

MapR

Key to Digital Transformation

MAPR.

Adrian Smolski
Principal Solutions Architect
May 2017



What differentiates these business from others?



Logos source: Individual companies corporate websites.

MAPR

MapR Confidential

© 2017 MapR Technologies

2

Talk about them that they don't have legacy systems and can deploy a lot of new technology which is more cost effective



Why are Businesses Transforming?



traditional approaches require an application first approach.

You start with the application and determine what are the data requirements. Then you prepare the data into specialized schemas, or cubes to serve the application. The keys to success are understanding all the requirements so you can have the right data model. Extract Transform and Get the data loaded.....



Application & Data Model Has Radically Changed



Each application solved one problem and created its own data type



Diverse data assets must be accessible from anywhere by microservices



Applications Dictate Data – Results in Silos



This results in application silos...According to Gartner the major data management issue facing organizations is the proliferation of data silos....with the typical organization having to deal with hundreds of separate data silos.

Advances in technology....



Next-gen Applications Have Complex Requirements



Open Source Runs Better with Scale, Speed & Reliability



HDFS API POSIX, NFS HBase API JSON API Kafka API

Web-Scale Storage	Database	Event Streaming
MAPR-FS	MAPR-DB	MAPRSTREAMS
High Availability	Real Time	Unified Security
	Multi-tenancy	Disaster Recovery
		Global Namespace

