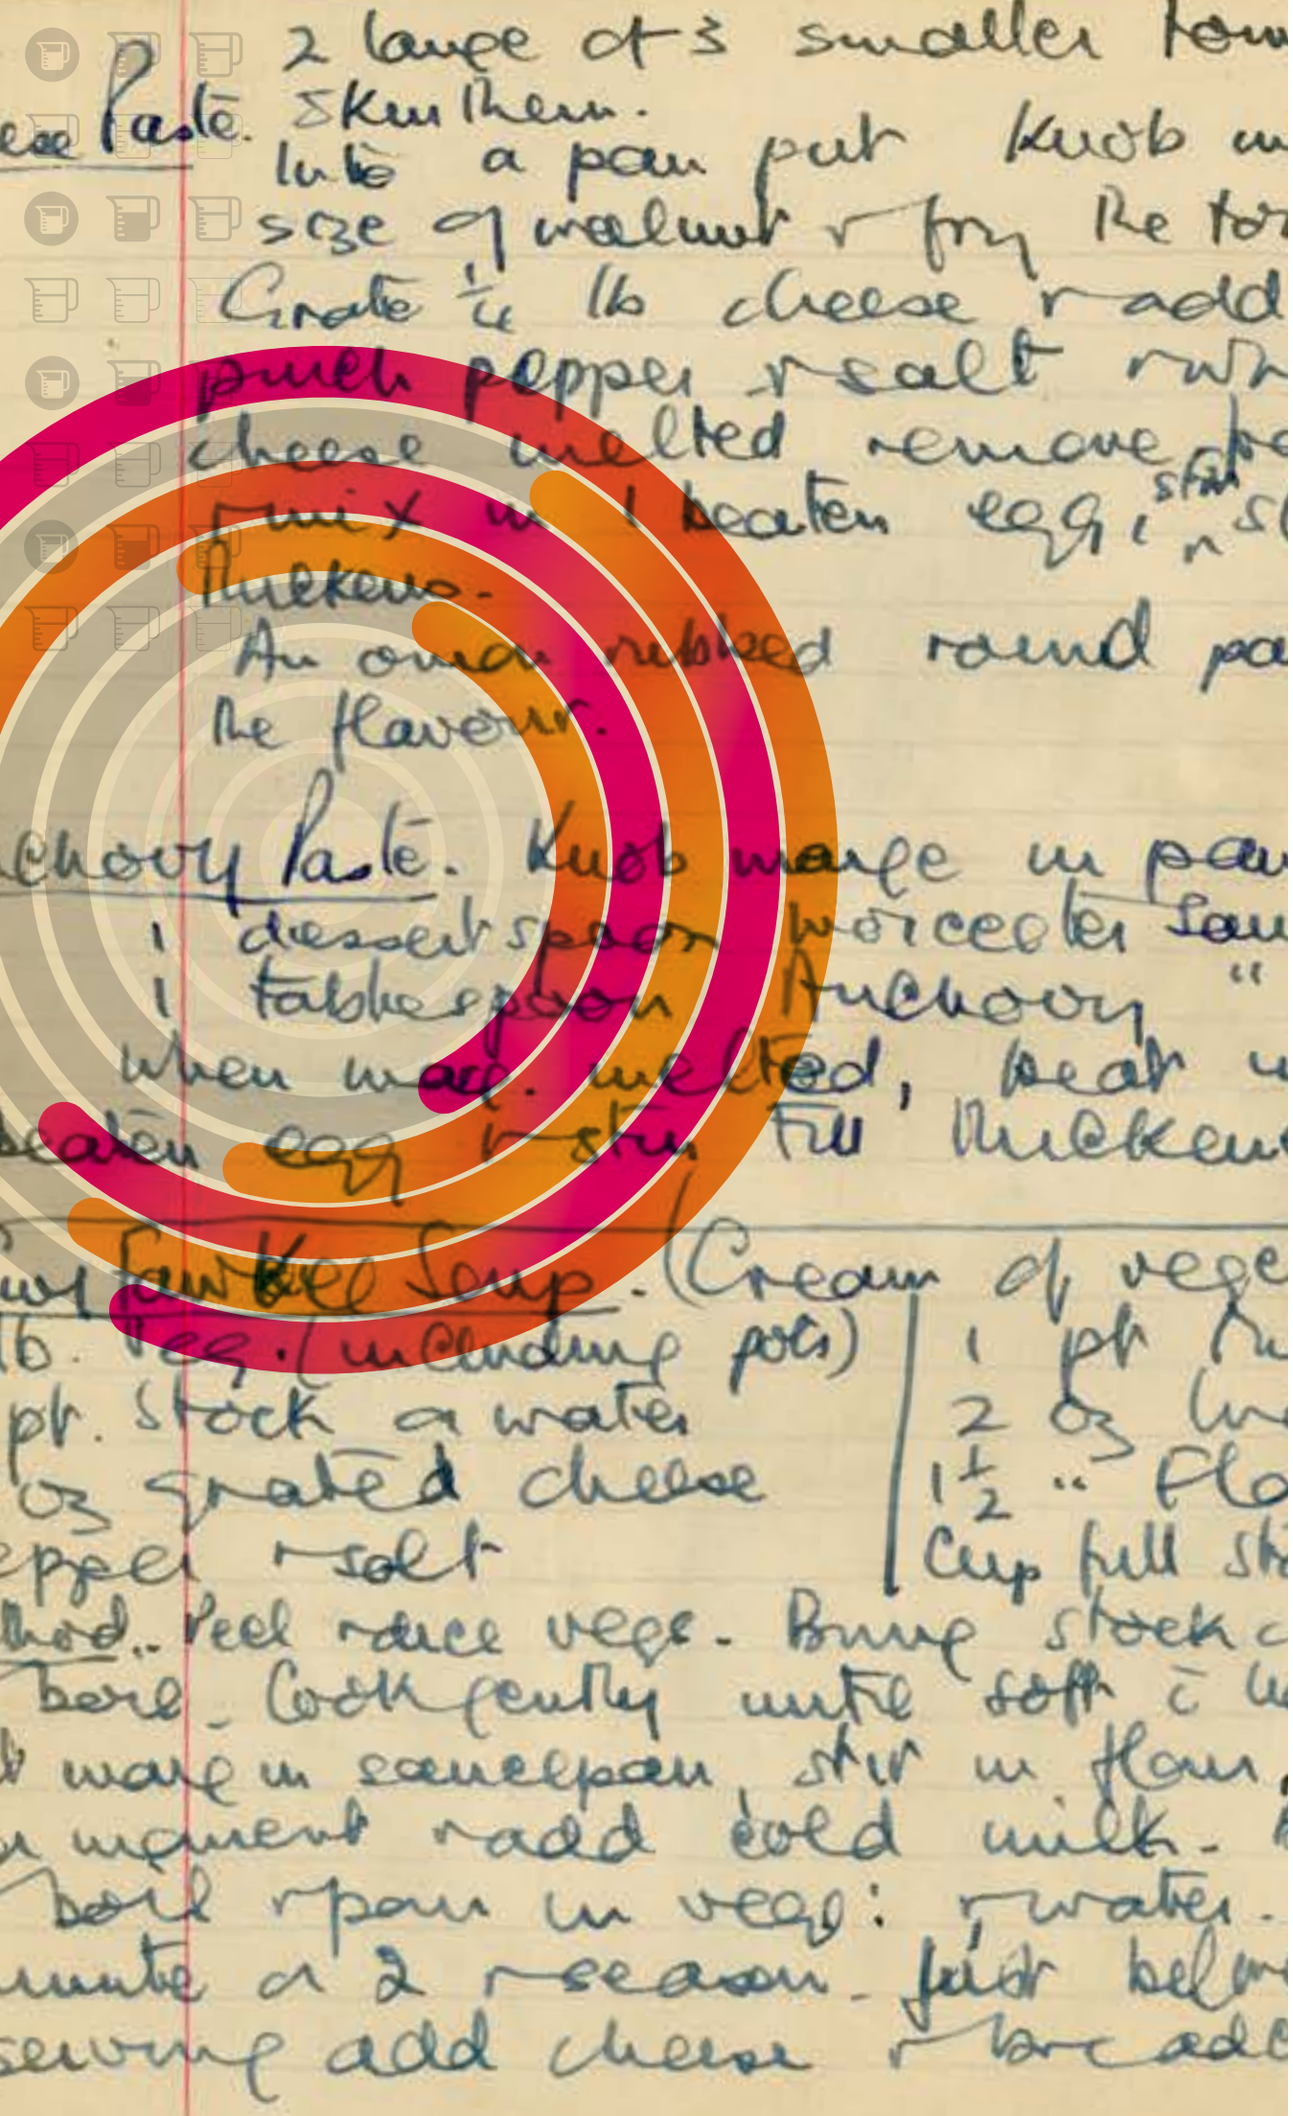

What's Cooking?

Recipes for bringing data to everything

The Secret Ingredient Is Data

"The Fannie Farmer Cookbook" was first published in 1896 and revolutionized home cooking by introducing something previously missing in cookbooks, or at least not written down: data. Not only did Farmer tell people what ingredients to use, she gave specific measures in tablespoons and cups, explained the chemical process of cooking and shared contemporary theories about nutrition. "Scientific cooking," Farmer said, "means the elevation of the human race." She also included a recipe called "Clam Frappé," but we'll forgive her for that.

Splunk customers have been finding cool and innovative ways to use the platform for years. We talk about people Splunking their homes and Splunking their commutes and Splunking their kids. In this book, you'll see a collection of new ways to use Splunk® tools to make your life easier. I think you'll find some of them surprising, but they all make sense. The fact that Splunk users have found ways to enhance and refine even centuries-old processes is not only fascinating, but it shows how technology has become an integral part of human life. It also shows how powerful the Splunk platform has become in helping drive nearly any outcome.

This book reflects the extraordinary time in which it was written. As people spend more time at home, they naturally turn their attention inward and think about how to make day-to-day activities not only more efficient, but more pleasant.

In this book, we'll showcase bakers who used a Raspberry Pi camera connected to a Splunk dashboard to monitor their sourdough starter instead of wrapping rubber bands around the jar. A music lover who looked to data to visualize their musical life. A lifelong surfer who built a dashboard to know the best time to go surfing. A dad who brought observability to his refrigerator when he noticed his pickles were going missing overnight.

Even if you don't surf, don't eat bread and have never detected inexplicable incursions into your pickle jar, the recipes in this book are built on sound principles and provide relevant examples of how data, analytics and visualization can make any process easier, more efficient, more precise and more effective.

I hope you enjoy this recipe book as much as I have, and have as much fun with it as the people who contributed the recipes and put it together. Maybe it will give you new insight into your work. Or maybe it will encourage you to finally begin your own sourdough starter.

Tim Tully
CTO, Splunk

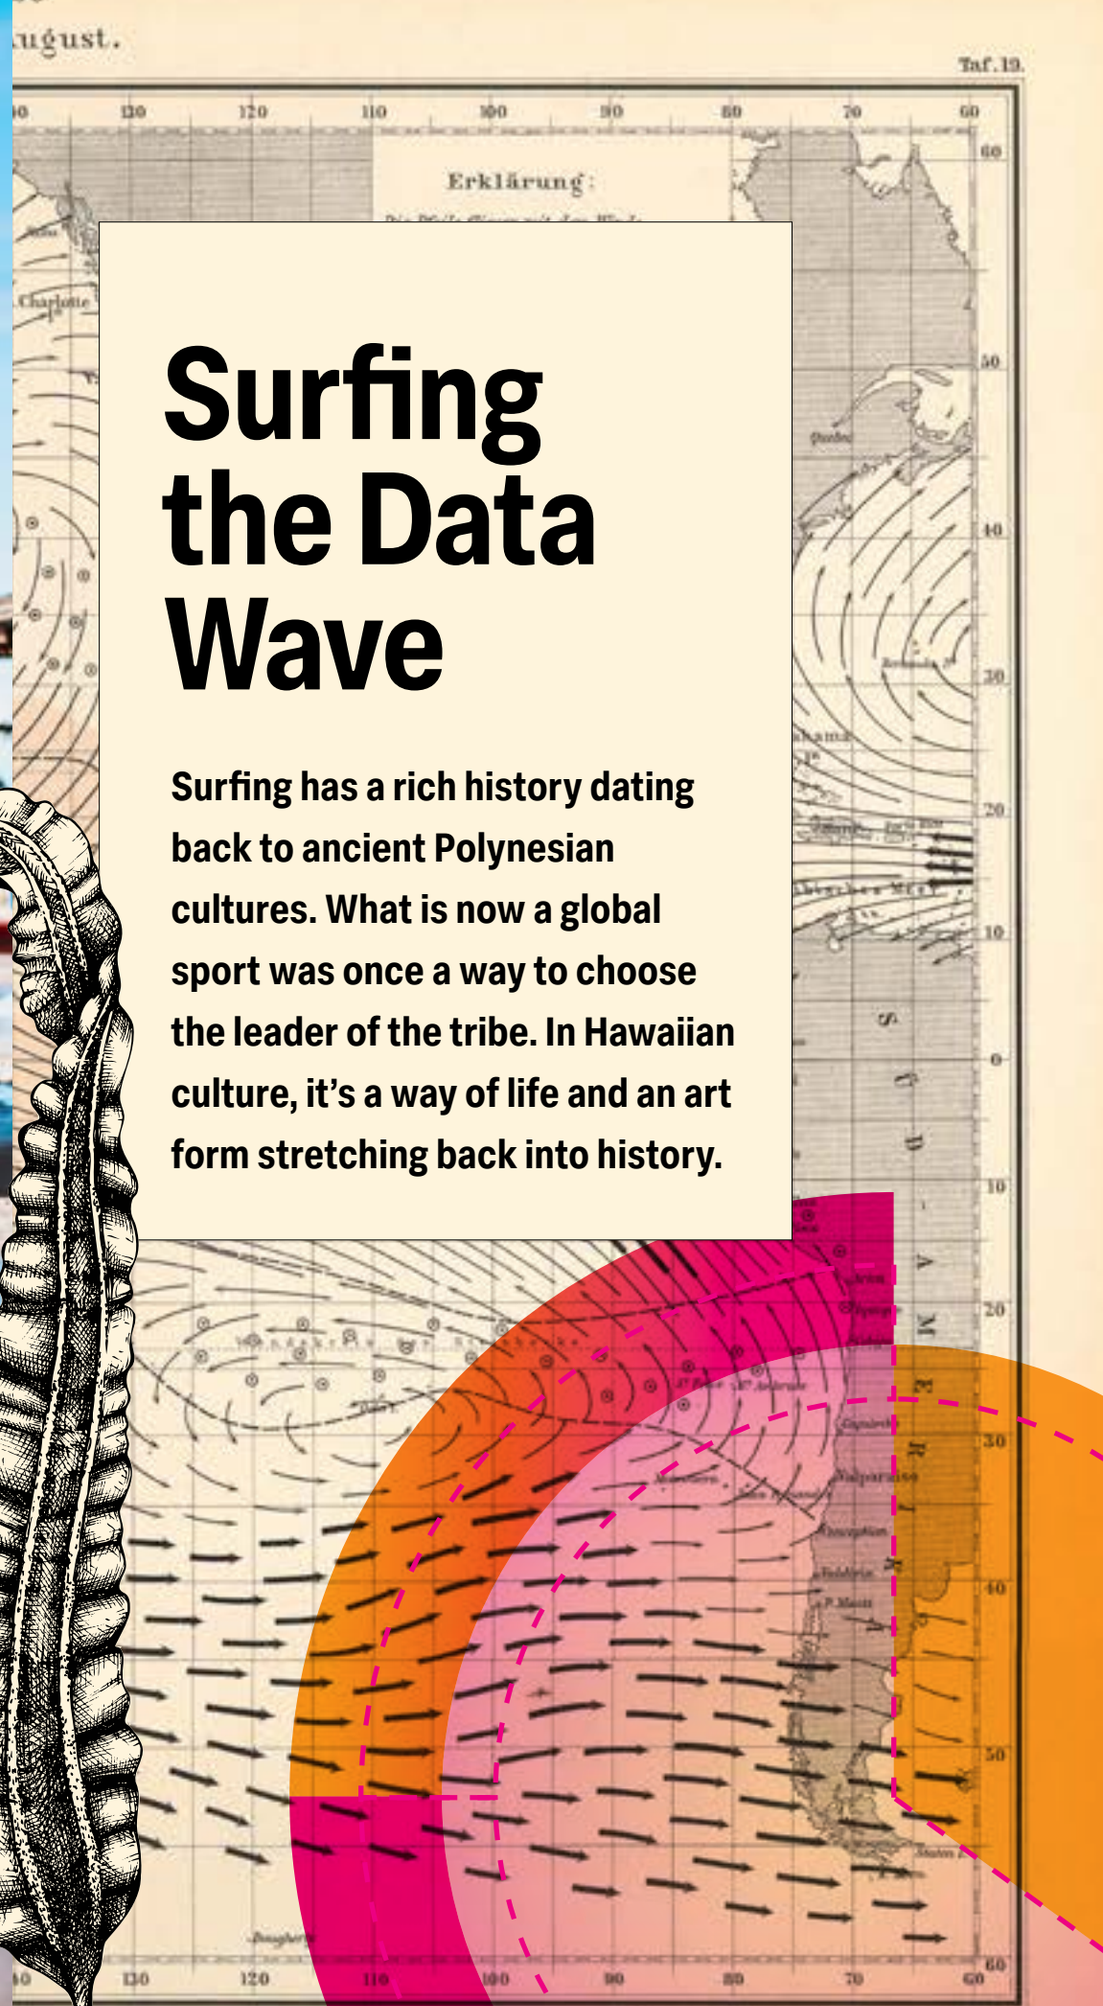

Surfing the Data Wave

Surfing has a rich history dating back to ancient Polynesian cultures. What is now a global sport was once a way to choose the leader of the tribe. In Hawaiian culture, it's a way of life and an art form stretching back into history.

TODAY, SURFING IN ALL ITS FORMS — longboard, shortboard and big wave surfing — is done on coasts around the globe, from Australia to Portugal to California.

Avoiding Kooks In Your Area

Surfing is continuing to evolve, with human-made wave pools, competitions featuring 100-foot waves and more. Data has been key in the evolution of the sport. Whether it's to train for bigger waves, engineer the perfect board or simply better understand the weather and surf, data is at the center.

Splunker Giovanni (Gio) Mola, a lifelong surfer, saw an opportunity in the increasing availability of surf- and surfing-related data.

The surfer's relentless hunt for the perfect wave inspired Gio to create the perfect surf report. One report, with the right information, available on all his devices. The information was out there but he could never get a single, clear view of where the waves would be best. He had to open three or four different apps — each offering one or two metrics he could trust — and still have to create his own mental forecast from the different sources.

Enter Splunk. The initial set-up only took him a few stints of work in his spare time. After a couple of weeks, he was fine-tuning his data sources and alerts. Now it's become a reliable way to predict the perfect surf conditions. In fact, it's so good, a few of his buddies use his dashboard on their mobile devices.

The report improves his surf sessions, allowing him to choose the best spots — and get there before the beach is crowded. He's even catching himself surfing on public cams that monitor the beach by correlating his surf data with the data coming from the cameras and downloading the images. The results have inspired him to look into building something similar for the ski/snowboard season.

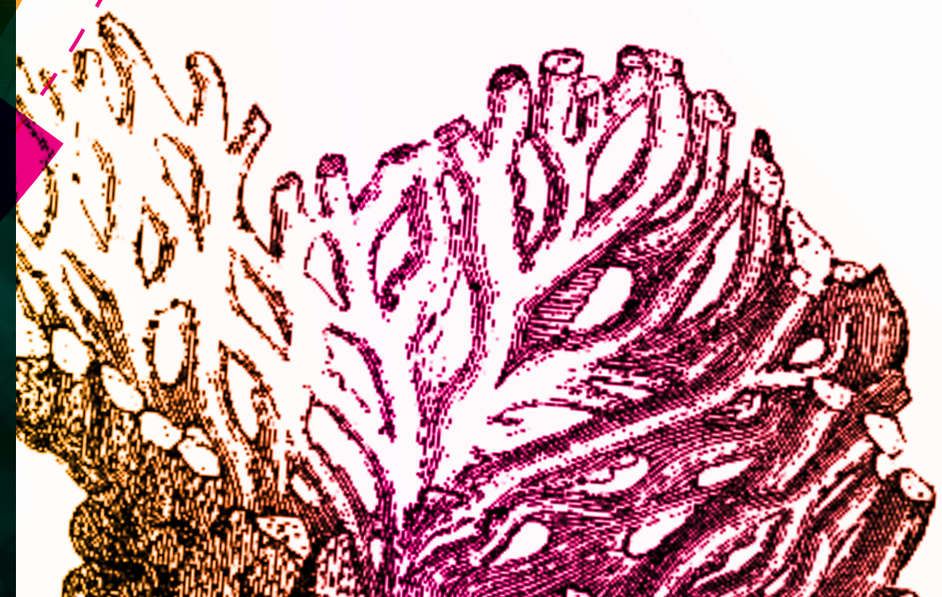

Get Tubular, On Time, Every Time

How to get the most reliable surf data from multiple sources in one dashboard, anywhere

Level	Time to create	Makes
●●○	10 hours	1 killer surf report

Ingredients

- A Splunk instance
- Basic Python scripting skills
- Surfline Premium subscription
- NOAA access

Cooking Instructions

1. Start with your favorite surf forecasting tool, understand the API and write a Python script to schedule your spot's forecast to be ingested once a day into Splunk using a scripted input.
   a. For Surfline forecast, you can log in to your preferred region's forecast, open the Google Chrome debugger, click the XHR tab and use the APIs labeled waves, tide, etc. Refer to the appendix for an example of the JSON endpoint for the wave information.

