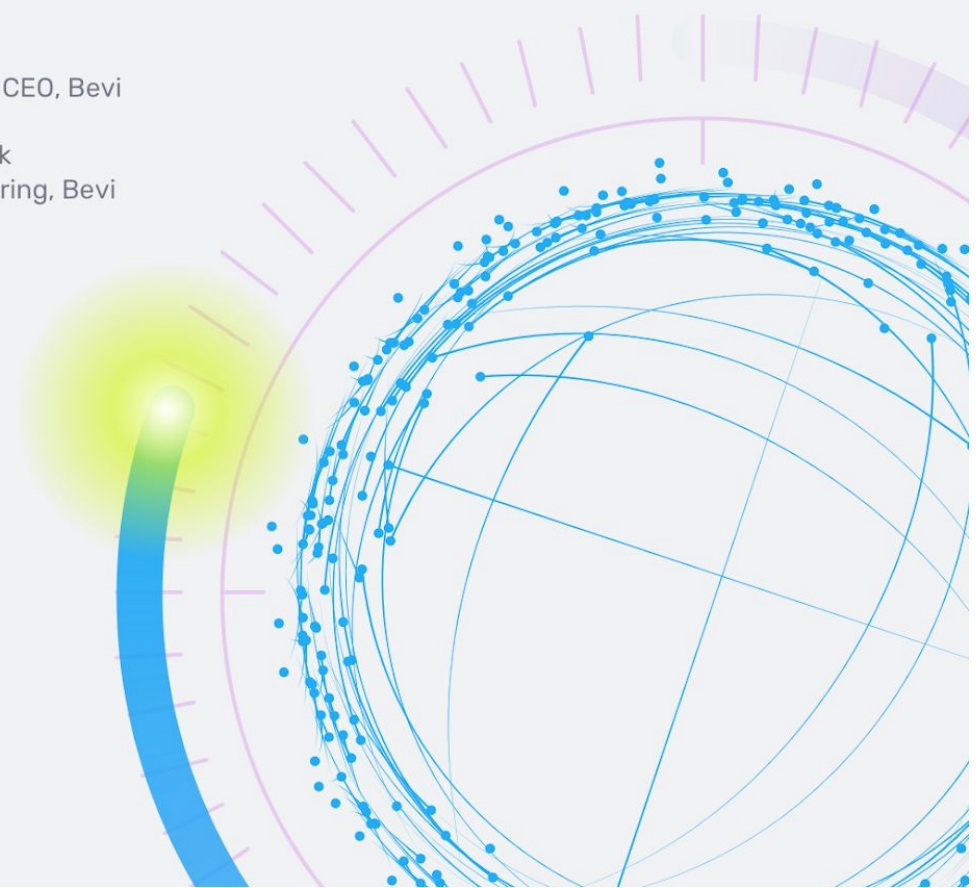*influxdata*

How the Bevi Smart Water Cooler Uses InfluxDB to Give You Less Bottles and More Bubbles

AN INFLUXDATA CASE STUDY

Sean Grundy
Co-founder & CEO, Bevi

Yvan De Boeck
VP of Engineering, Bevi

October 2017

Company in brief

Bevi, founded in 2013, makes internet-connected "smart" office water coolers that provide still, sparkling, and flavored water on-demand. Bevi's coolers, called Bevis, purify water directly from the tap, and then through a touchscreen interface, let users create customized drinks, from regular purified water to sparkling water to cucumber water, coconut water, and a wide variety of other healthy flavored drinks. The Bevi touchscreen is the data capture tool. At its core, Bevi is an environmental company that set out to entirely eliminate disposable bottles from the beverage supply chain. Bevi's goal is to take all the drinks that one can only get in a bottle today and provide them, or provide even better drinks, without the need for disposable bottles or cans. Available in Classic Standup format as well as all-new smaller Countertop format for small spaces, Bevi is the still and sparkling, multi-flavored, eco-friendly, healthy water cooler that never runs out.