Enterprise-Grade Platform Services

On-Premise, In the Cloud, Hybrid



A concrete Foundation with MapR Enterprise

Web-Scale Storage

MapR-FS

Database

MapR-DB

Event Streaming

MapR-Streams



Real Time



Data Protection



Disaster Recovery



Mirroring



Snapshots



Multi Tenancy



Performance



Replication



Self Healing



Scalability



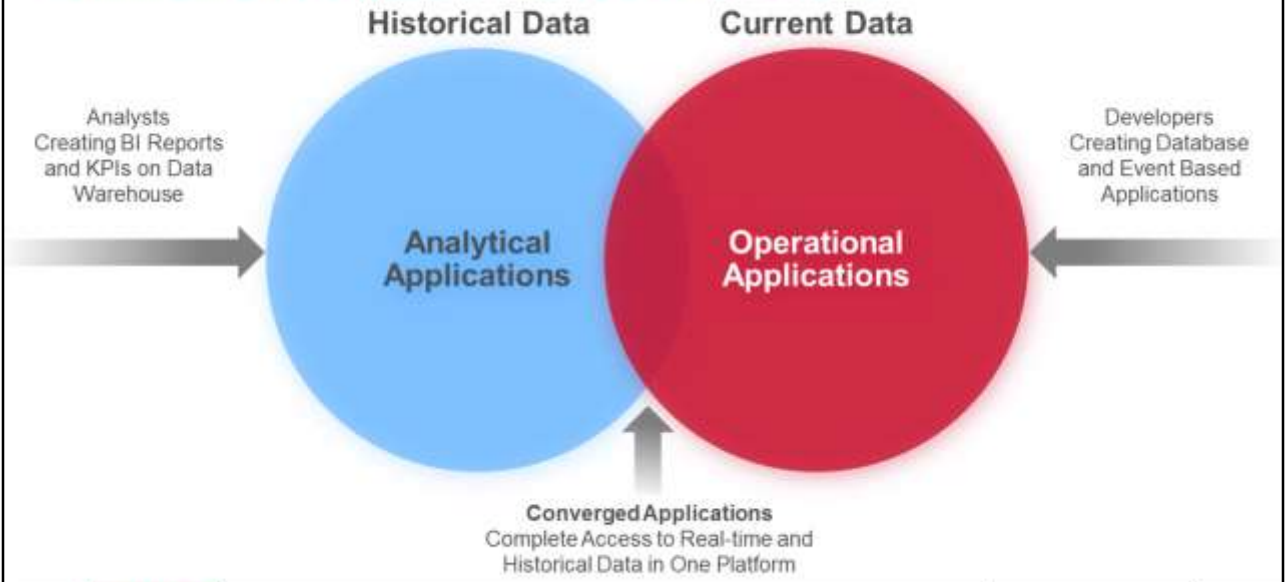
High Availability



Security



Traditional approach vs. Converged Approach



SAS with MapR

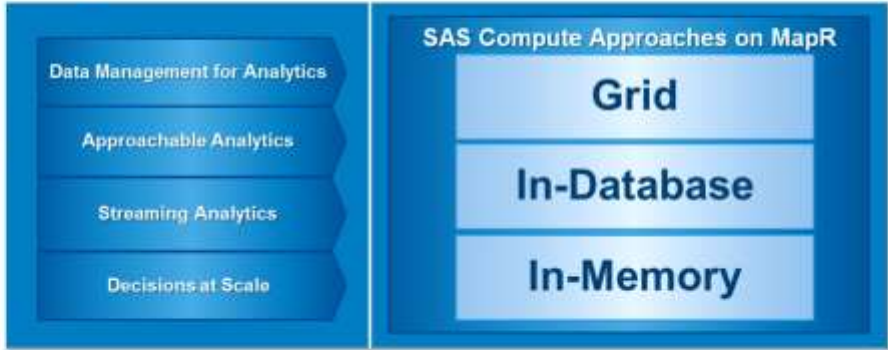


SAS® Data Loader
for Hadoop

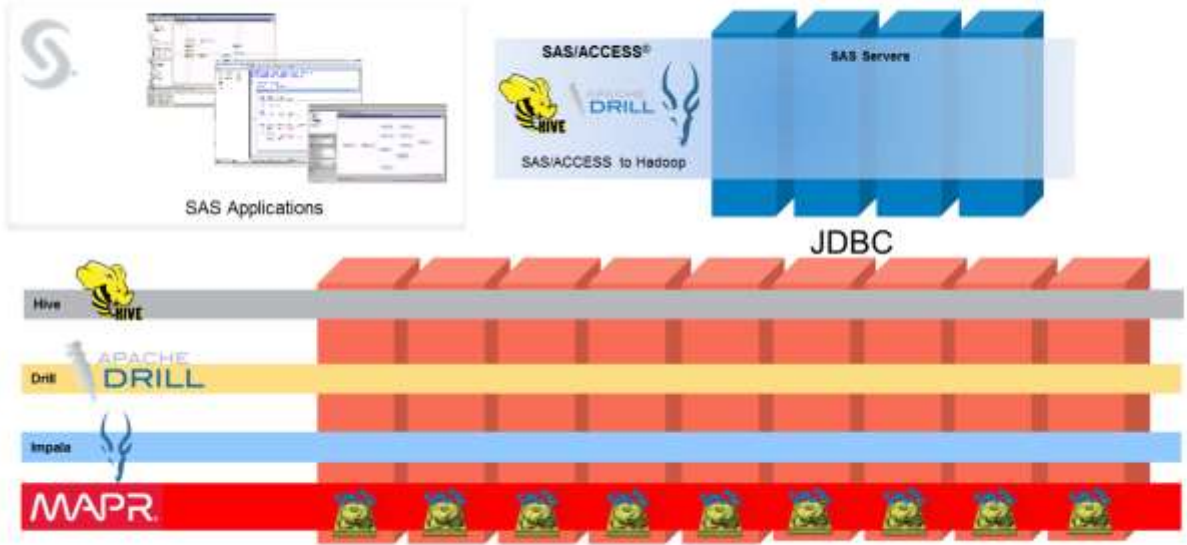
SAS® Visual Analytics

SAS® Event
Stream Processing

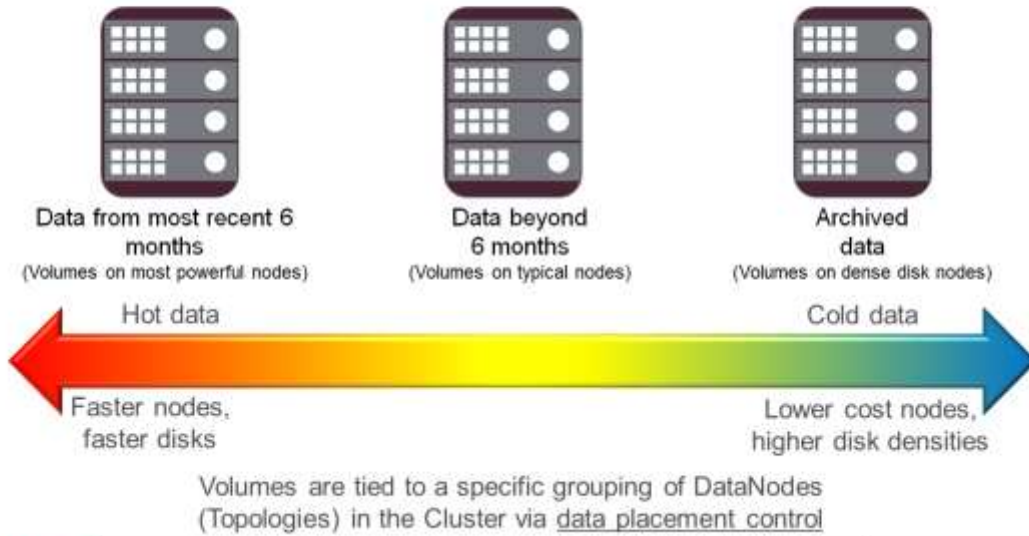
SAS® Analytics
SAS® Decision Manager



SAS & MapR Architecture




Data Placement Control for Multi-Temperature Data

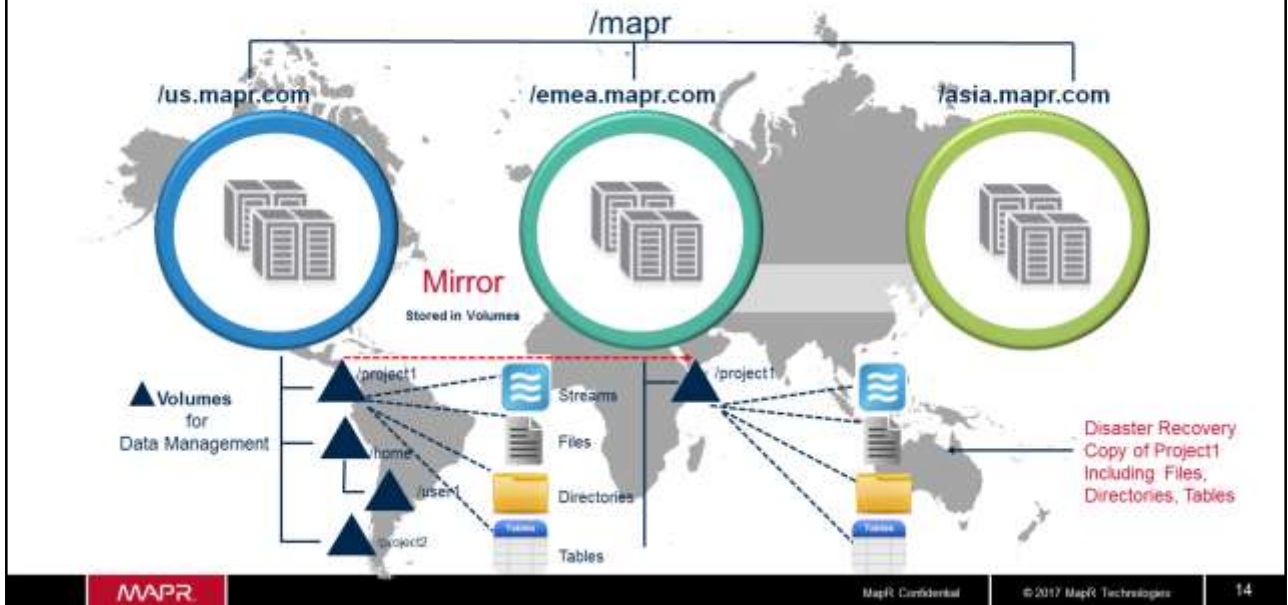




Industry's first Schema-free SQL query for Hadoop & NoSQL



MapR Global Namespace + Volume Mirroring



A consistent data tier leads to consistent, performant, and reliable functionality. Exhibit A: reliable, efficient, and secure data replication across the globe, for files, tables, documents, and streams.

Snapshots freeze your data at a moment in time. Instantaneous – no additional storage overhead unless data is modified. Provide backups against user error, and provide consistent, immutable views for higher-latency operations (such as mirroring). Applies to files, tables, documents and streams.

Mirrors maintain local or remote copies of a MapR volume. Based on a snapshot. Extremely efficient data transfer – only modified data is sent, as small as 8K blocks. Applies to files, tables, documents and streams.

MapR-DB also provides Active-Active table replication. MapR-Streams supports replication of event streams across the globe.

