

# Migration Guide

This guide provides instructions for migrating business-critical data and applications from an Apache Hadoop cluster to a MapR cluster.

The MapR distribution is 100% API-compatible with Apache Hadoop, and migration is a relatively straight-forward process. The additional features available in MapR provide new ways to interact with your data. In particular, MapR provides a fully read/write storage layer that can be mounted as a filesystem via NFS, allowing existing processes, legacy workflows, and desktop applications full access to the entire cluster.

Migration consists of the following steps:

- [Planning the Migration](#) — Identify the goals of the migration, understand the differences between your current cluster and the MapR cluster, and identify potential gotchas.
- [Initial MapR Deployment](#) — Install, configure, and test the MapR cluster.
- [Component Migration](#) — Migrate your customized components to the MapR cluster.
- [Application Migration](#) — Migrate your applications to the MapR cluster and test using a small set of data.
- [Data Migration](#) — Migrate your data to the MapR cluster and test the cluster against performance benchmarks.
- [Node Migration](#) — Take down old nodes from the previous cluster and install them as MapR nodes.