WHITE PAPER

Hitachi Unified Compute Platform for the SAP HANA Platform Using Hortonworks for Business Analytics

Reference Architecture Guide

By David Pascuzzi, Mark Walker

July 2017

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision History

Revision	Changes	Date
AS-607-00	Initial release	July, 2017

Table of Contents

Hitachi Unified Compute Platform for the SAP HANA Platform Using Hortonworks for Business Analytics

Reference Architecture Guide

Use this reference architecture guide to see how to configure Hitachi Unified Compute Platform for the SAP HANA Platform Using Hortonworks for Business Analytics. This shows an example environment to deploy big data infrastructure for advanced analytics.

This integrated big data infrastructure uses rack-optimized servers for Pentaho data integration. There are guidelines to create the best infrastructure for your big data projects with these hardware products:

- **Rack-optimized server for solutions, 2U single node** — This is a flexible server designed for optimal performance across multiple applications.
- **Brocade ICX 7250-48 switch** — This 48-port 1 GbE switch provides a management network.
- **Brocade ICX 7750-48 switch** — This 48-port switch provides 10 GbE connectivity for intra-rack networks.
- **Brocade VDX 6940-36 switch** — This 36-port switch provides 40 GbE connectivity for inter-rack networks and up-link connections to other racks and your network.

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Key Solution Elements

These are the key hardware and software components used for this integrated infrastructure. This example environment for Hitachi Unified Compute Platform for SAP HANA Using Hortonworks for Business Analytics implements Hortonwork.

Hardware Elements

These are the key hardware elements to power this big data solution. It is possible to create a scale-out configuration to power your Hadoop needs.

Rack Optimized Server for Solutions, 2U Single Node

The rack optimized server for solutions, 2U single node, is a rack mounted server designed for optimal performance and power efficiency. It supports up to 756 GB highly scalable memory capacity. It is powered by the Intel Xeon E5-2600 v4 processor product family for complex and demanding workloads. It supports flexible OCP and PCIe I/O expansion card options.

This solution supports up to 756 GB of memory in the rack optimized server. The Intel Xeon E5-2600 v4 processor provides power for this complex and demanding workload. It supports flexible PCIe I/O expansion card options. Figure 1 shows the front and back of this server.

Figure 1

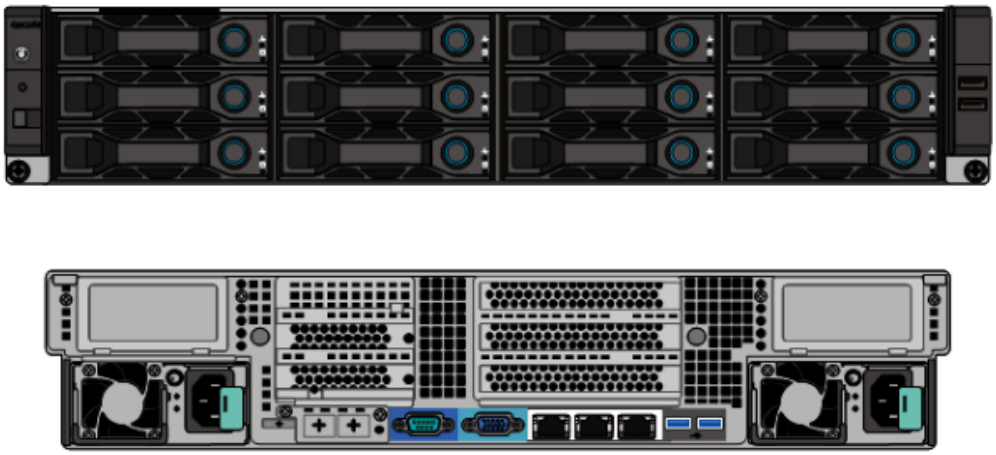

Brocade Switches

Brocade and Hitachi Data Systems partners to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

This solution includes the following Brocade products to provide Ethernet connectivity:

- Brocade VDX 6940-36 port switch
- Brocade ICX 7750-48 port switch
- Brocade ICX 7250-48 port switch

Software Elements

Big Data is a generic term to cover a set of components that are used with very large data sets to provide advanced data analytics. Big data usually refers to large volumes of unstructured or semi-structured data.

