



HP Reference Architecture for Hortonworks Data Platform on HP ProLiant SL4540 Gen8 Server

HP Converged Infrastructure with the Hortonworks Data Platform for Apache Hadoop

Table of contents

- Executive summary 2
 - HP SL4500 Series 3
- Hortonworks Data Platform overview..... 3
- High-availability considerations..... 5
- Pre-deployment considerations 6
 - Operating system..... 6
 - Computation 6
 - Memory 6
 - Storage..... 7
 - Network..... 7
 - Switches 8
- HP Insight Cluster Management Utility..... 8
- Server selection 10
 - Management nodes..... 10
 - Worker nodes..... 14
- Reference Architectures 17
 - Single Rack configuration..... 17
 - Multi-Rack Cluster configuration..... 20
- HP Vertica and Hadoop 22
- Summary..... 22
- For more information 23

Executive summary

HP and Apache Hadoop allow you to derive new business insights from Big Data by providing a platform to store, manage and process data at scale. However, Apache Hadoop is complex to deploy, configure, manage and monitor. This white paper provides several performance optimized configurations for deploying Hortonworks Data Platform clusters of varying sizes on HP infrastructure that provide a significant reduction in complexity and increase in value and performance.

The configurations are based on the Hortonworks Data Platform (HDP) and the HP ProLiant SL4540 Gen8 server platform. The configurations reflected in this document have been jointly designed and developed by HP and Hortonworks to provide a Hadoop solution that efficiently balances performance, storage and cost.

The configuration is built to support varying Big Data workloads. HP ProLiant SL4540 servers pack maximum performance, productivity and cost effectiveness in ultra-dense configurations. In fact, recent testing in April 2013 shows that HP ProLiant SL4540 holds the world's fastest 10TB Hadoop TeraSort result.¹

The HP ProLiant SL4540 servers offer breakthrough economics, density, and simplicity – and support up to a maximum of 27 servers with a total of 1.62 PB raw capacity in a single expansion 42U rack. In effect, the SL4540 Gen8 server's 3x15 configuration provides customers 67% more storage capacity and 33% more compute density at a 42U rack level when compared to 2U rack mount servers deployed today.

Most 2U 12 LFF drive rack mount servers are deployed for Hadoop with 6 core processors in order to maintain the 1:1 core to spindle ratio. Because the SL4540 Gen8 server's 3x15 configuration can provide 15 LFF drives per server, customers can take advantage of the 8 core processors to drive Hadoop workloads and maintain the 1:1 core to spindle ratio needed. As a consequence, SL4540 Gen8 servers deployed as worker nodes provide the following savings compared to 2U 12-drive rack mount worker nodes:

- 40% less Space
- 47% less Energy
- 41% less Cables

HP Big Data solutions provide best-in-class performance and availability, with integrated software, services, infrastructure, and management – all delivered as one proven solution as described at hp.com/go/hadoop. In addition to the benefits described above, the solution in this white paper also includes the following features that are unique to HP:

- **For analytics database**, the HP Vertica connectors for Hadoop allow seamless integration of both structured and unstructured data providing end-to-end analytics thereby simplifying bi-directional data movement for Hadoop and reducing customer integration costs. HP Vertica is a leading real-time, scalable, analytical platform for structured data.
- **For networking**, the HP 5900 Switch Series is a family of high-density ultra-low latency top-of-rack (ToR) data center switches. HP 5900 switches are suited for deployment at the server access layer of large enterprise data centers or for deployment at the data center core layer of medium-sized enterprises. The HP 5900AF-48XG-4QSFP+ Switch has 48x 10-Gigabit SFP+ ports with four QSFP+ 40-Gigabit ports for ultra-high capacity connections. The HP 5900AF enables customers to scale their server-edge 10GbE ToR deployments to new heights with high-density 48x10GbE ports delivered in a 1U design. The high server port density is backed by 4x40GbE uplinks to ensure availability of needed bandwidth for demanding applications. Redundant power and replaceable fan trays deliver highest availability.
- **For servers**, the HP ProLiant DL360p Gen8 and SL4540 Gen8 servers include:
 - The HP Smart Array P420i controller which provides increased² I/O throughput performance resulting in a significant performance increase for I/O bound Hadoop workloads (a common use case) and the flexibility for the customer to choose the desired amount of resilience in the Hadoop Cluster with either JBOD or various RAID configurations.
 - DL360p Gen8 – Two sockets with fast 6 core processors (Intel® Xeon® E5-2667) and the Intel C600 Series Chipset, providing the performance required for fastest time to completion for CPU bound Hadoop workloads. The HP FlexibleLOM network technologies provide customers the flexibility to easily move between 1GbE and 10GbE network interconnects on the server.
 - SL4540 Gen8 (3x15) – 3 node configuration of the HP ProLiant SL4540 Gen8 with two sockets, 8 core processors and the Intel C600 Series Chipset consists of three compute nodes and a total of forty-five (45) large form factor (LFF) 3.5" hard disk drives (HDD) in the chassis. The HP ProLiant SL4540 is a dual socket Gen8 server, with a choice of five different Intel Xeon processors, up to 192GB of memory and one PCIe slot for expansion per node. Every compute node also has its own dedicated networking ports. The product also supports InfiniBand interconnect. For more details see the [quickspecs](#) for the product.

¹ The HP ProLiant SL4540 Cluster had 23% faster throughput per U than the previous benchmark HP achieved using an HP ProLiant DL380p Cluster. Both tests were performed within a single 42U rack.

² Compared to the previous generation of Smart Array controllers

- The HP iLO Management Engine on the servers is a comprehensive set of embedded management features, including HP Integrated Lights-Out 4 (iLO 4), Agentless Management, Active Health System, and Intelligent Provisioning which reduces node and cluster level administration costs for Hadoop.
- **For management**, HP Insight Cluster Management Utility (CMU) provides push-button scale out and provisioning with industry leading provisioning performance (deployment of 800 nodes in 30 minutes), reducing deployments from days to hours. In addition, CMU provides real-time and historical infrastructure and Hadoop monitoring with 3D visualizations allowing customers to easily characterize Hadoop workloads and cluster performance reducing complexity and improving system optimization leading to improved performance and reduced cost. HP Insight Management and HP Service Pack for ProLiant, allow for easy management of firmware and the server.

All of these features reflect HP's balanced building blocks of servers, storage and networking, along with integrated management software and bundled support.

In addition, this white paper has been created to assist in the rapid design and deployment of the Hortonworks Data Platform software on HP infrastructure for clusters of various sizes. It is also intended to concretely identify the software and hardware components required in a solution in order to simplify the procurement process. The recommended HP Software, HP ProLiant servers, and HP Networking switches and their respective configurations have been carefully tested with a variety of I/O, CPU, network, and memory bound workloads. The configurations included provide the best value for optimum MapReduce and HBase computational performance.

HP SL4500 Series

The HP SL4500 Series is an excellent choice for running Hadoop and includes the HP ProLiant SL4540 Gen8 server. This recommended configuration document will focus on the 3 node configuration of HP ProLiant SL4540 Gen8. It consists of three compute nodes and a total of forty-five (45) large form factor (LFF) 3.5" hard disk drives (HDD) in the chassis. Each server has exclusive access to fifteen (15) large form factor (LFF) 3.5" hard disk drives (HDD) in the storage section. The HP ProLiant SL4540 is a dual socket Gen8 server, with a choice of five different Intel Xeon processors, up to 192GB of memory and one PCIe slot for expansion per node. Every compute node also has its own dedicated networking ports.

The HP ProLiant SL4540 Gen8 server offers the advancements and benefits of the Gen8 platform:

- **HP SmartMemory:** memory performance improvements and enhanced functionality as listed for Gen8 in the HP SmartMemory QuickSpecs at: http://h18000.www1.hp.com/products/quickspecs/14225_div/14225_div.html.
- **HP Smart Array RAID controllers:** An HP Dynamic Smart Array B120i SATA RAID Controller is located on the server node and controls the two small form factor (SFF/2.5") hard drives on the server node itself, and an HP Smart Array P420i controller is located in the I/O Module at the back of the SL4540 enclosure and controls the large form factor (LFF/ 3.5") hard drives in the chassis.
- **HP Agentless Management:** With HP iLO Management Engine in every HP ProLiant Gen8 server, the base hardware monitoring and alerting capability is built into the system (running on the HP iLO chipset) and starts working the moment that a power cord and an Ethernet cable is connected to the server.
- **HP Intelligent Provisioning** includes everything needed for system setup and deployment. Firmware, drivers, and tools to get systems online are embedded on a NAND flash chip on the server motherboard (thus SmartStart CDs or firmware DVDs are not necessarily needed).

