

CLOUDERA ENTERPRISE

Cloudera Enterprise is the modern platform for machine learning and analytics optimized for the cloud.

One Platform. Many Applications.

DATA SCIENCE & ENGINEERING

Process, develop, and serve predictive models

DATA WAREHOUSE

The modern data warehouse for today, tomorrow, and beyond

OPERATIONAL DATABASE

Real-time insights for modern data-driven business

Deploy and Run Anywhere

MULTI-CLOUD PROVISIONING

Deploy and manage Cloudera Enterprise across AWS, Google Cloud Platform, Microsoft Azure, and private networks

HIGH-PERFORMANCE ANALYTICS

Run your analytic tool of choice against cloud-native object stores like Amazon S3

ELASTIC AND FLEXIBLE

Support transient clusters and grow and shrink as needed, or permanent clusters for long-running BI and operational jobs

AUTOMATED METERING AND BILLING

Spin up and terminate clusters, and only pay for what you need, when you need it

Many of the world's largest companies rely on Cloudera's multi-function, multi-environment platform to provide the foundation for their critical business value drivers - growing business, connecting products and services, and protecting business. Find out what makes Cloudera Enterprise different from other data platforms.

Enterprise grade

The scale and performance required for today's modern data workloads meets the security and governance demands of today's IT. Cloudera's modern platform makes it easy to bring more users - thousands of them - to petabytes of diverse data and provides industry-leading engines to process and query data, as well as develop and serve models quickly. The platform also provides several layers of fine-grained security and complete audibility for companies to prevent unauthorized data access and demonstrate accountability for actions taken.

Shared data experience

Eliminate costly and unnecessary application silos by bringing your data warehouse, data science, data engineering, and operational database workloads together on a single, integrated data platform. Cloudera SDX enables these diverse analytic processes to operate against a shared data catalog that preserves business context like security and governance policies and schema. This common services framework persists even in transient cloud environments and makes it easier for IT to set and enforce policies while enabling business access to self-service analytics.

Hybrid deployment

Work where and how it's most convenient, affordable, and effective. Cloudera Enterprise can read direct from and write direct to cloud object stores like Amazon S3 and Azure Data Lake (ADLS) as well as on-premises storage environments, or HDFS and Kudu on IaaS. This provides flexibility to work on the data that you want wherever it lives, with zero copies or moves. Cloudera also provides the most popular data warehouse and machine learning engines that can run on any compute resource for ultimate deployment flexibility. Our hybrid control means users can self-service via PaaS offering, or opt for more configurability and management via IaaS, private cloud, or on-premises.

Powerful open source

Cloudera develops and validates the best of open source innovations into one seamless, rock-solid platform. Key features include:

- _ In-Memory Data Processing: The longest and deepest experience with Apache Spark
- _ Fast Analytic SQL: The lowest latency and best concurrency for BI with Apache Impala
- _ Updatable Analytic Storage: The only storage for fast analytics on fast changing data with Apache Kudu
- _ Open Source Leadership: Constant open source development and curation, with the most rigorous testing, for trusted innovation

One Platform. Many Uses. Designed for Your Needs.

Cloudera Enterprise is available on a subscription basis in five editions, each designed for your specific needs. Essentials provides superior support and advanced management for core Apache Hadoop. We also offer editions designed around how you're using the platform: Data Science and Engineering for programmatic preparation and predictive modeling; Operational DB for online

applications and real-time serving; and Data Warehouse for BI and SQL analytics. The Enterprise Data Hub gives you everything you need to become information-driven, with complete use of the platform. All editions are available in your environment of choice, whether it's cloud, on-premises, or a hybrid deployment.

	FREE up to 100 nodes	ANNUAL SUBSCRIPTION (PER NODE) ELASTIC CLOUD PRICING ALSO AVAILABLE				
	Express	Essentials	Data Science & Engineering	Operational DB	Data Warehouse	Enterprise Data Hub
Open Source Apache Hadoop Distribution CDH (100% open source data platform, including Apache Hadoop)	•	•	•	•	•	•
Automated Cluster Management (Cloudera Manager) Core Features: Multi-Cluster Deployment & Management, Service Configuration & Management (including HA, cluster templates) Service, Host, and Job Monitoring (including health tests, health history, & charting); Kerberos (including Active Directory); Diagnostic Tools & Alerting; & a Comprehensive API Advanced Features: Fine-grained user roles and permissions for administrators, automated wire encryption (TLS) setup for CDH+CM; Operational Reporting; Multi-Tenant Quota Management; Cluster Utilization Reporting; Configuration History & Rollbacks; Rolling Updates & Service Restarts; SNMP Support; Support integration (including scheduled diagnostics & proactive maintenance); External Authentication (LDAP/SAML) & Automated Backup & Disaster Recovery	•	•	•	•	•	•
Hybrid Deployment & Management Cloudera Altus Director: Flexible deployment across cloud environments; On-demand cluster creation/termination; Elastic cluster sizing; Cluster templates & cloning; Kerberos authentication & HA workflows; Rollback support; Multi-environment cluster management & monitoring; Comprehensive API & clients; & Customization	•	•	•	•	•	•
Components & Services Covered by Cloudera Support Basic Hadoop (HDFS, YARN, MapReduce, Hive, Pig, HUE, Sentry, Flume, Sqoop, Cloudera Manager, Cloudera Altus Director)		•	•	•	•	•
Flexible data and stream processing (Apache Kafka and Apache Spark including Spark Streaming, MLlib, and Spark SQL)			•	•	• (live-on-Spark only)	•
Analytic SQL (Apache Impala)					•	•
Real-time analytics (Apache Kudu)				•	•	•
Cloudera Search (Apache Solr)			•	•		•
Online NoSQL (Apache HBase)				•		•
Active Data Optimization (Cloudera Navigator Optimizer)* *Optionally included at time of purchase or may be purchased separately					•	•
Governance & Data Management (Cloudera Navigator including auditing, lineage, discovery, & policy lifecycle management)			•	•	•	•
Encryption & Key Management (Cloudera Navigator Encrypt & Key Trustee)			•	•	•	•
Support Features Only with Cloudera: Dedicated global support team, proactive technical guidance, predictive issue analysis, comprehensive knowledge base, production solution guides, open source community advocacy		•	•	•	•	•
Commercial License Warranty		•	•	•	•	•
Indemnification: Get protection from litigation stemming from use of open source technology			•	•	•	•
8x5 or 24x7 Expert Support: Get direct access to Cloudera's dedicated team of experts to help you resolve issues quickly and optimize your environment with the latest best practices, straight from the source		•	•	•	•	•
Premium Support: 15-minute time-to-first response for critical issues available as additional purchase option			•	•	•	•
Maximum nodes across all customer environments	100	unlimited	unlimited	unlimited	unlimited	unlimited

A software license is required for the following components: Cloudera Altus Director, Cloudera Manager (advanced features), Cloudera Navigator, Cloudera Navigator Encrypt and Key Trustee, Cloudera Navigator Optimizer.



Cloudera, Inc. 395 Page Mill Road Palo Alto, CA 94306 USA cloudera.com

© 2018 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice. Cloudera_Datasheet_Enterprise_131