



Partner Brief

Virtualization with VMware Brings Huge Benefits to Hortonworks Data Platform



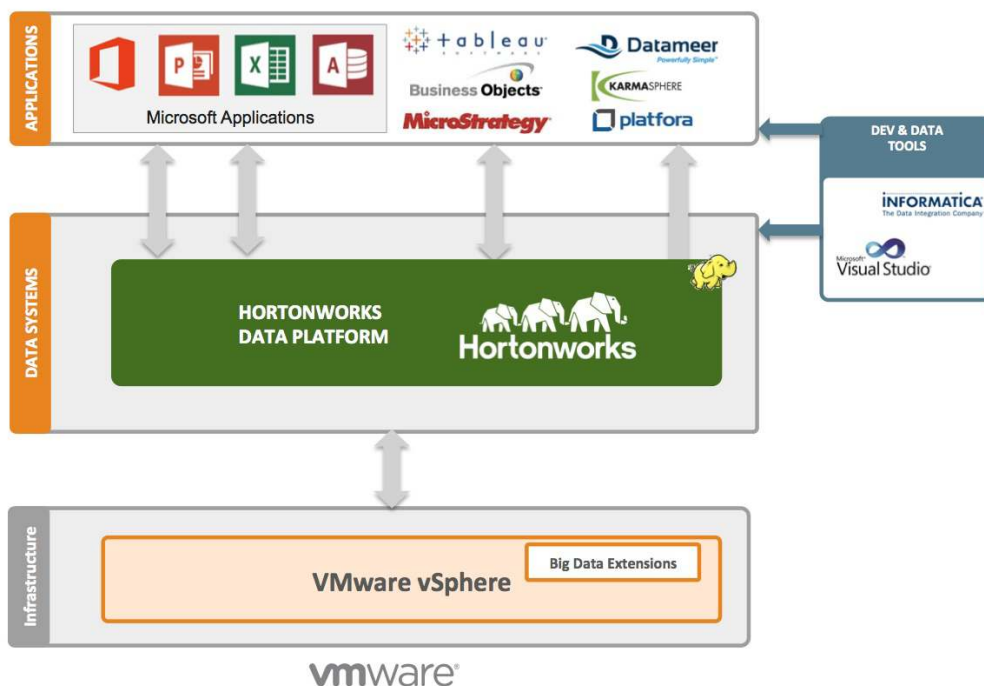
Two big forces are at work in the software industry today: the adoption of Big Data and the move to virtualizing everything in the data center, including every application and infrastructure component.

With the momentum behind Big Data and the use of the Hadoop architecture rapidly increasing, users are now seeking the best approaches for deploying Hadoop application clusters into their data centers. We believe the best way to do this is using the already-proven VMware vSphere virtualization platform as a basis for your Hadoop clusters.

Many customers have already used VMware to optimize and get the best utilization from their infrastructure and further move towards a software defined data center. Enterprises are also deploying Hadoop in the cloud, and virtualizing is a key step along that path also. This brief summarizes the main benefits you gain from deploying your Hortonworks Modern Data Architecture on VMware vSphere.

The Hortonworks and VMware teams have worked together to validate the HDP platform on the vSphere environment. The HDP platform is deployed onto virtual machines using a combination of the Hortonworks Ambari and the VMware vSphere technologies. Performance tests have shown that the behavior of Hadoop applications comes close to that seen on a native platform – and with suitable configuration changes, can in some cases exceed it. VMware's Big Data Extensions add specific optimizations for deploying on the virtual infrastructure.

Modern Data Architecture



Features & Benefits of Combined Solution

- **Rapid provisioning of your HDP clusters on to virtual machines for speed to market**
- **Optimal utilization of your infrastructure – shared with other workloads. Multiple versions running in parallel**
- **Multi-tenancy at the workload level – separation into contained resource pools**
- **Scaling of the compute layer independently of the data layer**
- **Best practices for node placement and replica handling in a virtualized environment**

VMware vSphere is the deployment platform for your HDP architecture, giving the flexibility to deploy multiple clusters containing different types of workload together onto a set of expandable, shared resources. vSphere provides the infrastructure for hosting virtual machines with Hortonworks Data Platform.

What Business Challenges are Solved?

- **Utilization and Costs**

Enterprises have recognized the cost benefits of virtualizing in their data center for consolidation reasons. However, many dedicated Hadoop clusters operate today on infrastructure that is not fully utilized. By virtualizing your HDP and Hadoop applications you can make better use of available hardware by sharing it across workloads in a controlled fashion.

- **Time to Market**

Provisioning a Hadoop cluster of many nodes for business applications use can be a time-consuming and error-prone task. By templating your configuration in a virtual machine and stamping out images of that template in an automated way, you gain time-to-market advantages and correctness at once.

- **Economy of Scale**

Scaling up or down is a key benefit with Hadoop and virtualization. Doing so with existing infrastructure presents a significant opportunity to businesses.

Industry Successes

VMware delivers virtualization solutions to over 300,000 customers worldwide. VMware's vSphere suite helps customers make best use of their data infrastructure and enables customers to adopt a cloud-based approach to deploying applications. VMware has been chosen by some of its largest customers for deployment of Hadoop workloads.

VMware is a Certified Technology Partner

VMware is a Certified Technology Partner on Hortonworks Data Platform (HDP). 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.

For additional questions, contact:



VMware Inc.,
Sales Support:
+1 (888) 506-2247
www.vmware.com
Hortonworks
USA: (855) 8-HORTON (1 for sales)
www.hortonworks.com

About Hortonworks

Hortonworks develops, distributes and supports the only 100-percent open source distribution of Apache Hadoop explicitly architected, built and tested for enterprise-grade deployments.



3460 West Bayshore Road
Palo Alto, CA 94303 USA

US: 1.855.846.7866
International: 1.408.916.4121
www.hortonworks.com