EMULEX
An Avago Technologies Company

Deploying Cloudera CDH (Cloudera Distribution Including Apache Hadoop) with Emulex OneConnect OCe14000 Network Adapters

Table of Contents

Introduction

The rapid growth of social media, cellular advances and requirements for data analytics has challenged the traditional methods of data storage and data processing for many large business and government entities. To solve the data storage and processing challenges, organizations are starting to deploy large clusters of Apache Hadoop – a solution that helps manage the vast amount of what is commonly referred to as big data. The Emulex OneConnect® family of OCe14000 10Gb Ethernet (10GbE) Network Adapters plays an important role in the Hadoop cluster, to move the data efficiently across the cluster.

This lab guide describes the necessary hardware, software and configuration steps needed to deploy the Cloudera CDH 5.x (Cloudera Distribution Including Apache Hadoop) with the Emulex family of OCe14000 10GbE Network Adapters. A brief introduction to centralized management of OCe14000 Adapters using Emulex OneCommand® Manager is also being reviewed.

Intended audience: System and network architects and administrators

Hardware requirements

For this lab guide we implemented a five-node cluster, but this is scalable as required. Adding a new DataNode to a Hadoop cluster is a very simple process. However, NameNode's RAM and disk space must be taken into consideration before adding additional DataNodes.

NameNode is the most important part of a Hadoop Distributed File System (HDFS). It keeps track of the directory tree of all the files in the file system, and tracks where in the cluster the file data is kept. It also is the single point of failure in a Hadoop cluster. With Hadoop 2.0, this issue has been addressed with the HDFS High Availability (HA) feature (refer to Apache's "HDFS High Availability Guide"). It is always recommended to implement high availability in the cluster.

A DataNode stores data in the HDFS. A cluster will always have more than one DataNode, with data replicated across all of them.

The required hardware for implementing and testing Hadoop with OCe14000 10GbE Adapters is listed below:

Hardware components	Quantity	Description	Comments
Server	4 or more	Any server with Intel/AMD processors, which support Linux	1 NameNode (MasterNode); 3 or more DataNode (SlaveNode, JobHistoryNode, ResourceManager)
Hard drives	2 or more per server	Any SAS or SATA drive	
RAID controller	4 or more	Anyserver, which supports Linux	
RAM	48GB+ per server		
Emulex OCe14000 Network Adapters	4 or more	10GbE network adapters	The OCe14102-UM adapter was used in this configuration
Switch	1	10Gbps switch	
Cables	4 or more	10Gbps optical SFP+ cables	

Figure 1. List of hardware required.