For more information see <http://h17007.www1.hp.com/us/en/whatsnew/proliantgen8/index.aspx>.

Target audience: This document is intended for decision makers, system and solution architects, system administrators and experienced users who are interested in reducing the time to design or purchase an HP and Hortonworks solution. An intermediate knowledge of the Hortonworks Data Platform and scale out infrastructure is recommended. Those already possessing expert knowledge about these topics may proceed directly to [Pre-deployment considerations](#).

Hortonworks Data Platform overview

Hortonworks is a major contributor to Apache Hadoop, the world's most popular big data platform. Hortonworks focuses on further accelerating the development and adoption of Apache Hadoop by making the software more robust and easier to consume for enterprises and more open and extensible for solution providers.

The Hortonworks Data Platform (HDP), powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner.

The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including Hortonworks MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, ZooKeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process; and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of their code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. Hortonworks sells only expert technical support, training and partner-enablement services. All of the technology is, and will remain, free and open source.

Hortonworks Data Platform (HDP) is the foundation for the next generation enterprise data architecture, one that deals with the volume and complexity of today's data. With HDP, you can now store, process and analyze data in any format and at any scale. Comprised of the essential Apache Hadoop components, HDP is 100% open source and includes everything you need to begin uncovering business insights from the quickly growing streams of data flowing into and throughout your business.

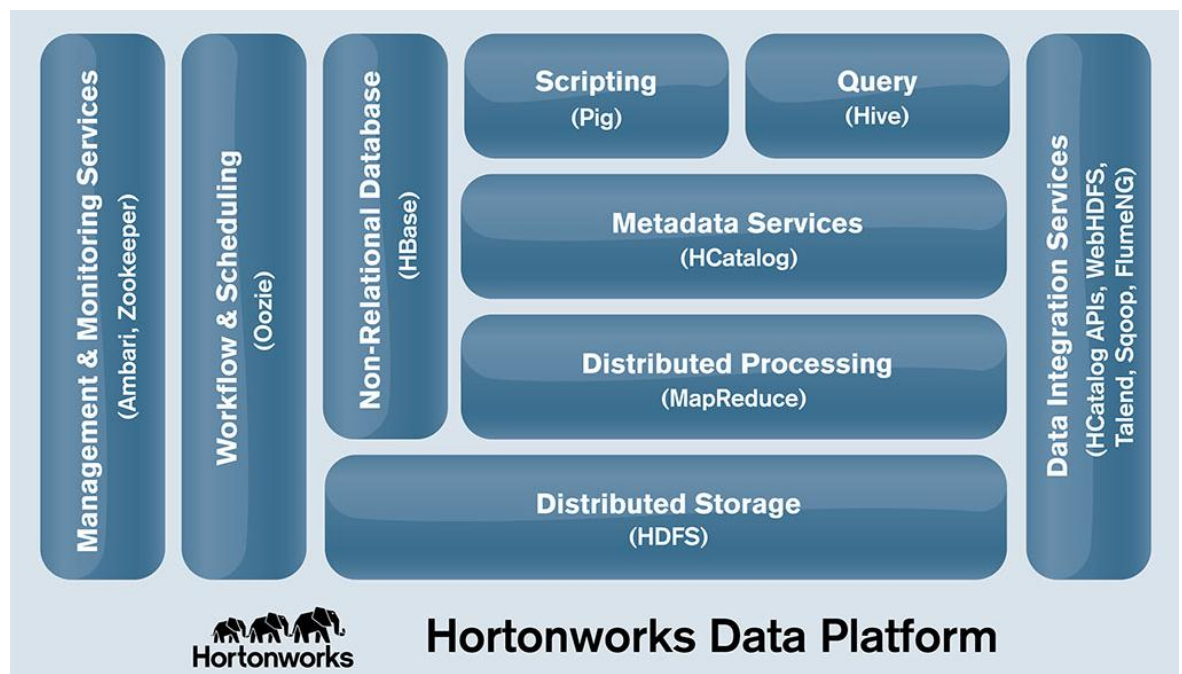
HDP combines the power and cost-effectiveness of Apache Hadoop with a number of other features critical for successfully deploying and managing Hadoop in enterprise environments, including:

- **Integrated and Tested Package** – HDP removes complexity and uncertainty by providing a comprehensive suite of stable and current versions of all essential Apache Hadoop components in an integrated and tested package.
- **Easy Installation** – HDP includes a simple, wizard-driven installation tool that makes Hadoop easy to install and provision across clusters of machines.
- **Management & Monitoring Services** – HDP includes Hortonworks Management Center, an open source and extensible management and monitoring tool with web-based dashboards that make it easy to monitor your clusters and create alerts.
- **Data Integration Services** – HDP integrates Talend Open Studio for Big Data, the leading open source data integration platform for Apache Hadoop. Included is a visual development environment and hundreds of pre-built connectors to leading applications that allow you to connect to any data source without writing code.
- **Centralized Metadata Services** – HDP includes HCatalog, a metadata and table management system that simplifies data sharing both between Hadoop applications running on the platform and between Hadoop and other enterprise data systems. HDP's open metadata infrastructure also enables deep integration with third-party tools.

Hortonworks Data Platform components

Hortonworks Data Platform is comprised of the essential and stable Apache Hadoop components needed for your big data projects, in an integrated and tested environment, and augmented with tools that help you to install, configure, provision and monitor your deployments. HDP includes the following components:

Figure 1. The Hortonworks Data Platform Overview



Typically, Hadoop clusters are used for batch MapReduce analysis of data or they are used to run HBase, which is an online distributed store for reading and writing structured data. The Hadoop Master services (NameNode, JobTracker and HBase Master) are installed and configured on HP DL360p head nodes. The Hadoop Worker services (DataNode, TaskTracker and HBaseRegionServer) are installed and configured on HP SL4540 Gen8 worker nodes so that they are co-located or co-deployed for optimal data locality.

The platform functions within the Hortonworks Data Platform (HDP) are provided by two key groups of services, namely the Management and Worker Services. Management Services manage the cluster and coordinate the jobs whereas Worker Services are responsible for the actual execution of work on the individual scale out nodes. The two tables below specify which services are management services and which services are workers services. Each table contains two columns. The first column is the description of the service and the second column specifies the number of nodes the service can be distributed to. The Reference Architectures (RAs) we provide in this document will map the Management and Worker Services onto HP infrastructure for clusters of varying sizes. The RAs factor in the scalability requirements for each service so this is not something you will need to manage.

Table 1. Hortonworks Data Platform Management Services

Service	Maximum Distribution across Nodes
Hortonworks Management Center	1
JobTracker	1
HBase Master	Varies
NameNode	1
Secondary NameNode	1

Table 2. Hortonworks Data Platform Worker Services

Service	Maximum Distribution across Nodes
DataNode	Most or all nodes
TaskTracker	Most or all nodes
HBaseRegionServer	Varies

High-availability considerations

The following are some of the high availability features considered in this reference configuration:

- OS availability and reliability – For the reliability of the server, the OS disk is configured in RAID 0+1 configuration thus preventing failure of the system from OS hard disk failures.
- Network reliability – The reference configuration uses two HP 5900 switches for redundancy, resiliency and scalability in using Intelligent Resilient Framework (IRF) bonding. We recommend using redundant power supplies.
- Power supply – To ensure the servers and racks have adequate power redundancy we recommend that each server have a backup power supply, and each rack have at least two Power Distribution Units (PDUs).

Pre-deployment considerations

There are a number of key factors you should consider prior to designing and deploying a Hadoop Cluster. The following subsections articulate the design decisions in creating a balanced baseline configuration for the reference architectures. The rationale provided includes the necessary information for you to take the configurations and modify them to suit a particular custom scenario.

