



Enterprise Analytics Using EMC Isilon And Hortonworks



Partner Brief



Hortonworks Data Platform (HDP) together with EMC Isilon storage running the OneFS operating system provide a complete, tested, and robust solution encompassing enterprise grade Apache Hadoop and related projects to form a Data Lake for maximizing the value of your data. The testing for OneFS with Hortonworks is amongst the most rigorous in the industry, which helps ensure that EMC customers can seamlessly deploy an Enterprise Hadoop distribution for advanced analytics and diverse scale-out data processing workloads. The combination of HDP and Isilon shared storage helps organizations speed time to insights, improve storage utilization, eliminate islands or silos of storage, and lower storage management costs of migration, security and protection.

EMC Isilon running OneFS is the first and only scale-out NAS platform with native Hadoop Distributed File System (HDFS) support in addition to traditional Server Message Block (SMB), network file system (NFS), HTTP, and FTP protocols. This enables organizations to deploy one shared storage system that works across traditional and new and emerging workloads.

The Hortonworks Data Platform (HDP) enables Enterprise Hadoop by providing a complete set of essential Hadoop capabilities required by any enterprise in a completely open, integrated and tested platform that is ready for enterprise usage—and most important, with no proprietary holdbacks aimed at locking customers into HDP.

Together Isilon OneFS and HDP provide a powerful and highly efficient solution that combines massively scalable big data storage with robust analytics capabilities to accelerate business insight.

What Business Challenges are Solved?

Cost efficiency

EMC Isilon scale-out NAS systems running OneFS allows enterprises to achieve great cost efficiencies by investing in the storage capacity and capability required today—smaller than traditional or Hadoop requirements; scaled in smaller steps proportionate to data growth; data tiered according to value, and greatly simplified management.

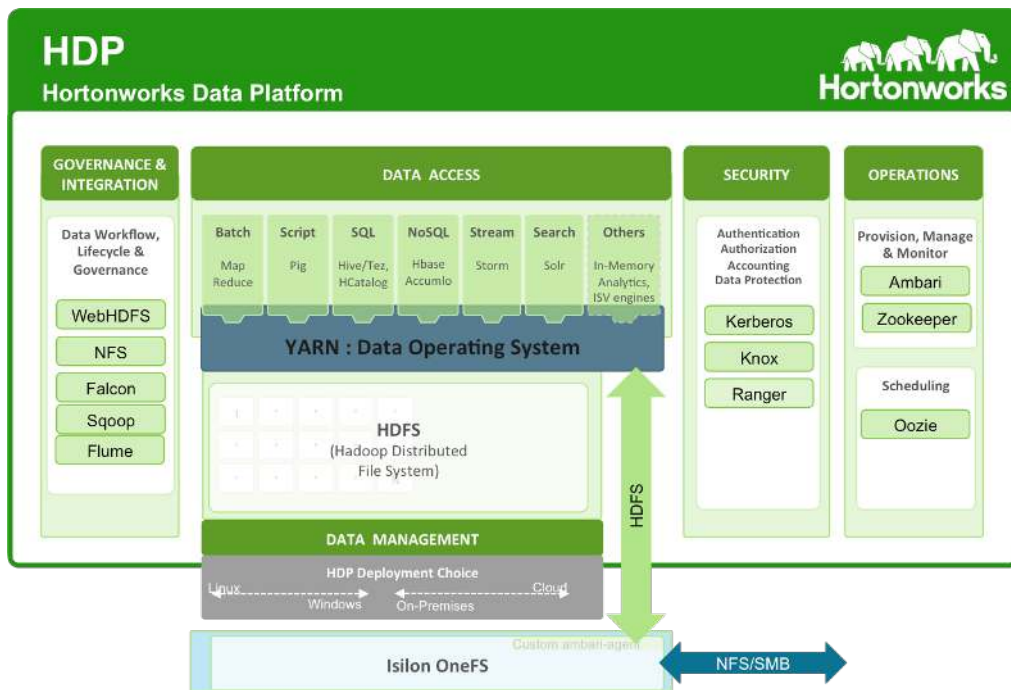
Protect data assets

The Isilon OneFS solution provides highly resilient storage with unmatched data protection and availability. Isilon OneFS uses the proven Reed-Solomon erasure encoding algorithm rather than RAID to provide a level of data protection that goes far beyond traditional storage systems.

Faster time to insights

Using Isilon NAS, enterprises can consolidate data from multiple islands into one single storage system based on Isilon OneFS. This is made possible through the multiple access mechanisms at the protocol level where data can be ingested into the Isilon store from a wide variety of sources and surfaced for use by others.

EMC Isilon in the Modern Data Architecture



Features and Benefits of the Combined Solution

- Accelerates insights and reduces costs through in-place analytics
- Secures and safeguards data with the appropriate level of protection
- Scales to meet the demands of future consolidation and growth
- Provides a tiering ability that enables organizations to manage their costs
- Maintains simplicity, even at the petabyte scale
- Enables multi-workload re-use of data
- Provides a data lake solution

EMC Isilon Scale-Out Data Lake

EMC Isilon combined with the Hortonworks Data Platform provides a complete and tested solution that helps serve as a rich repository of valuable insight that can help transform businesses so they can achieve significant advantages. The joint solution provides a powerful and highly efficient scale-out Data Lake that combines massive and robust enterprise-grade analytics capabilities.

EMC is a Hortonworks Certified Technology Partner

The Hortonworks Certified Technology Program reviews and certifies technologies for architectural best practices, validated against a comprehensive suite of integration test cases, benchmarked for scale under varied workloads and comprehensively documented.



Hortonworks. We do Hadoop.

Hortonworks is a leading commercial vendor of Apache Hadoop, the open source platform for storing, managing and analyzing Big Data. Hortonworks Data Platform, our distribution of Apache Hadoop, provides an open and stable foundation for enterprises and a growing ecosystem to build and deploy Big Data solutions.

Hortonworks is the trusted source for information on Hadoop, and together with the Apache community, Hortonworks is making Hadoop an enterprise data platform. Hortonworks provides unmatched technical support, training and certification programs for enterprises, systems integrators and technology vendors.

For additional questions, contact:

- EMC
www.emc.com
www.emc.com/big-data
- Hortonworks
www.hortonworks.com
(855) 8-HORTON