

M7 EDITION

ENTERPRISE-GRADE PLATFORM FOR NOSQL AND HADOOP

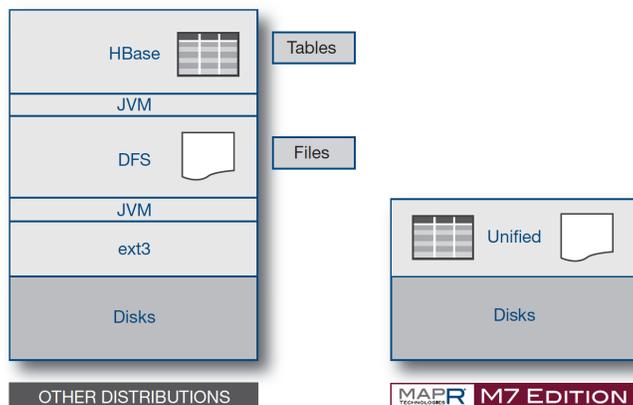
MapR M7 provides ease of use, dependability and performance advantages for NoSQL and Apache™ Hadoop® applications. MapR M7 has removed the trade-offs organizations face when looking to deploy a NoSQL solution. M7 provides scale, strong consistency, reliability and continuous low latency.

MapR M7 is architected from the ground up to deliver reliability and performance without requiring compactions or background consistency checks to work smoothly. M7 delivers performance of over one million operations/sec with a ten-node cluster that can scale linearly. Compared to Apache HBase™ applications running on other distributions, M7 delivers dramatically better performance with 4-10x more throughput across various read-write workloads and elimination of latency spikes.

M7 provides dramatic scalability advantages with support for up to one trillion tables across thousands of nodes. M7 also provides instant recovery from failures, ensuring 99.999% availability for Apache HBase and Hadoop applications.

With M7, there are no region servers, additional processes, or any redundant layer between the application and the data residing in the cluster. M7's zero administration approach includes automatic region splits and self-tuning with no downtime required for any operation including schema changes.

SIMPLIFIED ARCHITECTURE



M7 has a purpose-built architecture specifically designed to optimize the storage and processing of files as well as tables with a unified platform.

MAKING APACHE HBASE APPLICATIONS ENTERPRISE GRADE

- Patent pending architecture for a unified data platform for tables and files
- No RegionServers eliminates complexity and unreliability
- No manual administrative tasks such as table merges or splits
- Instant recovery from failures
- Full data protection and disaster recovery
- Consistent performance with no compactions
- High throughput and very low latency and no spikes
- Unprecedented scale

APACHE HBASE AND HADOOP MADE EASY

- Automated region splits
- Performance self-tuning with no compactions
- Automated data compression
- User & group usage tracking and quotas
- Online table management through UI, CLI, REST API
- Easy cluster management with the MapR Control System

APACHE HBASE AND HADOOP MADE DEPENDABLE

- Instant recovery from failures
- Snapshots for point-in-time recovery for tables and files
- Rolling upgrades for tables and files
- Mirroring for disaster recovery and multi-cluster synchronization of all data
- No single points of failure with HA and self-healing
- No NameNode HA™
- Comprehensive security with strong native authentication

APACHE HBASE AND HADOOP MADE FAST

- Consistent throughput and low latency for database operations
- Over 1 million ops/sec from 10 node cluster
- 4-10x more throughput and very low latency with no latency spikes compared to HBase applications running on other Hadoop distributions
- Increased throughput in MB/sec
- Direct block device I/O
- 100% data locality
- Automatic and transparent client-side compression

FEATURE SPOTLIGHT SPECIFIC TO HBASE

NO DOWNTIME WITH INSTANT RECOVERY

M7 delivers database high availability. Instead of requiring thirty minutes or more of downtime to reassign the region when a node goes down that HBase running on other Hadoop distributions requires, M7 enables instant recovery from data replicated in the cluster.

ZERO NOSQL ADMINISTRATION

M7's simplified architecture eliminates manual operations related to HBase administration with automated operations including region splits and self-tuning.

CONTINUOUS LOW LATENCY

MapR M7 provides consistent low latency by avoiding garbage collections or compactions that affect performance. Low disk I/O coupled with smaller disk footprint makes database operations on disk fast and predictable.

FULL DATA PROTECTION

M7 delivers full data protection for NoSQL applications, including Apache HBase applications. Snapshots enable point-in-time recovery of tables to protect against user or application errors. M7 expands snapshots to include all data - both files and tables. With M7, tables can be read directly from Snapshots and recovered directly without the downtime required to restore HBase tables in other distributions.

BUSINESS CONTINUITY WITH MIRRORING

Mirroring allows users to automatically replicate differential data in real-time across clusters. This could be employed to create disaster recovery solutions for databases or leveraged to provide read-only access to data from multiple locations. Because M7 does not require RegionServers to be reconstructed, databases can be brought up instantly on the mirrored site if the active site goes down.

M7: UNIQUE ADVANTAGES FOR HBASE APPS

	MAPR	OTHERS
99.999% High Availability	✓	X
Instant Recovery from Failures	✓	X
Continuous Low Latency (No compactions)	✓	X
High Performance with Strong Data Consistency	✓	X
Zero Administration (Automated splits, self-tuning)	✓	X
Instant Recovery, Snapshots, Mirrors	✓	X
Scalability (Number of tables supported)	1 Trillion	Hundreds

FEATURE SPOTLIGHT APPLICABLE TO ALL OPERATIONS

COMPREHENSIVE INTEGRATED SECURITY

MapR brings native authentication to Hadoop users via easy-to-use strong-wire authentication, so you can easily meet stringent security requirements and regulations. Secure all of the Hadoop ecosystem components through a simple, fast and self-contained security model.

SHARE THE CLUSTER SAFELY WITH MULTIPLE JOBS

MapR protects the core system by isolating it from user jobs so runaway jobs can't bring down your entire cluster. Features such as ExpressLane, data placement and job placement control mean you can run multiple jobs with different requirements without conflict.

REVOLUTIONARY DIRECT ACCESS NFS™

MapR enables you to simply mount the cluster as an NFS volume giving Hadoop a full read/write storage system to support multiple readers and writers as well as full random read and write.

MAPR M7 EDITION SUPPORT

- 24x7 phone and email support
- On-demand patches as needed
- Online incident submission and response

» Please refer to the MapR M5 Edition datasheet for more feature descriptions.

REGISTER TODAY FOR M7 TRIAL TO EXPERIENCE THE POWER OF ENTERPRISE-GRADE NOSQL AND HADOOP

You can also visit **THE MAPR ACADEMY** for free training videos and documentation to help you get the most out of your Hadoop investment.

Visit us at: www.mapr.com



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.

For more information, please visit www.mapr.com