Usually, a big data solution is part of the Apache Hadoop project. However, big data can include components from many different software companies.

This reference architecture uses Red Hat Enterprise Linux and Hortonworks Data Platform.

Figure 2 on page 3 is a sample showing some of the possible software in a Hadoop ecosystem.

Figure 2

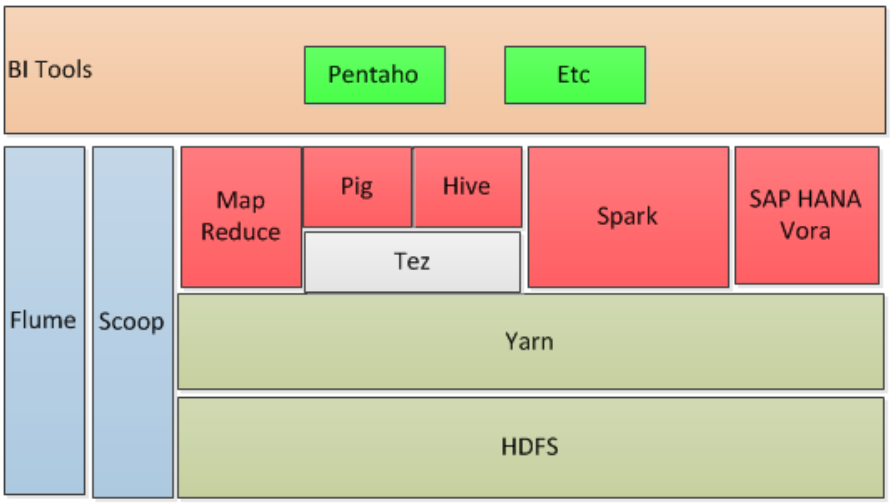

A partial list of the software components and modules used in this solution are the following:

- **Apache Hadoop Distributed File System**

  Hadoop Distributed File System (HDFS) is a distributed high-performance file system designed to run on commodity hardware.

- **Hadoop Common**

  These common utilities support the other Hadoop modules. This programming framework supports the distributive processing of large data sets.

- **Apache Hadoop YARN**

  Apache Hadoop YARN is a framework for job scheduling and cluster resource management. This splits into separate daemons the functionalities of the following:

  - **ResourceManager** interfaces with the client to track tasks and assign tasks to **NodeMangers** management
  - **NodeManager** launches and tracks execution on the worker nodes

- **Apache Spark**

  Apache Spark is a fast and general engine for large-scale data processing.

  - **Spark Master Node**

    In a Spark cluster, the master node oversees assigning tasks for the worker nodes to execute. It checks the status of those tasks and retrieves the results.

    Also, the master node can be used as a worker node, if necessary. In this solution, the master node will assign tasks to itself, as well as to worker nodes. This can be useful if there are few overall nodes in the cluster.

  - **Spark Worker Node or Nodes**

    The worker node or nodes in a Spark cluster do the work assigned to them by the master node. They connect to the master node, are assigned tasks, and execute those tasks. These nodes use the CPU and available storage.

- **SAP HANA Vora**

  SAP HANA Vora is a distributed computing solution for business. It leverages and extends the Apache Spark execution framework to provide enriched interactive analytics on enterprise and Apache Hadoop data.

- **Pentaho**

  Pentaho is a comprehensive big data integration and business analytics platform.

- **Apache Pig**

  Apache Pig is a platform for analyzing large data sets that consists of coupling the following:

  - A high-level language for expressing data analysis programs
  - An infrastructure for evaluating these data analysis programs

- **Apache Hive**

  Apache Hive is data warehouse software that facilitates reading, writing, and managing large datasets residing in distributed storage using SQL.

- **Apache Flume**

  Apache Flume is a distributed, reliable, and available service for efficiently collecting, aggregating, and moving large amounts of log data.

