

CLOUDERA MANAGER

End-To-End Administration for Your Enterprise Data Hub



Simplify Administration and Reduce Costs

MANAGE

_ Easily deploy, configure, and operate clusters with centralized, intuitive administration for all services, hosts, and workflows

MONITOR

_ Maintain a central view of all activity in the cluster through heatmaps, proactive health checks, and alerts

DIAGNOSE

_ Easily diagnose and resolve issues with operational reports and dashboards, events, intuitive log viewing and search, audit trails, and integration with Cloudera Support

INTEGRATE

_ Directly manage key third-party applications with Cloudera Manager and integrate with existing enterprise tools through SNMP, SMTP, and a comprehensive API



“We use Cloudera Manager on a daily basis to manage and get insights into our cluster. It’s been a game-changer for us.”

Jim Benedetto, Chief Technology Officer, Gravity

Cloudera Manager is the best-in-class holistic interface that provides end-to-end system management and key enterprise features to deliver granular visibility into and control over every part of an enterprise data hub. It is the only enterprise-grade Hadoop management application available - empowering operators to improve cluster performance, enhance quality of service, increase compliance, and reduce administrative costs.

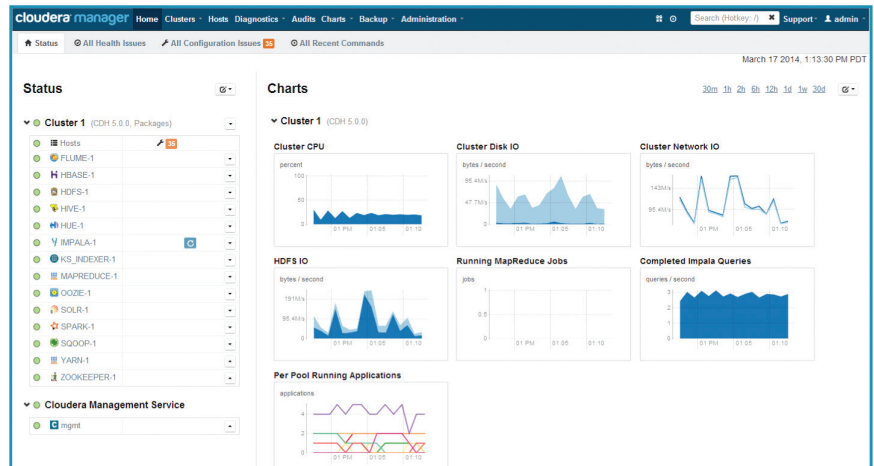
Manage Hadoop Clusters Easily and Effectively

As with any distributed computing or storage platform, deployment and ongoing administration of a Hadoop cluster can be difficult and time-consuming. Additionally, Hadoop consists of dozens of components and versions that can be deployed based on use cases; assigned roles for nodes; effectively configured and managed across the cluster; and diagnosed to optimize cluster performance. This can require significant expertise and constant attention without a holistic management tool. Cloudera Manager is designed make administration of Hadoop simple and straightforward, so you can focus on what’s most important to your business. With Cloudera Manager, you get:

- _ **Best-in-class Interface** - The only enterprise-grade Hadoop management application, featuring key enterprise features like zero-downtime rolling upgrades, backup and disaster recovery, and direct integration with Cloudera Support for faster time to resolution
- _ **Simple Efficiency** - One interface to easily manage the complexity of dozens of Hadoop components, which provides context around your data and the Hadoop system, and simplifies workflows to make administrators more productive
- _ **Broad Third-Party Integration** - Broadest network of partners allows you to easily integrate Hadoop and Cloudera into your existing systems, while removing the complexity of having to manage them all separately, to ensure the full functionality of an enterprise data hub

Cloudera Manager is available as part of both Cloudera Express and Cloudera Enterprise.