Table 3. Overview of Functional Components and Configurable Value

Functional Component	Value
Operating System	Improves Availability and Reliability
Computation	Ability to balance Price with Performance
Memory	Ability to balance Price with Capacity and Performance
Storage	Ability to balance Price with Capacity and Performance
Network	Ability to balance Price with Performance and Availability

Operating system

Hortonworks supports 64-bit Red Hat Enterprise Linux (RHEL) 5.x and 6.x and CentOS 5.x and 6.x as choices for the operating system.

Note

HP recommends using a 64-bit operating system to avoid constraining the amount of memory that can be used on worker nodes. 64-bit Red Hat Enterprise Linux 6.1 or greater is recommended due to better ecosystem support, more comprehensive functionality for components such as RAID controllers and compatibility with HP Insight CMU. The Reference Architectures listed in this document were tested with 64-bit Red Hat Enterprise Linux 6.2.

Computation

The processing or computational capacity of a Hortonworks Data Platform (HDP) cluster is determined by the aggregate number of MapReduce slots available across all nodes. MapReduce slots are configured on a per server basis. Employing Hyper-Threading improves process scheduling, allowing you to configure more MapReduce slots. Refer to the *Storage* section to see how I/O performance issues arise from sub-optimal disk to core ratios (too many slots and too few disks). For CPU bound workloads we recommend buying processors with faster clock speeds to remove the bottleneck.

Note

Oracle Java JDK 6 (not JRE) is required to execute MapReduce tasks.

Memory

Use of Error Correcting Memory (ECC) is a practical requirement for Hortonworks Data Platform (HDP) and is standard on all HP ProLiant servers. Memory requirements differ between the management services and the worker services. For the worker services, sufficient memory is needed to manage the TaskTracker and DataNode services in addition to the sum of all the memory assigned to each of the MapReduce slots. If you have a memory bound MapReduce Job we recommend that you increase the amount of memory on all the nodes running worker services.

Best practice

It is important to saturate all the memory channels available to ensure optimal use of the memory bandwidth. For example, on a two socket processor with three memory channels that supports two DIMMs each for a total of six (6) DIMMs per installed processor or a grand total of twelve (12) DIMMs for the server, one would typically fully populate the channels with 8GB DIMMs resulting in a configuration of 96GB of memory per server.

In addition, the Hortonworks HDP cluster can also be used for HBase which is very memory intensive, which is why we recommend 96GB of RAM per server.

Storage

Fundamentally, Hortonworks Data Platform (HDP) is designed to achieve performance and scalability by moving the compute activity to the data. It does this by distributing the Hadoop jobs to nodes close to their data, ideally running the tasks against data on local disks.

Best practice

Given the architecture of Hadoop, the data storage requirements for the worker nodes are best met by direct attached storage (DAS) in a Just a Bunch of Disks (JBOD) configuration and not as DAS with RAID or Network Attached Storage (NAS).

There are several factors to consider and balance when determining the number of disks a node requires.

- **Storage capacity** – The number of disks and their corresponding storage capacity determines the total amount of the DataNode storage capacity for your cluster. We recommend Large Form Factor (3.5”) disks due to their lower cost and larger storage capacity compared with the Small Form Factor (2.5”) disks. See the best practice below for alternate recommendations for I/O bound workloads.
- **Redundancy** – Hortonworks Data Platform (HDP) ensures that a certain number of block copies are consistently available. This number is configurable as a volume property, which is typically set to three. If a Hortonworks Data Platform (HDP) node goes down, HDP will replicate the blocks that had been on that server onto other servers in the cluster to maintain the consistency of the number of block copies. For example, if the NICs (Network Interface Cards) on a server with 12 TB of data fail, that 12 TB of data will be replicated onto other servers, generating 12 TB of traffic within the cluster. The failure of a non-redundant ToR (Top of Rack) switch will generate even more replication traffic. Make sure your network, whether 1GbE or 10GbE, has the capacity to handle block replication for your server configurations in case of failure.
- **I/O performance** – Each node has a certain number of MapReduce slots available for processing Hadoop tasks. Each slot operates on one block of data at a time. The more disks you have, the less likely it is that you will have multiple tasks accessing a given disk at the same time and it will also result in more tasks being able to run against node-local data.

Best practice

Customers can choose to use SATA or SAS MDL (Nearline 7200 RPM) disks in a Hortonworks Data Platform (HDP) cluster. SAS MDL disks are preferred by customers who are looking for disks that are more efficient at driving higher I/O throughput across a larger amount of disks delivering to customers a higher performing cluster. SAS MDL disks are recommended over SAS ENT (10K and 15K SAS drives) disks in a Hortonworks Data Platform (HDP) cluster because the SAS ENT disks deliver only marginally higher performance at a higher cost with lower storage capacity.

Network

Configuring only a single Top of Rack (ToR) switch per rack introduces a single point of failure for each rack. In a multi-rack system such a failure will result in a flood of network traffic as Hadoop rebalances storage, and in a single-rack system such a failure brings down the whole cluster. Consequently, configuring two ToR switches per rack is recommended for all production configurations as it provides an additional measure of redundancy. This can be further improved by configuring link aggregation between the switches. The most desirable way to configure link aggregation is by bonding one link from a server to one ToR switch with a second link from the same server to the other ToR switch in its rack. When done properly, this allows the bandwidth of both links to be used. If either of the switches fail, the servers will still have full network functionality, but with the performance of only a single link. Not all switches have the ability to do link aggregation from individual servers to multiple switches, however, the HP 5900AF-48XG-4QSFP+ switch supports this through HP's Intelligent Resilient Framework (IRF) technology. In addition, switch failures can be further mitigated by incorporating dual power supplies for the switches. The Hortonworks Data Platform can be configured to be aware of the server and rack topology within the cluster. This is beneficial in that it allows Hadoop to take rack affinity into account when placing block data on Hadoop Worker nodes. Hortonworks recommends 1-2 10GbE NICs per node to satisfy the disk throughput requirements.

Larger clusters with three or more racks will benefit from having ToR switches connected by 10GbE uplinks to core aggregation switches. During the map phase of Hadoop jobs that utilize the HDFS, the majority of tasks reference data on the server that executes the task (node-local). For those tasks that must access data remotely, the data is usually on other servers in the same rack (rack-local). Only a small percentage of tasks need to access data from remote racks. Although the

amount of remote-rack accesses increases for larger clusters, it is expected to put a relatively small load on the ToR and core switches.

The section on [Staging data](#) within the *Reference Architectures* section of the document provides a number of options for network configurations for ingress and egress.

Best practice

Each reduce task can concurrently request data from a default of five map output files, so there is the possibility that servers will be delivered more data than their network connections can handle. This will result in dropped packets and can lead to a collapse in traffic throughput. This is why we recommend ToR switches with deep buffering.

Switches

Hortonworks HDP clusters contain two types of switches, namely Aggregation switches and Top of Rack switches. Top of Rack switches route the traffic between the nodes in each rack and Aggregation switches route the traffic between the racks.

Top of Rack (ToR) switches (HP 5900AF-48XG-4QSFP+)

The HP 5900AF-48XG-4QSFP+ 10GbE high-density, ultra-low latency, top-of-rack (ToR) switch provides IRF Bonding and sFlow which simplifies the management, monitoring and resiliency of the customer's Hadoop network. This model has 48x 10-Gigabit / Gigabit SFP+ ports with four QSFP+ 40-Gigabit ports for ultra-high capacity connections. The high performance 10 GbE networking provides cut-through and nonblocking architecture that delivers industry-leading low latency (~1 microsecond) for very demanding enterprise applications; the switch delivers a 1.28 Tbps switching capacity and 952.32 Mpps packet forwarding rate in addition to incorporating 9 MB of packet buffers.

For more information on the HP 5900AF-48XG-4QSFP+ switch, please see http://h17007.www1.hp.com/us/en/products/switches/HP_5900_Switch_Series/index.aspx

The configuration for the HP 5900AF-48XG-4QSFP+ switch is provided below. The switch can also be used for connecting the SL4540 chassis iLOs from the rack, which will require eight HP X120 1G SFP RJ45 T Transceiver (JD089B) 1 GbE transceivers for full rack and nine transceivers for expansion rack.