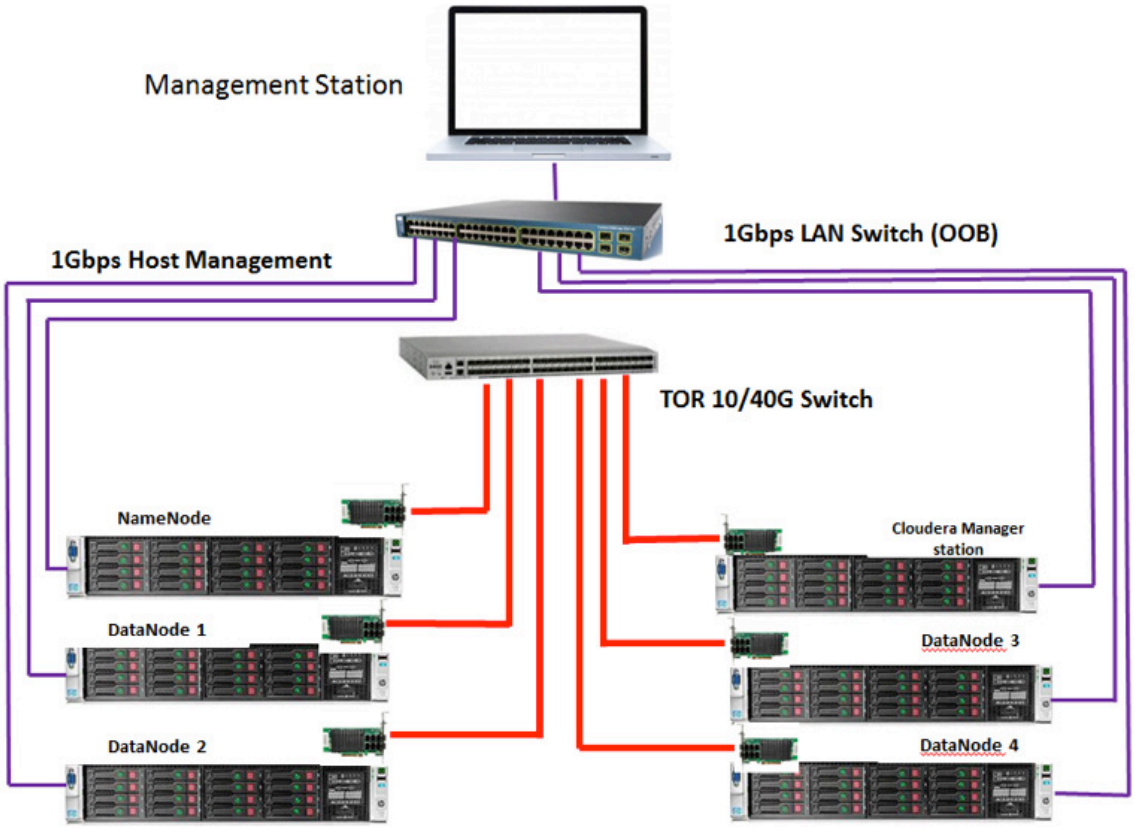

Figure 2. Sample configuration.

Recommended Hadoop cluster hardware specification

There are a lot of factors affecting the choice of hardware for a new Hadoop cluster. Hadoop runs on industry-standard hardware. However, selecting the right mix of CPU, hard drives or RAM for any workload can help you make the Hadoop cluster more efficient. The recommendations below are formulated from Cloudera. Please consider your organization's workload before selecting the hardware components.

Node	Hardware Components	Specifications
NameNode	CPU	2 quad/hex/oct core CPUs, running at least 2GHz
	Hard Drive	2 or more 1TB hard disks in a RAID configuration
	RAM	48 – 128GB
DataNode	CPU	2 quad/hex/oct core CPUs, running at least 2GHz
	Hard Drive	2 or more 1TB hard disks in a JBOD configuration
	RAM	64 – 512GB

Figure 3. Recommended server specifications for a NameNode and DataNode.

Software requirements

Software Components	Quantity	Description	Application Note Components
OS	5 or more	Any supported Linux or Windows OS	CentOS 6.4
Java	5 or more	Java 1.6.x or higher	Java 1.7.0
Cloudera Manager	1	Centralized management for Cloudera's CDH	Cloudera Manager 5.3.3
CDH 5	1	CDH 4.x or higher for implementing Hadoop 2.x	CDH 5.0.6
OCe14102 Firmware	5 or more	Download the latest firmware from the Emulex website.	10.2.370.19
OCe14102 Driver	5 or more	Download the latest driver from the Emulex website.	10.2.363.0
Emulex OneCommand Manager	5 or more	Download the latest version of OneCommand Manager from the Emulex website.	10.2.370.16

Figure 4. Software requirements.

Installation and configuration of servers

Install CentOS 6.4 on five servers. For this lab guide, five different names were assigned to the servers. Essential services like ResourceManager and JobHistory Server were split across the SlaveNode. The names along with the roles are listed below:

1. Elephant : NameNode Server (MasterNode)
2. Monkey : JobHistory Server, DataNode (SlaveNode)
3. Horse : ResourceManager, DataNode (SlaveNode)
4. Tiger : DataNode (SlaveNode)
5. Lion : DataNode (SlaveNode)
6. Workstation-centos: Centralized Cloudera Manager Server

Connect the OCe14102 Adapter to the PCI Express (PCIe) slots. Upgrade the adapter with the latest firmware, driver and version of OneCommand Manager.

