

Take full advantage of your Cloudera platform investment

# SMARTHEALTH: PLATFORM HEALTH CHECK FOR OPTIMAL PERFORMANCE



## Why choose Cloudera Professional Services?

- Sizeable Pool of Technical Experts -**  
 Cloudera’s services team is composed of over three hundred technical experts adept at designing, architecting and implementing big data solutions, including comprehensive migration planning and execution.
- Deep and Wide Knowledge Base -**  
 Cloudera has amassed an unparalleled knowledge base and delivery practices to deliver projects faster and more cost effectively.
- Experienced Delivering Over 1000 Big Data Projects -** Cloudera has delivered more than 1 000 projects and brings to bear valuable experience to help lower project risk and achieve success.
- Broad Range of Services -** Cloudera offers a broad range of services to help at every step of the journey, on all infrastructures, from on-premises to the cloud, and ranging from solution design, to implementation and production readiness
- Driven By Customer Success -** Cloudera is driven and measured by our customers’ success and will use every opportunity available to achieve their goals and objectives



Cloudera’s SMARTHEALTH offering delivers a comprehensive platform deployment and use case implementation health check, recommends optimizations and implements them to ensure peak performance and maximum value from enterprise data platform deployments.

### Exploit the Full Potential of Your Enterprise Data Platform

Our goal is to maximize the value of our customers’ data by aligning their information architecture to specific use cases and ensuring infrastructure outperforms standards at every stage of the data lifecycle. Enterprise data platforms deployed based on initial use case assumptions tend to perform sub-optimally as more use cases are added. Increased data volumes, additional processing, new users and security requires calibrating cluster configurations at periodic intervals for optimal performance and meeting SLAs. To remedy this, Cloudera offers a comprehensive cluster health check, developing recommendations and implementing them to ensure customers get the best value out of their deployments.

### Accelerate Time to Value for Your Big Data Solution

Cloudera offers the most capable technical experts and widest real-world experience to help you plan, pilot and migrate your big data solution to production quickly, painlessly, and with peak performance. No one knows the big data ecosystem better than Cloudera Solution Architects or has more skills configuring, securing, benchmarking, and deploying large-scale enterprise data hubs. Cloudera platforms, configured according to our own best practices, stand up faster, with less risk, and at a lower total cost.

### For more information

This document does not constitute an offer. For more information about Migration or any other Cloudera Professional Services offering, contact your Cloudera representative.

Learn more at [cloudera.com](https://cloudera.com)

### About Cloudera

At Cloudera, we believe that data can make what is impossible today, possible tomorrow. We empower people to transform complex data into clear and actionable insights. Cloudera delivers an enterprise data cloud for any data, anywhere, from the Edge to AI. Powered by the relentless innovation of the open source community, Cloudera advances digital transformation for the world's largest enterprises.

Learn more at [cloudera.com](https://cloudera.com)

## Benefit from Our Experience and Reduce Risk

Collectively, Cloudera Solutions Architects draw on the industry's most significant knowledge base, built over the course of delivering more than 1000 big data projects. Cloudera's combined knowledge consists of documented best practices and delivery tools aimed at reducing implementation time and costs to our customers. Cloudera has experience across all major industries, configuring our platforms to specific use cases in an optimal fashion to avoid downstream issues and potential adverse impacts.



A typical platform health check takes between one and three weeks, depending on the exact scope, and covers the following four areas:

#### Cluster and Component Review

- Current use case implementations and future plans
- Cluster hardware, OS, network and physical topology configurations
- Cluster design and component configurations including security
- Cluster resources, YARN queue/resource utilization and NameNode usage analysis
- Known cluster issues, log analysis and available patches/hotfixes
- Capacity planning, high availability, DR and backup/restore configurations

#### Data and Applications Review

- Overall solution design and application architecture for use cases
- Data lake schema, including data organization, data format and layout
- Data ingestion/ETL, data engineering, egress and analytics processes
- Sub optimal Hive/impala queries and analysis of problem jobs
- Existing SDLC processes including CI/CD and automation

#### Security and Governance

- Current design and implementation of data security access policies
- Data encryption with both data in motion and data at rest
- Data governance implementation including lineage and data tagging

#### Cluster Operations

- Review day-to-day cluster operation processes (incl. monitoring, alerting, log rotation)
- Enterprise Operations Integration related to escalation and access to knowledge base
- Assess operations team skills for a training/coaching and knowledge transfer plan
- DR, backup and recovery plan, capacity planning and patch/upgrade cadence
- Understand current issues to help build troubleshooting best practices
- Analysis of cross-team relationships, interactions and skills/readiness evaluation