Figure 3. HP 5900AF-48XG-4QSFP+ Top of Rack (ToR) switch



Table 4. HP 5900AF-48XG-4QSFP+ Single Switch options

Qty	Description
1	HP 5900AF-48XG-4QSFP+ Switch
2	HP 58x0AF 650W AC Power Supply
2	HP 5830AF-48G Back(power)-Front(prt) Fan Tray
1	HP X240 40G QSFP+ QSFP+ 1m DAC Cable

HP Insight Cluster Management Utility

HP Insight Cluster Management Utility (CMU) is an efficient and robust hyperscale cluster lifecycle management framework and suite of tools for large Linux clusters such as those found in High Performance Computing (HPC) and Big Data environments. A simple graphical interface enables an “at-a-glance” view of the entire cluster across multiple metrics, provides frictionless scalable remote management and analysis, and allows rapid provisioning of software to all the nodes of the system. Insight CMU makes the management of a cluster more user friendly, efficient, and error free than if it were being managed by scripts, or on a node-by-node basis. Insight CMU offers full support for iLO 2, iLO 3, iLO 4 and LO100i adapters on all HP ProLiant servers in the cluster.

Note

HP Insight CMU allows you to easily correlate Hortonworks Data Platform (HDP) metrics with cluster infrastructure metrics, such as CPU Utilization, Network Transmit/Receive, Memory Utilization, Power Consumption, Ambient Node and Rack temperature and I/O Read/Write. This allows you to characterize your Hortonworks Data Platform (HDP) workloads and optimize the system thereby improving the performance of the Hortonworks Data Platform (HDP) Cluster. CMU TimeView Metric Visualizations will help you understand, based on your workloads, whether your cluster needs more memory, a faster network or processors with faster clock speeds. In addition, Insight CMU also greatly simplifies the deployment of Hortonworks Data Platform (HDP), with its ability to create a Golden Image from a Node and then deploy that image to up to 4000 nodes. Insight CMU is able to deploy 800 nodes in 30 minutes.

Insight CMU is highly flexible and customizable, offers both a graphical user interface (GUI) and a command-line interface (CLI) interface, and is being used to deploy a range of software environments, from simple compute farms to highly customized, application-specific configurations. Insight CMU is available for HP ProLiant and HP BladeSystem servers with Linux operating systems, including Red Hat Enterprise Linux, SUSE Linux Enterprise, CentOS, and Ubuntu. Insight CMU also includes options for monitoring graphical processing units (GPUs) and for installing GPU drivers and software.

For more information, please see hp.com/go/cmu.

Table 6. HP Insight CMU options

Qty	Description
1	HP Insight CMU 1yr 24x7 Flex Lic
1	HP Insight CMU 1yr 24x7 Flex E-LTU
1	HP Insight CMU 3yr 24x7 Flex Lic
1	HP Insight CMU 3yr 24x7 Flex E-LTU
1	HP Insight CMU Media

Figure 4. HP Insight CMU Interface – real-time view

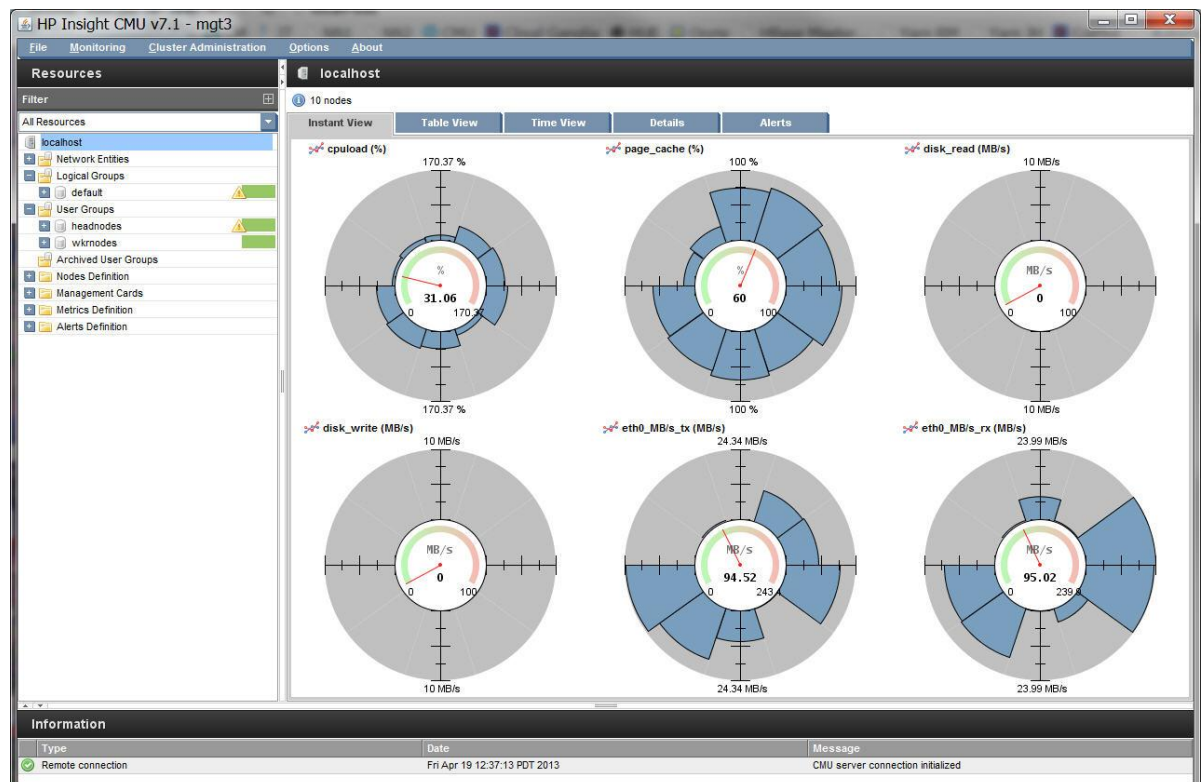
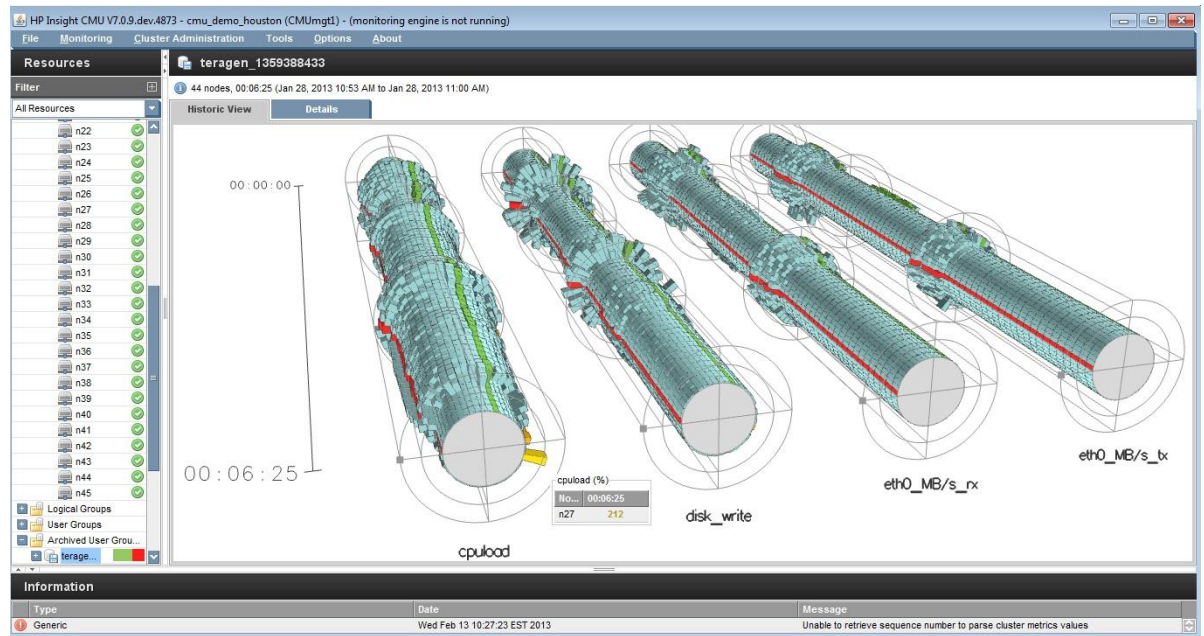


Figure 5. HP Insight CMU Interface – Time View



Server selection

Depending on the size of the cluster, a Hadoop deployment consists of one or more nodes running management services and a quantity of worker nodes. We have designed these reference architectures so that regardless of the size of the cluster, the server used for the management nodes and the server used for the worker nodes remains consistent. This section specifies which servers to use and the rationale behind it. The *Reference Architectures* section will provide topologies for the deployment of management and worker nodes for single and multi-rack clusters.