With Cloudera Express, you get a full set of functionality to deploy, configure, monitor, diagnose, safeguard, and scale your cluster—the most comprehensive and advanced set of management capabilities available from any vendor. When you upgrade to Cloudera Enterprise, you’ll get additional enterprise capabilities for integration with Cloudera’s partner network, zero-downtime rolling upgrades, backup and disaster recovery, and direct access to Cloudera Support, necessary for mission-critical applications.



Key Features of Cloudera Manager

FEATURE	DESCRIPTION	CLOUDERA EXPRESS	CLOUDERA ENTERPRISE
Manage			
Automated Deployment and Hadoop Readiness Checks	Install the complete CDH stack in minutes and ensure optimal settings	•	•
Service Management	Configure and manage all CDH services, including Impala and Search, from a central interface	•	•
Security Management	Configure and manage security across the cluster – including Kerberos authentication and role-based (administrator and read-only) administration	•	•
Resource Management	Allocate cluster resources by workload or by user/group/application to eliminate contention and ensure Quality-of-Service (QoS)	•	•
High Availability	Easily configure and manage High Availability for various services like HDFS, MapReduce, Oozie, YARN, HBase	•	•
Client Configuration Management	Centrally configure all client access to the cluster	•	•
Node and Cluster Templates	Easily deploy and expand heterogeneous clusters by creating templates for node roles and clusters	•	•
Comprehensive Workflows	Perform end-to-end tasks such as start/stop/restart clusters, services and roles, add/delete hosts, decommission nodes, etc.	•	•
Multi-Cluster Management	Manage multiple CDH clusters from a single instance of Cloudera Manager	•	•
Monitor			
Service, Host, and Activity Monitoring	Get a consolidated, real-time view of the state of all services, hosts, and activities running in the cluster	•	•
Events and Alerts	Create, aggregate, and receive alerts on relevant Hadoop events pertaining to system health, log messages, user actions, and activities Set thresholds, and create custom alerts for metrics collected by Cloudera Manager	•	•
Diagnose			
Global Time Control	Correlate all views along a configurable timeline to simplify diagnosis	•	•
Proactive Health Checks	Monitor dozens of service performance metrics and get alerts when you approach critical thresholds	•	•
Heatmaps	Visualize health status and metrics across the cluster to quickly identify problem nodes and take action	•	•
Customizable Charts	Report and visualize key time-series metrics about services, roles, and hosts	•	•
Intelligent Log Management	Gather, view, and search Hadoop logs collected from across the cluster	•	•
Integrate			
Comprehensive API	Easily integrate Cloudera Manager with your existing enterprise-wide management and monitoring tools, including Isilon as an alternative storage backend	•	•
Third-Party Application Management	Deploy, manage, and monitor services for third-party applications running on the cluster (e.g. data integration tools, math/machine learning applications, non-CDH services, etc.)	•	•
Advanced Management Features (Enabled by subscription)			
Operational Report and Quota Management	Visualize current and historical disk usage; set user and group-based quotas; and track MapReduce, Impala, YARN, and HBase usage		•
Configuration History and Rollbacks	Maintain a trail of all actions and a complete record of configuration changes, including the ability to roll back to previous states		•
Cluster Utilization Reporting	Visibility into historical usage and efficiency across users, tenants, and applications		•
Rolling Updates	Stage service updates and restarts to portions of the cluster sequentially to minimize downtime when upgrading or updating your cluster		•
AD Kerberos Integration	Integrate directly with Active Directory to get started easily with Kerberos		•
Kerberos Wizard	Easily configure Kerberos and trigger automated workflows to secure clusters		•
Hadoop SSL Related Configs	Simplify configs and eliminate need for safety valves		•
LDAP/SAML Integration	Integrate user credentials with Active Directory and enable single sign-on (SSO) capabilities		•
SNMP Support	Send Hadoop-specific events and alerts to global monitoring tools as SNMP traps		•
Scheduled Diagnostics	Take a snapshot of the cluster state and automatically send it to Cloudera Support to assist with optimization and issue resolution		•
Automated Backup and Disaster Recovery	Centrally configure and manage snapshotting and replication workflows for HDFS, Hive, and HBase		•