Connect port 0 of every server to the top-of-rack (ToR) switch and ensure that the link is up.

Note: There will be a system reboot required for upgrading that firmware.

1. Verification of NIC profile/firmware/driver versions for all servers.

Verify the Network Interface Card (NIC) profile in the following steps:

a. Start OneCommand Manager.

b. Select the OCe14102-UM Adapter and go to the Adapter Configuration tab.

c. Ensure that the personality is set to NIC.

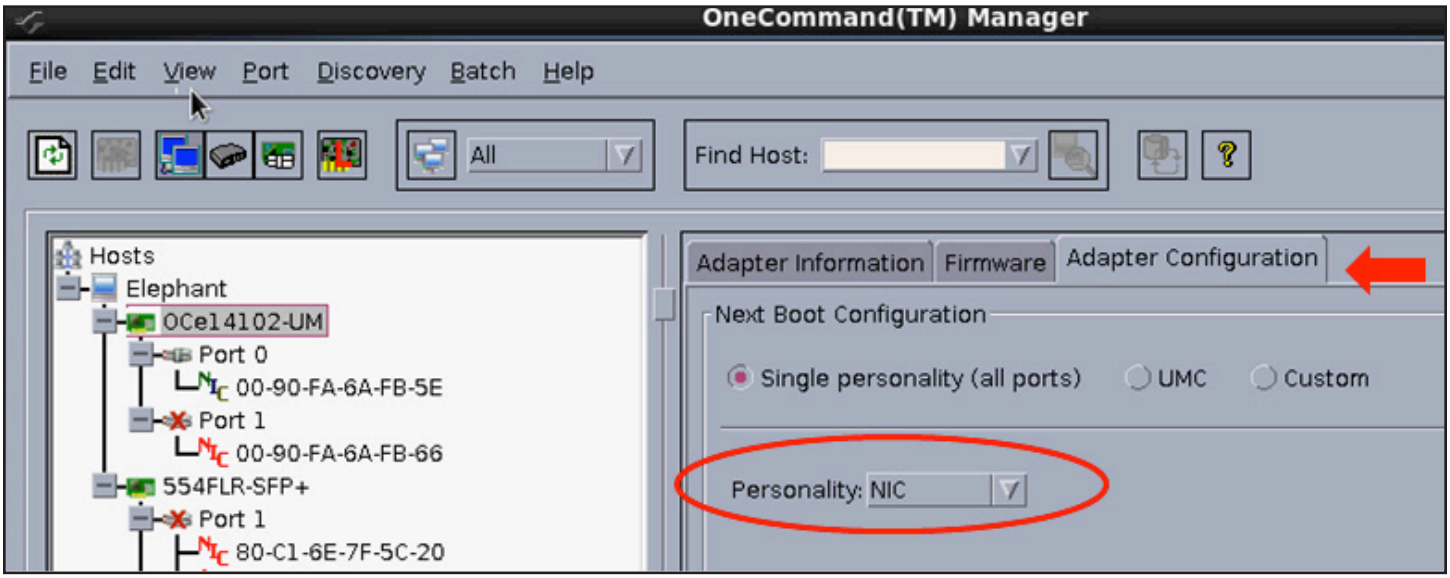

Figure 5. NIC personality verification.

Verify the firmware version in the following steps:

a. Select the OCe14102-UM Adapter and go to the Firmware tab.

b. Ensure that the firmware version is same as the one in which you have upgraded.