   ```
   https://services.surfline.com/kbyg/spots/forecasts/
   wave?spotId=5842041f4e65fad6a770880a&days=
   6&intervalHours=1&maxHeights=false
   ```

2. Follow up with finding the closest NOAA buoys within your region and schedule the hourly delivery of those buoy readings like swell height, swell period, wind wave height, wind wave period and mean wave direction into your Splunk surf index.
3. At the NOAA website you can access all real-time buoy-related data and choose the buoy closest to your region. You could, for instance, get tabular format spectral buoy data from the San Francisco Bar buoy.
4. Choose the weather forecasting tool of your choice to decipher wind and weather conditions in your local area and schedule those inputs at your given preference.
5. Finally, create one dashboard that combines all of the preferred sources of surf, tide and weather forecasts into one simplified view, and let the data tell you where and when the surf in your local region will be best.

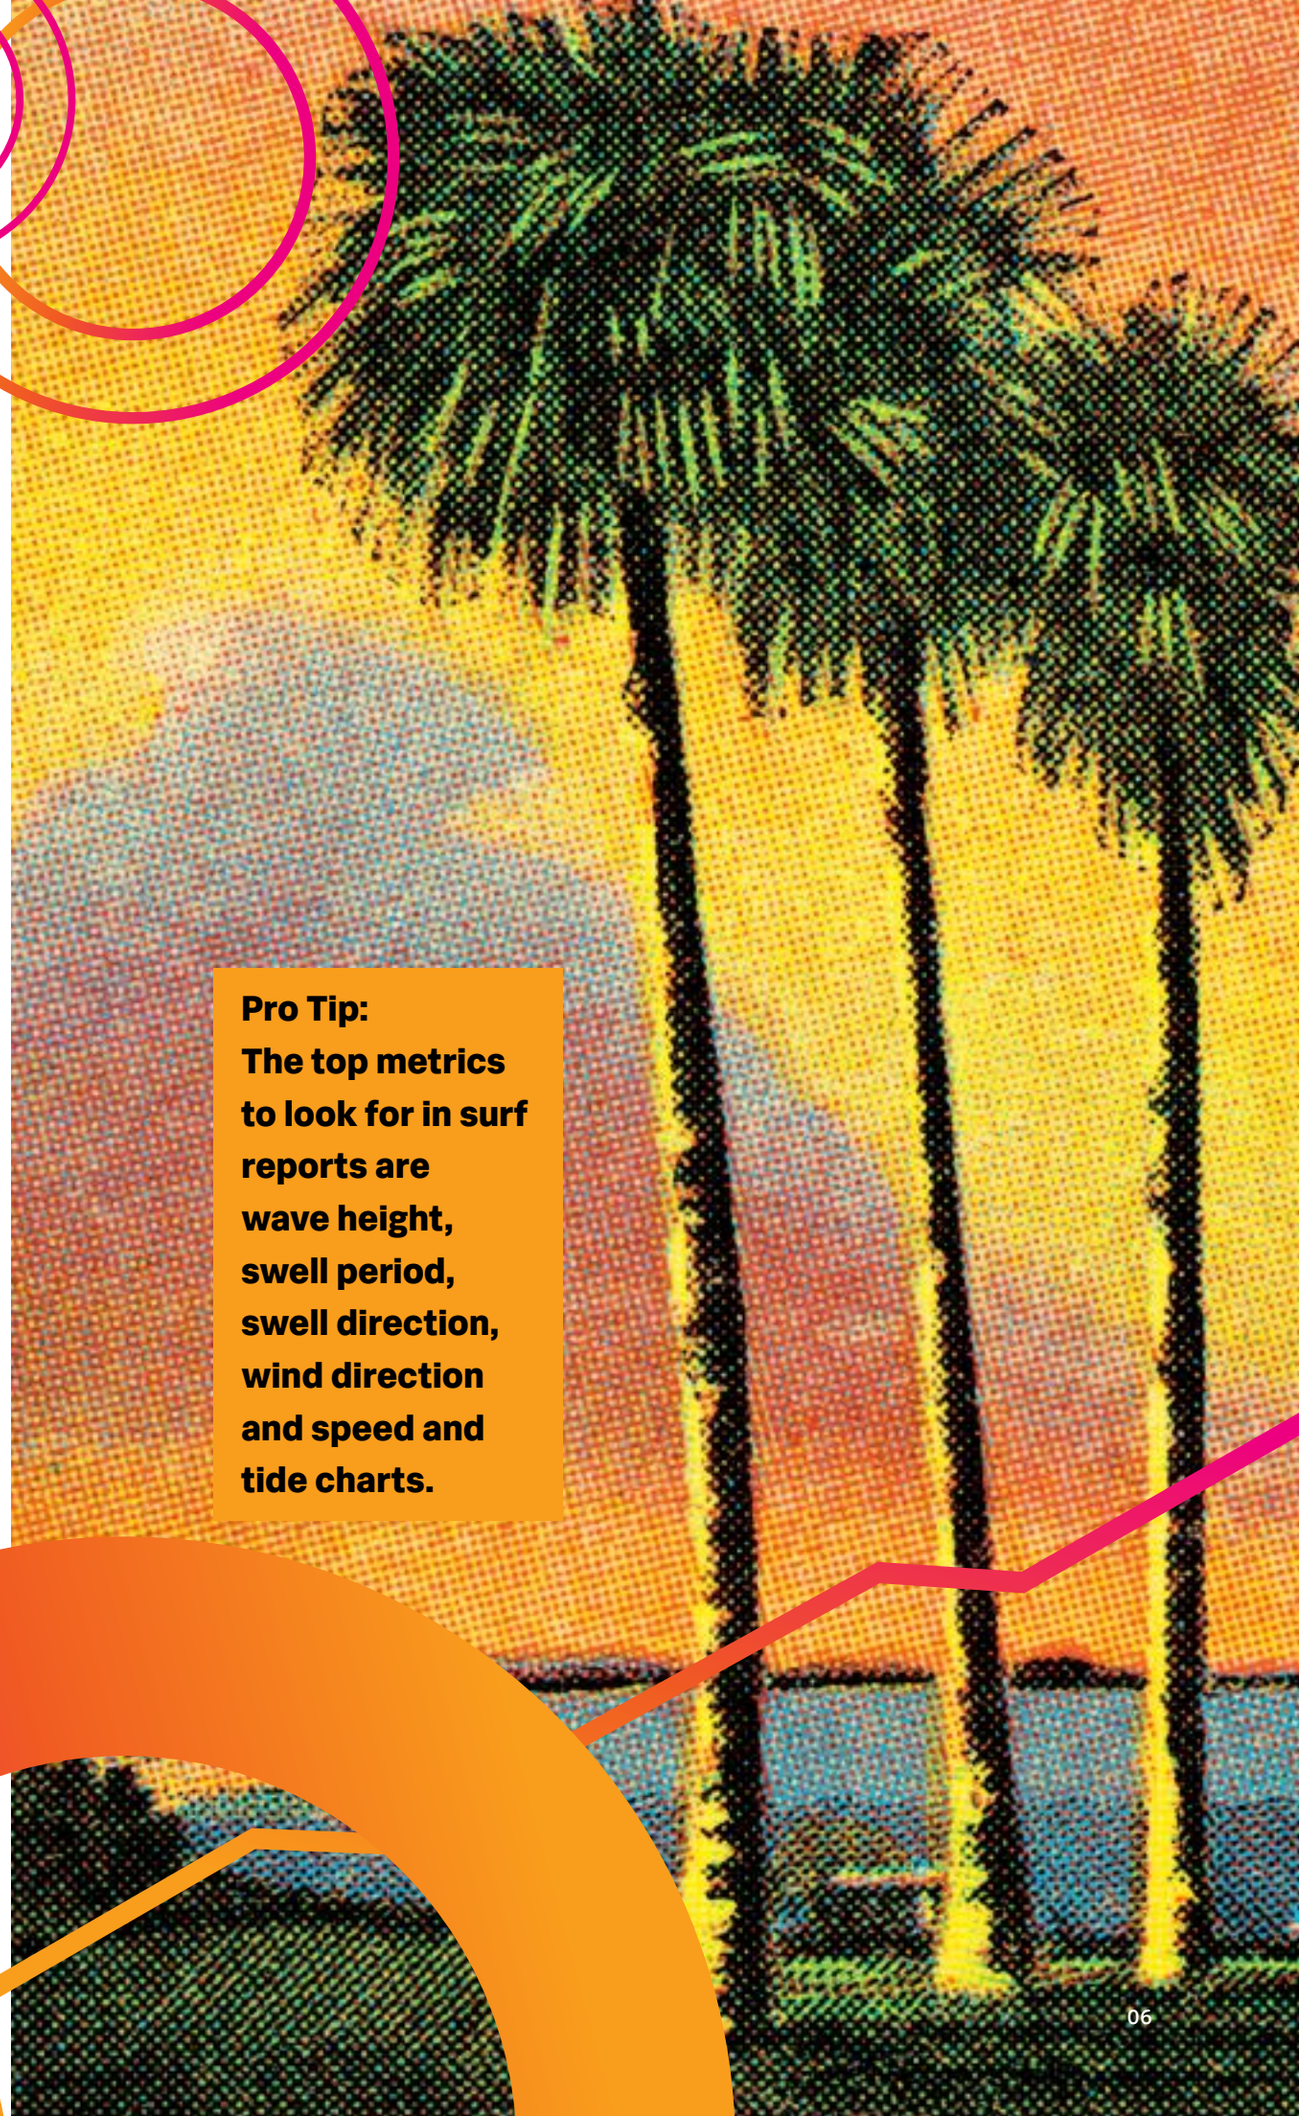

Pro Tip:
The top metrics to look for in surf reports are wave height, swell period, swell direction, wind direction and speed and tide charts.

Musical Journeys Through Data

Music is as old as rocks. The origin of music predates recorded history itself, but experts believe that the first music might have been invented in Africa, before ancestral populations first dispersed around the world.

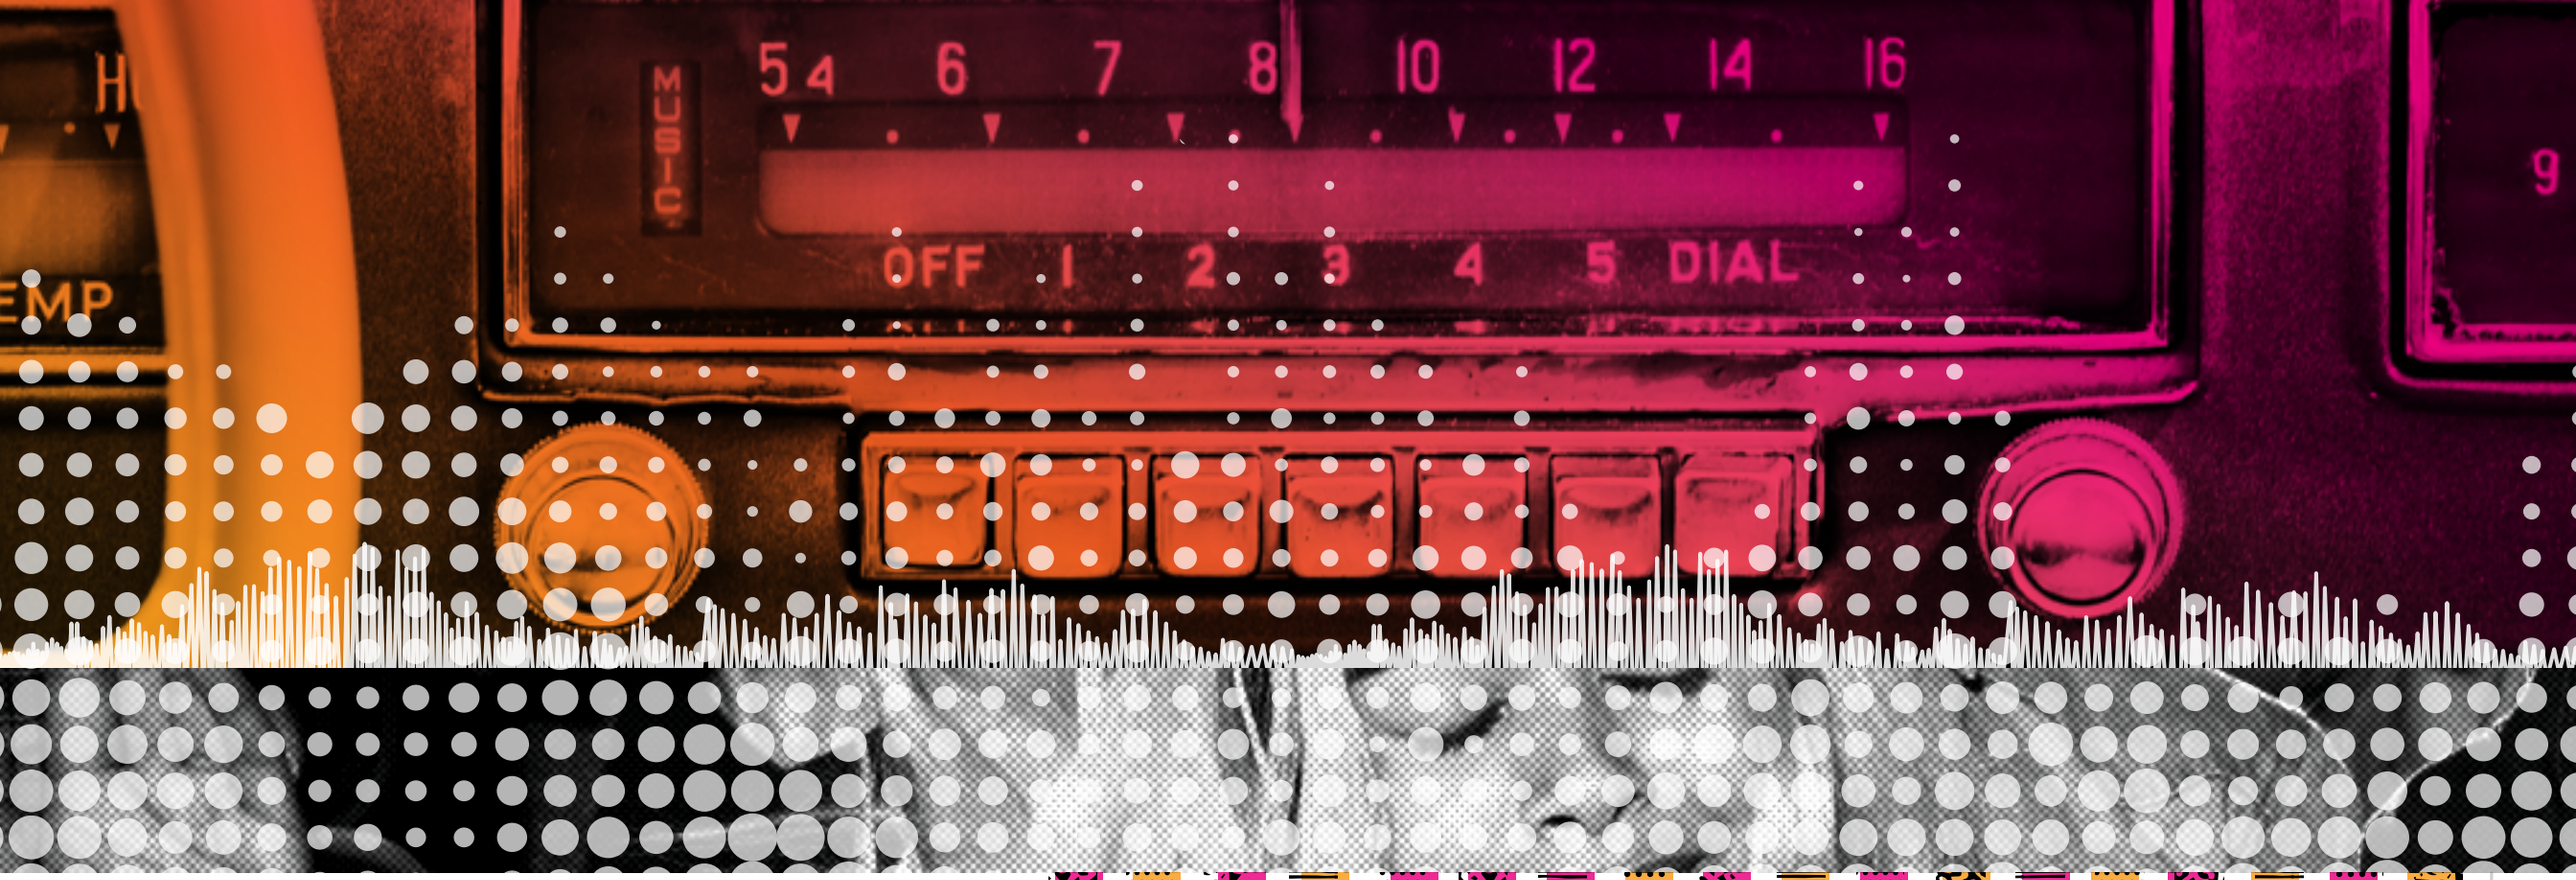

ACROSS THE CENTURIES, music has evolved within every culture and every community, and it's as much a part of our lives as ever before. We stream it while we're hard at work. We play it on road trips while we gawk at majestic sunsets. We feel it with every fiber of our beings when we're dancing on concert floors. And in today's digital age, music is only ever a click or tap away.

Listening to the Data in Your Music

Digitization of music consumption and production has dramatically increased the amount of accessible, music-related data. In most cases, listening to music today means listening across multiple devices and multiple services — Spotify, Pandora, SoundCloud and more — generating symphonies of untapped information.

Splunk's Sarah Moir, a music aficionado with a whopping 150+ concerts under her belt, saw an opportunity in the intersection of music and data. She had questions about her listening patterns that existing tools couldn't quite answer. Songkick told her when artists in her library were coming to town, but what if she wanted to figure out which of her recently discovered artists were playing locally? Spotify Wrapped identified her most played songs of the year, but how could she identify her personal one-hit wonders?

To get the answers she wanted, Sarah imported concert data and information from the music services she used into Splunk. Searches and visualizations uncovered patterns in her musical listening, identified her most-listened-to songs since moving to California, highlighted her most frequented concert locations and answered multitudes of other questions.

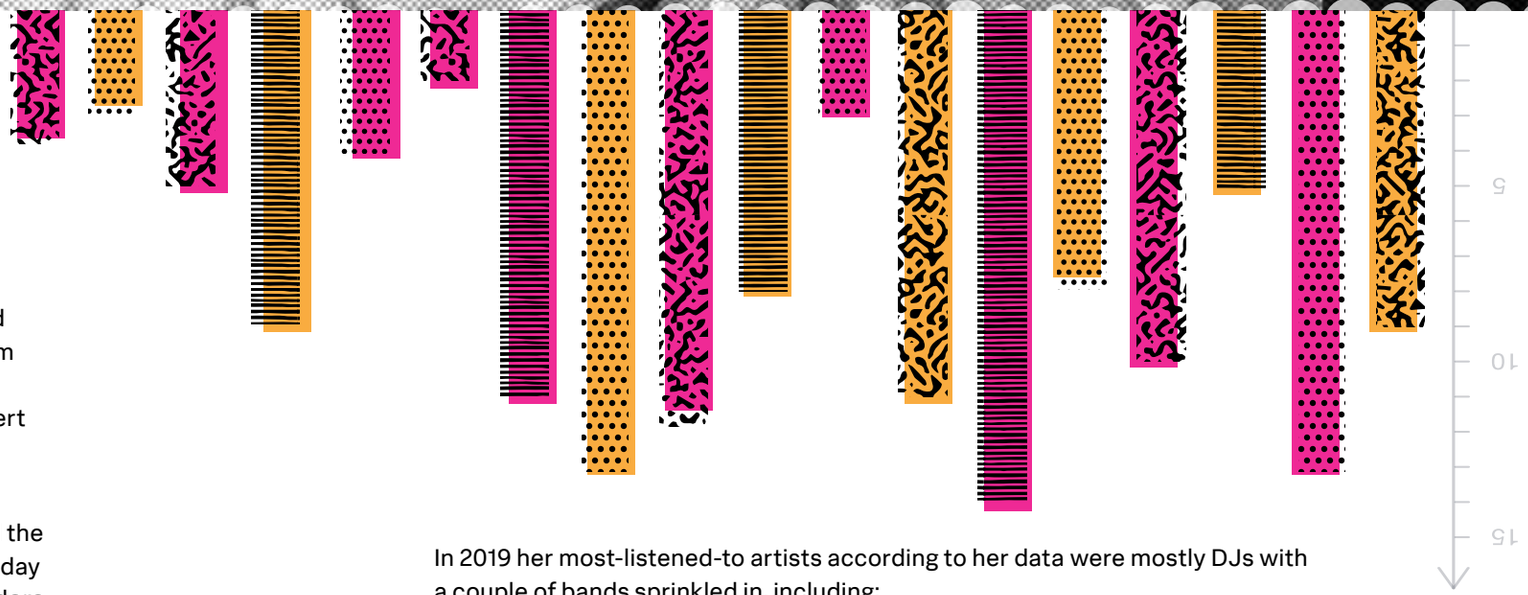

In 2019 her most-listened-to artists according to her data were mostly DJs with a couple of bands sprinkled in, including:

- Tourist
- Lane 8
- Benoit & Sergio
- The Vaccines
- Litany

Bringing all her music data into Splunk gave Sarah a deeper understanding of her musical life, including insights into how her listening habits have altered since she began sheltering in place. It came as no surprise that she's seen a marked rise in time spent listening to music, directly corresponding to key stages of the pandemic's unfolding.

