



# Cloudera® Enterprise 5 Reference Architecture

A PSSC Labs Reference Architecture Guide  
December 2016



## Introduction

PSSC Labs continues to bring innovative compute server and cluster platforms to market. Focusing on specific applications where performance and reliability are critical, highlights PSSC Labs strengths. The Apache Hadoop data framework requires substantial compute and storage capabilities coupled tightly together. Cloudera® Enterprise 5 adds layers of unification, security and management.

PSSC Labs is the first manufacturer to design a server platform specifically for Hadoop. Introducing the World's highest density, lowest power consuming, Enterprise Ready Big Data server platform designed specifically for Hadoop workloads. The CloudOOP 12000 offers the absolute highest possible compute and storage density combined with high performance Data IO throughput. PSSC Labs already delivers the CloudOOP 12000 for small POC to large production clusters spanning hundreds of nodes.



### Key Features:

- Cloudera Enterprise 5 Certified
- Reduce Data Center Footprint By 50%
- Reduce Power Consumption By 40%
- Nearly 50% Greater Data IO Rates
- Patent Pending Tool Free Maintenance

### Technical Specifications:

- Up to 12 x 3.5" SATAIII or 14 x SSDs in 1U Rack Space
  - 96 TB using twelve x 8 TB Hard Drives
- Supports UP to 2 x Intel® Xeon® E5 Series Processors
- Supports Up to 256 GB ECC Enterprise Memory
- 10 GigE, 40 GigE and Infiniband Network Support
- Red Hat®, CentOS, Ubuntu, MS Windows® Compatible



"Certification of the CloudOOP 12000 on Cloudera Enterprise 5 ensures customers have the most innovative data management technology available."

Tim Stevens  
Cloudera®, Vice President





## CloudOOP 12000 Sample Configurations

Every organization or use case requires different computing needs. The CloudOOP 12000 offers the greatest flexibility possible. Below are three different proposed architectures for different workloads: high density storage, high computational requirements and a balanced configuration.

### CloudOOP 12000 High Density

- 72 TB Total Storage
- 16 Xeon E5 Processor Cores
- 128 GB ECC Memory
- 2 x 10GigE Network Ports
- 2 x GigE Network Ports
- Remote IPMI Management
- CentOS Linux OS
- Power Draw Estimate  
215 Watt Idle / 275 Watt Full Load

### CloudOOP 12000 High Compute

- 24 TB Total Storage
- 28 Xeon E5 Processor Cores
- 256 GB ECC Memory
- 2 x 10GigE Network Ports
- 2 x GigE Network Ports
- Remote IPMI Management
- CentOS Linux OS
- Power Draw Estimate  
265 Watt Idle / 345 Watt Full Load

### CloudOOP 12000 Balanced

- 48 TB Total Storage
- 20 Xeon E5 Processor Cores
- 256 GB ECC Memory
- 2 x 10GigE Network Ports
- 2 x GigE Network Ports
- Remote IPMI Management
- CentOS Linux OS
- Power Draw Estimate  
230 Watt Idle / 315 Watt Full Load

A Sample of Organizations Currently Using PSSC Labs for Cloudera Enterprise 5





## CloudRAX : Cloudera Enterprise 5 Validated Turn Key Cluster Solutions

PSSC Labs offers a complete, turn-key cluster that is ready to run on delivery. With over 2000+ HPC Clusters delivered to date, PSSC Labs understands everything that is necessary for a successful deployment. All necessary hardware including servers, network equipment, power and infrastructure are included. PSSC Labs Cluster Engineers preconfigure network, storage, operating system and BIOS settings to the end user's specifications. Cloudera Enterprise 5 is installed at PSSC Labs factory. The final step is the running of sample data sets to ensure proper functionality and performance. Each CloudRax Turn Key Cluster is custom configured for specific needs and budget.