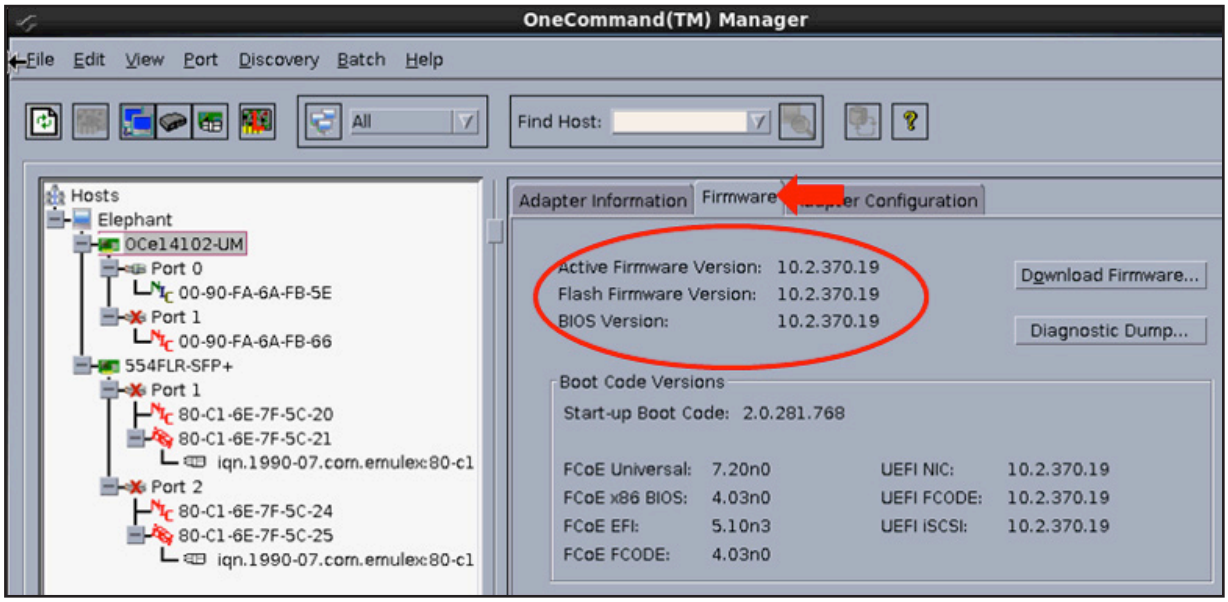

Figure 6. Firmware verification.

Verify the the driver version in the following steps:

a. Start OneCommand Manager.

b. Select NIC from Port 0 or Port 1 of the OCe14102-UM Adapter and go to the Port Information tab.

c. Ensure that the driver version is the same as the one in which you have upgraded.