But the insights didn't stop there. Sarah also realized the number of tracks she listened to from May 2019 to May 2020 increased by 25%, but the amount of time she spent listening had grown by 74%, or 150 hours, since the pandemic started.

Pro Tip:
Trellis layouts are great for splitting visualizations and showing you charts for each artist, while choropleth maps are handy for mapping your concert location data.

The Quantified, Musical Self

How to listen to your music data across all your devices and services

Level	Time to create	Makes
● ● ●	Around 3-4 hours	Any number of insight-packed reports

Ingredients

- A Splunk instance
- Basic Python scripting skills
- Regex knowledge (or friends with regex knowledge)
- A Last.fm account and configured scrobblers
- Splunk Add-On Builder
- Music App for Splunk (optional)
- Lookup File Editor (optional)
- Records of concerts you've attended, livestreams you've watched or tickets you've purchased (optional)

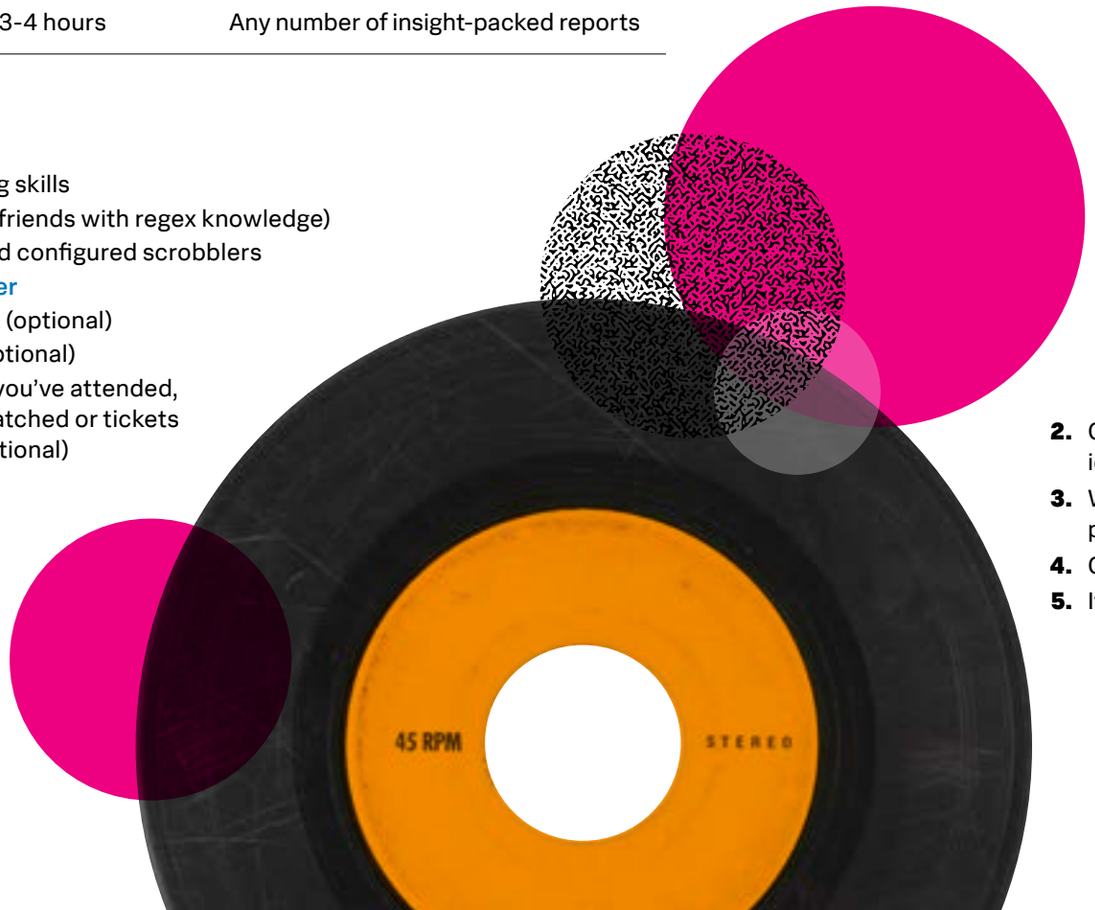

Cooking Instructions

1. Add all the relevant data into Splunk. There's a variety of methods, depending on the music data source.
   a. iTunes library for music library insights: upload files with the help of the iTunes_csv script, which converts XML files to CSV files. Alternatively, the Music App for Splunk can extract the tracks from an iTunes Library XML file. However, the latest version of Apple Music is not supported.
   b. Spotify library for more music library insights: Use a Python script to pull the Spotify library data. You can also request structured data about your playlists from Spotify and Shazam directly.
   c. Last.fm for listening insights: Use the Add-on Builder or your own Python skills to write a modular input to retrieve recent tracks from the Last.fm API.
   d. Concert and ticket data for event insights: Use the Lookup File Editor app for uploading and formatting data, because it lets you edit the data in Splunk itself.
2. Generate a list of questions that you have about your music listening, and identify the data necessary to answer them.
3. Write searches to answer your burning questions in Splunk's search processing language (SPL).
4. Create visualizations to show off the results.
5. Iterate on searches and visualization for precision.

The Future Is Buzzing With Data

Beekeeping and being a beekeeper, or apiarist, has been a practice since before the domestication of bees. In fact, in the Cueva de la Araña (Cave of the Spider) near Valencia, Spain, a painting dating back to 9000 B.C.E. depicts a brave (or foolish) human climbing a tree to stick their hand directly into a beehive.

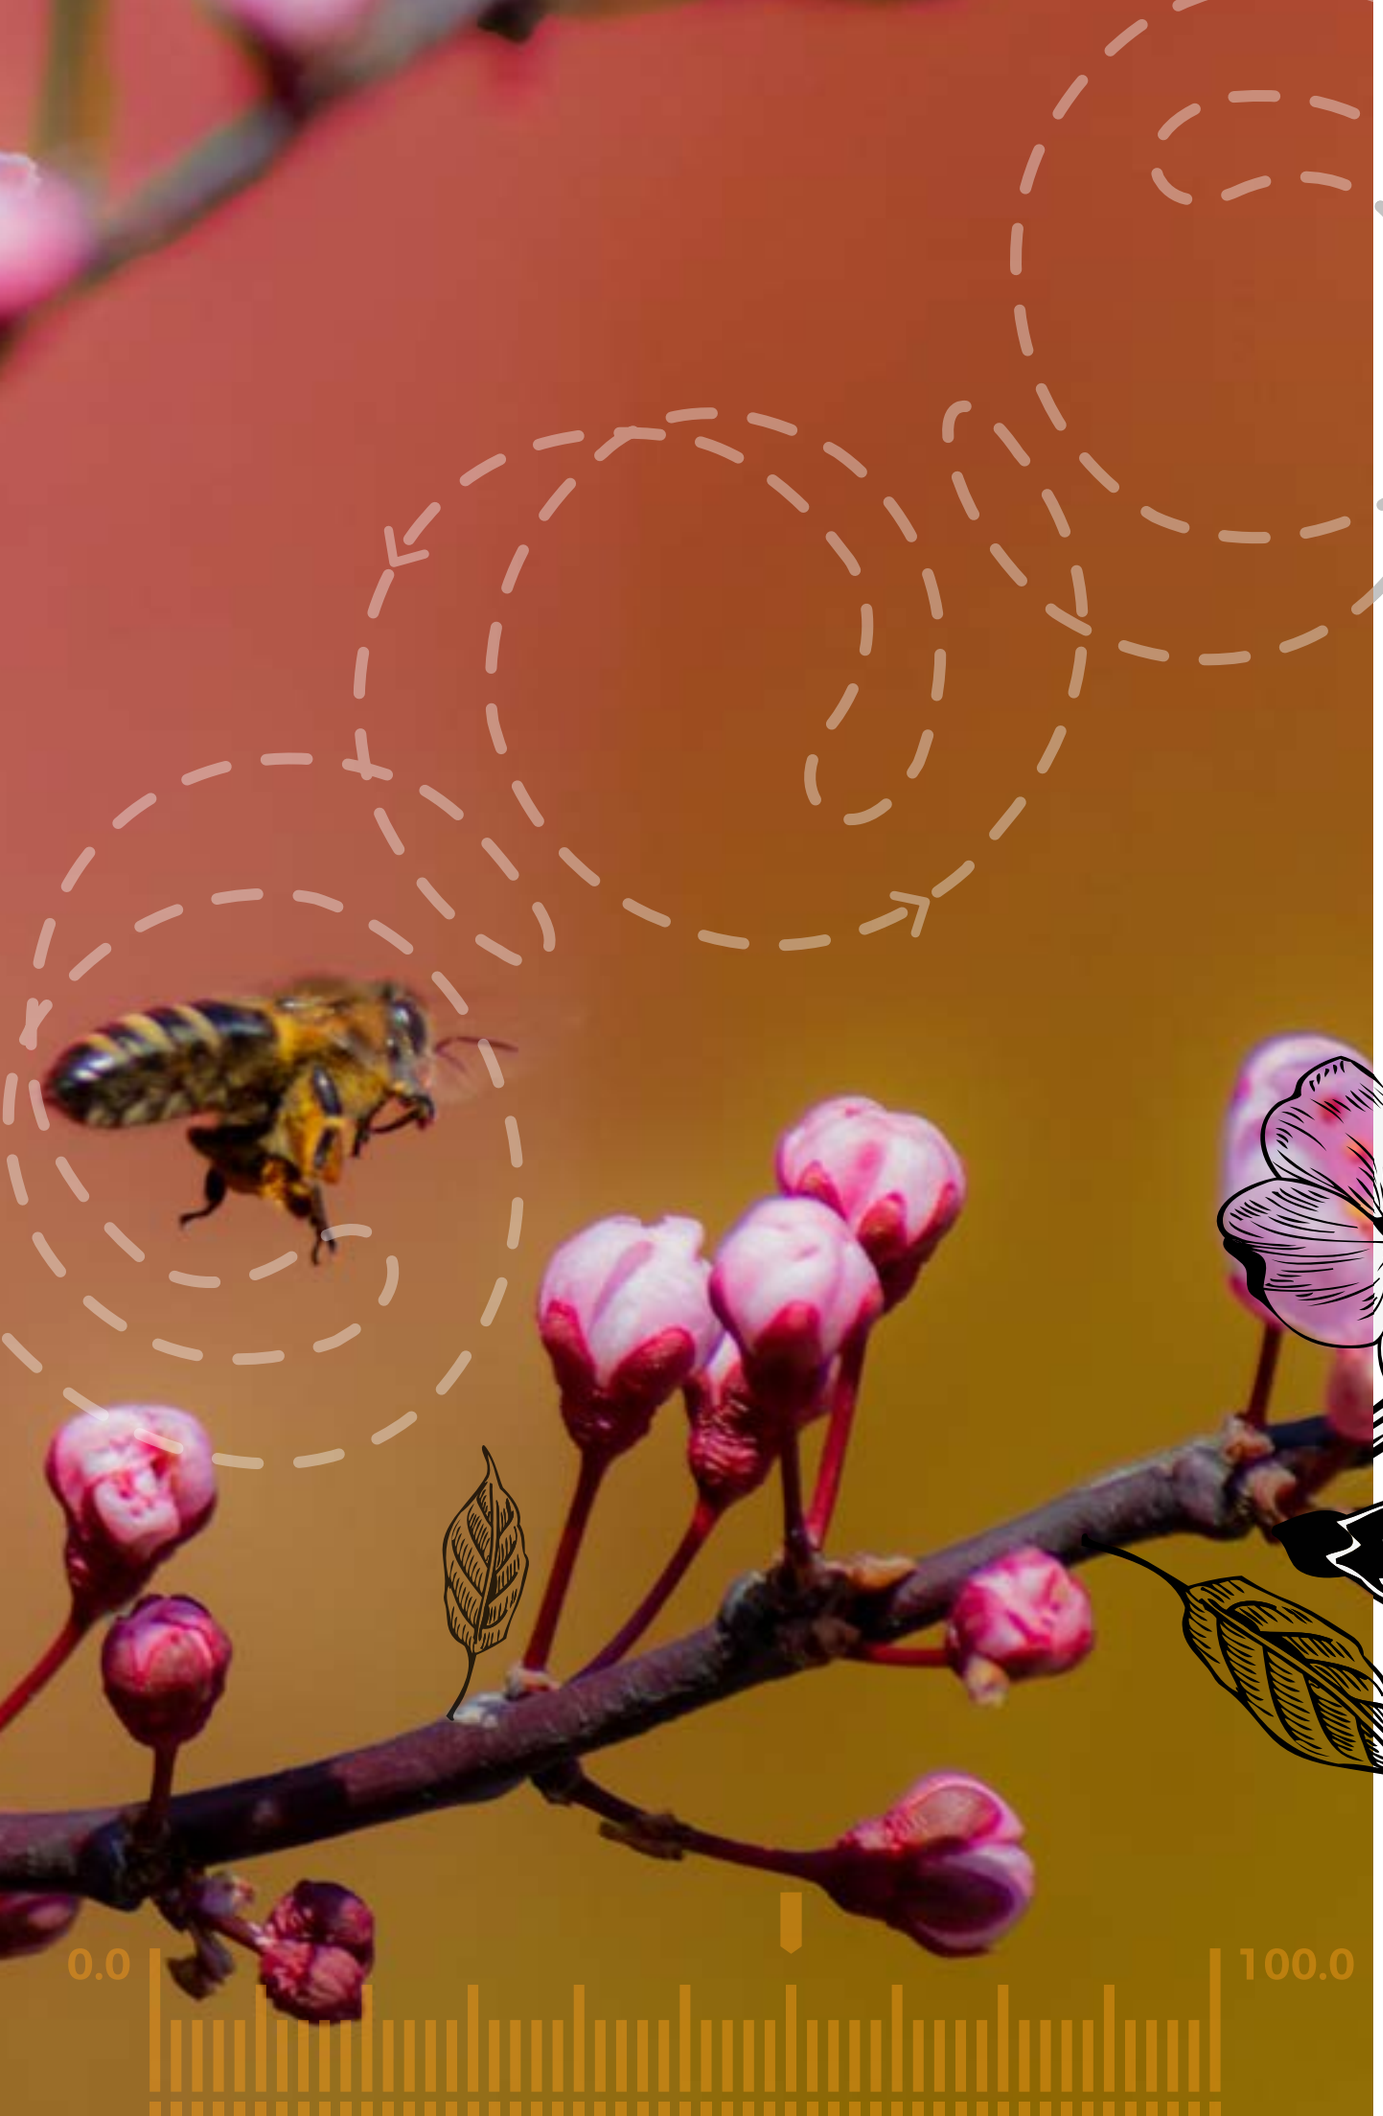

IN THE 18TH AND 19TH CENTURIES, apiarists started to explore sustainable ways of harvesting honey without harming colonies, leading to wooden hive boxes developed and improved by François Huber and Thomas Wildman.

In the 1850s, a bee hobbyist in Pennsylvania, Lorenzo Langstroth, observed that bees wouldn't build a honeycomb in a space tighter than 1 centimeter, or 3/8 of an inch. So Langstroth developed a hanging bar with removable frames that were spaced exactly 1 centimeter apart from each other and 1 centimeter away from the box walls. To this day, Langstroth hives are the most popular amongst hobbyist beekeepers and professional apiarists alike.

Bringing Data to Bees and the Environment

Imagine how quickly the apiarists of history could have innovated on beehive design with the amount of data available to us today. That's exactly what Splunk software engineer Nia Watts is doing. Nia connects SparkFun environmental sensors directly in and around her beehive to constantly track the health and output of her hive. She monitors environmental factors such as barometric pressure, temperature, sound, light, air quality, weight and more. The sensors are covered in waterproof casings to protect them from weather and hive activity. Every five minutes, the data from these sensors is collected by Nia's Raspberry Pi and forwarded into Splunk.

Nia began to care for her 60-pound beehive five years ago, but brought Splunk into the picture about two years ago when she wanted to learn more about the activity, health and output of her bees. Because Nia keeps her Langstroth beehive at a local botanical garden, she also uses Splunk to remotely receive alerts when hive activity seems out of the ordinary.

On top of using weight sensors and Splunk to detect the hive's potential readiness for harvest, Nia can use her data to track bee behavior due to environmental events, such as wildfires or thunderstorms. Nia has been able to detect nearby wildfires and impending thunderstorms by correlating bee activity — tracked by light sensors and microphones — with barometric pressure and air quality data.

Nia wants to track more data in the future, including UV indexes and wind speed, and use Splunk to detect events like wasp threats or unusual behavior due to weather. Bringing data to beekeeping can help apiarists maintain healthy hives and even bring insights to environmentalists looking to understand the effects of climate change on bees.

