

World-class expertise to get the most out of your Apache™ Hadoop® investment

We offer a wide range of services to help you plan, implement, deploy and manage your big data installation.

From implementation to data migration to tuning and optimization to data engineering and advanced analytics, MapR's Professional Services team will work with you every step of the way.

As the technology leader in Hadoop, MapR employs the best and brightest engineers and data scientists in the big data industry. MapR augments its in-house expertise by collaborating with industry experts and certified service providers to ensure you get the best Hadoop experience.

## Core Hadoop Services

### Implementation and Data Migration

We will assess your current environment and then install and configure MapR for your particular infrastructure. If you're currently using a different Hadoop

distribution, we'll move the data over to your new MapR installation, ensuring a smooth transition with zero data loss.

### Health Check and Tuning/Optimization

Maintaining and optimizing MapR for your current environment is critical to the long-term success of your Hadoop deployment. Our Professional Services team will review the optimization and

performance attributes of your environment, and help you understand how to diagnose, improve and maintain the performance of your MapR deployment.

## Data Engineering

### Solution Design/Implementation

Our Data Engineers will help you fully realize the benefits of your MapR investment by designing and implementing a comprehensive solution that will perform optimally within your particular

environment. All aspects of your implementation will be addressed, including hardware infrastructure, data sources, ecosystem software, and operations considerations.

### ETL

The first step in any Hadoop implementation is data ingestion and transformation. MapR will develop and implement a customized ingestion/ETL plan that includes identifying the multiple data

sources and file formats, transforming them to meet your needs, and loading the data into the data structures best suited to your needs for further analysis.

### Application Development

Our team of data engineering experts can help you design, build and optimize your big data applications. Whether it's building Java MapReduce jobs, implementing in Apache Pig, building distrib-

uted indexes, constructing Apache Hive (HQL) queries, or building complex data models, MapR provides the experience and expertise you need.

### Data Aggregation

Our Data Engineers will assist you with data aggregation so that you can quickly

and accurately consolidate and fuse data from various sources.

### Presentation Layer

From start to finish, MapR will create the most optimal data flow: from integrating

data sources, to Hadoop, to the presentation layer.



Maximize your investment in Hadoop with MapR Professional Services

## Advanced Analytics

<b>Use Case Discovery</b>	As part of our Advanced Analytical services, MapR will send a Data Scientist to your site who will conduct a use case discovery session in order to gain a clear understanding of business priorities as well as existing workflows/data sources available for analytics. From there, use cases will be identified, and a road map will be created for big data development and training so that you can become self-sufficient.
<b>Data Modeling</b>	Our Data Scientists will identify the best way to transform and structure raw data, and will create a custom data model for your particular use case.
<b>Statistical Methods</b>	MapR will create the right algorithm or statistical model in order to meet your speed, reliability and maintenance goals.
<b>Machine Learning</b>	MapR will assess your business needs, define the business case, and recommend the appropriate machine learning methods.

## M7/HBase Practice

MapR can schedule an M7/HBase Practice workshop on site to ensure that you are up to speed on the following:

<b>HBase Schema Design</b>	You'll learn about the basic underlying concepts of the storage layer in HBase, as well as how to build a schema, load, and query data.
<b>Application Analysis</b>	You'll take a close look at the HBase architecture, data model and Java API, as well as several advanced topics and best practices.
<b>Performance Tuning</b>	Performance tuning requires a deep understanding of both Hadoop and HBase. During this practice session, you'll learn how to fine-tune HBase to boost performance, based on projected workloads and data flows.

MapR delivers on the promise of Hadoop with a proven, enterprise-grade platform that supports a broad set of mission-critical and real-time production uses. MapR brings unprecedented dependability, ease-of-use and world-record speed to Hadoop, NoSQL, database and streaming applications in one unified Big Data platform. MapR is used across financial services, retail, media, healthcare, manufacturing, telecommunications and government organizations as well as by leading Fortune 100 and Web 2.0 companies. Amazon, Cisco, and Google are part of MapR's broad partner ecosystem. Investors include Lightspeed Venture Partners, Mayfield Fund, NEA, and Redpoint Ventures.

