



## OP<sup>U</sup>WER

### Opower Uses Big Data to Help Consumers Save \$320 Million in Utility Bills

#### Company Overview

Opower's mission is to motivate everyone on Earth to save energy. The Virginia-based company's pioneering approach – a white-label platform that pairs behavioral science with Cloudera-powered big data analytics – engages utilities customers with a **consumer engagement platform** that presents relevant insights into household energy use through informative dashboards, alerts, incentives, and household comparisons. The information delivered generates heightened awareness and incentivizes consumers to reduce energy consumption, which saves money and reduces the impact on our environment.

Behaviorally fine-tuned and decision-tree optimized, the innovative platform reaches customers through the mail, email, text messages, and web for Opower's 85-plus utility clients – including eight of the US' 10 largest. The platform reaches more than 18 million homes across North America, Europe, and Asia.

With raised consciousness and a data-driven view of how their energy use is impacting their wallets, customers are taking action – altering their behavior to use energy during off-peak times, to swap out older and inefficient appliances, and to anticipate factors such as weather in order to save energy and money.

The resulting efficiency is a reliable and sustainable energy resource that is helping Opower's utility clients offset the intensifying energy demands on their aging infrastructures.

#### Business Challenges Before Cloudera

“The obvious focus for utility companies is on their physical infrastructure,” notes Eric Chang, Opower's technical lead for data services. “How do I manage my flow of electricity? How do I make sure the lights are on? Managing the data that gets collected from that infrastructure is secondary to them.”

Utilities are not set up to manage that data, which continues to escalate in volume. Advanced metering infrastructure (AMI) technologies such as smart meters, thermostats, and other in-home devices generate and send providers frequent voltage and temperature readings (consider an example of 100 readings per day, per household, multiplied

## Key Highlights

### Industry

- Software-as-a-Service for utilities

### Locations

- Arlington, VA, US (headquarters)
- San Francisco, CA, US
- London, England
- Singapore

### Business Application Supported

- Enterprise data hub empowering customer engagement platform

### Impact

- 2.7 terawatt hours saved to date
- US\$320+ million in consumer savings to date

### Technologies in Use

- Hadoop Platform: Cloudera Enterprise + RTD
- Hadoop Components: Cloudera Manager, HBase, Hive
- Servers: Dell PowerEdge C5220, Dell PowerEdge C2100
- Data Warehouse: Hive

### Big Data Scale

- Platform supports data collection, processing, and analytics for 85+ global utilities providers
- Reports reach 18+ million homes across North America, Europe, and Asia

by thousands of customer households for just one utility provider). Other related data streams compound data size and include demographics data, smart appliances and sensors, weather data, consumer behavior information, and social media data.

But data is only valuable if you do something with it. Opower recognized that processing and analyzing this information would provide the utilities with much-needed insights into customer's energy consumption that could be used to drive energy savings, reduce demand at peak times, adopt new rate structures, support smart meter roll outs, and lower cost to serve.

Opower's first data management platform was a MySQL open source relational database management system (RDBMS). "With a small number of utilities clients, we could manage with just relational databases," notes Chang. As its client base and consumption data grew and Opower's products offered more functionality, though, the company needed to adopt a new architecture – specifically one that could scale to process and store large amounts of data without the cost and operational overhead of its relational databases. "For us," says Chang, "that new architecture included [Apache Hadoop](#) and [\[Apache\] HBase](#)."

In 2011, the company began moving to a distributed services-oriented architecture. Opower deployed [Cloudera Enterprise + RTD](#) as the enterprise data hub to expand its big data capabilities and to store, transform, and query time series and other data.

## Use Case

Drew Hylbert, Opower's vice president of technology and infrastructure, sees use cases for big data in the utilities sector taking a variety of forms – including grid optimization and production optimization. "For Opower, the value of big data in the utility space is to help providers understand how their customers use energy, and to engage those customers with insights that empower them to change their behaviors, reduce their utility bills, and save energy."

Opower's customer engagement platform is delivered as software-as-a-service (SaaS) which offers multiple benefits for utilities providers. "Utilities simply can't cope with the vast volumes of smart meter data – not just with storing the data, but being able to analyze it and put it to use," says Hylbert. Opower's service model delivers that scalability – and the cost scales as well, based on use. Opower continuously evolves its products so that as devices, communication channels and preferences, and customer behaviors change, its clients can respond effectively.

Key components of Opower's architecture include [CDH](#) and [HBase](#). [Apache Hive](#) serves as Opower's data warehouse. The Cloudera-based enterprise data hub delivers numerous improvements to Opower's infrastructure and capabilities – including processing speed and efficiency, and the ability to easily pull data in parallel from multiple sources.

HBase serves as a highly scalable data store that is built on top of Hadoop and offers fast, random read/write access to users and applications. Raw energy consumption data is batched regularly and imported into HBase, analytics are performed, and the results are stored there. Immediate consumers of the data stored in HBase include

- Web front ends that directly forward queries from end users
- Other internal processes, for example, that generate customer emails or paper mailings

"HBase is the means by which we are able to deliver our insights at scale, in real time, to any user," says Chang.