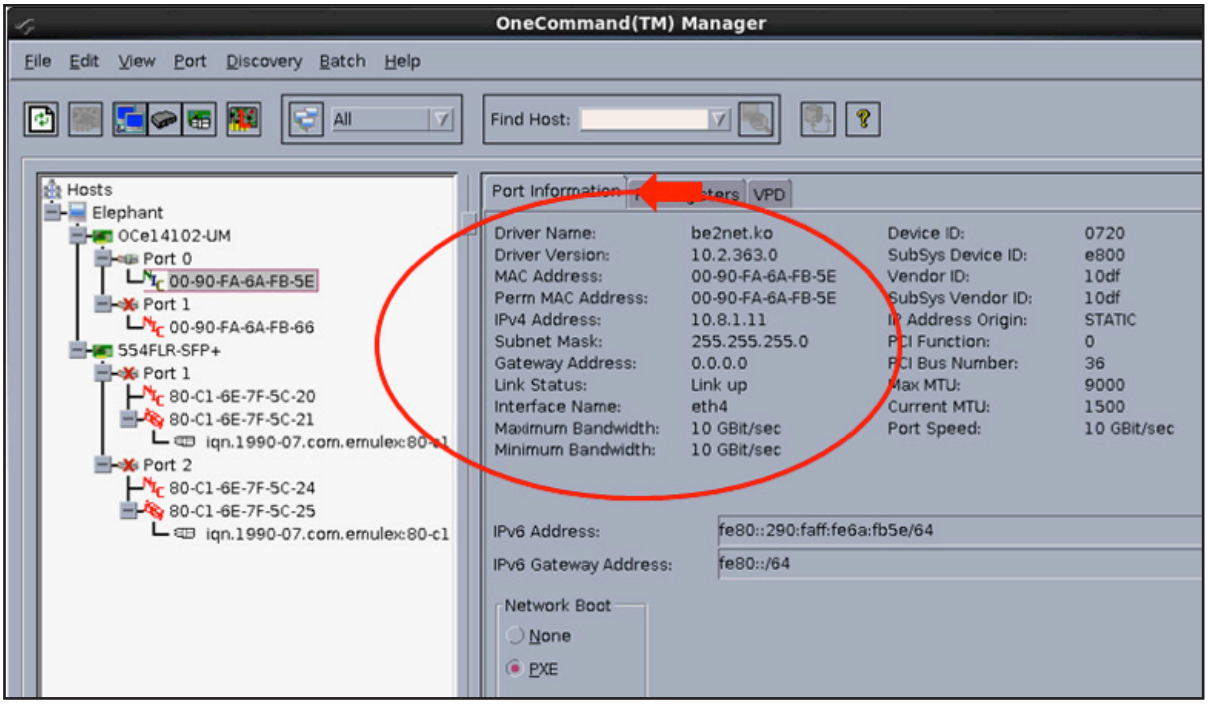

Figure 7. Driver verification.

2. Assign IPv4 address to port 0 on every server.

a. Assign the following IP addresses to port 0 of the OCe14102 Adapter using Network Manager or ipconfig.

b. Verify the IP address assignment using OneCommand Manager.