- **Apache Sqoop**

  Apache Sqoop is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured data stores, such as relational databases.

- **Apache ZooKeeper**

  Apache ZooKeeper is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services.

  - **ZooKeeper Master Node**

    ZooKeeper is a high-availability system, whereby two or more nodes can connect to a ZooKeeper master node. The ZooKeeper master node controls the nodes to provide high availability.

  - **ZooKeeper Standby Master Node**

    When ZooKeeper runs in a highly-available setup, there can be several nodes configured as ZooKeeper master nodes. Only one of these configured nodes is active as a master node at any time. The others are standby active nodes.

    If the currently-active master node fails, then the ZooKeeper cluster itself promotes one of the standby master nodes to be active as the master node.

- **Apache Oozie**

  Apache Oozie is a workflow scheduler system to manage Apache Hadoop jobs.

- **Apache Ambari**

  Apache Ambari makes Apache Hadoop management simpler by developing software for provisioning, managing, and monitoring Hadoop clusters.

- **HUE**

  HUE (Hadoop User Experience) is a web interface for analyzing data with Apache Hadoop.

  This is a Cloudera product, not an Apache product.

- **Apache Zeppelin**

  Apache Zeppelin is a web-based notebook that enables interactive data analytics.

- **Apache Tez**

  Apache Tez is a framework for batch and interactive data processing applications.

Solution Design

This is the detailed design for an integrated infrastructure for big data from Hitachi to implement big data and business analytics.

Rack Optimized Server for Solutions, 2U Single Node

This solution uses multiple rack-optimized servers for solutions, 2U single node. The architecture supports using one-node servers in multiple configurations.

For ease of maintenance and flexibility in deployment, all nodes use the same configuration.

There are three types of nodes:

Edge Node

An edge node resides on both the client network and the data network. Often these nodes are management or master nodes.

This type of node can be anything that needs to access both networks. This includes the following types of edge nodes:

- **Hitachi Unified Compute Platform for SAP HANA node**

  Applications on a Unified Compute Platform for SAP HANA node access SAP HANA and can access data in a big data product.

- **Pentaho node**

  Pentaho does the following:

  - Transfer data from existing databases into Hadoop
  - Access data from Apache Hive and any relational database management system

- **Gateway node**

  Users go to a gateway node before accessing any data node.

Master Node

Master nodes are nodes that control other processes on the network. This includes the following types of master nodes:

- Apache Ambari node
- NameNode for the Apache Hadoop Distributed File System
- Apache Hive Manager node
- System management node

These master nodes can be in the same configuration as the worker nodes. These nodes are usually edge nodes, also.

Table 1 shows the configuration of the master nodes.

TABLE 1. MASTER NODE CONFIGURATION

Component	Description
CPU	■ 2 Intel Xeon E5-2600 v4 processors
Memory Options	■ 64 GB ■ 128 GB
Network Connections	■ 2 port 10 GbE on board connections ■ 1 GbE LOM management port ■ 1 × 10 GbE OCP mezzanine card (optional)
Disk Controllers	■ LSI 3108 RAID controller mezzanine card for data disks

TABLE 1. MASTER NODE CONFIGURATION (CONTINUED)

Component	Description
Disk	■ 6 HDD data disks ■ 2 operating system disks
Racks	■ 42 U rack
Number of Servers	■ Depends on implementation.

Worker Node

Worker nodes process the data. These nodes have very diverse needs and can have many different configurations.

Table 2 shows the worker node configuration options.

TABLE 2. WORKER NODE CONFIGURATION OPTIONS

Component	Description
CPU	■ 1 or 2 Intel Xeon E5-2600 v4 processors
Memory Options	■ 64 GB ■ 128 GB ■ 256 GB ■ 384 GB ■ 512 GB ■ 768 GB
Network Connections	■ 2 port 10 GbE on board connections ■ 1 GbE LOM management port
Disk Controllers	■ LSI 3108 RAID controller mezzanine card for data disks
Disks	■ Up to 12 data disks ■ 2 operating system disks
Racks	■ 42 U rack
Number of Servers	■ 9 to 18 data nodes per rack

Edge or Master Rack Configuration

Install the master nodes and edge nodes in the same rack. Also, you can use this rack for spine switches. The empty units can have worker nodes installed.