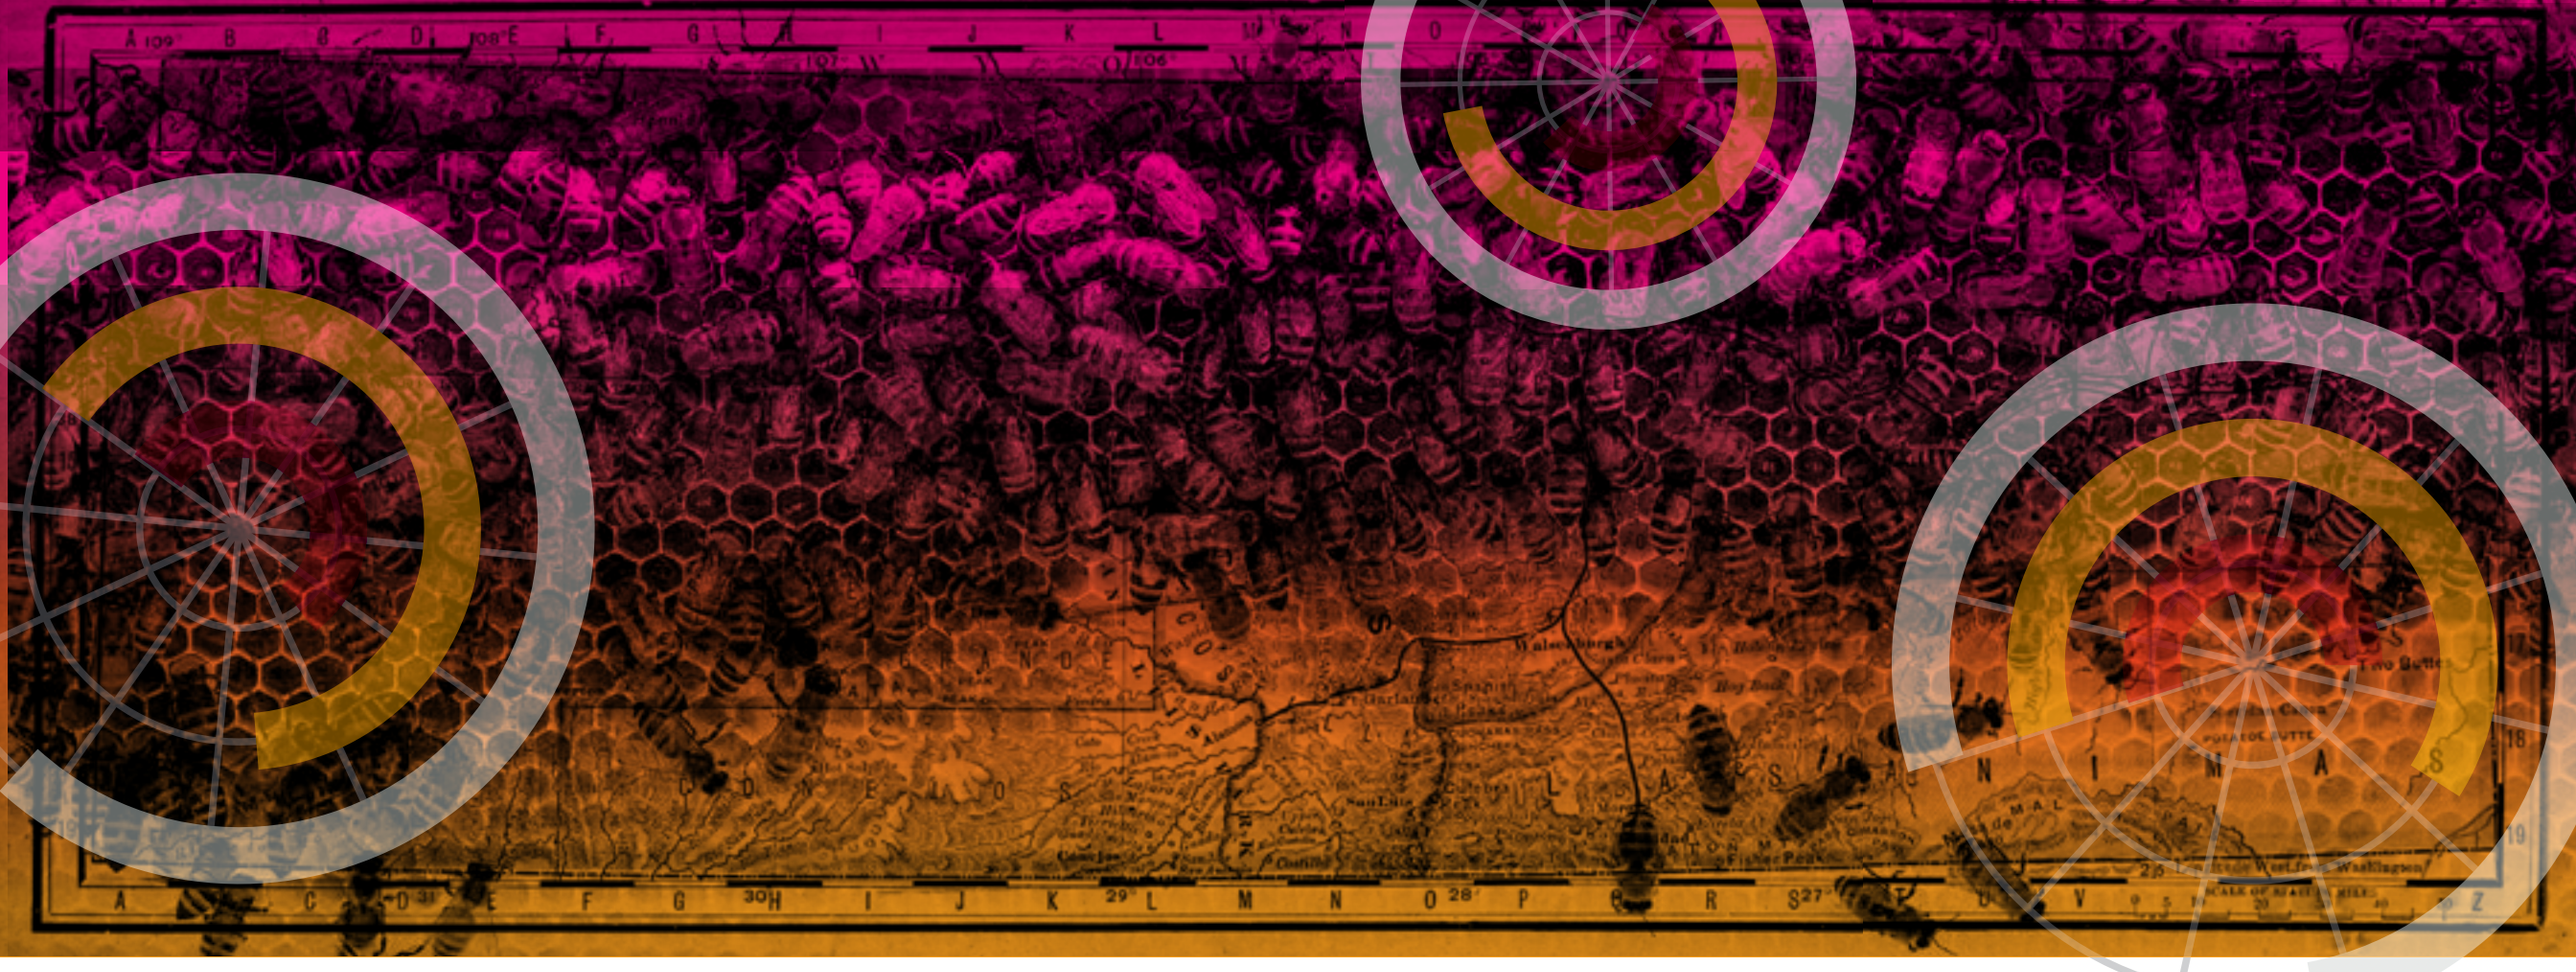

Become the Hivemind

Bringing data to bee-havior

Level	Time to create	Makes
● ● ●	10 hours	1 happy, healthy beehive and a honey of a visualization

Ingredients

- A Splunk instance
- A Beehive
- 3-5 SparkFun Environmental Combo Breakout Sensors
  - Including waterproof sensor containers
- 1 microphone (optional)
- 1 Qwiic pHAT for Raspberry Pi
  - To connect breakout sensors to Raspberry Pi
- 1 Raspberry Pi

Cooking Instructions

1. Set up the Raspberry Pi.
2. Plug the Qwiic pHAT into your Raspberry Pi.
3. Set up the sensors (put them in the containers and drill tiny holes into them for the wires).
   a. 2-3 sensors inside the hive (internal conditions)
   b. 1-2 sensors on the bottom of the hive (external conditions)
   c. Tip: Run the wires directly out of the hive to the Raspberry Pi, not through the hive, to protect the hive and your hardware. The data from the sensors will go straight to the pHAT and then into the Raspberry Pi.
4. Set up the microphone (if you want)
5. Enable I^2C (Inter-Integrated Circuit) pins on Raspberry Pi.
   a. Need to use the Raspberry Pi config tool (raspi-config) to enable the I2C pins.
   b. `i2c_arm_baudrate` needs to be set to 10000 (setting the speed at which data is fed through the system).
6. Reboot Raspberry Pi.
7. Install Python packages.
   a. splunk-sdk
   b. Sparkfun_qwiic
8. Configure your Splunk Enterprise credentials and location in the code (example below).
   a. HOST = “localhost”
   b. PORT = 8089
   c. USERNAME = “admin”
   d. PASSWORD = “yourpassword”
9. Run the code at Nia’s GitHub repository (sending data into Splunk every 5 minutes).
10. Set up dashboards in Splunk for temperature, humidity, air quality, total volatile organic compounds (TVOC), CO2 and barometric pressure data.
    a. Create time-series dashboards for all of the above.
11. Iterate on your data needs, use cases and visualizations (e.g., light, sound, etc.) to learn more about how a bee’s environment and the weather affects the health of your hive.

Your Health Is in the Data

Ancient philosopher Hippocrates once said, “If we could give every individual the right amount of nourishment and exercise, not too little and not too much, we would have found the safest way to health.”

TODAY, HEALTH is less of a guessing game. Nutritionists and athletes are continually looking to data to learn how diet and exercise impact overall performance. New technology is making real-time, granular fitness data readily available. Accessible apps like Strava make it possible to get data from all sorts of activities like cycling, running, hiking and skiing.

And it can go even deeper. Connecting Strava to Splunk allows individualized analysis of fitness progress, consistency and overall health. Everyone from professional athletes to work-from-home accountants can use Splunk and Strava to bring data to their health.

Improving Workouts Without Breaking a Sweat

When Patrick Peeters first started working at Splunk, he also embarked on a journey to improve his overall fitness. So naturally, he immediately sat down for a few hours in front of a computer. Patrick decided it would be a great idea to connect Splunk to his Strava app. However, Strava had recently implemented OAuth 2.0 for API access, requiring additional inputs to access the data. So Patrick took matters into his own hands, creating the Strava for Splunk Add-On by using Python and Splunk's Add-On Builder.

Patrick uses Splunk to track a number of different fitness data points, including bike rides or running performance when jogging alone, or when he jogs with one or two of his kids in a stroller — a particular partner favorite. Patrick created dashboards in Splunk to track the number of "kudos" (likes) he gets for his Strava activity, total calories burned, elevation gain, longest activity streaks and money saved in gas or tolls due to cycling to work. Additionally, Patrick uses the Strava data in Splunk and overlays it on a map to create a visualization of everywhere he runs and rides his bike.

Others have been using the Splunk for Strava Add-On to track not only their own progress but, because it supports multiple users, also that of friends, family and colleagues for some extra motivation and healthy competition.

Working on Your Workouts

Sweating the small stuff

Level	Time to create	Makes
●○○	2 hours	1 powerful data platform for Strava + Splunk

Ingredients

- A Splunk instance
- A Strava account
  - Strava API Access
- Strava for Splunk Add-On
- Python knowledge (optional, only for integrating heatmap)

Cooking Instructions

1. Read the Strava for Splunk Add-On's details page to complete the initial setup and authorize the app via OAuth 2.0.
2. Install the Strava for Splunk Add-On.
   a. For the required fields, use the values you got from Strava (Pro tip: The starting timestamp is optional if you only want to get data for a certain time.)
3. You should now have access to your Strava events in Splunk. (The add-on has a sample dashboard that can help to get you started.)
4. The Strava API provides a wealth of details on each activity, allowing you to now build dashboards for the data most pertinent to your needs.
   a. E.g., total cycling distance, total running distance, total activity duration, weekly run distance, etc.
   b. Code example for "Most Kudo'd Activity":

```
<panel>
   <single>
   <title>Most kudo'd activity</title>
   <search>
         <finalized>
         <set token="kudos_name">$result.name$</set>
         <set token="kudos_date">$result.date$</set>
         </finalized>
         <query>
         index=strava sourcetype=strava:activities type=*
         | eval date = strftime(_time,"%d-%m-%Y")
         | stats max(kudos_count) as kudos by name, id, date
         | head 1
         | fields kudos name date id
         </query>
         <earliest>$top_timepicker.earliest$</earliest>
         <latest>$top_timepicker.latest$</latest>
   </search>
   <option name="underLabel">$kudos_date$: $kudos_name$</option>
   <option name="unit">kudos</option>
   </single>
</panel>
```

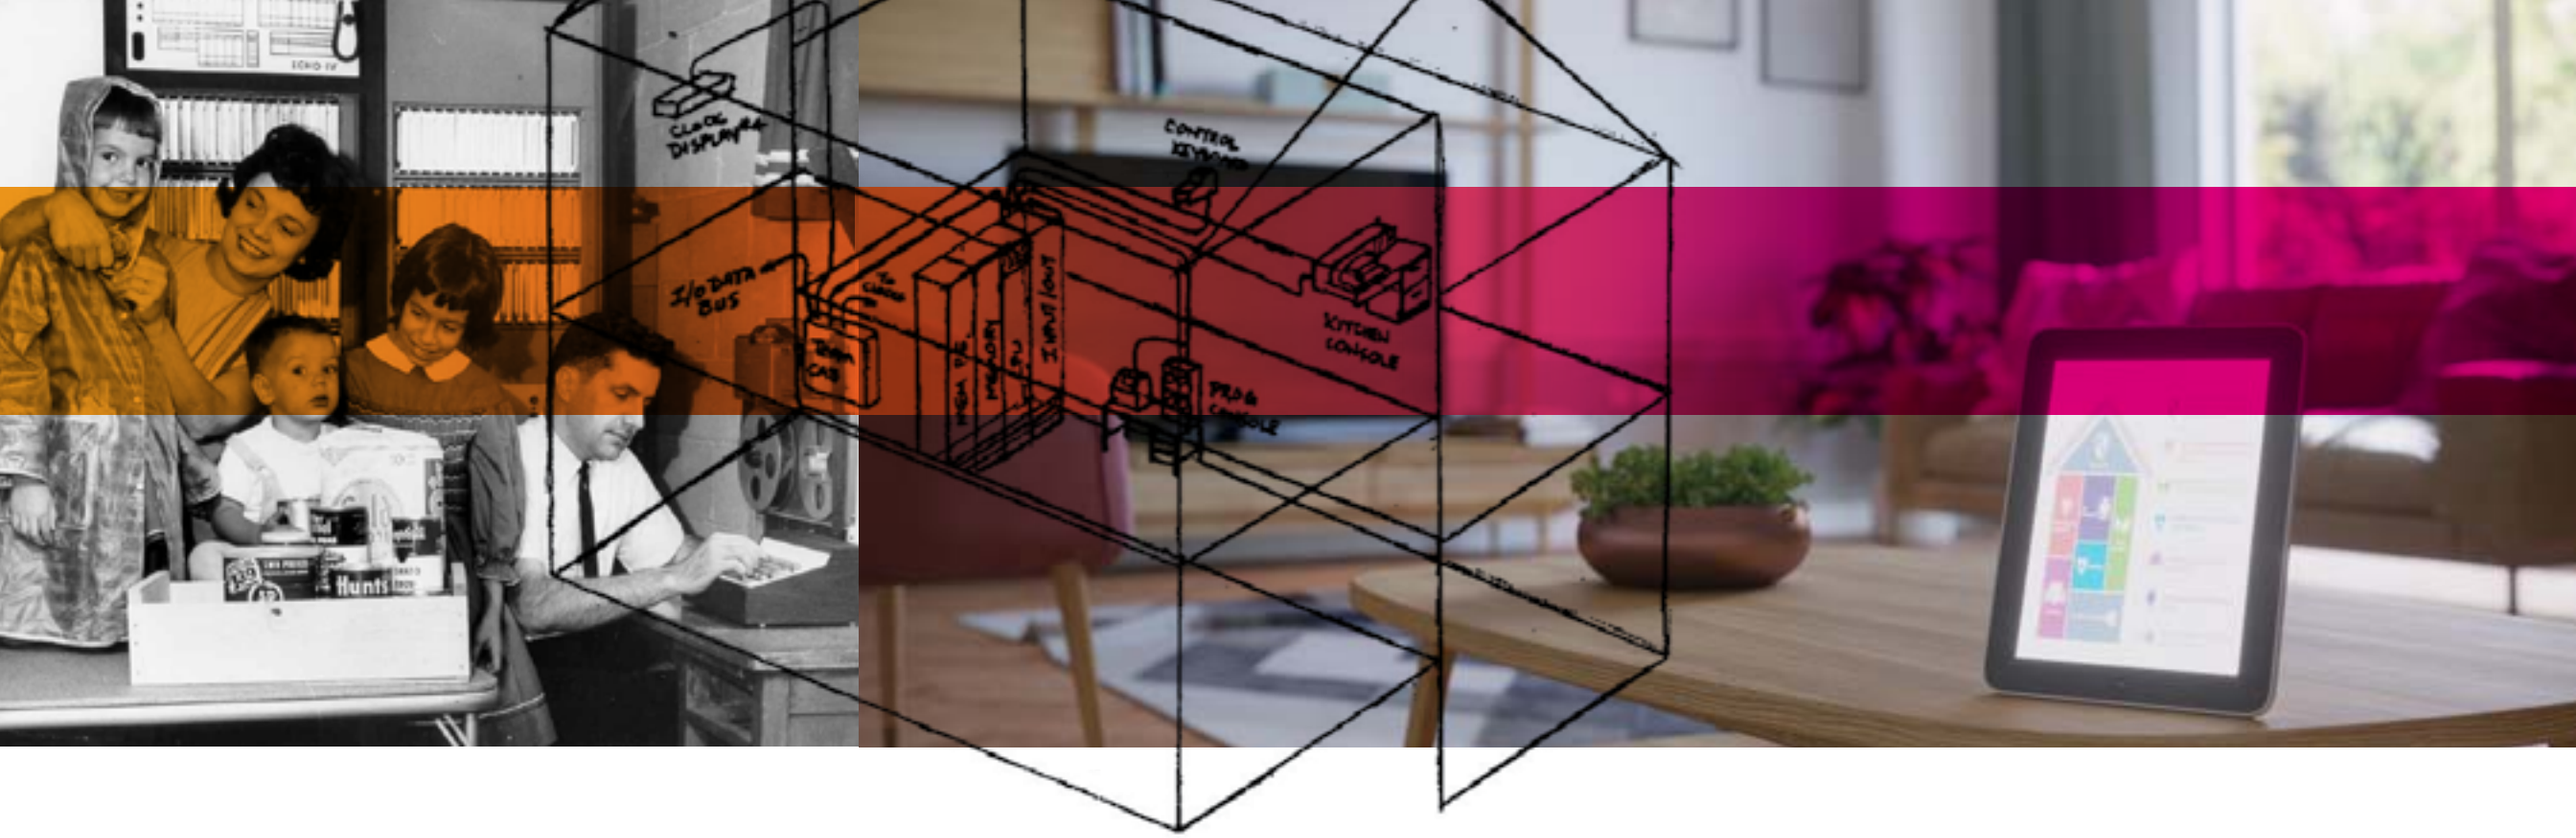

Smartest House on the Block

Smart homes have gone from trend to reality for many homeowners. Whether it's smart lighting, thermostats or security cameras, smart devices are increasingly present in our home life.

THE DRIVE TO AUTOMATE and better understand the day-to-day functions of our home has a history that dates back to before the internet. In 1966, computerized home management took its first huge leap forward when Jim Sutherland, an engineer at Westinghouse Electric, created the ECHO IV, or Electronic Computing Home Operator, with surplus parts from his job. For a decade, Jim's ECHO IV computed shopping lists, controlled the air conditioning, maintained a calendar and clock and even locked down the TV and antenna so the Sutherland kids couldn't watch on school nights without answering a series of questions.

Since then, IoT devices, sensors and widespread connectivity have made it possible to do all that Jim did and more. Some people fear that IoT could turn into HAL 9000 (the spaceship computer in "2001: A Space Odyssey" that malfunctions in a creepily murderous fashion), but that didn't stop Splunker and data geek Stephen Luedtke. He wanted to see if there were a better way to monitor his utilities, save on monthly bills and prevent problems before they happened by applying Splunk's end-to-end visibility to his entire house.

What was he looking to understand? First, Stephen wanted to monitor electricity usage by individual circuit breaker to understand energy consumption across devices in his home. This would help determine whether lights were left on or an appliance was running less efficiently than it should, and also discover unknown energy consumers. He wanted to monitor water usage for spikes that could immediately identify costly leaks. Plus, someone in the house was taking insanely long showers. Finally, he wanted to monitor air conditioning and heating usage, and compare it to past bills to see if usage and cost could be lowered.

The experiment was a success. Stephen can now observe everything in detail from a single Splunk dashboard. Setup took some time, but the Splunk Add-On Builder sped up the process of connecting to REST APIs for his utilities, and now he has a NOC in his home.

Tracking Your Home's Life

Tap into the ins and outlets of your home

Level	Time to create	Makes
●●○	Home operations center: 4-5 hours Energy consumption: 1-2 hours Heating/AC: 2-3 hours Sensors and API calls: ~1 week	1 smart home

Ingredients

- A Splunk instance
- For energy monitoring, Curb energy sensor or Sense energy monitor
- For water monitoring, Flume Water Sensor and app
- For AC/heat monitoring, try an energy meter and a Nest thermostat

Cooking Instructions

1. **Energy**
   - **a.** Install Curb Energy unit or Sense following the product instruction manuals. This can take a few hours.
   - **b.** Connect to the Curb API (you may need to contact CURB support to obtain an access token) using the Splunk Add-On Builder (or build your own modular input) and the following REST URL:

     ```
     https://app.energycurb.com/api/v3/latest/<yourdeviceid>
     ```
2. **Water**
   - **a.** Install the Flume Water Sensor using the product instruction manuals.
   - **b.** Connect to Flume API using the Splunk Add-On built by Splunk's Antoine Toulme.
3. **Nest**
   - **a.** Try Splunk Add-On for Nest found on Splunkbase for more granular insights. Curb readings from the AC/heater tell you if it's running or not.