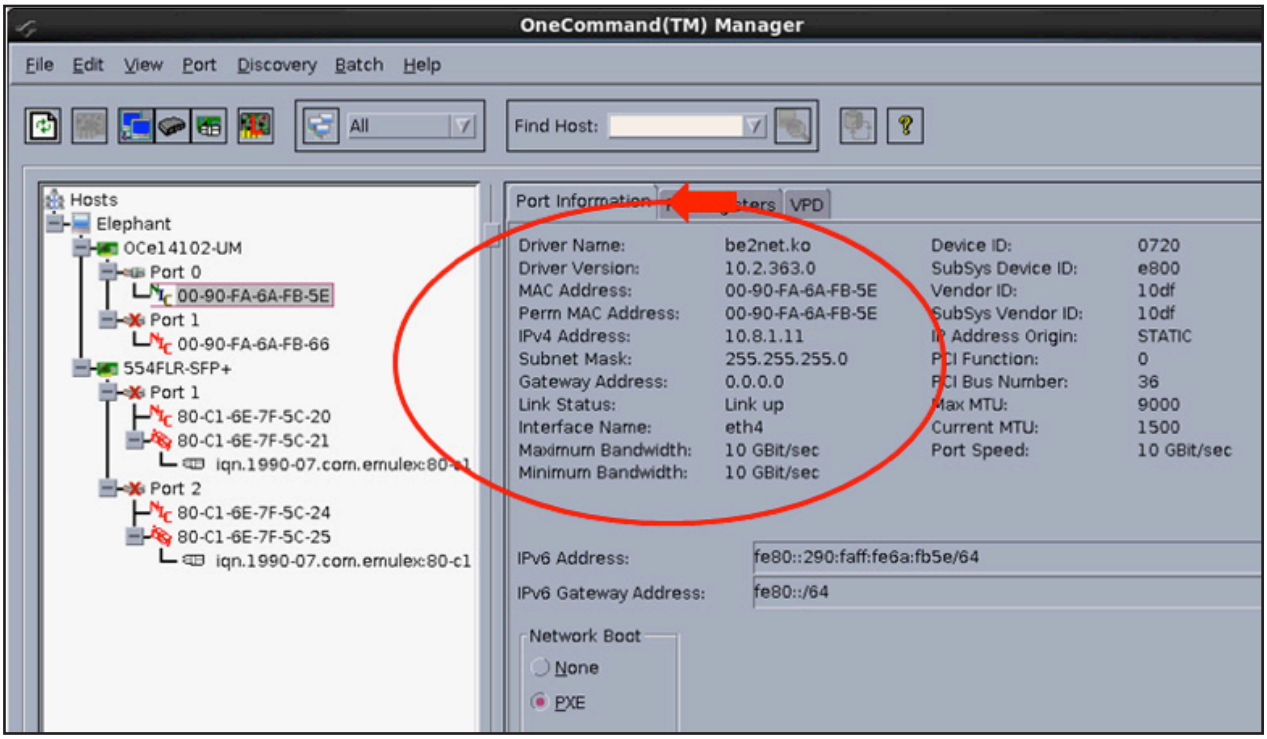

Figure 8. IP address verification.

3. Setup passwordless SSH on all servers.

a. Generate authentication Secure Shell (SSH)-keygen keys on Elephant using ssh-keygen-t rsa.

```
[root@Elephant ~]# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
/root/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
e4:0d:55:25:e0:10:e3:c9:cc:3e:0a:20:79:8e:5c:a1 root@Elephant
The key's randomart image is:
+--[ RSA 2048]----+
|    .    +.ooo..  |
| .. .  = *   .    |
|oEo.    O .       |
|.=..   + o        |
|... .   S .       |
|     . . .        |
|      .           |
|                  |
|                  |
+-----------------+
```

Figure 9. ssh-keygen output.

Note: logged in user should have privileges to run the ssh-keygen command or run the sudo command before running the ssh-keygen command.

b. Run ssh-copy-id -i ~/.ssh/id_rsa.pub root@HOSTNAME for Monkey, Lion, Tiger and Horse.

```
[root@Elephant ~]# ssh-copy-id -i ~/.ssh/id_rsa.pub root@monkey
root@monkey's password:
Now try logging into the machine, with "ssh 'root@monkey'", and check in:

  .ssh/authorized_keys

to make sure we haven't added extra keys that you weren't expecting.
```

Figure 10. ssh-copy sample output.

Note: ssh-copy-id command must be run for all the servers in the cluster to enable the login to SSH without a password (passwordless login).

c. Verify that the password ssh is working.

```
[root@Elephant ~]# ssh horse
Last login: Tue Oct  7 09:53:56 2014 from 10.193.253.33
[root@Horse ~]# exit
logout
Connection to horse closed.
[root@Elephant ~]#
```

Figure 11. Verification of passwordless ssh.

4. Configure/etc/hosts file.

The host names of the servers in the cluster along with the corresponding IP address need to be added to the /etc/hosts file. This is used for the operating system to map host names to IP addresses.

a. The Port 0 IP assigned to the OCe14000 series adapter corresponding to each host should be added to the /etc/hosts file. Add the lines listed below to the /etc/hosts file on Elephant.

10.8.1.11 Elephant
10.8.1.15 Horse
10.8.1.12 Monkey
10.8.1.12 Tiger
10.8.1.14 Lion
10.8.1.45 workstation-cento6u4

b. Copy the /etc/hosts file from Elephant to Horse, Monkey, Tiger, Lion and workstation-cento6u4 using the scp command

```
[root@Elephant ~]# scp /etc/hosts root@horse:/etc/hosts
hosts                                         100%   99     0.1KB/s   00:00
```

Figure 12. scp/etc/hosts.

Note: /etc/hosts should be copied to all the servers in the Hadoop cluster.

5. Install Java
Download and install Java on all of the hosts in the Hadoop cluster.

6. Disable Firewall
For this Application Note, the firewall was disabled. Please consult your network administrator for allowing the necessary services by the firewall to make the Hadoop cluster work.