Below is an overview of each different server platform PSSC Labs offers for Cloudera® Enterprise 5 turn-key deployments. Depending on the complexity of the environment, some software resources can be installed on different server platforms.

<b>CLOUDOOP 12000 DATA NODE</b>		
<b>Tech Specs</b>	<b>Key Features</b>	<b>Software Resource</b>
<ul style="list-style-type: none"> <li>○ 1U High Density Form Factor</li> <li>○ 2 x Intel® Xeon® E5 Processors</li> <li>○ 12 x SATAIII or SAS Hard Drives or 14 x SSDs</li> <li>○ 12 TB to 72 TB Storage Capacity</li> <li>○ 32 GB to 256 GB ECC Memory</li> <li>○ 2 x GigE Network Adapters</li> <li>○ Optional 10 GigE, 40 GigE, Infiniband Support</li> <li>○ Dedicated IPMI / iKVM</li> </ul>	<ul style="list-style-type: none"> <li>○ Enterprise Platform</li> <li>○ Redundant Power Supply</li> <li>○ Improved Data IO Throughput</li> <li>○ 40% Reduction in Power Consumption</li> <li>○ 2 x the Density of Standard Server</li> <li>○ Flexible Configuration Options</li> <li>○ 3 Year Warranty Included (24 x 7 x 365 NBD Available)</li> </ul>	<ul style="list-style-type: none"> <li>○ TaskTracker Daemon</li> <li>○ DataNode Daemon</li> <li>○ Zookeeper</li> <li>○ Hbase Region Server</li> <li>○ Impala Daemon</li> <li>○ Solr Servers</li> </ul>

<b>SURESTORE U2000 MANAGEMENT NODE (NAME NODE &amp; SECONDARY NAME NODE)</b>		
<b>Tech Specs</b>	<b>Key Features</b>	<b>Software Resources</b>
<ul style="list-style-type: none"> <li>○ 2 x Intel® Xeon® E5 Processors</li> <li>○ 1 TB to 24 TB SATA III, SAS, SSD Hard Drives</li> <li>○ 32 GB to 256 GB ECC Memory</li> <li>○ 2 x GigE Network Adapters</li> <li>○ Optional 10GigE, 40GigE, Infiniband Support</li> <li>○ Dedicated IPMI / iKVM</li> </ul>	<ul style="list-style-type: none"> <li>○ Enterprise Platform</li> <li>○ Redundant Power Supply</li> <li>○ Redundant Storage</li> <li>○ Raid Levels 0,1,5,6,10,50</li> <li>○ Flexible Configuration Options</li> <li>○ 3 Year Warranty Included (24 x 7 x 365 NBD Available)</li> </ul>	<ul style="list-style-type: none"> <li>○ Cloudera Manager</li> <li>○ Hbase Master</li> <li>○ JobTracker</li> <li>○ Standby JobTracker</li> <li>○ NameNode</li> <li>○ Standby NameNode</li> <li>○ JournalNodes</li> <li>○ Zookeeper</li> <li>○ Oozie</li> </ul>

<b>SURESTORE U1000 EDGE NODE</b>		
<b>Tech Specs</b>	<b>Key Features</b>	<b>Resources</b>
<ul style="list-style-type: none"> <li>○ 2 x Intel Xeon E5 Processors</li> <li>○ 1 TB to 12 TB SATAIII, SAS, SSD Hard Drives</li> <li>○ 32 GB to 256 GB ECC Memory</li> <li>○ 2 x GigE Network Adapters</li> <li>○ Optional 10GigE, 40GigE, Infiniband Support</li> <li>○ Dedicated IPMI / iKVM</li> </ul>	<ul style="list-style-type: none"> <li>○ Enterprise Platform</li> <li>○ Redundant Power Supply</li> <li>○ Redundant Storage</li> <li>○ Raid Levels 0,1,5,6,10,50</li> <li>○ Flexible Configuration Options</li> <li>○ 3 Year Warranty Included (24 x 7 x 365 NBD Available)</li> </ul>	<ul style="list-style-type: none"> <li>○ Hive command line client</li> <li>○ Hadoop command line client</li> <li>○ Impala command line client</li> <li>○ Flume Agents</li> <li>○ Hue Server</li> </ul>