Figure 3 shows a typical edge or master rack configuration:

- All racks in this solution have ToR switches.
- Place optional spine switches in the management rack.
- Use Ambari to manage the Apache Hadoop cluster.
- This solution uses five master nodes to control and manage processing for Hadoop components.
- There is one gateway node.
- There is one Pentaho server node for accessing data.
- There is one Hitachi Unified Compute Platform for SAP HANA node

Figure 3

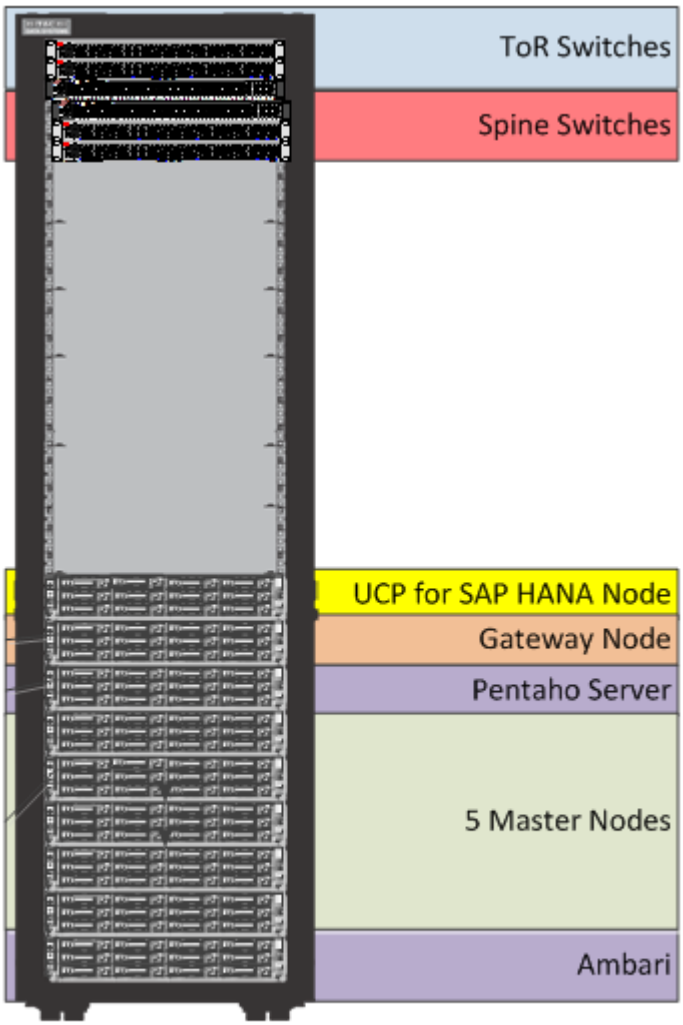

Worker Rack Configuration

Rather than pre-define the number of nodes, this solution allows many different configurations. Table 3 has the examples of base rack configurations options for the worker nodes.

TABLE 3. BASE RACK CONFIGURATIONS OPTIONS

Number of Servers	Description

9	■ Use for this type of system:
	  ■ Starter system
	  ■ Test or development system
	■ Can be used as a spine switch rack
	■ Allows for future growth without adding a rack
	■ Leaves space for other node types, including but not limited to the following:
	  ■ Edge nodes
	  ■ Application nodes
	  ■ Infrastructure nodes
	  ■ System management nodes
	  ■ Future growth

18	Leaves space for 6U of the following: ■ 2 leaf data switches ■ 1 leaf management switch ■ 2 spine data switches ■ 1 spine management switch ■ 1 management server

Figure 4 shows a rack with the following:

- 18 × 2U 1-node big data servers
- 2 ToR leaf data switches
- 1 ToR leaf management switch