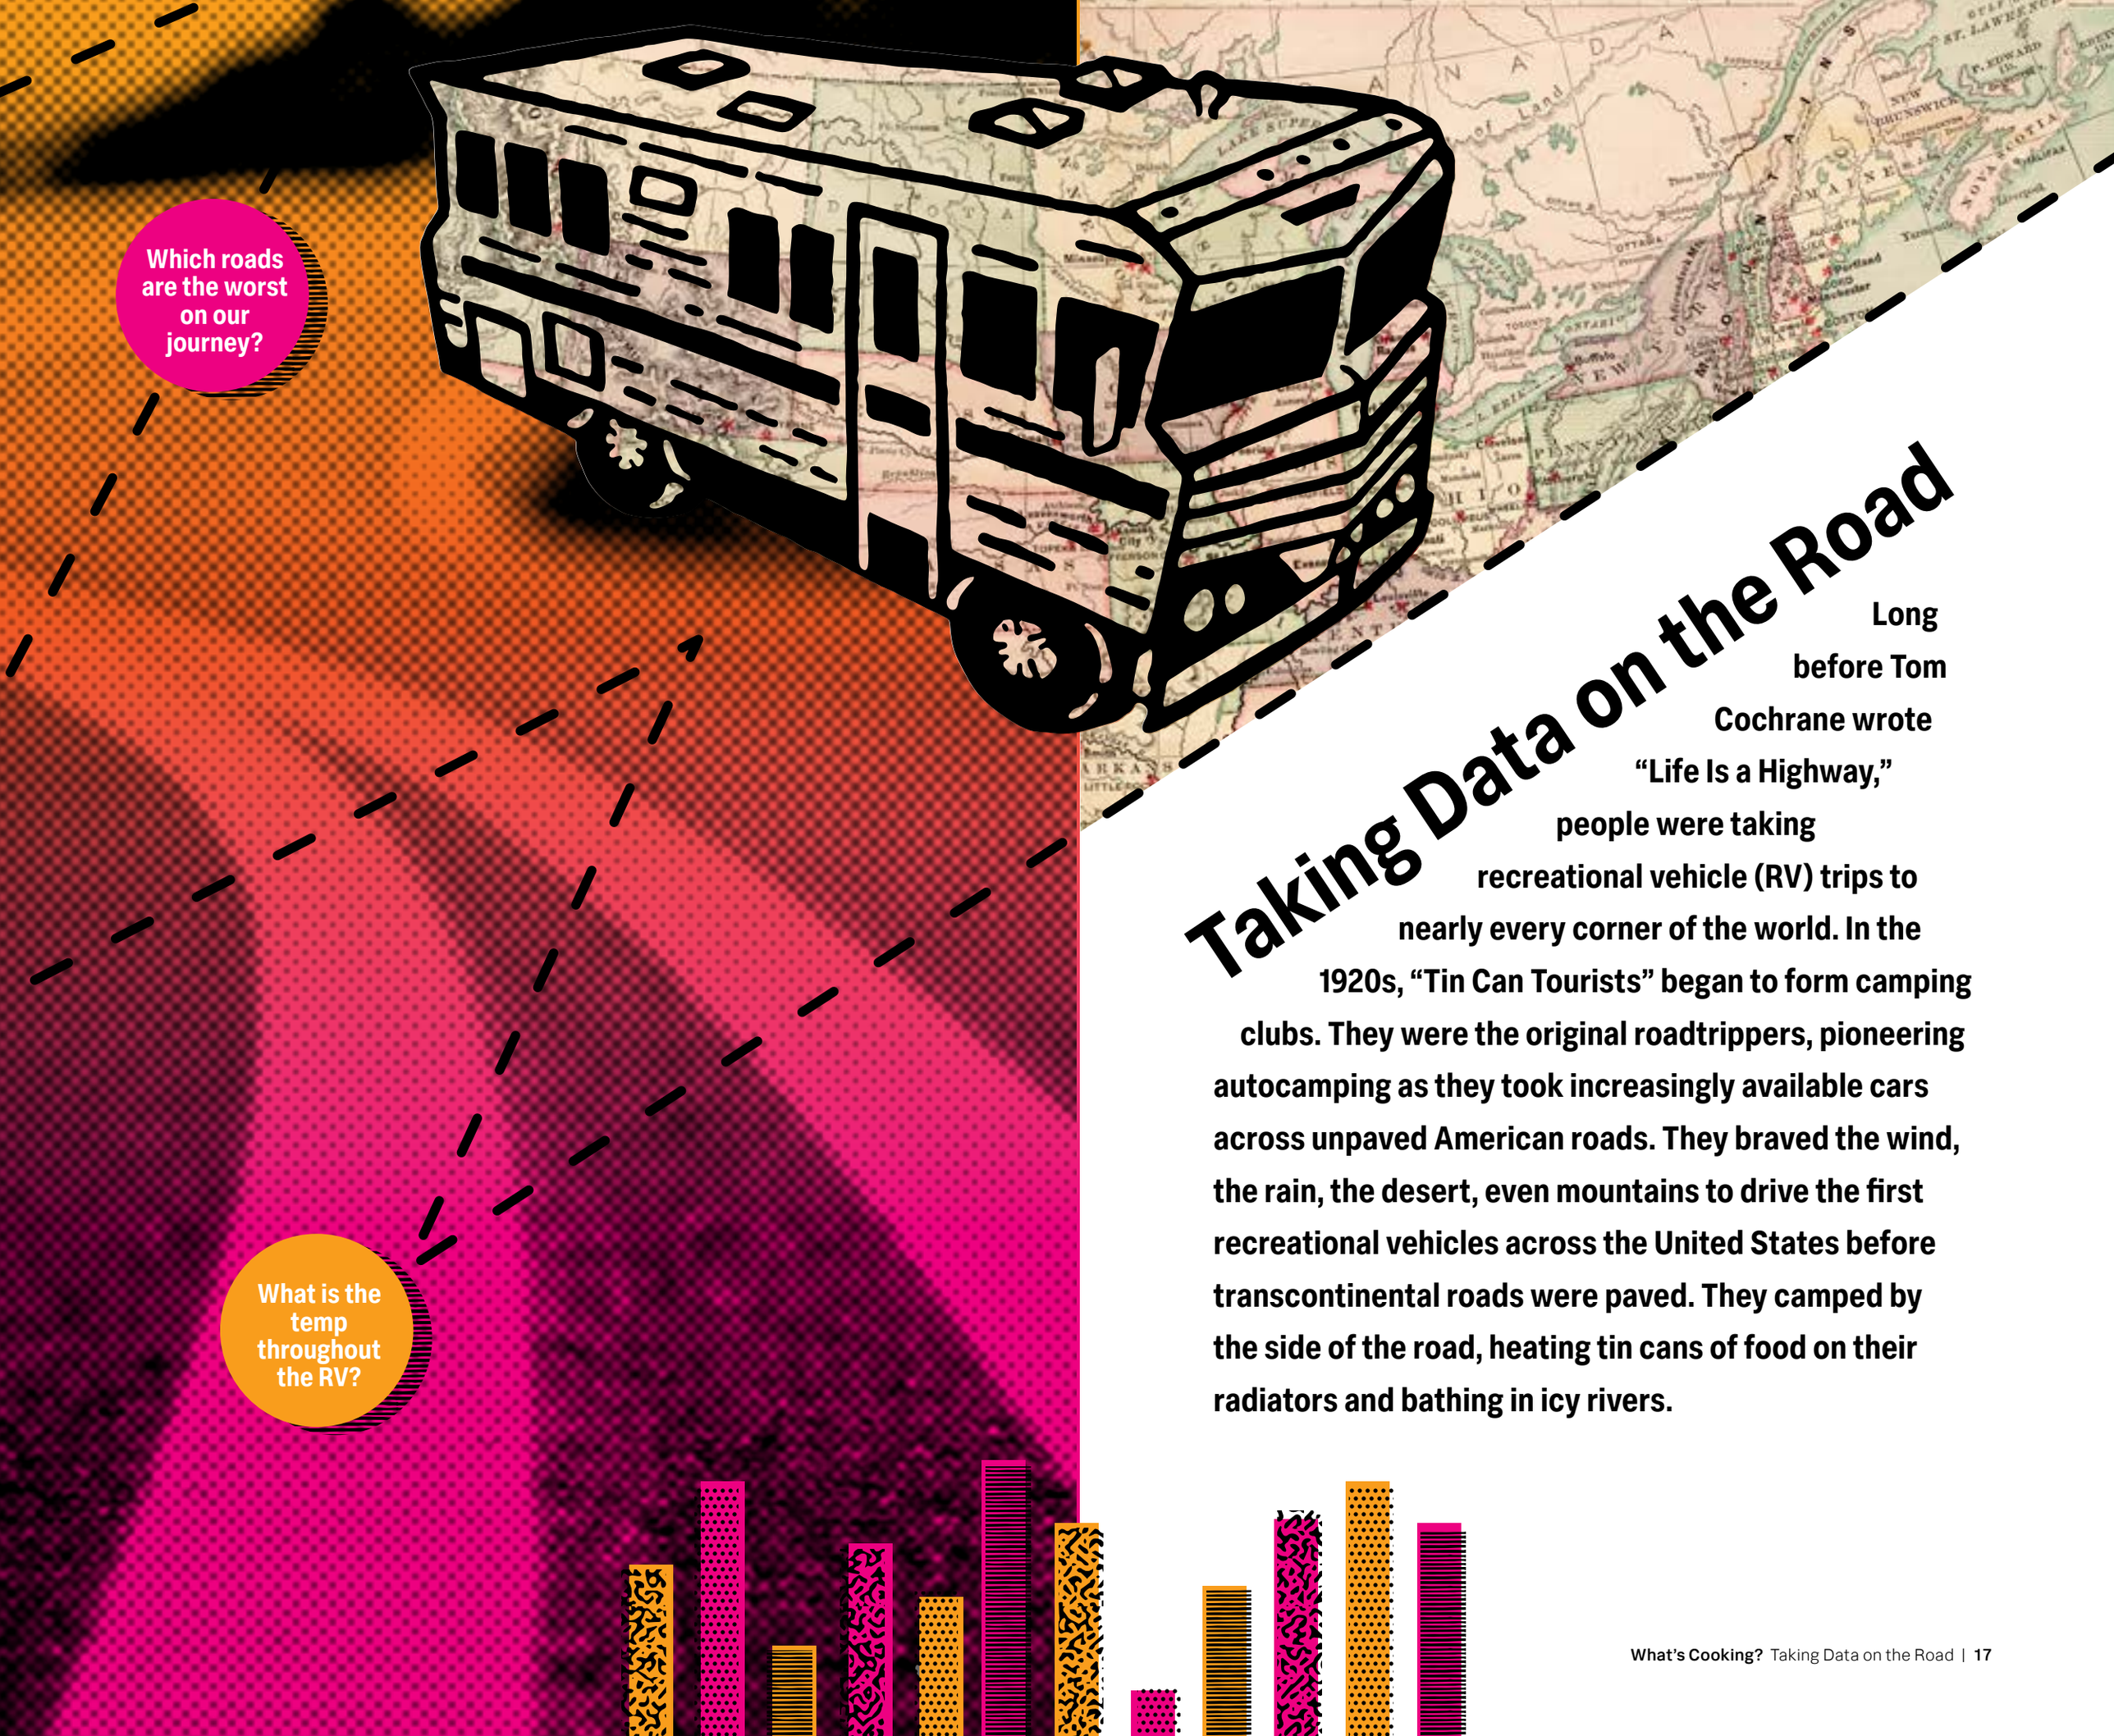

Taking Data on the Road

Long before Tom Cochrane wrote "Life Is a Highway," people were taking recreational vehicle (RV) trips to nearly every corner of the world. In the 1920s, "Tin Can Tourists" began to form camping clubs. They were the original roadtrippers, pioneering autocamping as they took increasingly available cars across unpaved American roads. They braved the wind, the rain, the desert, even mountains to drive the first recreational vehicles across the United States before transcontinental roads were paved. They camped by the side of the road, heating tin cans of food on their radiators and bathing in icy rivers.

NOW, WE HAVE LAND BOATS navigating cross-country highways. Generations after the first recreational vehicles began traveling coast-to-coast, we're still obsessed with life on the road. We've evolved to Airstreams, Winnebagos, vans and school buses converted into campers — the spirit of adventure endures.

Country Roads, DataNodes

With that spirit, Brett Roberts, Kyle Prins and Cory Minton of the Big Data Beard podcast decided to grab an RV, set up some IoT sensors, build an edge-to-cloud computing environment and pack up the recording equipment for an epic road trip across the country from Boston to Las Vegas.

Over the span of their two-week journey, they traveled across 15 states, driving more than 3,889 miles and stopping in 13 different cities. There were many data challenges along the way.

The road trippers wanted to pull data insights around location, RV comfort, music and even bathroom usage. They still haven't figured out who wasn't washing their hands at 3:30 in the morning. Who would have guessed that the threat of being called out by Splunk for not washing your hands was still not enough to get 100% compliance? Along with their location, their listening habits were also being broadcast, and watchers discovered Cory's penchant for bluegrass.

There were a few bumps in the road. Networking wasn't consistent — 4G was especially patchy in Utah. It turns out that the Jersey Turnpike is indeed the worst road in America. But through it all, they continued to come up with new ways to use data on their journey.

Big Beard Road Trip

How to broadcast your favorite tunes, monitor your bathroom and track your location when a podcast isn't enough

Level	Time to create	Makes
●●●	Around 3-4 hours	Any number of insight-packed reports

Ingredients

- A Splunk instance
- Basic Python knowledge
- Dell Gateway 3001
- 3 Raspberry Pi 3 B+
- Splunk Add-On Builder
- Spotify Add-On (optional)
- What3words API (optional)
- Opencage Geocoder API (optional)
- Dark Sky API

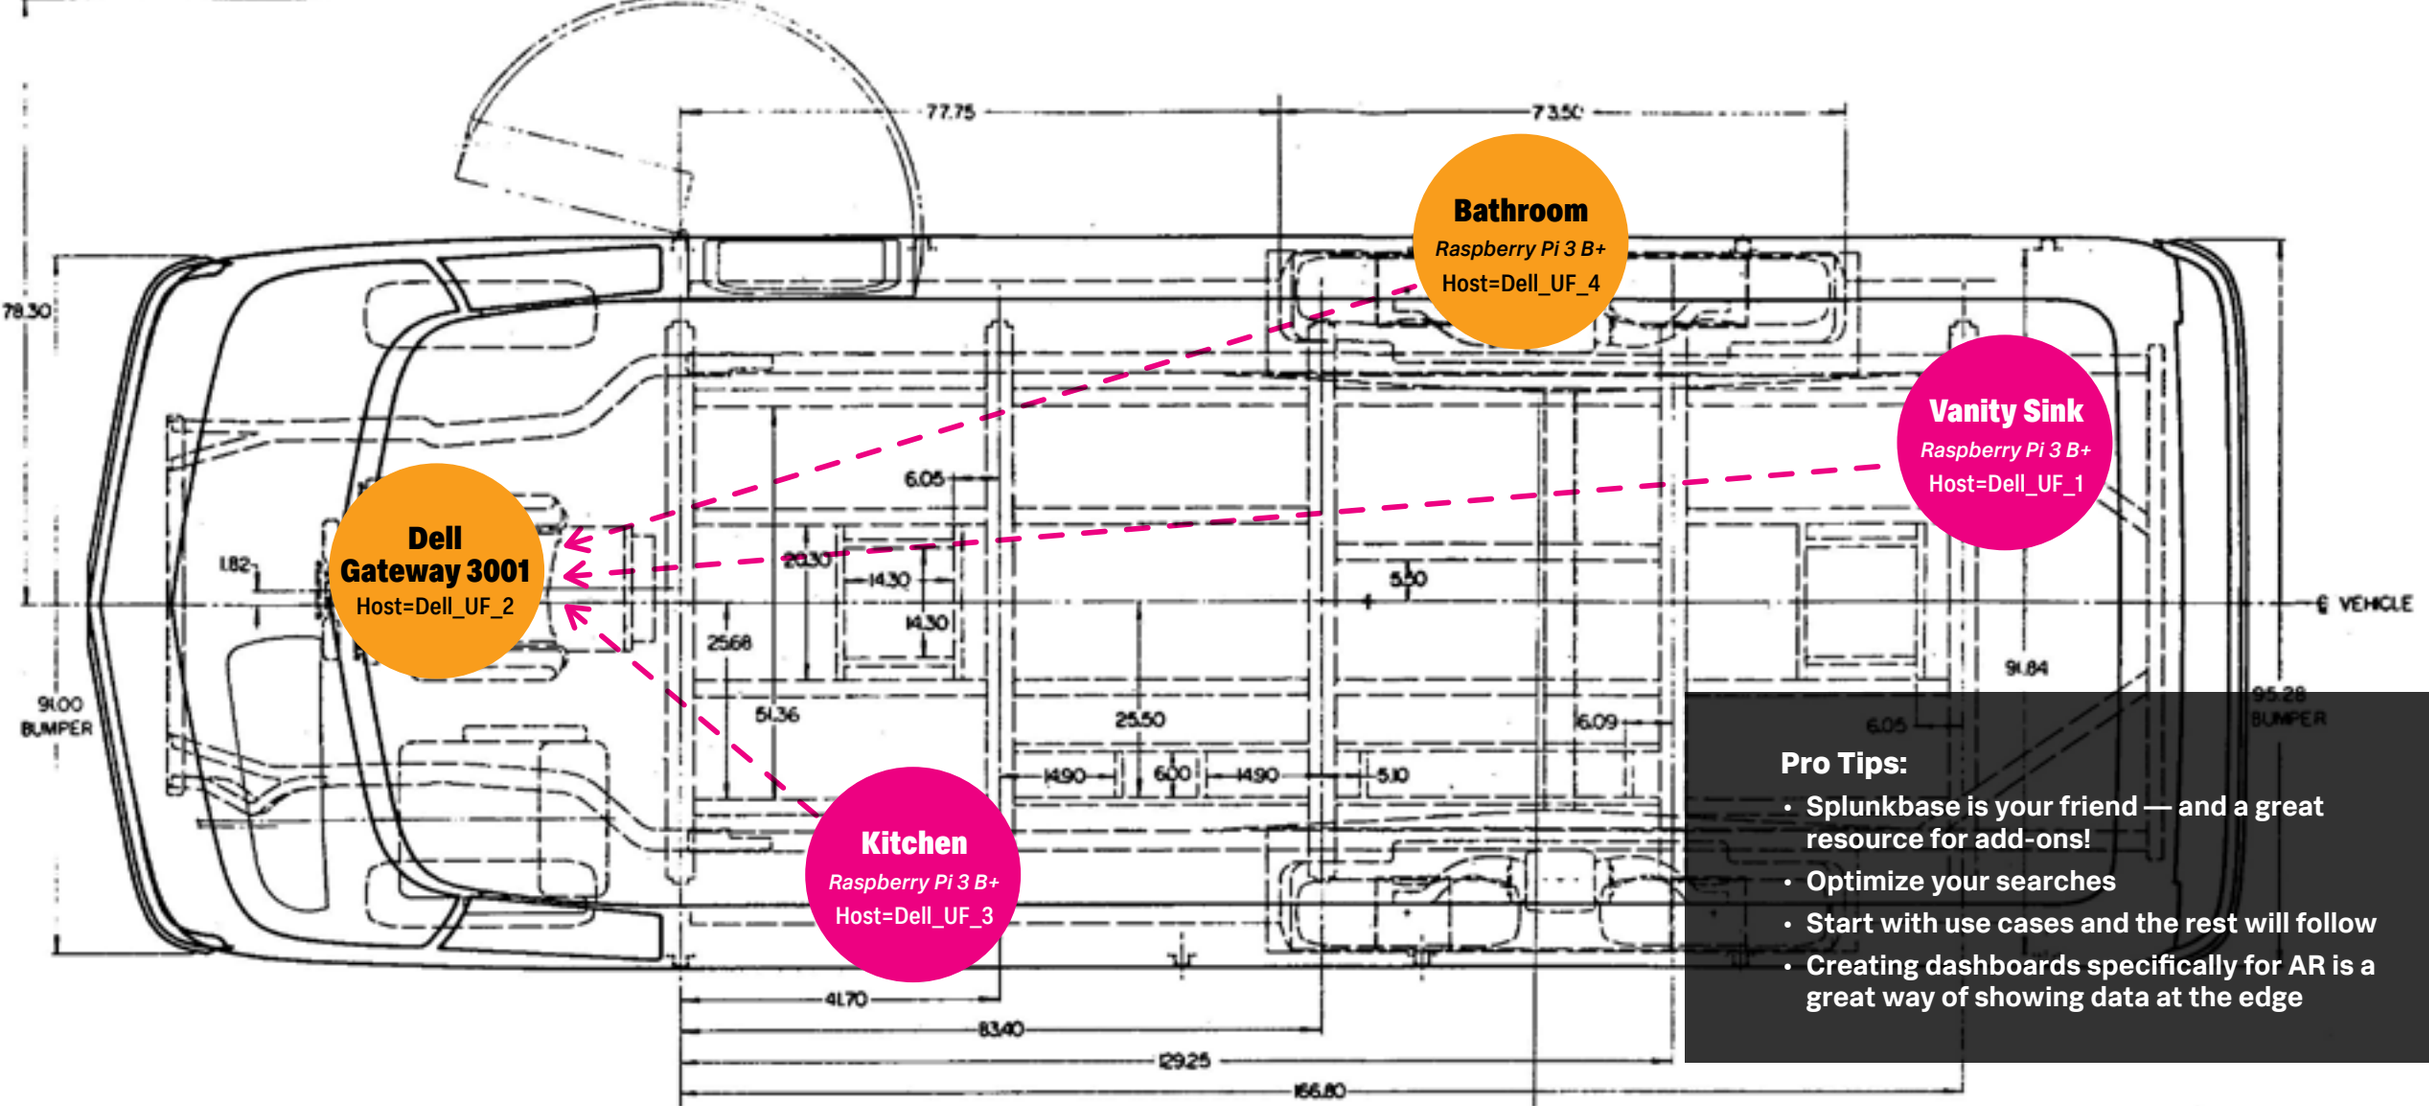

Pro Tips:

- Splunkbase is your friend — and a great resource for add-ons!
- Optimize your searches
- Start with use cases and the rest will follow
- Creating dashboards specifically for AR is a great way of showing data at the edge

Cooking Instructions

1. Install edge devices. Check out the diagram above for how the Big Beard team did it.
2. Choose how to measure your data: temperature, vibration, buttons, sound, GPS.
3. Forward data from the Raspberry Pis into a heavy forwarder.
4. Add all the relevant data into Splunk.
5. Keep in mind, when you're moving at speeds of 60 mph, the 4G may go in and out. That's why it's important to forward your data to the cloud. Otherwise, your measurements won't be precise.
6. Everything can be written in a log file. However, the Big Beard team recommends using Rest API calls.