Opower manages its Hadoop cluster through **Cloudera Manager**. “Cloudera Manager provides a lot of tools that enable us to know exactly what the cluster is doing at a given time,” says Tim Luo, a software engineer on the Opower data services team. The tool allows Opower’s team to respond to alerts quickly, manage and tweak configurations, restart servers, and ensure that overall cluster health is good.

While Opower is focused primarily on analyzing energy usage data and delivering insights to end users, Chang explains the closed loop system that makes it easy for Opower to correlate the impact of those efforts. “The utilities provide us with meter reads, we perform analysis, we target a certain number of customers to deliver that analysis to, the customers act on that information and change their behavior, the utilities provide new meter reads, we compare those to the original reads, and we measure the actual impact in energy savings in that targeted population.” Even minute changes in behavior can make a big difference when they are driven on a large scale. Chang stresses, “A two to three percent energy savings across many millions of households really adds up.”

### **Impact: More than \$320 Million Saved**

Placing the impact of Opower’s customer engagement efforts in terms of real numbers, the company has helped its utility partners save energy consumers more than US\$320 million in utility bills – and generated nearly three terawatt hours in energy savings. That’s enough energy to power every household in Salt Lake City and St. Louis for an entire year.

### **Impact: Driving a Greener World from the Inside Out**

“There is amazing potential for leveraging big data in the energy industry,” says Ogi Kavazovic, Opower’s vice president of marketing and strategy. Today, Opower is using big data to engage customers and motivate them to save energy and money. Increasingly, utilities will use big data to fuel insights that will drive new opportunities for innovation and efficiency in their energy delivery.

Not only is Opower influencing consumer energy savings, it’s a mindfulness that the company also imposes on itself. “The Hadoop platform infrastructure allows us to be much smarter about how we utilize our resources across the platform,” notes Chang. “We also worked closely with **Dell** on our latest acquisition of hardware to make sure we are using lower power consumption drives.”

That energy consciousness is one of the main reasons Alexandre Normand, a software engineer at Opower, wanted to work there. “Opower is one of the few companies that helps people save energy and makes an impact on the environment. As a software engineer, there just are not that many opportunities to be involved in that.”

### **Impact: Poised for Rapid Innovation**

“We’re constantly innovating, constantly measuring the outcomes, and using that information to improve our platform,” says Kavazovic. “When utilities work with us, not only do they get the most advanced data platform in the industry, but they also get a continuously improving customer engagement platform that’s proven to deliver their data outcomes.”

“Cloudera Manager provides a lot of tools that enable us to know exactly what the cluster is doing at a given time.”

Tim Luo,  
Software Engineer, Opower

“We chose Cloudera because of their understanding of Hadoop and HBase – and they were willing to engage with us on the kinds of problems that we were solving. They are an excellent partner both in keeping the systems going and delivering new technology on top of HBase. As we continue to scale, we know we can rely on Cloudera to help us understand how to grow – the impact on our cluster, how we need to provision, and what configurations need to be tweaked to accommodate that additional scale.”

Drew Hylbert,  
Vice President of Technology and Infrastructure  
Opower

Cloudera plays an important role in enabling these innovations. Hylbert explains, “We chose Cloudera because of their understanding of Hadoop and HBase – and they were willing to engage with us on the kinds of problems that we were solving. They are an excellent partner both in keeping the systems going and delivering new technology on top of HBase. As we continue to scale, we know we can rely on Cloudera to help us understand how to grow – the impact on our cluster, how we need to provision, and what configurations need to be tweaked to accommodate that additional scale.”

Hylbert also likes the **open source** aspect of Cloudera, and the new flight and functionality that comes out of that. “I’m impressed with Cloudera’s **ability to support** our use of open source technology in ways that not many others are using it,” he says. “I think that’s a particularly valuable piece of Cloudera’s offering – not just understanding the software, but understanding how to use it and how to interact with the open source community to get the most value from it.”

## About Cloudera

Cloudera is revolutionizing enterprise data management by offering the first unified Platform for Big Data, an enterprise data hub built on Apache Hadoop. Cloudera offers enterprises one place to store, process and analyze all their data, empowering them to extend the value of existing investments while enabling fundamental new ways to derive value from their data. Only Cloudera offers everything needed on a journey to an enterprise data hub, including software for business critical data challenges such as storage, access, management, analysis, security and search. As the leading educator of Hadoop professionals, Cloudera has trained over 40,000 individuals worldwide. Over 1200 partners and a seasoned professional services team help deliver greater time to value. Finally, only Cloudera provides proactive and predictive support to run an enterprise data hub with confidence. Leading organizations in every industry plus top public sector organizations globally run Cloudera in production. [www.cloudera.com](http://www.cloudera.com).

---

[cloudera.com](http://cloudera.com)

1-888-789-1488 or 1-650-362-0488  
Cloudera, Inc. 1001 Page Mill Road, Palo Alto, CA 94304, USA

© 2015 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice.

cloudera-casestudy-opower-102