Installing and Configuring CDH Hadoop

Notes:

- Please follow the Cloudera instructions on installing CDH Hadoop according to your environment's need.
- For this guide we have implemented CDH 5.3 using the package method using Cloudera Manager.
- All config file changes are made using Cloudera Manager and are deployed on the clients using Cloudera Manager.
- Config files for Hadoop are present under /usr/lib/hadoop/etc/hadoop or /etc/hadoop.
- Scripts are located in the Appendix.

1. **Download Cloudera Manager on the centralized management server.** Follow the instructions listed on the page to install the Cloudera Manager.
2. **Open a browser and type the IP address and login.** Select Cloudera Express from the start page.

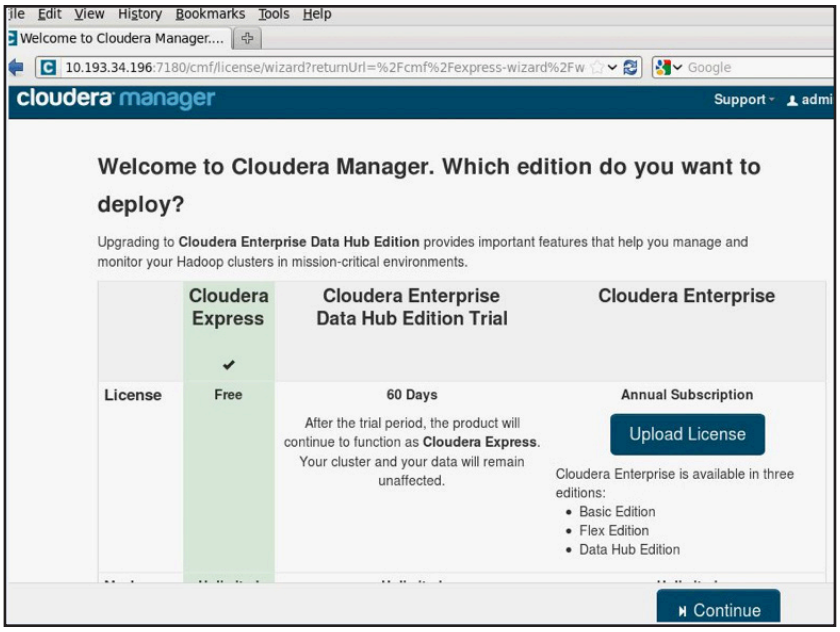

Figure 13. Cloudera Manager.