Figure 4

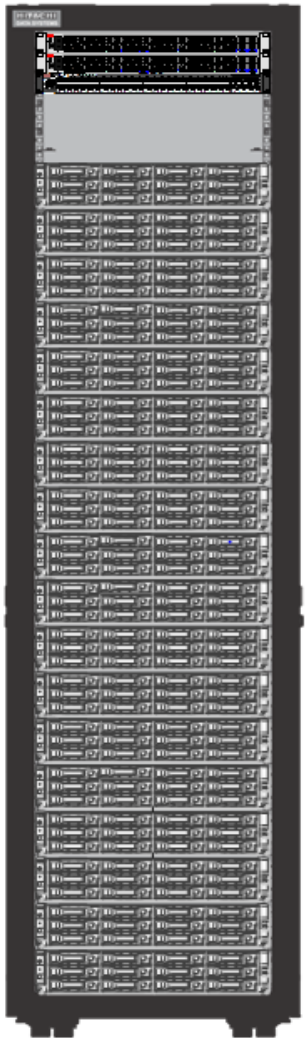

With 18 nodes, a rack can have the following:

- 36 CPUs
- 13 TB of memory
- 1296 TB of data storage

Worker Node Storage Architecture

This solution uses internal storage drives. As every deployment has different needs, this solution allows you to customize the storage.

Table 4 lists the storage configuration options.

TABLE 4. STORAGE CONFIGURATION OPTIONS

Use	Option	Number of Disks	Configuration
Operating System Disk	■ 500 GB SATA HDD, 7200 RPM	2	RAID-1
Data Disk	■ Multiple disk types, sizes, and speeds from 300 GB to 6 TB	0 - 12	Just a bunch of disks (JBOD)

Network Architecture

This solution uses three logical networks. There can be multiple network configurations, depending on the Apache Hadoop deployment.

Figure 5 on page 12 shows one of the possible configurations.

- **Client Network** — Client access to edge nodes and management servers
- **Data Network** — Communication between nodes
- **Management Network** — Management of hardware

