



Solution Brief



cloudera®



Dell In-Memory Appliance for Cloudera Enterprise

Accelerate time to insights with interactive analytics

Key benefits

- Gain fast time-to-value with an appliance-based solution. Focus on driving innovation, not system setup and tuning.
- Leverage a preconfigured hardware and software stack designed for next-generation workloads that move beyond slower, disk-intensive MapReduce batch processing.
- Accelerate time-to-insights. The Dell solution incorporates Spark to run jobs up to 100X faster than MapReduce and up to 10X faster in-memory compared to disk.
- Operationalize analytics. Take your analysis and move to production with fast processing using Spark Streaming.
- Get a turnkey, flexible solution. Start now with an open, ready-to-use, secure solution that scales to meet business needs.

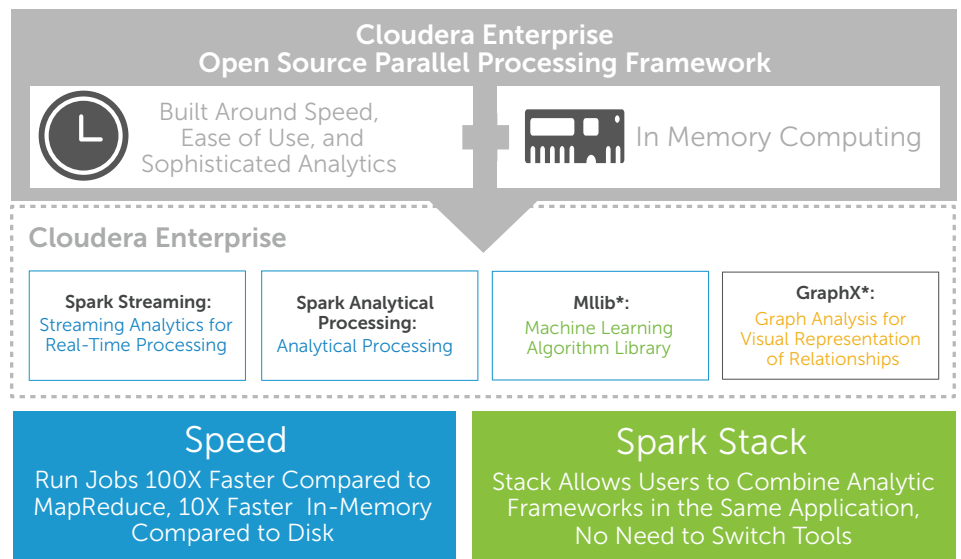
Across all industries and markets, data is the new currency and competitive differentiator. But data alone is never enough. To capitalize on this opportunity, and to protect your business and its customers, your organization needs solutions that allow you to quickly ingest, store, analyze, and build insights from Big Data.

That's the idea behind the Dell In-Memory Appliance for Cloudera Enterprise. Building on a deep engineering partnership among Dell, Intel, and Cloudera, this next-generation analytics solution solves the Big Data challenge with a turnkey, purpose-built, in-memory advanced analytics data platform.

Reduce analytics cycle from weeks and days to minutes and seconds

To enable fast analytics and stream processing, the Dell In-Memory Appliance for Cloudera Enterprise is bundled with Cloudera Enterprise, which includes Apache Spark. Cloudera Enterprise allows your business to implement powerful end-to-end analytic workflows, comprising batch data processing, interactive query, navigated search, deep data mining, and stream processing, all from a single common platform.

With a single common platform, there is no need to maintain separate systems—with separate data, metadata, security and management—that drive up complexity and cost.



*Cloudera Enterprise only Supports Spark Streaming

Figure 1. The Dell In-Memory Appliance is bundled with Apache Spark, an open source, parallel data processing framework that complements Apache Hadoop. Spark makes it easy to develop fast, unified applications that combine batch, streaming, and interactive analytics on all your data.

Deliver faster time-to-value with an appliance approach

With its appliance-based approach, the Dell solution simplifies and accelerates the otherwise complex process of creating large cluster deployments. Rather than focusing on building and deploying an analytics platform, your IT team can now spend more time helping the business gain fast, critical insights from huge amounts of data.