Case overview

Bevi wanted to connect its "smart" water coolers to the internet to fulfil its mission of reinventing the wasteful beverage supply chain. To achieve that, Bevi needed internet-connected coolers that would provide data to uncover beverage trends, alert when stocks are running low in the machine, and enable proactive machine servicing before drinks run out. Bevi uses InfluxData to gather metrics and events to remotely conduct real-time data analysis so that it can understand how its water coolers are performing. Remote monitoring frees the Bevi team from having to push upgrades manually and visit every cooler in person. Using InfluxData, Bevi was able to adopt a distributed supply chain and thereby help eliminate plastic bottles and make the best drinks instantly available, through a touchscreen, using purified tap water and natural ingredients.

"We couldn't push updates remotely. We had to go machine to machine. So Yvan and I had to visit every single machine and plug into the tablets and try to test different fixes and see what was going on."

Sean Grundy, *co-founder and CEO*
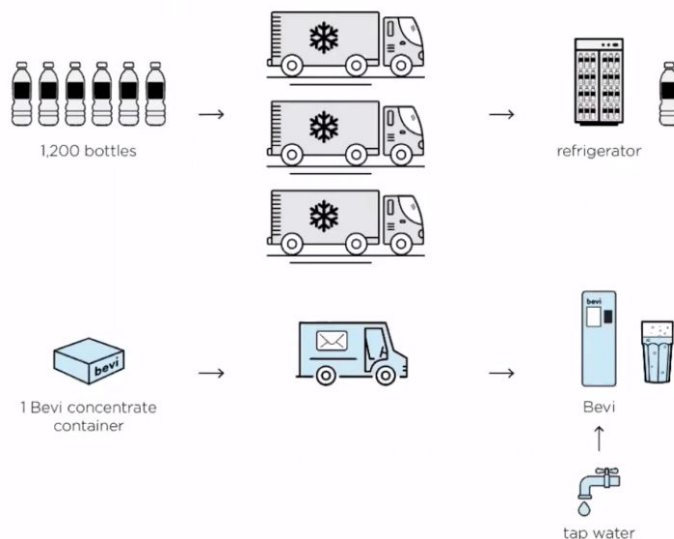
The business problem

Problem: Meeting growing market demand for eco-friendly beverages
Bevi found that despite the massive market for bottled water, millennials have a strong desire for more environmental solutions and that, if given the choice and despite being addicted to various brands of bottled sparkling water, they prefer environmental and healthy solutions. Millennials also want their water to be pure, bubbly, flavored to their personal tastes, without sugar, sweeteners, or BPAs and without wasting a plastic bottle.

Problem: Building a distributed resource-saving supply chain
Bevi needed to move away from the traditional beverage industry model because in that model, more than 30 billion plastic bottles end up in a landfill each year, 17 million barrels of oil are used to make plastic bottles annually in the US alone, and 3 bottles' worth of tap water are used in the manufacturing process of 1 water bottle.

Moving from a centralized to a distributed supply chain

Bevi is building the new beverage supply chain

The technical problem

To solve its business problems, Bevi needed an advanced cooler that would:

- Let people create personalized drinks, set flavor strength, and choose carbonated or non-carbonated options through a touchscreen (whereby they would actually be controlling pumps and valves inside the cooler).
- Be able to create drinks using extremely ultra-filtered water and then mixed to very precise ratios.
- Make all drinks come out of a single nozzle without ever having a crossover of one flavor to another.

Problem: Becoming an internet-connected product

Bevi needed to transition from a hardware-driven to a software-driven company. The first generation of Bevi coolers had off-the-shelf tablets powered by an actual battery, which would often run out and not fully charge overnight. The team has to visit every single machine and plug into the tablets to test different fixes since they could not yet monitor or push upgrades remotely.

Software enables the customer experience

Becoming an internet-connected product was absolutely essential for Bevi because remote monitoring would enable:

- Providing proactive cooler maintenance, improve system stability and uptime, track frequency of use, and handle inventory management
- Improving the end quality of beverages and service by providing real-time data on when ingredients and filters need to be replaced, as some of the best flavorings have low shelf lives and can't be used in a traditional bottle or traditional soda fountain

For Bevi, software is the key to collecting an enormous amount of data. Every time users touch a Bevi touchscreen, Bevi can discover what is popular where and then correlate flavor preferences.