Figure 5

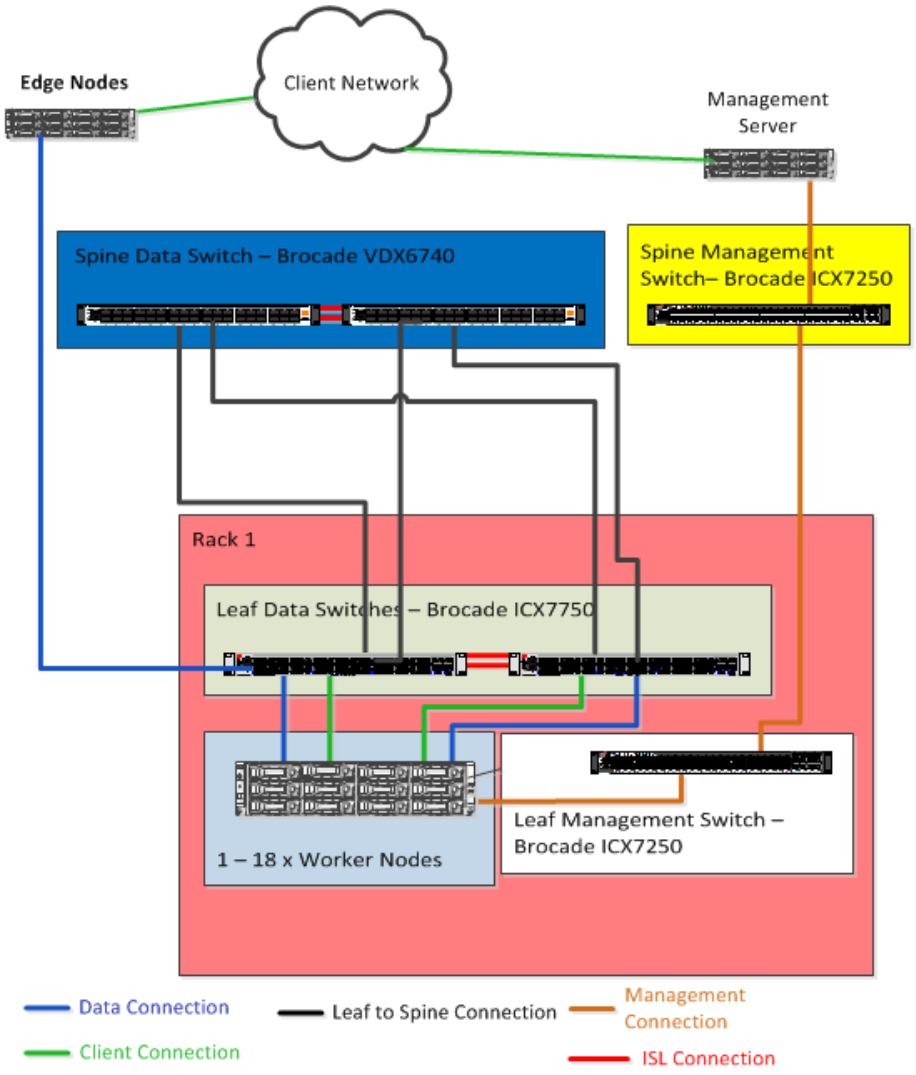

The network architecture has these components:

- “Switches” on page 13
- “Data Network” on page 13
- “Client Network” on page 13
- “Management Network” on page 13

Switches

This solution requires the following three types of switches:

- **Spine Data Switches — Brocade VDX 6940-36**

  These spine data switches interconnect leaf switches from different racks.

  Connect two switches together using an inter-switch link (ISL). This lets both switches act together as a single logical switch. If one switch fails, there still is a path to the hosts.

- **Leaf Data Switches — Brocade ICX 7750-48**

  These leaf data switches connect all nodes in a rack together. Then, the leaf switches are uplinked to the spine data switches.

  Connect two switches together using stacking. This lets both switches act as one single logical switch. If one switch fails, there still is a path to the hosts.

- **Leaf and Spine Management Switches — Brocade ICX 7250-48**

  These leaf and spine switches connect the management ports of the hardware to the management server. When there is more than one rack, use a spine switch to connect all the management leaf switches together.

This reference architecture uses Brocade switches with a 40 GB uplink network. However, you may be able to use different switches.

Data Network

Use the data network for communications between the nodes. It uses the 2 × 10 GbE on motherboard ports.

Provide redundancy with two network interfaces configured at the operating system level to use the **active-active** network-bonding mode.

Client Network

The client network is an optional network. It is usually used on edge nodes and management servers.

Management Network

The management network allows for access to the nodes using the 1 GbE LAN on motherboard (LOM) interface. This network provides out of band monitoring and management of the servers.

Sample System Configurations

You can use many different configurations for Hitachi Unified Compute Platform for SAP HANA as a foundation for business analytics solution.

- “Basic Apache Hadoop Dense Storage Configuration ” on page 14
- “Basic Apache Hadoop Multiple Rack Configuration” on page 14
- “Apache Hadoop with Apache Spark Cluster Configuration” on page 16

Basic Apache Hadoop Dense Storage Configuration

Table 5 lists the components for the sample single rack solution in Figure 4 on page 10. It uses single node units with following:

- 12 large disks per unit
- Separate operating system disks
- One work node rack
- One master node rack

This system configuration is for low usage with high storage requirements. This solution provides 1.3 PB per rack. With a standard duplication factor of three, there is approximately 400 TB of data. This configuration uses the sample master node rack.

