

## **Introduction** ▶

Hadoop/Spark has emerged as an open standard Big Data Operating System with wide community support. It allows the usage of in-cloud or on-premise commodity hardware to effectively address the “3 V’s” of Big Data: volume, variety and velocity. Moreover, it brings a new generation of intelligent functions as standard components including machine learning, real-time streaming and graphical analysis. This combination makes it an ideal operating system for data processing.

Data Operating Systems, like regular operating systems, provide powerful base functions. These functions are essential for programmers, but business users can only benefit from them through applications created by programmers. This is analogous to how business users need an app like Microsoft Office to benefit from Microsoft Windows. InetSoft’s data intelligence app is the application for data analytics, visualization, dashboards, and reporting that’s specifically designed for this Big Data Operating System.

## **Don’t Many Analytic Tools Work with Hadoop?** ▶

The answer is a little bit of ‘yes’ and a big bit of ‘no.’ Hadoop, recognized as a de facto standard, enjoys widespread analytic tool support. That is the small bit of ‘yes.’

The big bit of ‘no’ is that traditional tools just “connect” to Hadoop. To these tools, Hadoop is just another data store or database to get data from. Besides this data link, these tools reside and execute in a totally different environment. This vastly limits how much the tools can utilize the Big Data OS. The PC analogy is how non-Windows software can connect to Windows to access files stored on it, yet it is not an application running on Windows.

## **InetSoft’s App Is a Native Big Data OS Analytic Engine** ▶

InetSoft’s analytic engine natively executes on and inside Hadoop/Spark Big Data OS. It takes advantage of Hadoop/Spark’s power to the fullest extent. Instead of retrieving data out of the Hadoop environment in order to process it, InetSoft brings the engine into the Big Data OS. This allows the analytic engine’s mash up, visual analytics, data instructions, machine learning processing, and many others functions to execute with the full power of Hadoop.

Beyond the native engine, InetSoft also brings data into the Big Data OS. Enterprise data, from traditional databases, data warehouses, API-accessed data, and NoSQL data, can all be dynamically brought into the Big Data OS as analytics demand. A data lake, with little or no design for any particular purpose, normally resides inside a Big Data platform. InetSoft’s analytic engine can tap into and mash up such sources without any data movement.

## A Small Learning Curve, Big-Data-In-A-Box ▶

InetSoft's analytic engine is designed for easy and gradual learning and scaling. As a first step, it can even be deployed without Hadoop/Spark. Even though this limits accessibility to advanced Big Data functions, it offers good analytic dashboarding and reporting.

Big-Data-In-A-Box is InetSoft's answer to the complexity of the Big Data OS. Install InetSoft's pre-packaged Hadoop/Spark Docker Container on each cluster node, and InetSoft's visualization server manages the environment. No more Big Data technology knowledge is required. This means organizations with no Big Data experts can deploy the InetSoft Big Data solution. Visual analytics and reporting will automatically take advantage of the Big Data OS.

For organizations with more mature Big Data implementations, the InetSoft engine can be dropped into an existing Hadoop/Spark installation. Existing data, machine learning models and other assets will all become available to users for visual analysis and mashup. Data scientists, developers and business users can seamlessly cooperate on a single Big Data platform.

## An Example Visualization ▶

### Marketing Lead Trends