The Rise of Data and Dough

From slender baguettes to chewy focaccia, buttery brioche to golden challah, the most widely consumed food in the world has been around for as many as 30,000 years. The first iteration was prehistoric flatbread, made with starchy roots ground into flour and mixed with water, then cooked on heated rocks.

LEAVENED BREADS — the lighter, fluffier counterparts to unleavened varieties like pita and naan — made their initial appearance later on. Ancient Egyptian bakers pioneered the commercial production of leavened bread in 300 B.C.E., and bread has continued to be a universal dietary staple ever since, especially once mechanized slicing arrived and companies like Wonder Bread ensured their pre-sliced loaves were top of mind in every household during the 1930s.

Breadmaking's Unexpected Rise

We've established that bread eating has been popular for a very long time, but breadmaking in particular has been on the rise since shelter-in-place began in early 2020. The increase in time spent at home, coupled with what many bakers describe as the stress-relieving nature of the activity, has germinated a rise in home bakers. Two such bakers are Skye Lowry and his partner Byeong Kim, who not only learned to make sourdough, but also figured out how to Splunk their sourdough starter.

The upkeep of sourdough starter is a somewhat tricky, painstaking business, an ordeal that involves discarding most of the starter every few days and feeding the remainder with more flour and water. It needs constant care, like a baby, a plant or a Tamagotchi.

The typical baker uses rubber bands to track the height of the starter and to know when it's ready for additional feeding. Skye and Byeong opted for a different method: After hearing about others' experiences using computer vision for monitoring the dough fermentation process, they decided to Splunk their sourdough starter.

The two used a Raspberry Pi camera that took photos of the starter. Byeong wrote a Python script to analyze the photos, ascertain the exact height of the starter and output the height of the starter with timestamps. All this logged data was then picked up by Splunk's universal forwarder for Raspberry Pi. Alerts were generated within Splunk every time the height of the dough indicated the starter needed feeding. Visualizations in Splunk came in handy for tracking the fluctuations in height as the starter grew.

The entire endeavor — from the technical setup to the breadmaking process itself — took some trial-and-error experimentation, but Skye and Byeong achieved delicious results.

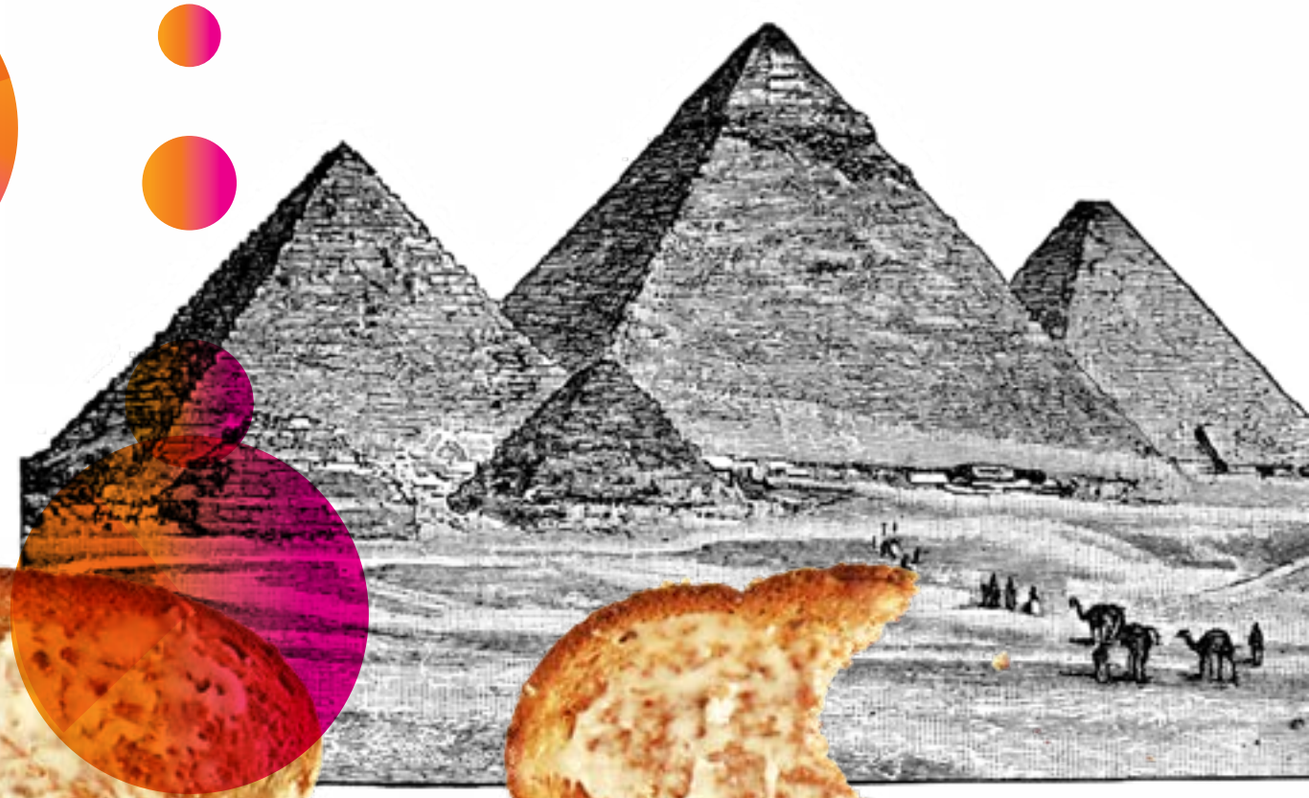

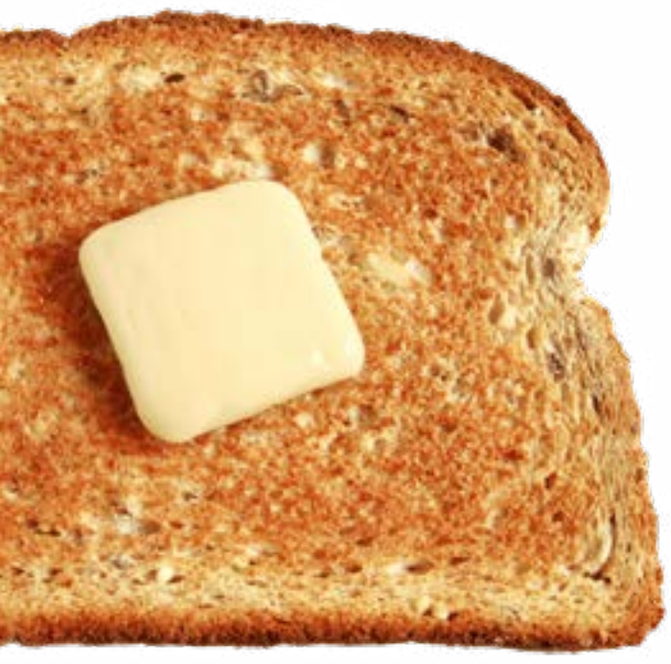

The Best Thing Since Sliced Bread

No loafing when it comes to data

Level	Time to create	Makes
●●○	2 days	A well-baked report ... and a sourdough loaf

Ingredients

- A Splunk instance
- A Raspberry Pi
- Raspberry Pi camera
- An IFTTT (free) account
- Ability to expose a webhook endpoint to the internet
- Splunk Universal Forwarder
- Philips Hue lights (optional)

Cooking Instructions

1. Make your sourdough starter. Come back in a couple of weeks.
2. Set up a Raspberry Pi camera. Use the webhook connector “Make Event” API and IFTTT, in conjunction with Philips Hue lights, to turn on the light only when photos are taken every 10-30 minutes. Use a shell script for the timing of the photos.
3. Write a Python script to process the images and analyze them for logging the sourdough starter height. For an example script, refer to Skye’s GitHub repository.
4. Use Splunk’s universal forwarder for Raspberry Pi to transfer all the data into Splunk, as well as the actual photos.
5. Set up height alerts within Splunk so that you can be notified when the starter needs feeding. The photos taken can be uploaded to an AWS S3 bucket and included within the actual dashboards.

When Life Gives You Pickles

The very first pickles were created 4,000 years ago in ancient Mesopotamia and have been beloved by many over the years. Aristotle praised their purported healing effects. Julius Caesar fed them to his troops. Cleopatra ate them as part of her beauty regimen.

The Report That Gets You in a Pickle

How to monitor activity around your house, your fridge and even your pickle jar

Level	Time to create	Makes
● ○ ○	2 hours	1 juicy report and some annoyed kids

Ingredients:

- A Splunk instance
- Arlo wireless cameras
- Splunk Stream (optional)
- Ability to expose a webhook endpoint to the Internet
- An IFTTT (free) account

Cooking Instructions:

1. Consider ingesting three potential data sources: proximity card reader data, local DHCP server logs and camera activity data. Combine any/all of them and you can start to get a very good picture of who is in your space, or who is trying to access your space.
2. Install a camera system that logs its activity when it senses motion, the more info the better. Arlo cameras can be placed almost anywhere in your home and can also identify moving objects like animals, people or vehicles.
3. Ensure that the activity data is in a time-stamped format that can be consumed in Splunk — a flat file written to a server, an email notification that Splunk Phantom could parse, data retrievable via API call, or, in Arlo's case (and as James did), an IFTTT applet that triggers an HTTPS post via webhook to Splunk's HTTP Event Collector.

TODAY, THE DEPARTMENT OF AGRICULTURE estimates the average American consumes 8.5 pounds of pickles a year, which amounts to a staggering 2.5 million pounds consumed as a nation. However, it's Canada that takes first place as top pickle purchaser, racking up $72 million in imports in 2019.

Observability Worth Preserving

When he started working from home, Splunk's James Brodsky realized his household consumes groceries with astonishing speed, particularly his favorite pickles.

So James set out to monitor his household's grocery consumption rate, particularly the frequency of the pickle culprit's visits, with Splunk. He had been running Splunk in his house for seven years, gathering data of all sorts from his firewall, home theater system, thermostat and more. The next step was to tap into his home's wireless cameras, which could inhabit pretty much any corner of the house as long as they were close to the system's base station.

One of these cameras was placed atop a pickle jar in the Brodsky fridge. James's instance of Splunk recorded how often the fridge was opened and by whom, since the camera detected motion and video recorded each time someone raided the fridge. And his suspicions were supported: His 11-year-old son was the 1 a.m. pickle thief.

Of course, it didn't have to stop with the pickles. James had cameras both inside and outside his house and started using Splunk to monitor activity on his doorstep and in his driveway. It's an experience akin to omniscience: James would observe a spike of activity when his kids chalked messages on the driveway such as "We Support Our Healthcare Workers" and "Save Us, Daddy is Creepily Monitoring Our Food Intake!" He even started using Splunk for his DHCP server, so when his kids finally do get phones, he will know exactly how many minutes past curfew his kids return home (based on their phone acquiring an IP address).

Pro Tip:
Applying these ideas to a commercial or corporate environment is easy. Splunk has customers that use camera data, DHCP data, proximity card reader data, even visitor registration log data to do things from understanding human traffic patterns in retail stores to managing access in corporate environments.

To Splunka Songbird

According to the fossil record, birds have been around for 160 million years, since dinosaurs roamed the earth during the Jurassic period. While all birds are vertebrates and have anatomies designed for flight, the songbird has especially developed vocal organs and belongs to the particular suborder Passeri, of the order Passeriformes.

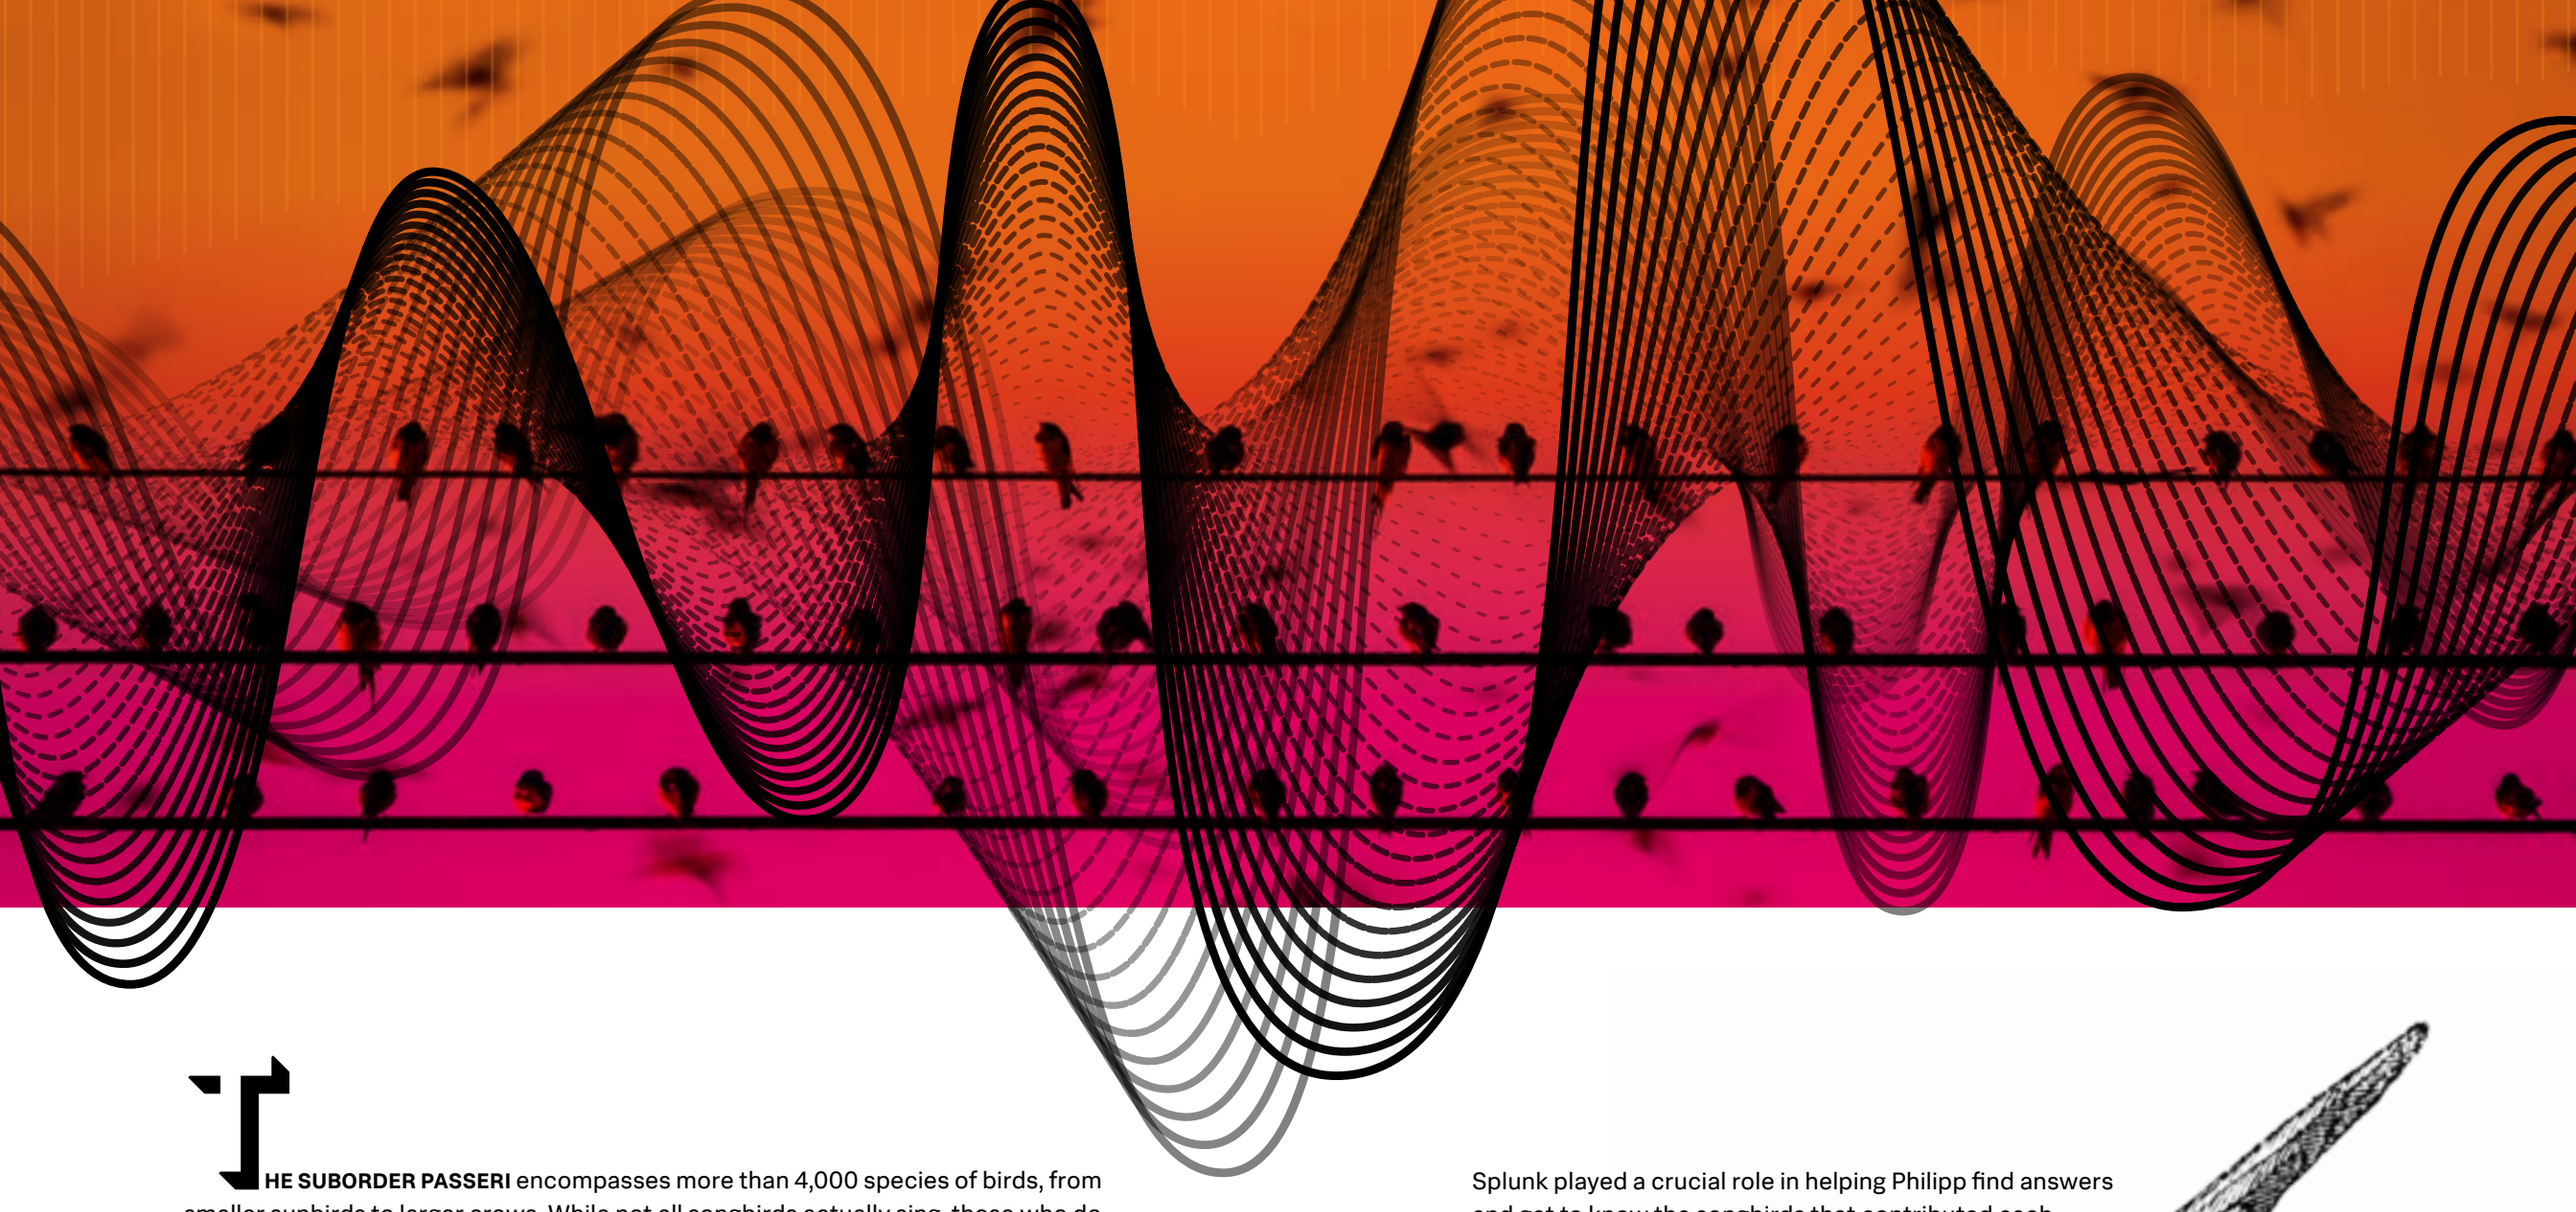

THE SUBORDER PASSERI encompasses more than 4,000 species of birds, from smaller sunbirds to larger crows. While not all songbirds actually sing, those who do primarily sing for the sake of social communication. During courtship rituals, males spout melodies to serenade females, and at times pairs even perform duets.