Better still, with the Dell In-Memory Appliance for Cloudera Enterprise you gain the agility and flexibility to scale up seamlessly as your analytics needs grow. The solution, built on the Dell™ PowerEdge™ two-socket Intel-based servers and 10G network connectivity, scales economically from an entry-level system to 48 nodes, all without the need to rip and replace technology.

Use cases

The Dell In-Memory Appliance for Cloudera Enterprise supports a wide range of use cases that help you accelerate time to insights to protect your business and its customers, capitalize on market opportunities, and maintain the availability of the critical systems.

Use cases.

- **Log aggregation and analysis:** Continually ingest log data from different data sources and process and analyze the data in a short window of time to accelerate insights and drive faster decision making.
- **Software application management and monitoring:** Provide end-to-end visibility across distributed infrastructures by troubleshooting application environments to monitor for performance degradation by tracing transactions across systems.
- **Security and compliance:** Provide rapid incident response with near real time correlation and in-depth monitoring across data sources. Conduct statistical analysis for advance pattern detection and threat defense.

- **Infrastructure and operations management:** Proactively monitor IT to ensure uptime by rapidly pinpointing and resolving problems. Identify infrastructure relationships, establish baselines and create analytics to report on service-level agreements (SLAs) or to track the SLAs of service providers.

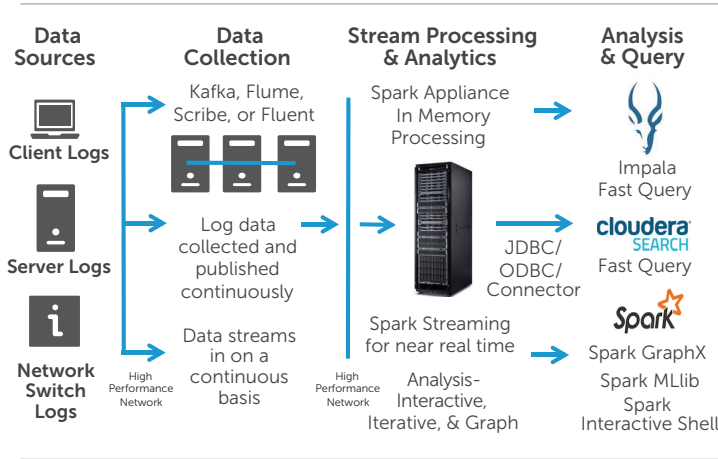


Figure 2. The Dell solution helps you gain an understanding of all the transactions in an environment by quickly aggregating and analyzing log data from diverse sources.

Get there today with the Dell solution

To increase your competitive edge, reduce costs, and protect your business and its customers, your organization needs to capitalize on Big Data in a manner that accelerates time to insights. Dell, together with Cloudera and Intel, helps you achieve this goal with a turnkey, purpose-built, in-memory advanced analytics data platform.

Dell In-Memory Appliance for Cloudera Enterprise Configurations at a glance		
Starter configuration	Mid-size configuration	Small enterprise configuration
<ul style="list-style-type: none">• 8 node cluster• 4 infrastructure nodes with passport on PoweEdge 2-socket servers• 4 data nodes with ProSupport on PowerEdge 2-socket servers• Cloudera Enterprise• 10G Force10 networking• Dell rack 42U• ~ 176 TB (disk raw space)	<ul style="list-style-type: none">• 16 node cluster• 4 infrastructure nodes with passport on PowerEdge 2-socket servers• 12 data nodes with ProSupport on PowerEdge 2-socket servers• Cloudera Enterprise• 10G Force10 networking• Dell rack 42U• ~ 528 TB (disk raw space)	<ul style="list-style-type: none">• 24 node cluster• 4 infrastructure nodes with passport on PowerEdge 2-socket servers• 20 data nodes with ProSupport on PowerEdge 2-socket servers• Cloudera Enterprise• 10G Force10 networking• Dell rack 42U• ~ 880 TB (disk raw space)

* Expansion Unit - PowerEdge - 4 data nodes with ProSupport, Cloudera Enterprise, Scale in Block

For a closer look at the Dell In-Memory Appliance for Cloudera Enterprise, visit DellBigData.com or Dell.com/Hadoop

Have a question? Please send it to us at Hadoop@Dell.com.