Problem: Gathering metrics and events in real time

Bevi wanted to visualize the data captured from the touchscreens and started looking for a time series database. The team learned from a friend about InfluxDB, InfluxData's time series database. Then, upon seeing an InfoQ presentation by InfluxData co-founder Paul Dix, Bevi adapted the presentation examples to its use case. The team found it easy to start the Docker container with InfluxDB, hook up Grafana to that, and visualize the data, and had their proof-of-concept implemented very rapidly.

Solution

> *"Sean went to get a Bevi drink, and by the time he was back at his desk, I could tell him what kind of drink he was drinking. So, that was very impressive and very easy without any frame points. From then on, that moment on, we continued using InfluxDB."*
>
> **Yvan De Boeck**, *VP of engineering*

Bevi gives users a worry-free, set-it-and-forget-it solution:

- Bevi's single 3-gallon concentrate container has enough concentrate to produce 1200 bottles worth of flavored drinks, which enables much lower shipping costs and much less fuel usage than the traditional beverage supply chain, thereby competing both on sustainability and cost.
- The Bevi coolers, by connecting to the Internet and having a remote monitoring backend, eliminate the need to constantly restock heavy beverage cases.
- Each Bevi machine sends diagnostic data over the internet to a Central Command, and the Bevi team or their service partners monitor remotely and refill the machines proactively.

Why InfluxDB?

Bevi chose InfluxDB because:

- Bevi's team knew from previous experience that using another database, such as PostgreSQL, would have been much more time-intensive and labor-intensive for time series data, and found InfluxDB to be spot-on for their use case.
- All the data is persistent in InfluxDB, so Bevi could store data on user response to new flavors, adherence to new flavors, and form data entered by their service people documenting ingredient and filter replacement schedules.

> *"The most characteristic thing about InfluxDB is that it was so painless to set up and to grow with our needs. So we're still not at all at limit of how we can use InfluxDB."*
>
> **Yvan De Boeck**, *VP of Engineering*

Technical architecture

InfluxData powers Bevi's architecture
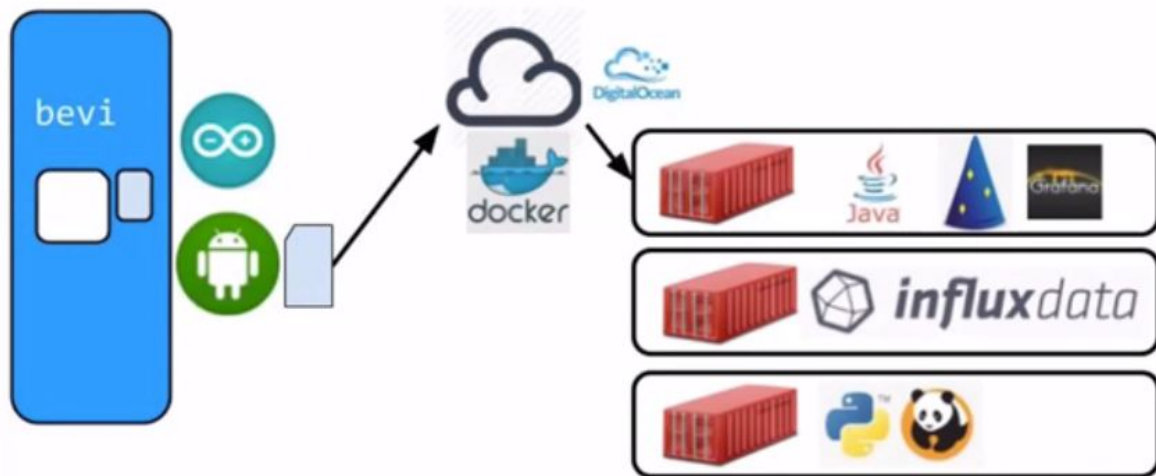
- The Bevi smart water coolers have an android tablet that is the main interaction point with end users. Bevi updates and pushes its apps regularly. The android tablet also functions as an IoT device that relays all events to Bevi's backend.
- The Bevis in the field accumulate events, which are written to a log file.
- Events are sent to the Cloud (Bevi's service are on DigitalOcean with Docker containers).
- Every event has four shared properties (timestamp, units, sequence, and type) which are used to track performance and enable proactive maintenance.
- InfluxDB is the only database in Bevi's architecture, and it stores the history of all Bevi's machines and all service data.
- A post-processing layer (involving data science and Python) uses a .CSV file that Bevi gets from InfluxData.
- Bevi's business logic is in Java using Dropwizard, and Grafana, in the beginning, was used to visualize the time series stored in InfluxDB.
- Java with Dropwizard is used to handle streams of data coming from the machines and create derived streams to compute the status of consumables and flag any abnormal behavior. The data is continuously used to optimize the Bevis' operation and user experience.