TABLE 5. SAMPLE SINGLE RACK CONFIGURATION

Component	Number	Description
Rack	1	42U Rack
Single node server	1 rack × 18 servers	2U Apache Hadoop server, 18 servers per rack
Brocade ICX7750 switch	2	2 TOR switches per rack
Brocade ICX7250 switch	1	1 TOR management switch per rack
Spine switches	0	In a single rack configuration, you do not need spine switches.
Memory per node	128 GB per node	
Storage	12 × 6 TB SAS 7200 RPM HDD per node	72 TB per node

Basic Apache Hadoop Multiple Rack Configuration

Table 6, “Sample Three Rack, One Data Center Configuration,” on page 15 lists the components for a three rack, one data center configuration. Figure 6 on page 15 shows the layout of this sample solution. This sample solution uses the following:

- Three worker node racks
- One master node rack

Each rack is one-half filled with data nodes. This provides space to add more nodes. You can locate the racks as follows:

- Locate the racks in one data center
- Locate the racks in multiple data centers, to provide disaster recovery capability

TABLE 6. SAMPLE THREE RACK, ONE DATA CENTER CONFIGURATION

Component	Number	Description
Rack	3	42U Rack
Single node server	3 racks × 9 servers	2U Apache Hadoop server, 9 per rack
Hadoop Leaf, Brocade ICX7750 switch	3 × 2 per rack	2 TOR switches per rack
Management Leaf, Brocade ICX7250 switches	3 × 1 per rack	1 TOR management switch per rack
Hadoop Spine, Brocade VDX6940 switch	2	Redundant spine switches shared across multiple racks
Management Spine, Brocade ICX7250 switch	1	Spine switch, shared across multiple racks
Memory per node	128 GB	
Storage	12 × 2 TB SAS 7200 RPM HDD per node	24 TB per node

Figure 6

Apache Hadoop with Apache Spark Cluster Configuration

When designing a cluster to run applications using the Apache Spark computing engine, several extra node types are necessary. These extra node types are the following:

- Apache Spark master node
- Apache Spark worker node
- Apache ZooKeeper master node
- Apache ZooKeeper standby master node

The different types of nodes have different requirements. Some can be run on the same node, such as the Spark master node and standby master node, for example.

This reference architecture has Spark worker nodes running on the same node as the data nodes. In this sample configuration, the worker nodes are in the edge or master rack. Table 7 shows this sample configuration.

TABLE 7. SAMPLE APACHE HADOOP WITH APACHE SPARK WORKER NODES CONFIGURATION

Component	Number	Description
Rack	1	42U Rack
Single node server	1 rack × 18 servers: ■ 11 worker nodes ■ 7 Edge or Master Nodes	2U Apache Hadoop server, 18 servers per rack
Brocade ICX7750 switch	2	2 ToR switches per rack
Brocade ICX7250 switch	1	1 ToR management switch per rack
Spine switches	0	In a single rack configuration, you do not need spine switches.
Memory per worker node	512 GB per node	
Storage	10 × 1.2 TB SAS 7200 RPM HDD per node for data 2 x 800 TB SATA SSD for Spark	12 TB HDD for data 1.6 TB SDD for Spark storage

Figure 7 on page 17 shows the rack layout for this sample solution. There can be multiple components installed on each node. Some of the key components that are part of this solution are the following:

- Ambari
- Apache ZooKeeper master node
- Apache Spark master node or ZooKeeper standby node
- NameNode for Apache Hadoop Distributed File System
- NameNode standby
- A node for other master process
- Pentaho node
- Apache Spark worker nodes and Apache Hadoop worker nodes
- ToR switches

When using Spark with Hadoop, it is recommended to collocate the Apache Spark worker nodes with the Apache Hadoop worker nodes.