Management nodes

Management services are not distributed across as many nodes as the services that run on the worker nodes and therefore benefit from a server that contains redundant fans and power supplies, as well as an array controller supporting a variety of RAID schemes and SAS direct attached storage. In addition, the management services are memory and CPU intensive; therefore, a server capable of supporting a large amount of memory is also required. Management nodes do not participate in storing data for the HDFS and have much lower storage capacity requirements than worker nodes and thus a 2U server with a large amount of disks is not required.

This section contains the following subsections:

- High Availability
- Server platform
- Management node
- JobTracker server
- NameNode server
- Edge Nodes

High Availability

The Hortonworks Data Platform provides several options for High Availability that mitigate failure scenarios for the management nodes. For more information on setup and configuration, please see [hortonworks.com/products/HortonworksDataPlatform](http://hortonworks.com/products/ HortonworksDataPlatform)

Server platform: HP ProLiant DL360p Gen8

The HP ProLiant DL360p Gen8 (1U) is an excellent choice as the server platform for the management nodes.

Figure 6. HP ProLiant DL360p Gen8 Server



Processor configuration

The configuration features two sockets with fast 6 core processors (Intel Xeon E5-2667) and the Intel C600 Series chipset, which provide 12 physical cores and 24 Hyper-Threaded cores per server at the fastest clock speeds available. We recommend that Hyper-Threading be turned on.

We recommend processors with 6 cores for the management servers because the JobTracker, NameNode and Hortonworks Management Services are CPU intensive and multi-threaded and will fully utilize all the cores available. Furthermore, the configurations for these servers are designed to be able to handle an increasing load as your Hortonworks HDP cluster grows so one needs to ensure the right processing capacity is available to begin with.

Drive configuration

While the Hortonworks Data Platform provides High Availability solutions for the Management nodes in the cluster the solution can be made further redundant by employing RAID configurations for the drives within those nodes. The Smart Array P420i Controller is specified to drive eight 900GB 2.5" SAS disks on the Management node and four 900GB 2.5" SAS disks on the JobTracker and NameNode servers. The Management node has more disks than the JobTracker and NameNode servers due to the fact that it needs to have extra storage capacity for RAID Mirroring, the Hortonworks Management Console, and logs, as well as to act as a multi-homed staging server for data import and export out of the HDFS. Hot pluggable drives are specified so that drives can be replaced without restarting the server. Due to this design, one should configure the Gen8 P420i controller to apply the following RAID schemes:

- Management node: 4 Disks with RAID 1+0 for the OS, 4 Disks with RAID 1+0 for data staging
- JobTracker and NameNode Servers: 4 Disks with RAID 1+0 for OS

Best practice

The Gen8 P420i controller provides two port connectors per controller with each containing 4 SAS links. Each drive cage for the DL360p contains 8 disks and thus each disk has a dedicated SAS link which ensures the server provides the maximum throughput that each drive can give you. For a performance oriented solution, we recommend NL-SAS drives as they offer a significant read and write throughput performance enhancement over SATA disks.

Memory configuration

Servers running management services such as the HBaseMaster, JobTracker, NameNode and Hortonworks Management Center should have sufficient memory as they can be memory intensive. When configuring memory, one should always attempt to populate all the memory channels available to ensure optimum performance. The dual Intel Xeon E5-2667 2.9 GHz processors in the HP ProLiant DL360p Gen8 have 4 memory channels per socket which equates to 8 channels per server. The configurations for the management servers were tested with 64GB of RAM, which equated to eight 8GB DIMMS.

Network configuration

The HP ProLiant DL360p Gen8 is designed for network connectivity to be provided via a FlexibleLOM. The FlexibleLOM can be ordered in a 4 x 1GbE NIC configuration or a 2 x 10GbE NIC configuration. This Reference Architecture was tested using a 4 x 10GbE NIC configuration by using a 2 x 10 FlexibleLOM and a 2 x 10 PCI e card (as specified in the server configuration below).

Best practice

For each management server we recommend bonding and cabling two of the 10GbE NICs to create a single bonded pair which will provide 20GbE of throughput as well as a measure of NIC redundancy. The **Management node** in the cluster can be leveraged for ingress and staging data. In the reference configurations later in the document you will notice that we use two IRF Bonded switches. In order to ensure the best level of redundancy we recommend cabling NIC 1 to Switch 1 and NIC 2 to Switch 2.

*Bill of materials***Table 7.** The HP ProLiant DL360p Gen8 Server Configuration

Qty	Description
1	HP DL360p Gen8 8-SFF CTO Chassis
1	HP DL360p Gen8 E5-2667 FIO Kit
1	HP DL360p Gen8 E5-2667 Kit
8	HP 8GB 1Rx4 PC3-12800R-11 Kit
4	HP 900GB 6G SAS 10K 2.5in SC ENT HDD (Note: Management node needs 8)
1	HP Ethernet 10GbE 2 port 560FLR SFP+ adapter
1	HP Ethernet 10GbE 2 port 560 SFP+ adapter
1	HP 512MB FBWC for P-Series Smart Array
2	HP 460W CS Gold Hot Plug Power Supply Kit
1	HP 1U SFF BB Gen8 Rail Kit
1	ProLiant DL36x(p) HW Support

Management node

The Management node hosts the applications that submit jobs to the Hortonworks HDP cluster. We recommend that you install the following software components:

Table 8. Management node software

Software	Description
RHEL 6.2	Recommended Operating System
HP Insight CMU 7.1	Infrastructure Deployment, Management, and Monitoring
Oracle JDK 1.6.0_31	Java Development Kit
Hortonworks MC	Hortonworks Management Center
NFS Server	Provides an NFS Mount for the NameNode Edit Log
Apache Pig and/or Apache Hive	Analytical interfaces to the Hadoop Cluster
ZooKeeper	Synchronization service (Only if running HBase)

Please see the following link for the Hortonworks Management Center Installation Guide, http://docs.hortonworks.com/CURRENT/index.htm#Deploying_Hortonworks_Data_Platform/Using_HMC/Getting_Ready_To_Install/Getting_Ready_to_Install.htm

Note: The Hortonworks Management Center enables the deployment of the cluster as well as which additional components to install. Please see the following link for deploying Apache Pig and Apache Hive with HMC, http://docs.hortonworks.com/CURRENT/index.htm#Deploying_Hortonworks_Data_Platform/Using_HMC/Configuring_And_Deploying_The_Cluster/Configuring_and_Deploying_the_Cluster.htm

The Management node contains the following base configuration:

- Dual Six-Core Intel E5-2667 2.9 GHz Processors
- Smart Array P420i Controller
- Eight 900GB SFF SAS 10K RPM disks
- 64 GB DDR3 Memory
- 4 x 10GbE Ethernet NICs

JobTracker server

The JobTracker server contains the following software. Please see the following link for more information on installing and configuring the JobTracker and secondary NameNode,

http://docs.hortonworks.com/CURRENT/index.htm#Deploying_Hortonworks_Data_Platform/Using_HMC/Configuring_And_Deploying_The_Cluster/Configuring_and_Deploying_the_Cluster.htm

Table 9. JobTracker Server Software

Software	Description
RHEL 6.2	Recommended Operating System
Oracle JDK 1.6.0_31	Java Development Kit
JobTracker	The JobTracker for the Hadoop Cluster
<i>HBaseMaster</i>	<i>The HBase Master for the Hadoop Cluster (Only if running HBase)</i>
<i>ZooKeeper</i>	<i>Synchronization service (Only if running HBase)</i>
Secondary NameNode	Process to handle Check pointing from the NameNode

