

# Sovrn Quickly and Automatically Scales Its Real-time Data Processing with DC/OS

Big Data Processor Adopts and Deploys Containers with Push-Button Ease

## Introduction

Effectively deriving insights from big data can be a huge challenge for today's organizations. Few companies have the resources to collect and analyze large volumes of data in real time. Sovrn has set out to change that dynamic for publishers and content creators.

Sovrn's goal is to arm its customers with the insights they need to make better decisions about growing audiences, engaging readers, and monetizing content. Sovrn achieves this by aggregating data from across its network and making it easy to consume through a proprietary ad technology stack with direct connections to every major buyer in the world.

Sovrn serves approximately 3,500 publishers and 250 million daily readers on over 50,000 websites. As one of the largest publisher networks in the world, Sovrn collects, combines, and analyzes an immense amount of data from across its network. By aggregating this data and sharing it through its meridian platform, Sovrn provides smaller publishers with advanced insights, regardless of the amount of site traffic.

## Sovrn Project At-a-Glance

### Customer

Sovrn – online advertising technology firm

### Challenge

Quickly scale infrastructure up and down to meet ever-changing levels of demand

### Solution

DC/OS + data and analytics technologies (Cassandra, Kafka, Spark and Apache)

### Impact

Running dynamic workloads and evolving engineering capabilities

## Challenge:

### Modernizing Sovrn's Infrastructure Capabilities

Half of Sovrn's revenue is driven through a real-time ad exchange where it supports billions of ad calls a day. Key to its success is building real-time capabilities that enable Sovrn to scale its system up and down in line with the cyclical trends of its business. Plus, it must do this while processing a growing amount of data faster and more frequently.

On any given day, Sovrn scales up and down an order of magnitude anywhere from four to 10X, depending upon the amount of traffic on the ad exchange. As part of that, Sovrn works with upstream and downstream partners and must optimize its operations to work effectively with this diverse ecosystem spanning global locations and multiple cloud-based and hosted data centers.

As Sovrn's VP of Platform Engineering, Kyle Gilliland, explains, "We need to ensure our system can handle the large volume of data that comes through our publisher network at the velocity at which the data is seen. This must all be done without slowing down the responsiveness of our system."

According to Gilliland, it takes a lot of horsepower and compute power to process big data effectively and efficiently. "We build machine learning — or artificial intelligence — algorithms to manipulate all the data we collect. Scalable technologies that allow us to do that fast and efficiently are crucial to our business." However, to scale up its infrastructure to meet ever-changing levels of demand, Sovrn was manually acquiring and provisioning systems — a process that took weeks. As a result, the business was not able to move as quickly as it wanted, nor ensure necessary infrastructure resiliency.

## Solution:

### DC/OS Eases the Transition to Containers

Gilliland knew that driving an as-a-service-based software architecture would allow Sovrn to isolate services in containers and easily scale those services.

However, though he wanted to get into containers, Gilliland knew Sovrn couldn't afford to jump into the complex world of Kubernetes. Moreover, he didn't want to Sovrn to develop its own open source capabilities using other open source tools.

As far as Gilliland was concerned, DC/OS represented the perfect marriage of the two, removing barriers to entry. "We saw that DC/OS could help orchestrate our workloads, enabling us to handle this efficiently in any location," says Gilliland.

Gilliland was impressed by the way D2iQ guided Sovrn. "D2iQ convened experts to discuss the technical and business problems we were facing. It then tailored a solution to our specific needs, accelerating our time to market," he explains.

**"It's a build-once/deploy-anywhere construct that has helped evolve our engineering organization. The push-button capabilities of DC/OS enable my engineers to develop and create more value."**

**—Kyle Gilliland, VP of Platform Engineering, Sovrn**

While Sovrn had successfully been able to deploy tools like Kafka and Cassandra, DC/OS enabled it to do so with a push-button capability. "DC/OS enabled our developers to do this easily and to automate that whole structure end-to-end, without any interruption of configuration or management in between."

Currently Sovrn is running 18 applications in a production state across its DC/OS cluster, including Java applications, Spark, and Apache, and is looking at moving ephemeral workloads to DC/OS.

## Impact:

### Running Dynamic Workloads in Production

Using customized open source technologies that specialize in delivering real-time data processing, Sovrn can enable analysis of nearly any event within 20 seconds of its occurrence. Plus, by scaling up new infrastructure within days or hours, it can more quickly and reliably deliver this capability.

DC/OS has given Sovrn the opportunity to put in place repeatable automated processes. As a result, Sovrn has matured through its container journey. Now it builds real-time data processing pipelines with continuous integration in mind from the start. And it deploys applications seamlessly in the same fashion from Development to QA to production.

“It’s a build-once/deploy-anywhere construct that has helped evolve our engineering organization. The push-button capabilities of DC/OS enable my engineers to develop and create more value,” says Gilliland.

At the same time, DC/OS positions Sovrn to further innovate. Over the years, Sovrn has developed Java apps to support its business, and these apps have become monoliths of sorts. Sovrn is on a journey to start breaking services off those monoliths and deploy those into containers where it can run them effectively. DC/OS makes this possible. “DC/OS has given us the capability to peel a service off the monolith and run it in an isolated state. This enables us to continually evolve those services and our platform,” Gilliland concludes.

“D2iQ convened experts to discuss the technical and business problems we were facing. It then tailored a solution to our specific needs, accelerating our time to market.”



Kyle Gilliland,  
VP of Platform Engineering  
Sovrn



[d2iq.com](http://d2iq.com) | [sales@d2iq.com](mailto:sales@d2iq.com)