Figure 7

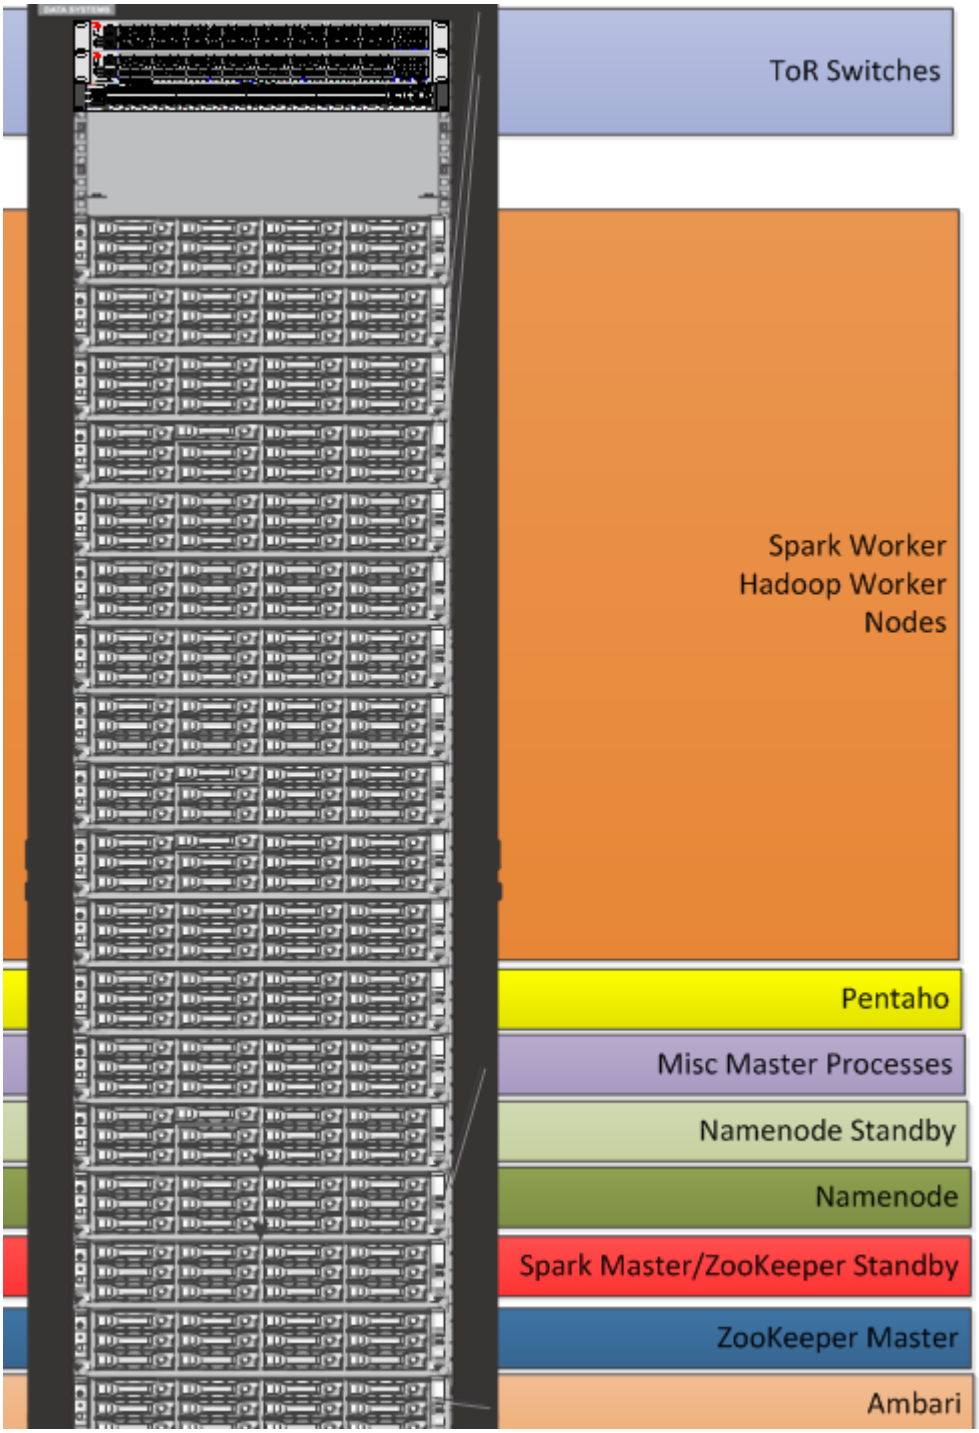

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Resources website.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Training and Certification website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-607-00, July 2017.