The JobTracker server contains the following base configuration:

- Dual Six-Core Intel E5-2667 2.9 GHz Processors
- Four 900GB SFF SAS 10K RPM disks
- 64 GB DDR3 Memory
- 4 x 10GbE Ethernet NICs
- 1 x Smart Array P420i Controller

NameNode server

The NameNode server contains the following software. Please see the following link for more information on installing and configuring the NameNode,
http://docs.hortonworks.com/CURRENT/index.htm#Deploying_Hortonworks_Data_Platform/Using_HMC/Configuring_And_Deploying_The_Cluster/Configuring_and_Deploying_the_Cluster.htm

Table 10. NameNode Server Software

Software	Description
RHEL 6.2	Recommended Operating System
Oracle JDK 1.6.0_31	Java Development Kit
NameNode	The NameNode for the Hadoop Cluster
NFS Client	Allows the NameNode to write logs to the Management node
ZooKeeper	<i>Synchronization service (Only if running HBase)</i>

The NameNode server contains the following base configuration:

- Dual Six-Core Intel E5-2667 2.9 GHz Processors
- Four 900GB SFF SAS 10K RPM disks
- 64 GB DDR3 Memory
- 4 x 10GbE Ethernet NICs
- 1 x Smart Array P420i Controller

Edge nodes

Edge nodes are servers that are typically multi-homed into both the Hortonworks Data Platform private network as well as a network accessible by users. Edge nodes are typically configured to run Application Clients to the Hortonworks Data Platform or are used to manage ingress and egress of data.

Edge nodes configured to run applications do not require a large amount of Direct Attached Storage and thus are well suited to the 1U HP ProLiant DL360p Gen8 server. When an Edge node is multi-homed it enables it to run applications that can submit jobs to the Hortonworks Data Platform cluster and have the application interfaces accessible to users within the user network. One example of how an Edge node can be leveraged would be to install and configure an application such as Datameer. This provides users with a tool that interacts with the Hortonworks Data Platform to allow users to analyze and visualize their data.

Edge nodes configured to explicitly support the staging of data and thus require a significant amount of storage capacity, are well suited to the HP ProLiant SL4540 Gen8 server.

Worker nodes

The worker nodes run the TaskTracker (or HBaseRegionServer) and DataNode processes and thus storage capacity and performance are important factors.

Server platform: HP ProLiant SL4540 Gen8 (3 x 15)

The HP ProLiant SL4540 Gen8 (3x15) (4.3U) is an excellent choice as the server platform for the worker nodes. For ease of management we recommend a homogenous server infrastructure for your worker nodes. The 3 node configuration of the HP ProLiant SL4540 Gen8 consists of three compute nodes and a total of forty-five (45) large form factor (LFF) 3.5" hard disk drives (HDD) in the chassis (15 per node) and six (6) small form factor (SFF) 2.5" hard drives in the front of the chassis (two per node). The HP ProLiant SL4540 is a dual socket Gen8 server, with a choice of five different Intel Xeon processors, up to 192GB of memory and one PCIe slot for expansion per node. Every compute node also has its own dedicated networking ports.

Figure 7. HP ProLiant SL4540 Gen8 (3 x 15) Server



Processor configuration

The configuration for each server node in the SL4540 features two sockets with 8 core processors and the Intel C600 Series chipset which provide 16 physical or 32 Hyper-Threaded cores per server. Hadoop manages the amount of work each server is able to undertake via the amount of MapReduce slots configured for that server. The more cores available to the server, the more MapReduce slots can be configured for the server (see the [Computation](#) section for more detail). We recommend that Hyper-Threading be turned on.

Drive configuration

The HP Smart Array P420i controller which provides increased³ I/O throughput performance resulting in a significant performance increase for I/O bound Hadoop workloads (a common use case). Redundancy is built into the Apache Hadoop architecture and thus there is no need for RAID schemes to improve redundancy on the worker nodes as it is all coordinated and managed by Hadoop. Drives should use a Just a Bunch of Disks (JBOD) configuration, which can be achieved with the HP Smart Array P420i controller by configuring each individual disk as a separate RAID 0 volume. Additionally array acceleration features on the P420i should be turned off for the RAID 0 data volumes. The worker node design includes a total of 17 drives per node, 2 are used for the operating system and 15 are used for data.

Best practice

The two 500GB SATA MDL disks in the front of the SL4540 are configured as RAID 1 using the HP Smart Array B120i Controller for mirrored OS and Hadoop runtime. This provides additional measures of redundancy on the worker nodes. We do not recommend sharing drives that contain the OS and Hadoop runtimes with drives that contain the temporary MapReduce data and the HDFS block data as it results in degraded I/O performance.

Performance

The HP Smart Array P420i controller provides two port connectors per controller with each containing 4 SAS links. The controller drives the 15 disks in the chassis per node. The 2 disks in the front per node are controlled by B120i SATA RAID controller. For a performance oriented solution, we recommend SAS MDL drives as they offer a significant read and write throughput performance enhancement over SATA disks.

Core to Disk Ratio

The more drives a server contains, the more efficiently it can service I/O requests because it reduces the likelihood of multiple threads contending for the same drive which can result in interleaved I/O and degraded performance.

Memory configuration

Servers running the worker node processes should have sufficient memory for either HBase or for the amount of MapReduce Slots configured on the server. The Intel Xeon E5-2470 has 3 memory channels per processor. When configuring memory, one should always attempt to populate all the memory channels available to ensure optimum performance.


³ Compared to the previous generation of Smart Array controllers

Best practice

It is important to saturate all the memory channels available to ensure optimal use of the memory bandwidth. For example, on a two socket server with three memory channels that supports two DIMMs each for a total of six (6) DIMMs per installed processor or a grand total of twelve (12) DIMMs for the server, one would typically fully populate the channels with 8GB DIMMs resulting in a configuration of 96GB of memory per server.

See Figure 8 for the memory slots to populate when using 8GB memory DIMMs.

Figure 8. SL4540 8GB DIMM memory configuration recommendation



Network configuration

For 10GbE networks we recommend that two 10GbE NICs be bonded to improve throughput performance to 20 Gb/s. In addition, in the reference configurations later on in the document you will notice that we use two IRF Bonded switches. In order to ensure the best level of redundancy we recommend cabling NIC 1 to Switch 1 and NIC 2 to Switch 2.

The worker node contains the following software. Please see the following link for more information on installing and configuring the TaskTracker (or HBaseRegionServer) and DataNode, http://docs.hortonworks.com/CURRENT/index.htm#Deploying_Hortonworks_Data_Platform/Using_HMC/Configuring_And_Deploying_The_Cluster/Configuring_and_Deploying_the_Cluster.htm

Table 11. Worker Node Software

Software	Description
RHEL 6.2	Recommended Operating System
Oracle JDK 1.6.0_31	Java Development Kit
TaskTracker	The TaskTracker process for MapReduce (Only if running MapReduce)
DataNode	The DataNode process for HDFS
HBaseRegionServer	The HBaseRegionServer for HBase (Only if running HBase)

The ProLiant SL4540 Gen8 configured as a worker node should have the following base configuration:

- Dual Eight-Core Xeon E5-2470 2.3 GHz Processors with Hyper-Threading
- 15 x 2TB 3.5" SAS SC MDL 7.2K RPM disks (higher density disks are also available)
- 96 GB DDR3 Memory
- 2 x 10GbE Ethernet NICs
- 2 x Smart Array P420i Controllers

Best practice

Customers have the option of purchasing a second power supply for additional power redundancy. This is especially appropriate for single rack clusters where the loss of a node represents a noticeable percentage of the cluster.