From the hooting of owls to the whistling of mockingbirds, the cheeping of sparrows to the warbling of finches, bird songs have been much adored by poets through the ages — and are still appreciated today by people who care to pause to listen to nature's choruses.

AI and Piping Birds

Splunker Philipp Drieger often wondered exactly which birds sang in his neighborhood. He would hear twitters and peeps, occasionally detect more complex songs and even wake up to vigorous chirping from outside his window.

While researching the kinds of birds that surrounded him, Philipp discovered BirdNET, an artificial neural network developed by the Cornell Lab of Ornithology and the Chemnitz University of Technology and trained to classify 984 species of birds.

But Philipp wanted to do more than just identify songbirds one at a time. He wanted to know which birds sang when, as well as how long they sang, and even if it's true that birds of a feather flock — and therefore sing — together.

Splunk played a crucial role in helping Philipp find answers and get to know the songbirds that contributed each trill and caw he heard around his neighborhood. All it took was running several searches, analyzing simple time charts and creating some Sankey diagram visualizations.

The wealth of possibilities when it comes to data collection and the power of AI tremendously excites Philipp, because what worked in his songbird project would work similarly for image recognition systems or any other AI system. Birds produce songs that can be analyzed, whereas other systems provide logs, metrics, events and more that can likewise be quickly analyzed to produce answers and actionable results.

The Words of Birds, Demystified

How to use AI to tell your prothonotary warblers from your red-breasted nuthatches

Level	Time to create	Makes
●●○	1-2 hours	Reports by the flock

Ingredients

- A Splunk instance
- BirdNET repository on GitHub
- Sankey Diagram (optional)
- 3D Graph Network Topology Visualization (optional)
- Container development knowledge

Cooking Instructions

1. Clone the repo BirdNET on GitHub.
2. Build a docker container.
3. Run BirdNET on a recorded sound file and save the results.
4. Define a data input in Splunk and point to the TSV files generated by BirdNET.
5. Write a few searches to populate a dashboard.
6. Define a drill down to show a thumbnail picture and hyperlinks for more information about a bird species.
7. To figure out which birds sing together, you can use a search like the one below:

```
index="birds" sourcetype="birds-tsv"
| rename "Begin Time _s" as time_start, "End Time _s" as time_end,"Common
  Name" as name, "Confidence" as confidence, "Species Code" as code
| eval _time=_time+time_start
| eval duration=time_end-time_start
| search confidence>0.9 NOT name="Human"
| sort 0 _time
| streamstats dc(name) as dc_names first(name) as src last(name) as dest
  time_window=10
| table _time src dest dc_names
| where dc_names==2
| stats count by src dest
```

8. Use a Sankey Diagram to visualize which birds sing together.
   a. Alternatively, for a more abstract visualization you can use the 3D Graph Network Topology Visualization built-in Graph Analysis Framework.
9. Increase the accuracy of your results by adding geoinformation. This lets BirdNET apply additional “knowledge” about bird species that are known to be domestic in certain regions.

Should I Be Breathing This?

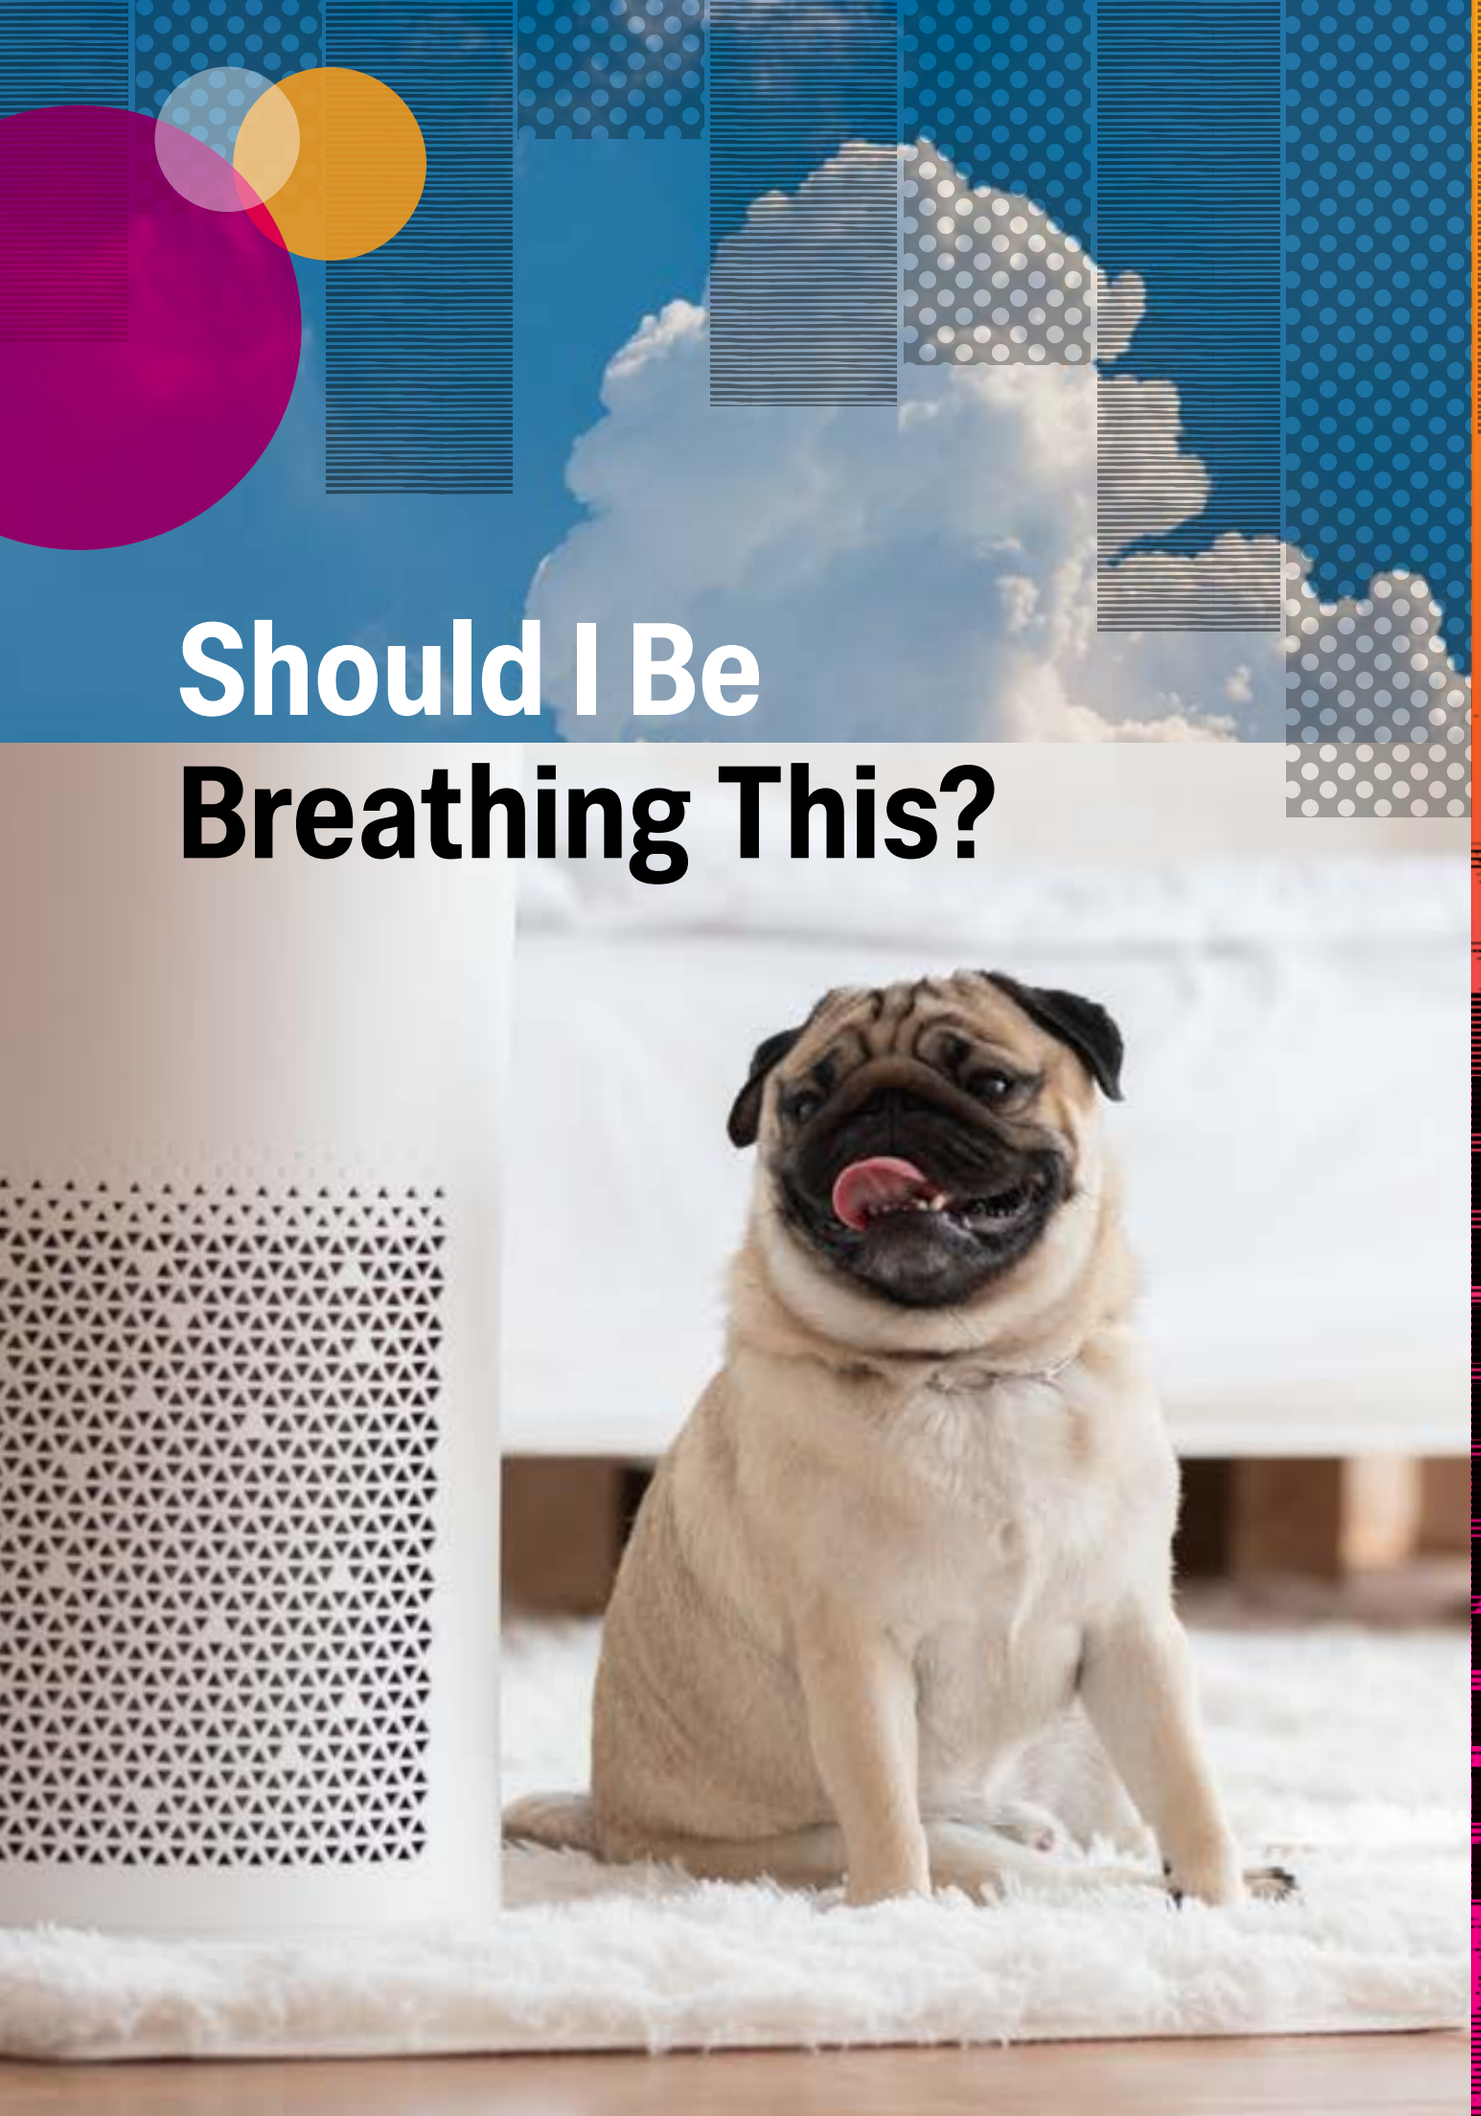

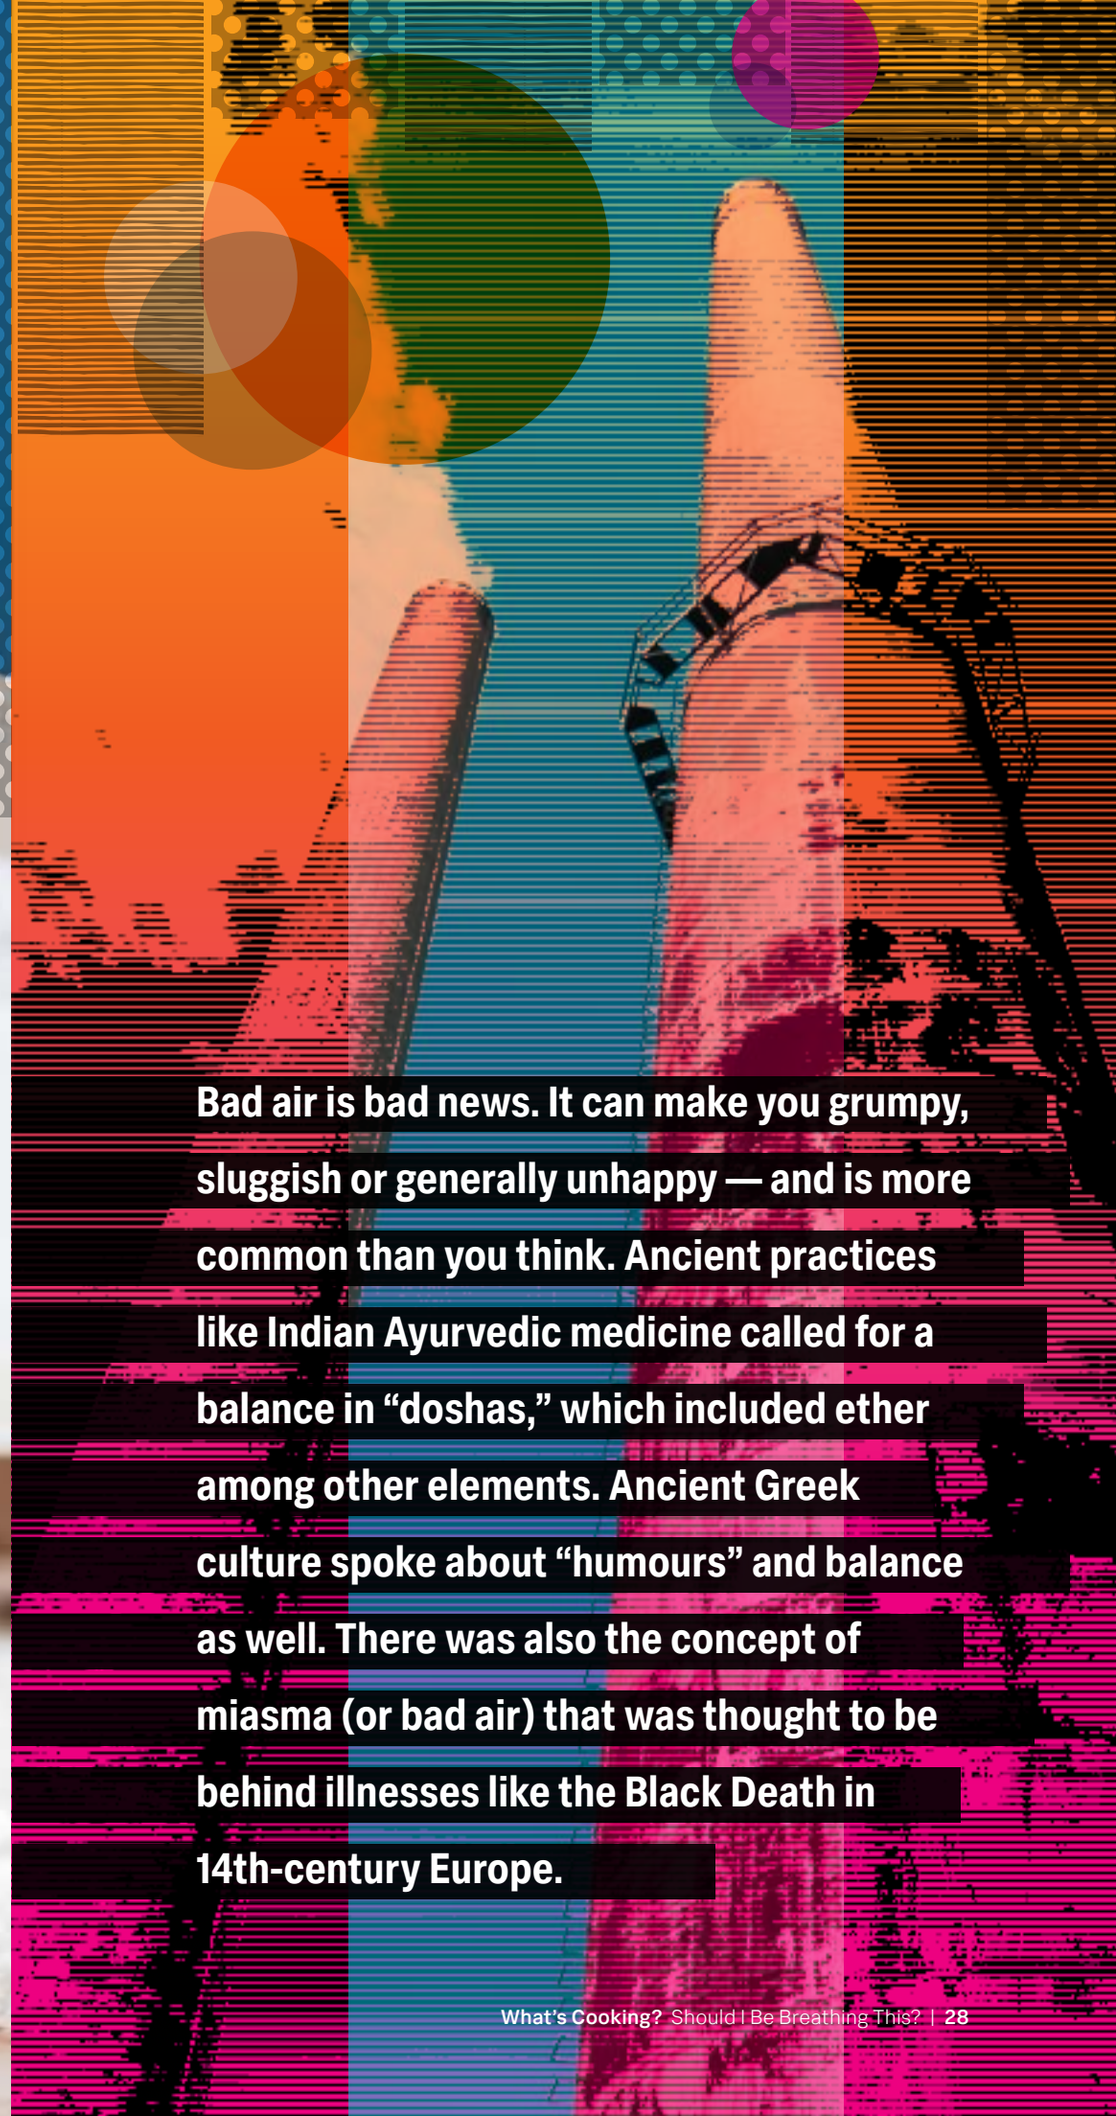

Bad air is bad news. It can make you grumpy, sluggish or generally unhappy — and is more common than you think. Ancient practices like Indian Ayurvedic medicine called for a balance in "doshas," which included ether among other elements. Ancient Greek culture spoke about "humours" and balance as well. There was also the concept of miasma (or bad air) that was thought to be behind illnesses like the Black Death in 14th-century Europe.