## CloudRAX Turn-Key Cluster Sample Configurations

### CloudOOP Rax 150

- 150 TB Total Storage
- 1 Name Node
- 3 Data Nodes
- 48 Xeon E5 Processor Cores
- 196 GB ECC Memory
- 10GigE Network Backplane
- Remote IPMI Management
- CentOS Linux OS
- Cloudera Installation Service
- Cloudera Validation Service

### CloudOOP Rax 500

- 500 TB Total Storage
- 2 Name Nodes
- 1 Edge Node
- 10 Data Nodes
- 160 Xeon E5 Processor Cores
- 1280 GB ECC Memory
- 10GigE Network Backplane
- Remote IPMI Management
- CentOS Linux OS
- Cloudera Installation Service
- Cloudera Validation Service
- Rack & Roll Service

### CloudOOP Rax 1500

- 1500 TB Total Storage
- 2 Name Nodes
- 1 Edge Node
- 30 Data Nodes
- 480 Xeon E5 Processor Cores
- 3840 GB ECC Memory
- 10GigE Network Backplane
- Remote IPMI Management
- CentOS Linux OS
- Cloudera Installation Service
- Cloudera Validation Service
- Rack & Roll Service



“We believe strongly in our ability to deliver the highest performance, highest reliability server platforms to Cloudera Enterprise end users. Our experience delivering clusters ranging from several hundred TBs to several dozen PBs ensures a successful Cloudera Enterprise deployment.”

**Larry Lesser**  
PSSC Labs, CTO



## Total Cost of Ownership Comparison

PSSC Labs goal is to offer solutions with the absolute lowest total cost of ownership. The below chart compares different server manufacturer's solution for a 1 Petabyte (raw) Hadoop environment. PSSC Labs CloudRax 1000 requires 50% less rack space and consumes 40% less power

	PSSC Labs CloudRAX 1000 for 1PB Total Storage Space	Dell Configuration for 1PB Total Storage	HP Configuration for 1PB Total Storage	Lenovo Configuration for 1PB Total Storage
Required Data Center Footprint	Single x 42U Rack	Two x 42U Rack	Two x 42U Rack	Two x 42U Rack
Power Consumption Estimate*	4300 Watts Total @ Idle 5500 Watts Total @ Load	5800 Watts Total @ Idle 8500 Watts Total @ Load	6000 Watts Total @ Idle 8800 Watts Total @ Load	5700 Watts Total @ Idle 8700 Watts Total @ Load
Required Power Circuits	Single x 30 Amp / 208V / Single Phase	Two x 30 Amp / 208V / Single Phase	Two x 30 Amp / 208V / Single Phase	Two x 30 Amp / 208V / Single Phase
Pre-installation and Validation of Cloudera Enterprise 5 at Factory	Yes. Cloudera Enterprise 5 preinstalled and tested.	No. Additional services and fees required.	No. Additional services and fees required.	No. Additional services and fees required.
Onsite Physical Installation	Yes. Cluster arrives pre-racked, cabled and labeled.	No. Additional services and fees required.	No. Additional services and fees required.	No. Additional services and fees required.
Cluster Management Training	Yes.	No. Additional services and fees required.	No. Additional services and fees required.	No. Additional services and fees required.
Dedicated Remote Monitoring Capabilities	Yes. IPMI 2.0 Network Standard	Yes.	Yes.	Yes.
Hardware Warranty	3 Year NBD Service Available.	3 Year NBD Service Available.	3 Year NBD Service Available.	3 Year NBD Service Available.

\*Dell, HP and Lenovo power estimates based on manufacturers website power draw estimates.