Table 12. The HP ProLiant SL4540 Gen8 (3 x 15) Server Configuration

Qty	Description
SL4540 base	
1	HP SL4540 (3 x15) Gen8 Chassis
4	HP 750W CS Gold Hot Plug Power Supply Kit
1	ProLiant SL4540 Gen8 HW Support
1 server node (order 3 nodes per SL4540 chassis)	
1	HP 3xSL4540 Gen8 Tray Node Svr
1	HP SL4540 Gen8 E5-2470 Kit FIO Kit
1	HP SL4540 Gen8 E5-2470 Kit
12	HP 8GB 2Rx4 PC3L-10600R-9 Kit
15	HP 2TB 6G SAS 7.2K 3.5in SC MDL HDD
2	HP 500GB 6G SATA 7.2k 2.5in SC MDL HDD
1	HP 10G IO Module
2	HP RAID Mez Ctrllr p420i FIO Kit
2	HP 12in Super Cap for Smart Array
2	HP 1GB FBWC for P-Series Smart Array

Reference Architectures

This section provides a progression of configurations from single rack to full scale out Hortonworks HDP cluster configurations. Best practices for each of the components within the configurations specified have been articulated earlier in this document.

Single Rack configuration

The Single Rack Hortonworks Data Platform configuration provides a starting point for an HDP cluster configuration. The configuration reflects the following components.

Single Rack network

As previously described in the [Network](#) section, two IRF Bonded HP 5900AF-48XG-4QSFP+ ToR switches are specified for performance and redundancy. Keep in mind that if IRF bonding is used, it requires 2 40GbE ports per switch, which would leave 2 40GbE ports on each switch for uplinks.

Cluster isolation and access configuration

It is important to isolate the Hortonworks HDP cluster on the network so that external network traffic does not affect the performance of the cluster. In addition, this also allows for the Hadoop cluster to be managed independently from that of its users, which ensures that the cluster administrator is the only one capable of making changes to the cluster configurations. To achieve this, we recommend isolating the JobTracker, NameNode and Worker nodes on their own private Hadoop cluster subnet.

Best practice

Once a Hortonworks HDP cluster is isolated, the users of the cluster will still need a way to access the cluster and submit jobs to it. To achieve this we recommend multi-homing the Management node so that it participates in both the cluster subnet and a subnet belonging to the users of the cluster. Hortonworks Management Center is a web application that runs on the Management node and allows users to be able to manage and configure the Hadoop cluster (including seeing the status of jobs) without being on the same subnet, provided the Management node is multi-homed. Furthermore, this allows users to be able to shell into the Management node and run the Apache Pig or Apache Hive command line interfaces and submit jobs to the cluster that way.

Staging data

In addition, once the Hortonworks HDP cluster is on its own private network one needs to think about how to be able to reach the HDFS in order to move data onto it. The HDFS client needs to potentially be able to reach every Hadoop DataNode in the cluster in order to stream blocks onto it to move data onto the HDFS. The Reference Architecture provides several ways to do this.

The first option is to use the already multi-homed Management node. This server has been configured with twice the amount of disk capacity (an additional 3.6 TB) compared to the other management servers in order to provide an immediate solution to move data into the Hadoop Cluster from another subnet.

Another option is to leverage WebHDFS which provides an HTTP proxy to securely read and write data to and from the Hadoop Distributed File System. For more information on WebHDFS, please see http://docs.hortonworks.com/CURRENT/index.htm#Using_WebHDFS_REST_API/Using_WebHDFS_REST_API.htm

Rack enclosure

The rack contains eight HP ProLiant SL4540 chassis with 24 server nodes, three HP ProLiant DL360p servers and two HP 5900AF-48XG-4QSFP+ switches within a 42U rack.

Network

As previously described in the [Switches](#) section, two HP 5900AF-48XG-4QSFP+ switches are specified for performance and redundancy. The HP 5900AF-48XG-4QSFP+ includes up to four 40GbE uplinks which can be used to connect the switches in the rack into the desired network. Keep in mind that if IRF bonding is used, it requires 2 40GbE ports per switch, which would leave 2 40GbE ports on each switch for uplinks.

Management nodes

ProLiant DL360p Gen8 management nodes are specified:

- The Management Node
- The JobTracker Node
- The NameNode

Detailed information on the hardware and software configurations for these three management nodes is available in the [Server selection](#) section of this document.

Worker nodes

After the redundant switches and management servers are installed in the rack there is space for eight ProLiant SL4540 Gen8 chassis with 24 worker nodes within the rack.

Best practice

One can have as few nodes as a single worker node, however starting with at least three worker nodes is recommended to provide the redundancy that comes with the default replication factor of 3. Performance improves with additional worker nodes as the JobTracker can leverage idle nodes to land jobs on servers that have the appropriate blocks, leveraging data locality rather than pulling data across the network. These servers are homogenous and run the DataNode and the TaskTracker (or HBaseRegionServer) processes.

Power and cooling

In planning for large clusters, it is important to properly manage power redundancy and distribution. To ensure the servers and racks have adequate power redundancy we recommend that each server have a backup power supply, and each rack have at least two Power Distribution Units (PDUs). There is an additional cost associated with procuring redundant power supplies. This is less important for larger clusters as the inherent failover redundancy within the Hortonworks Data Platform will ensure there is less impact.

Best practice

For each SL4540 chassis, there are four power supplies and we recommend that two power supplies be connected to two different PDUs. Furthermore, the PDUs in the rack can each be connected to a separate data center power line to protect the infrastructure from a data center power line failure.

Additionally, distributing the server power supply connections evenly to the in-rack PDUs, as well as distributing the PDU connections evenly to the data center power lines ensures an even power distribution in the data center and avoids overloading any single data center power line. When designing a cluster, check the maximum power and cooling that the data center can supply to each rack and ensure that the rack does not require more power and cooling than is available.

Figure 9. Single Rack Hortonworks HDP Configuration

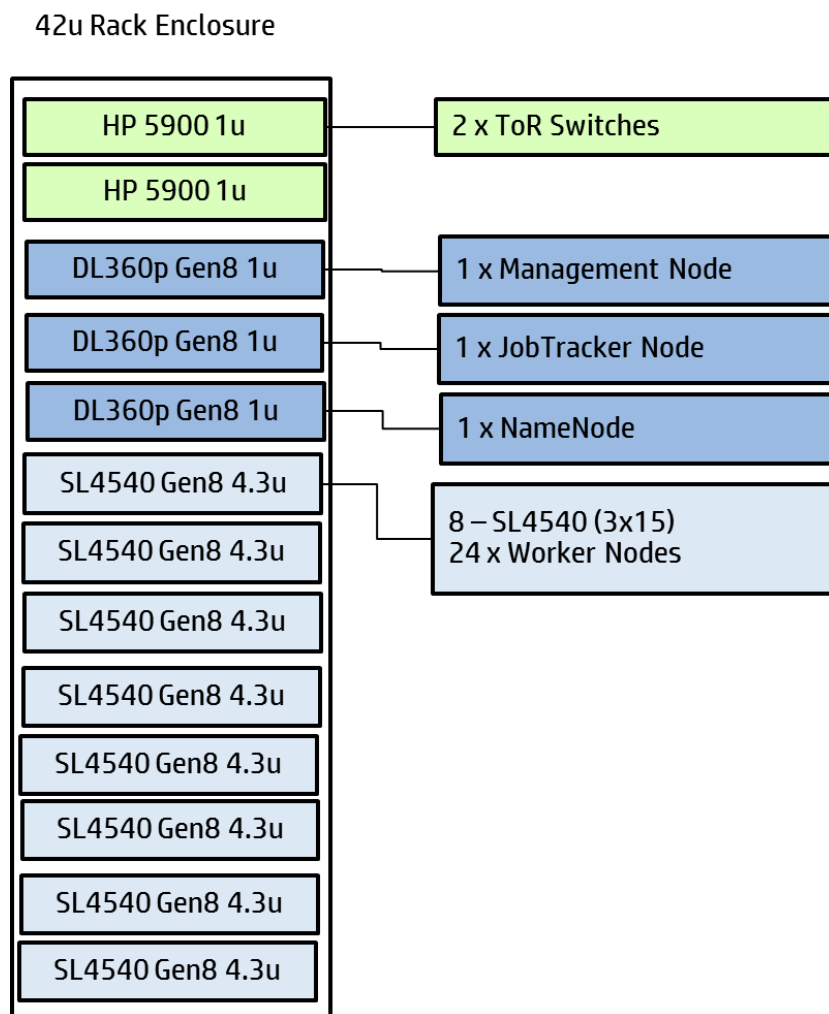


Figure 10. Single Rack Hortonworks HDP Configuration – Rack Level View

Management Node

1 x HP ProLiant DL360p Gen8
with 12 cores
2 x E5-2667 (Intel Xeon 2.9GHz, 15M Cache)
64 GB - 8 x HP 8GB 1Rx4 PC3-12800R
7.2 TB - 8 x HP 900GB 6G SAS 10K 2.5in SC ENT HDD
1 x HP 512MB FBWC for P-Series Smart Array
2 x HP dual port 10GbE