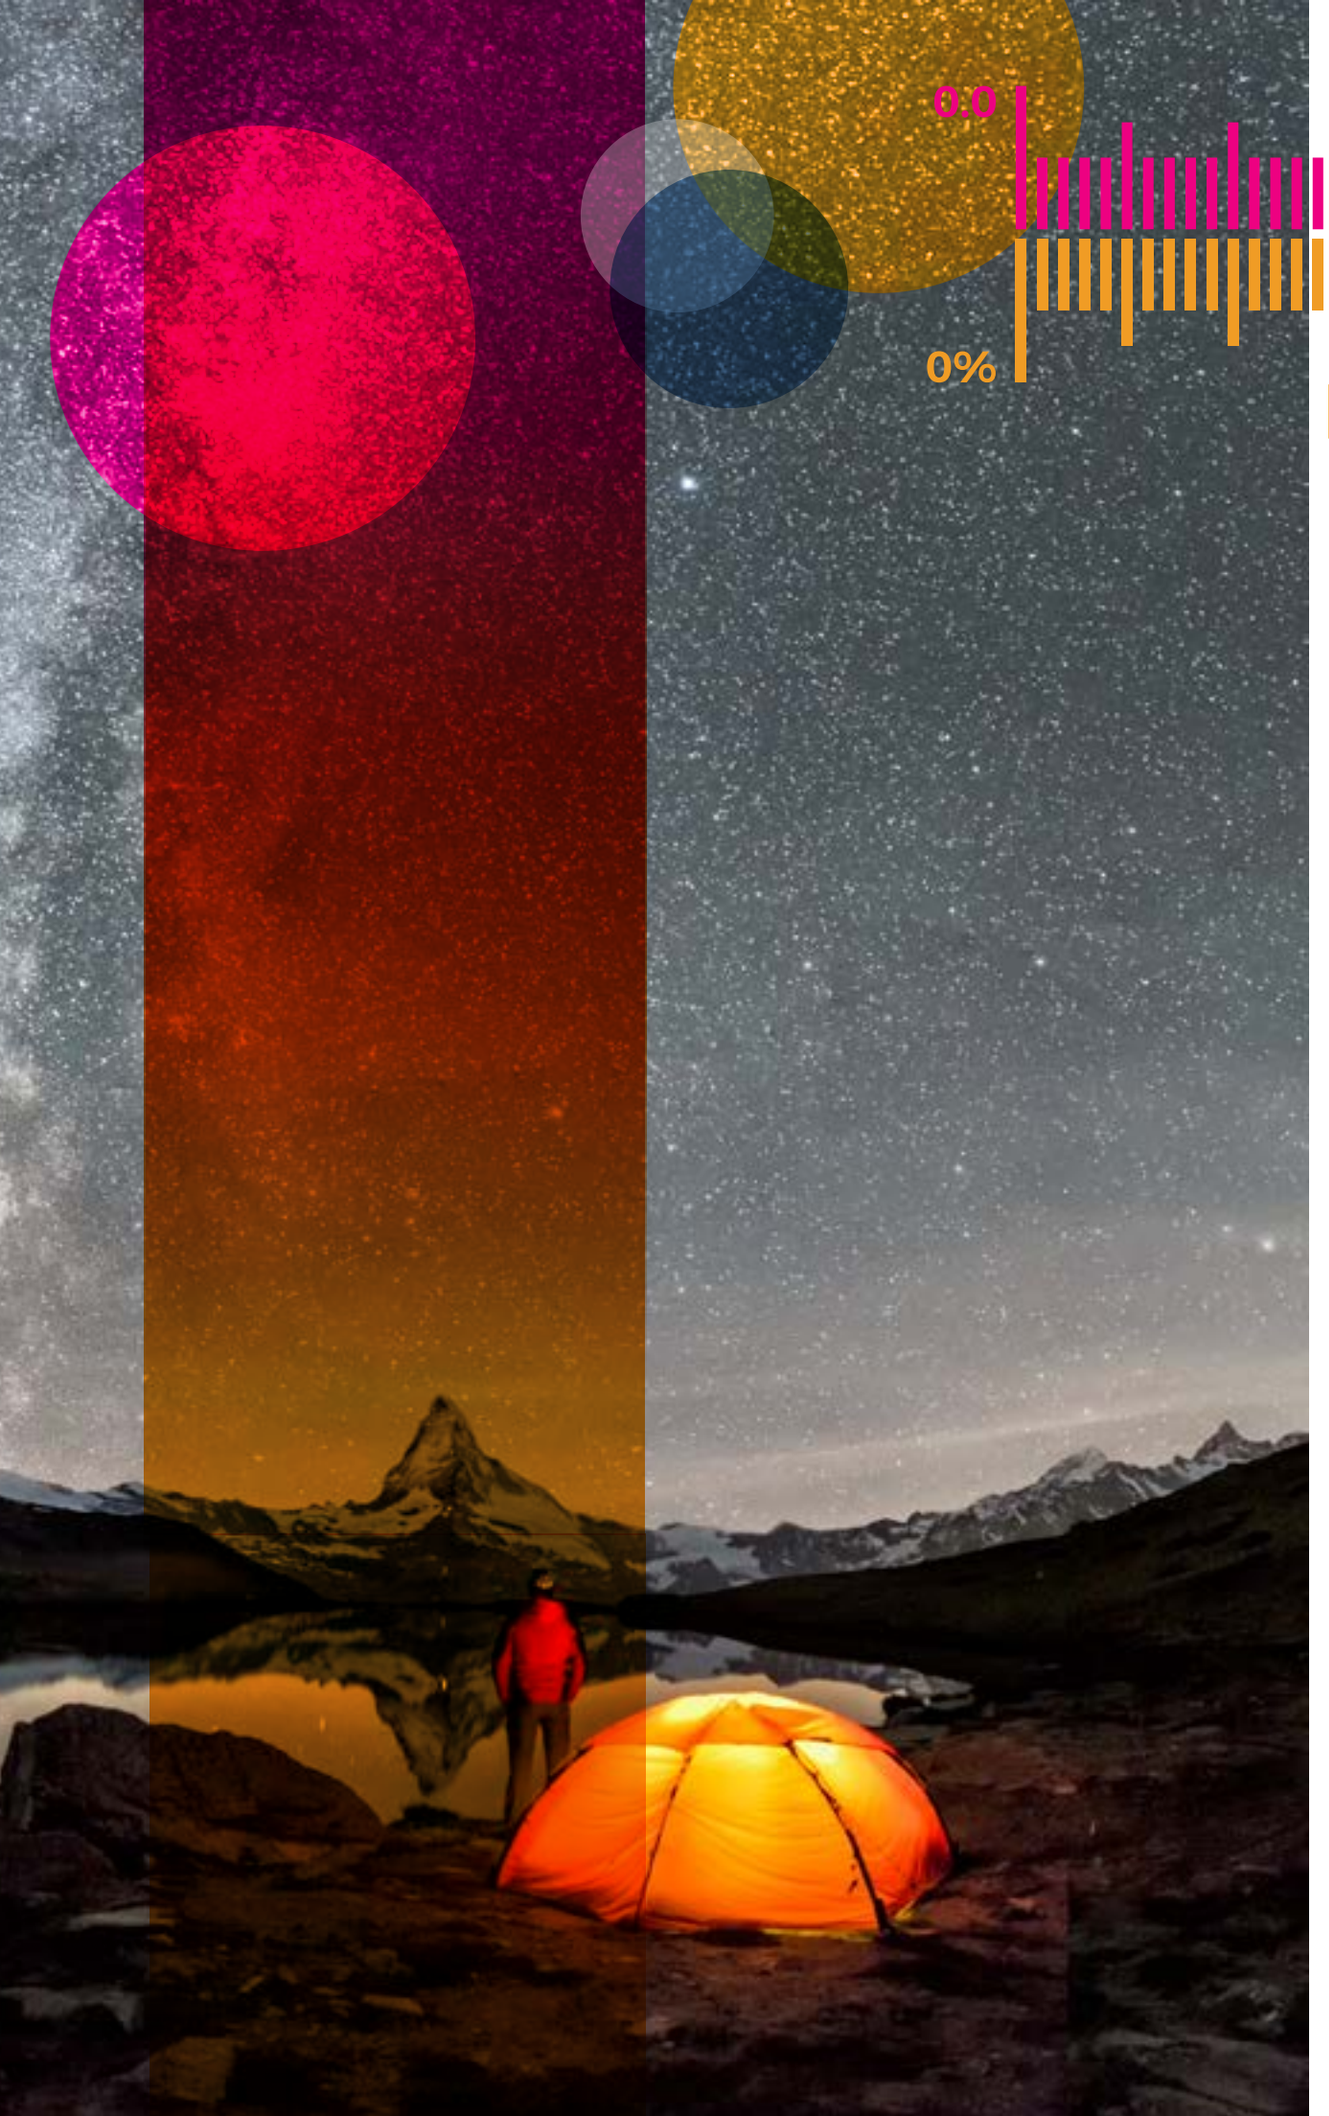

WE TAKE A MORE data-based approach to bad air today and use CO2 to clue in on what we're inhaling.

How are CO2 levels and getting sick connected? Aerosols — and not the kind you spray out of a can. Levels of atmospheric particulates in the air are a good indicator of the amount of air exhaled by a person — think moisture in the air or bad breath. Knowing CO2 concentration can reveal how much human-originated aerosols are in the air because people are usually the only source of CO2 in a room.

Airing rooms out with high CO2 (and accompanying aerosols) concentrations creates healthier spaces. CO2 data in offices and conference rooms would reduce the likelihood of fatigue and droplet-transmitted infectious diseases. By monitoring CO2 levels with Splunk, people can set healthy benchmarks for CO2 data (less than 1000 parts per million is a good threshold) and alert individuals when rooms have poor air quality.

Cleaner Data for Cleaner Air

Many CO2 measuring devices have a built-in traffic light, informing people when air quality is good, when you should take caution and when air quality is poor.

Splunker Tomas Baublys took it further by using Splunk with CO2 data to get even more insight. His inspiration might have been cleaner air, but it has also helped him keep weird body odors at bay with better ventilation. When done effectively, Splunk and CO2 trackers can work together to achieve the following applications:

1. Storing CO2 data over a longer period of time
2. Real-time alerting through e-mail and mobile devices
3. Predicting air quality
4. Monitoring and correlating CO2 data with external data such as weather data, local infection rates, number of people in a room and events, such as window or door openings

The more data ingested and dissected within Splunk, the more we can learn about the effects of air quality on health and productivity. And amid shelter-in-place, it could not be more important.

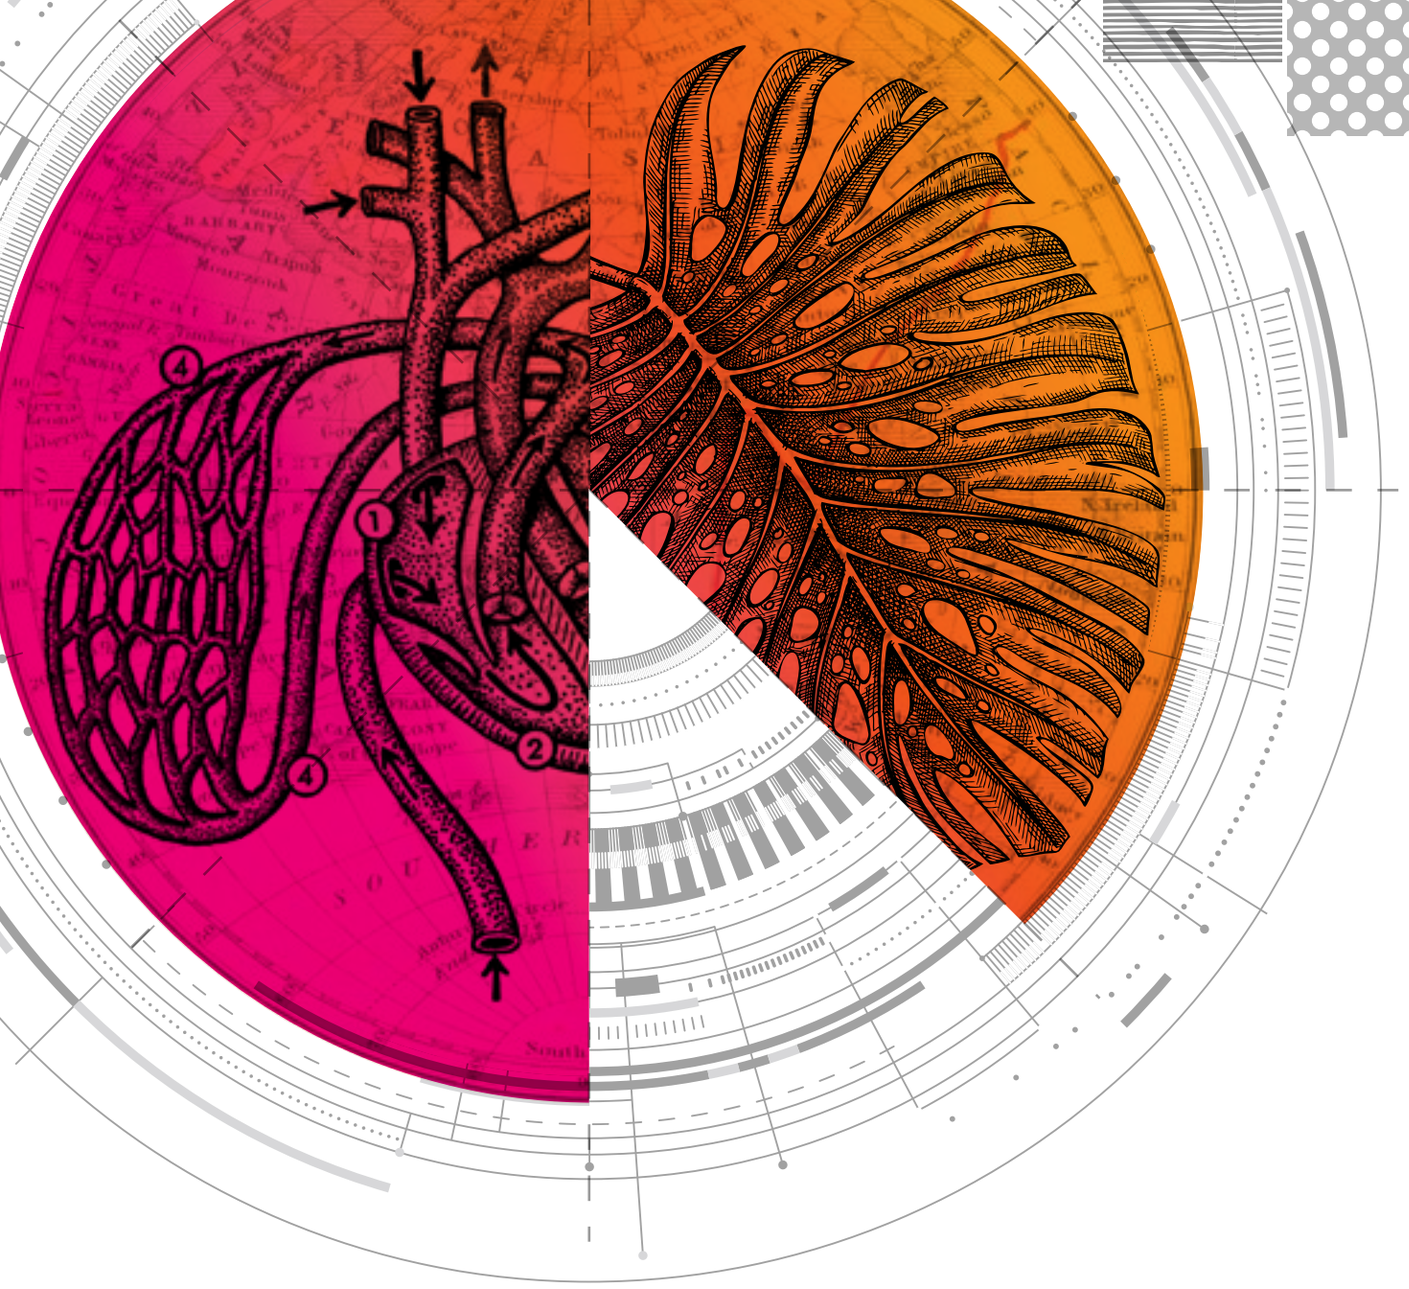

A Breath of Fresh Air

How data can improve air quality, focus and health

Level	Time to create	Makes
●●●	1 hour	1 clean, healthy meeting room

Ingredients

- A Splunk instance
- 1 Raspberry Pi
- USB CO2 Measuring Device(s) (at least 1 for each room measured, depending on size)
- Splunk Universal Forwarder
- Splunk Cloud Gateway Mobile App (optional)
- Linux knowledge
- Python knowledge

Cooking Instructions

Hardware

1. TFA Dostmann AirCO2ntrol Mini CO2 Monitor (or any other CO2 sensor that can be read via IC2 or USB). The TFA Dostmann tracker has the advantage that a display shows the data directly and a small warning LED shows the danger level in traffic light colors.
2. Set up your Raspberry Pi.

Software

1. Install Linux on the Raspberry Pi.
2. Refer to this Python script for querying data via USB.
   a. You will want to make this small adjustment to write the data to a logfile:

```
with open('/var/log/co2monitor.log', 'a') as out:
        out.write(eventtime + "," "TMP,%3.1f" % (tmp) + "," "M01" + '\n')
        out.write(eventtime + "," "CO2,%4i" %(co2) + "," "M01" + '\n')
```

3. Set up the Splunk Universal Forwarder.
   a. Alternatively, you could send the data directly via HEC but then you would have no buffer in case of a network failure.
4. Establish an accessible Splunk server.

Configuration

1. Install Splunk Universal Forwarder on Raspbian.
2. Script setup, customization and testing can be found here.
3. Connect the CO2 measuring device to Raspi via USB (with the original cable supplied).

   Start script:

```
sudo nohup /opt/jobs/monitor.py /dev/hidraw0 &
```

   Create Splunk input:

```
[monitor:///var/log/co2monitor.log]
disabled = false
sourcetype = metrics_csv
index = raspi
```

   On the server side:

   a. Create searches and dashboards.
   b. Use machine learning to create forecasts.
   c. Install Cloud Gateway and utilize mobile alerts, dashboards and AR.

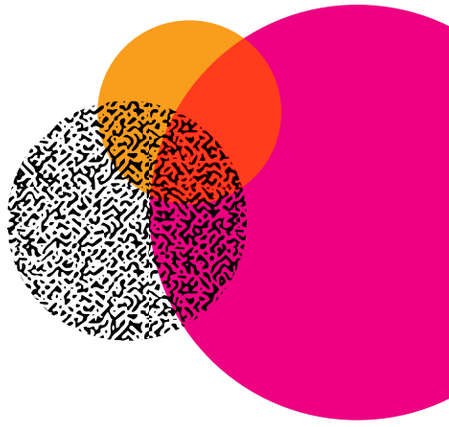

Let’s Bring Data-to-Everything™

Inventors rarely know how their products will be used. Did Henry Ford expect teenagers to make out in the back seats of cars? Were the creators of 4G just looking for an easier way to order delivery? Probably not. And when we created Splunk, none of our engineers had beehives, songbirds or surfboards in mind. But the off-label applications of Splunk have delighted us for years. We’ve always known that data can change our lives, our minds and our world. We just didn’t see the pickle thing coming.

Now that you’ve seen some of the cool things people are doing with Splunk — and maybe took a stab at following their recipes — it’s your turn. Learn more about the Splunk community, and maybe submit your own wild and wonderful Splunk data project. Visit splunk.com/whatscooking.

splunk> turn data into doing™

20-14965-SPLK-What's Cooking-120-digital