

# Accelerating Big Data Deployment in the Enterprise

## The Big Picture

BlueData Software, Inc. is the pioneer in Big Data private clouds. The company is democratizing Big Data by streamlining and simplifying Big Data infrastructure and eliminating complexity as a barrier to adoption. With its EPIC™ software platform, enterprises can now build agile, secure and cost-effective Big Data deployments that deliver value in days instead of months and at a cost savings of 50-75% compared to traditional approaches. With BlueData™, enterprises of all sizes can create a public cloud-like experience from their on-premise environments and get the same value out of their Big Data as companies like Google, Facebook and Yahoo at a fraction of the cost and with far fewer resources.

The BlueData platform incorporates many patent-pending virtualization enhancements for distributed data workloads that were previously barriers to the use of virtualization. BlueData delivers self service, speed and scale through:

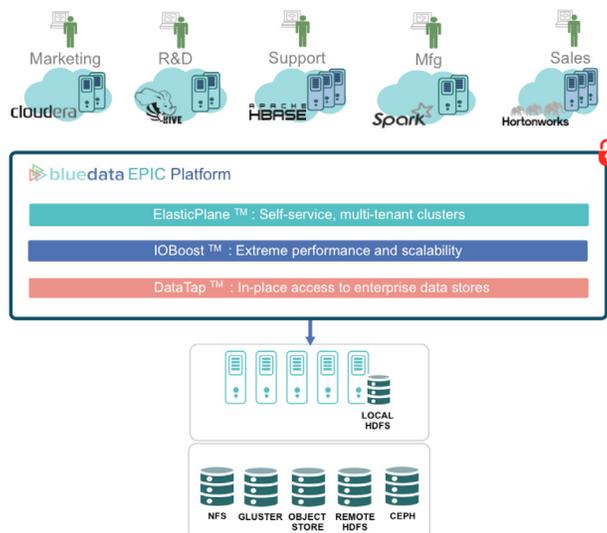
**IOBoost™** – application-aware caching service for extreme performance.

**DataTap™** – accelerates time to results by eliminating time-consuming data movement.

**ElasticPlane™** – virtualization and containers for self-service multi-tenant clusters.

## key benefits

- » **Increase business agility** by empowering users to create and manage Big Data clusters and nodes in matter of hours, not months.
- » **Maintain security and control** over multi-tenant Big Data infrastructure on-premise, integrating with your enterprise's security model (LDAP, Active Directory or Kerberos).
- » **Streamline operations** with automated provisioning, unified management, policy-based automation and push-button upgrades.
- » **Leverage automation and existing data center investments** using your existing shared storage infrastructure to run Big Data applications.
- » **Accelerate cost savings of up to 70%** compared to traditional hardware, storage, and operating costs while increasing server utilization by up to 350%.



BlueData lets organizations create secure private clouds for Big Data applications, with elasticity, multi-tenancy, enterprise security and management simplicity.

## The Private Cloud Infrastructure for Big Data

BlueData is a software solution that gives Hadoop Big Data deployments the benefits of virtualization and private clouds— efficiency, utilization, multi-tenancy, elasticity and policy based automation—without compromising on performance. BlueData dramatically reduces deployment complexity and improves business agility.

Using BlueData, enterprises can now derive value from their growing data sets within days instead of months, at a fraction of the cost of current methods. Virtualization increases server utilization by allowing hundreds of concurrent workloads with automated, policy-based scheduling.

## Instant Clusters for Big Data Apps

BlueData provides an elastic self-service infrastructure for Big Data applications. Data scientists and application developers can create clusters on demand and execute jobs without worrying about underlying infrastructure. BlueData effectively allows users and data scientists to consume Big Data as a service within secure, private cloud environments.

BlueData installs as a cluster virtualization layer between the underlying physical infrastructure and the Big Data distributions and applications, transforming physical infrastructure into an agile, flexible platform. It delivers all of the benefits of private-cloud computing, including multi-tenancy and data isolation, elasticity, enterprise security and control. A secure web-based portal gives IT administrators insight and control over all clusters, jobs and infrastructure status. BlueData also automates routine tasks such as provisioning, updates and monitoring.

### key features

#### Rapid, Self-service Provisioning

Use BlueData to deploy and test applications in minutes with rapid, self-service provisioning. Users can create instant, custom clusters in the ElasticPlane™ cluster virtualization layer. Data scientists and development engineers use BlueData's self-service user portal for creating and managing clusters and jobs, reducing the IT overhead of managing disparate and growing Big Data needs within the enterprise. The IT and DevOps teams can use the administrative portal to monitor underlying resource usage, and manage users and access controls for clusters, applications and data storage.

#### Secure Multi-tenancy

Different groups in the enterprise can run Big Data jobs simultaneously on the same physical infrastructure. BlueData provides multi-tenancy and data isolation at the virtualization layer. The platform integrates with enterprise security and authentication mechanisms such as LDAP, Active Directory and Kerberos.

#### Storage Efficiency & Data Virtualization

Because BlueData DataTap™ can separate compute and storage infrastructure, enterprises don't have to make multiple copies of data for Big Data analysis. Sensitive data can remain within shared enterprise storage such as NFS, GlusterFS, CEPH, Swift or HDFS, without the cost and risks of creating and maintaining multiple copies.

#### Policies and Resource Management

A policy engine defines service levels based on priority and automates resource management based on tenant and application needs. Administrators can define user groups and assign policies to restrict access to jobs, data or clusters based on departments or roles.

#### Run Multiple Apps on the Same Hardware

Users may create virtual clusters running multiple versions of a Hadoop distribution or different Hadoop distributions on the same physical cluster. Enterprises can then evaluate their options on an apples-to-apples basis, reducing the need for—and cost of—bare-metal resources.

#### Scalable Performance

BlueData applies IOBoost™, patented I/O optimization technologies to provide the benefits of virtualization while delivering near bare-metal performance. Application-aware caching and elastic resource management adapts dynamically to changing workload and application requirements to deliver the best performance on a virtualization platform.