Measuring CO2 levels and flavor-mixing trends using Derived Stream

Bevi coolers have sensors for CO2, the mechanism used to obtain sparkling water. To measure fullness of the coolers' CO2 tank, the coolers have a built-in load sensor that reports its measurement in millivolts. Every machine has a different format to derive from the actual weight in pounds. When users mix flavors, Bevi bundles the data from different dispensers as one mixed drink to extract behavioral knowledge of how users actually use the machines.

Derived Stream Example
Load Sensor mV (millivolts) => Lbs (pounds)

> *"Having InfluxDB from the get-go was really beneficial. In software, we could track errors and bugs when they were happening. In the field, we would have very fast turnaround times to see what's going on and if the tablet is well-behaving."*
>
> *Yvan De Boeck, VP of engineering*

What's next for Bevi?

Bevi's plans for the future include:

- Interesting ways to visualize plastic bottle reduction to show user bases the impact they're having
- Using InfluxData to achieve more customization from a software UI perspective as well as more ways to customize drinks
- Developing a way for individual users to track their own drinks in order to meet personal consumption goals and track nutrition and vitamin intake

- Utilizing machine learning to guide the introduction of new flavors

Results

> *"Beverage companies are interested in providing their beverages through our platform...because we collect all this information, in terms of geographies, industries, demographics, on where different drinks are popular. That's something that if you're a traditional beverage company, it's often very hard to find out."*
>
> **Sean Grundy**, *co-founder and CEO*

With real-time access to metrics and events, Bevi has been able to build on millennials' interest in eco-friendly products and let them guide its development on what kind of products to build for customers:

- Bevi has already saved over 10 million bottles and cans from entering the ecosystem and has saved a massive amount of fuel, putting technology in service of the environment.
- Bevi, empowered with remote real-time monitoring, has been able to use much higher quality filters that need more frequent replacement and much higher quality ingredients that have a low shelf life.
- Bevi's visibility into individual flavor preferences has enabled it to correlate favorite flavors, which has led to happier customers and higher revenue.
- Beverage companies are showing interest in providing their own beverages through Bevi's platform.

A one-stop-shop for still and sparkling water drinks in the office, Bevi is bringing the future of office water to smart offices today.

About InfluxData

InfluxData is the creator of InfluxDB, the open source time series database. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by IoT devices, applications, networks, containers and computers. We are on a mission to help developers and organizations, such as Cisco, IBM, PayPal, and Tesla, store and analyze real-time data, empowering

them to build transformative monitoring, analytics, and IoT applications quicker and to scale. InfluxData is headquartered in San Francisco with a workforce distributed throughout the U.S. and across Europe.

Learn more.

InfluxDB documentation, downloads & guide

Download InfluxDB

Get documentation

Additional case studies

Join the InfluxDB community

![influxdata logo]

799 Market Street
San Francisco, CA 94103
(415) 295-1901
www.InfluxData.com
Twitter: @InfluxDB
Facebook: @InfluxDB