

Turnkey Hunk™ with Pre-Integrated Hadoop

Simplify and streamline archival and search of historical Splunk data

SOLUTION HIGHLIGHTS

- ▶ **Turnkey solution for Hunk (Splunk Analytics for Hadoop):** Search and analyze your real-time data in Splunk and historical data in Hadoop using the same familiar Splunk Enterprise interface
- ▶ **Self-service, multi-tenant Hunk deployment:** Offers a purpose-built multi-tenant solution that provides rapid creation of multiple Hunk environments on shared infrastructure
- ▶ **Business agility and faster time-to-value:** Empowers users to self-service provision Hunk pre-integrated with an elastic Hadoop cluster in minutes versus weeks or months
- ▶ **Flexibility and choice:** Provides choice of major open source Hadoop distributions as well as remote shared storage such as NFS or HDFS for archival and search of Splunk Enterprise data
- ▶ **Lowest infrastructure TCO:** Delivers the elastic scaling and utilization benefits of containers with system performance comparable to that of bare-metal solutions at a fraction of the cost

Big Data Infrastructure Software

BlueData™ provides a software platform that makes deployment of Big Data infrastructure easier, faster, and more cost effective—eliminating complexity as a barrier to adoption.

The BlueData EPIC™ software platform solves the infrastructure challenges and limitations that can slow down and stall Big Data deployments. BlueData incorporates patent-pending innovations for running distributed data workloads in a virtualized environment, leveraging containers as nodes for Hadoop clusters. Using EPIC software, you can deploy Hadoop clusters in minutes rather than months, with the analytical and search tools that your users need.

BlueData delivers self-service, speed, and scale through:

- **ElasticPlane:** Enables users to deploy virtual clusters on demand in a secure, multi-tenant environment
- **IOBoost:** Provides application-aware caching for extreme performance and scalability with virtualization
- **DataTap:** Accelerates time-to-value by supporting in-place Hadoop compute processing on remote NFS, HDFS and Object storage

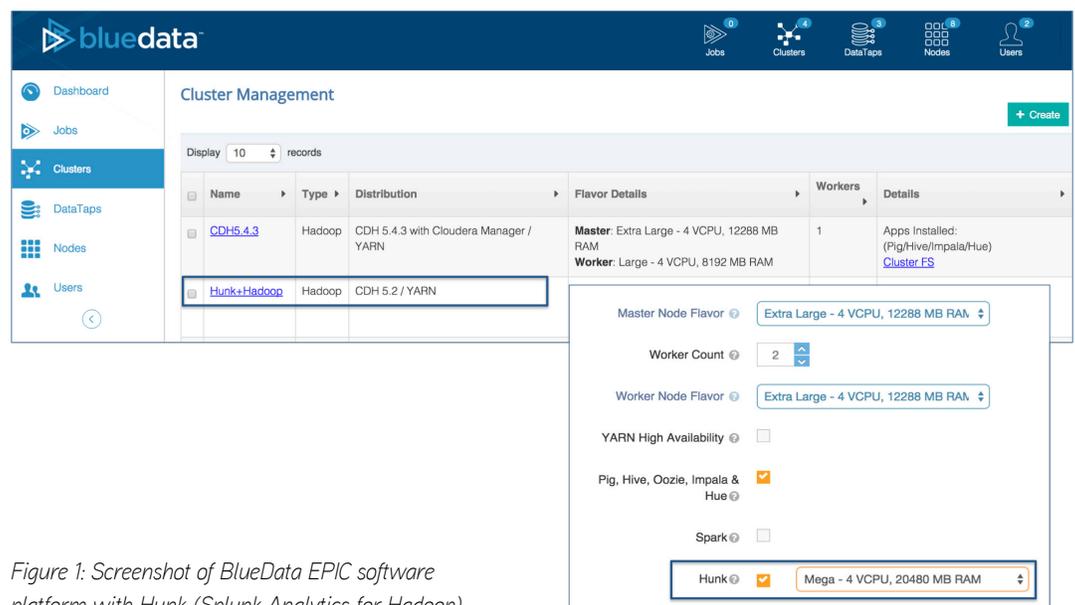


Figure 1: Screenshot of BlueData EPIC software platform with Hunk (Splunk Analytics for Hadoop)

Simplify Hadoop Deployment for Hunk

Today, most enterprise deployments for Hadoop and Big Data analytics require multiple, dedicated physical servers on-premises. This approach results in high management overhead, rigid infrastructure, and low hardware utilization. More importantly, it hinders the time to business value for applications such as Hunk (Splunk Analytics for Hadoop), which utilizes Hadoop to archive and search historical data.

With BlueData, enterprises can get a turnkey Hunk environment pre-integrated with a Hadoop cluster on shared infrastructure with exceptional simplicity, speed, efficiency, and performance:

- Fully automated deployment of Hunk with a Hadoop cluster in minutes, leveraging the power of containers
- Automatic, transparent integration of Hunk with the Hadoop cluster file system and resource manager as well as auto-creation of a default virtual index provider
- Flexibility for user groups or tenants to create multiple Hunk deployments with no administrative overhead
- Consolidate Big Data server, storage, and networking resources in a unified on-premises platform to improve data center utilization and reduce costs

The screenshot shows the 'Hunk with Hadoop' web interface. At the top, there are tabs for 'Node(s) Info', 'Job(s) Info', and 'Charts', along with a 'Cluster FS Browser' link. Below the tabs is a 'Node List' section with a 'Display 10 records' dropdown and a search box. A table lists nodes with columns for Node Name, Distribution, Role, Public IP, and Process List. The row for 'bluedata-4' is highlighted, showing it is a 'Hunk for CDH 5.2' edge node with IP '10.39.251.5' and processes 'Splunk Dashboard'. To the right, a 'Hadoop Cluster Information' panel shows configuration details: Hadoop Version (Hadoop 2.x, Yarn), File System (dtap://bluedata-defaultfs-1), Resource Manager Address (bluedata-1.bdlocal:8032), and Resource Scheduler Address (bluedata-1.bdlocal:8030). There is also a 'Splunk Settings' section with 'HDFS Working Directory' set to /user/root/hunk.

Figure 2: BlueData transparently provisions and configures Hunk with a Hadoop cluster, including a default virtual index provider

BENEFITS

Rapid, simplified deployment: BlueData pre-integrates Hunk software image for turnkey deployment with a Hadoop cluster in minutes vs. months (Figure 1)

Easy, ready-to-run for faster time-to-value: BlueData provides a link to the Hunk web app and pre-creates a virtual index against the cluster file system to search & archive (Figure 2)

Simple, point-and-click scalability: With BlueData, the Hunk and Hadoop clusters (as well as the underlying infrastructure) can be scaled to meet your search / archival SLAs without requiring downtimes or complex procedures

FEATURES

App-centric elastic cluster creation: Any authorized end user can get a fully configured Hunk environment with just a few inputs on a simple Web-based form (such as the Hadoop distribution and number/size of worker nodes)

Automated container orchestration and configuration: Hadoop settings includes the resource manager and file system in Hunk, which are auto-configured by BlueData as part of the fully automated container node creation and orchestration

Web-based GUI to scale horizontally: Expanding a running Hadoop cluster linked to Hunk is simple as increasing the worker count on the cluster creation form. BlueData CPU and memory can be increased by adding IP addresses of new machines

Try BlueData for Splunk at www.bluedata.com/splunk