Worker Nodes

8 x HP 3xSL4540 Gen8
with 16 cores each
2 x E5-2470 (Intel Xeon 2.3GHz, 20M Cache)
96 GB - 12 x HP 8GB 2Rx4 PC3L-10600R
30 TB - 15 x HP 2TB 6G SAS 7.2K 3.5in SC MDL HDD
1TB - 2 x HP 500GB 6G SATA 7.2K 2.5in SC MDL HDD
2 x HP 1GB P-series Smart Array FBWC
2 x HP Smart Array P420i Mez Ctrlr FIO Kit
1 x HP 10G IO Module (2x1GbE, 2x10GbE)

Ethernet Switch


2 x HP 5900AF-48XG-4QSFP+ Switch

Name Node & JobTracker

2 x HP ProLiant DL360p Gen8
with 12 cores each
2 x E5-2667 (Intel Xeon 2.9GHz, 15M Cache)
64 GB - 8 x HP 8GB 1Rx4 PC3-12800R
3.6 TB - 4 x HP 900GB 6G SAS 10K 2.5in SC ENT HDD
1 x HP 512MB FBWC for P-Series Smart Array
2 x HP dual port 10GbE

Software

OS – Linux (RHEL 6.2 x86_64)
Hortonworks Data Platform
HP Insight Cluster Management Utility (HP CMU) v7.1



Multi-Rack Cluster configuration

The Multi-Rack Cluster design assumes the Single Rack configuration is already in place and extends its scalability. The Single Rack configuration ensures the required amount of management services are in place for large scale out. For Multi-Rack clusters, one simply adds more racks of the configuration provided below to the Single Rack configuration. This section reflects the design of those racks.

Rack enclosure

The rack contains nine HP ProLiant SL4540 Gen8 chassis with 27 server nodes and two HP 5900AF-48XG-4QSFP+ switches within a 42U rack.

Network

As previously described in the [Switches](#) section, two HP 5900AF-48XG-4QSFP+ switches are specified for performance and redundancy. The HP 5900AF-48XG-4QSFP+ includes up to four 40GbE uplinks which can be used to connect the switches in the rack into the desired network. Keep in mind that if IRF bonding is used, it requires 2 40GbE ports per switch, which would leave 2 40GbE ports on each switch for uplinks.

Software

The ProLiant SL4540 Gen8 (3 x 15) chassis with 27 server nodes in the rack all are configured as Worker nodes in the cluster as all required management processes are already configured in the Single Rack RA. Aside from the OS, the following worker processes are typically present.

- DataNode
- TaskTracker (or HBaseRegionServer if you are using HBase)

Figure 11. Large Cluster Hortonworks HDP Configuration

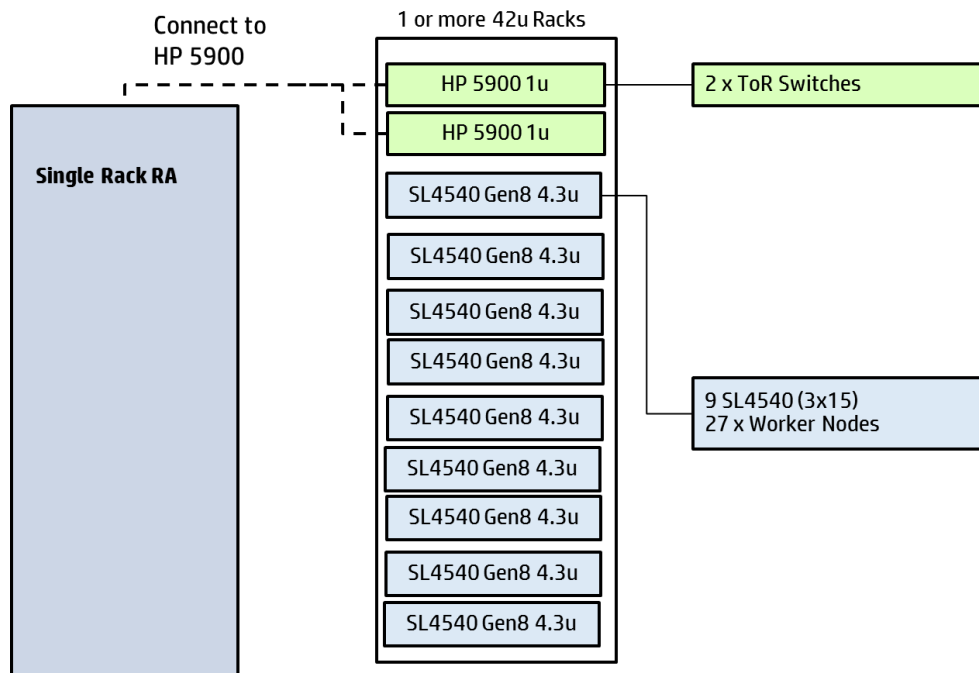



Figure 12. Large Cluster Hortonworks HDP Configuration – Rack Level View



HP Vertica and Hadoop

Relational database management systems such as HP Vertica excel at analytic processing for big volumes of structured data including call detail records, financial tick streams and parsed weblog data. HP Vertica is designed for high speed load and query when the database schema and relationships are well defined. Hortonworks Data Platform (HDP), built on the popular open source Apache Software Foundation project, addresses the need for large-scale batch processing of unstructured or semi-structured data. When the schema or relationships are not well defined, Hadoop can be used to employ massive MapReduce style processing to derive structure out of data. The Hortonworks Data Platform (HDP) simplifies installation, configuration, deployment and management of the powerful Hadoop framework for enterprise users.

Each can be used standalone – HP Vertica for high-speed loads and ad-hoc queries over relational data, Hortonworks Data Platform (HDP) for general-purpose batch processing, for example from log files. Combining Hadoop and Vertica creates a nearly infinitely scalable platform for tackling the challenges of big data.

Benefit of HP Vertica

HP Vertica was the first analytic database company to deliver a bi-directional Hadoop Connector enabling seamless integration and job scheduling between the two distributed environments. With HP Vertica's Hadoop and Pig Connectors, users have unprecedented flexibility and speed in loading data from Hadoop to Vertica and querying data from Vertica in Hadoop as part of MapReduce jobs for example. The HP Vertica Hadoop and Pig Connectors are supported by HP Vertica, and available for download.

For more information, please see vertica.com/the-analytics-platform/native-bi-etl-and-hadoop-mapreduce-integration/

Summary

HP and Hortonworks allow one to derive new business insights from Big Data by providing a platform to store, manage and process data at scale. However, designing and ordering Hadoop Clusters can be both complex and time consuming. This white paper provides several reference configurations for deploying clusters of varying sizes of the Hortonworks Data Platform (HDP) on HP infrastructure. These configurations leverage HP's balanced building blocks of servers, storage and networking, along with integrated management software and bundled support. In addition, this white paper has been created in order to assist in the rapid design and deployment of the Hortonworks Data Platform (HDP) on HP infrastructure for clusters of various sizes.

For more information

Hortonworks, hortonworks.com

HP Solutions for Apache Hadoop, hp.com/go/hadoop

HP Insight Cluster Management Utility (CMU), hp.com/go/cmu

HP ProLiant SL4540 Gen8, hp.com/servers/sl4540

HP ProLiant DL380e Gen8, hp.com/servers/dl380e

HP ProLiant servers, hp.com/go/proliant

HP Enterprise Software, hp.com/go/software

HP Networking, hp.com/go/networking

HP iLO Management Engine, hp.com/servers/ilo

HP Product Bulletin (QuickSpecs), hp.com/go/quickspecs

HP Services, hp.com/go/services

HP Support and Drivers, hp.com/go/support

HP Systems Insight Manager (HP SIM), hp.com/go/hpsim

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.



Sign up for updates

hp.com/go/getupdated

© Copyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates.

4AA4-7057ENW, May 2013, Rev. 1