3. Follow the instructions to install CDH 5.3.3

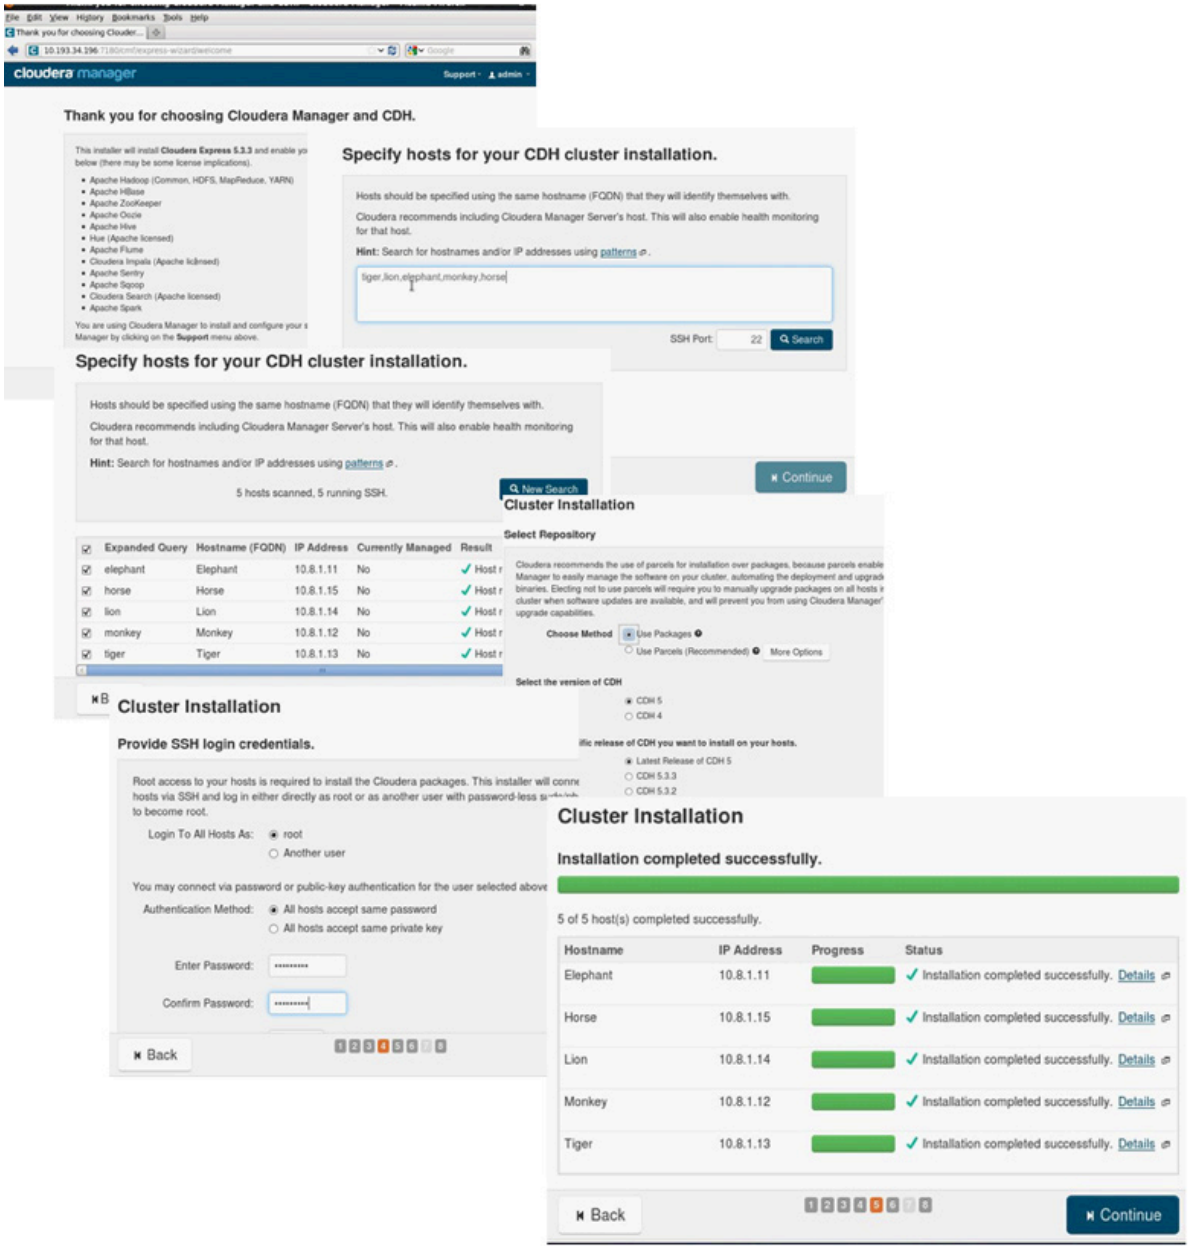

Figure 14. CDH 5.3.3 install.

4. **Install the services required for your environment.** For this application note, under Custom Services, HDFS and YARN services were chosen. Assign roles to the host in the cluster.

Figure 15. CDH 5.3.3 install.

5. **Set up the slaves, *.xml config files and hadoop_env.sh according to the needs of your environment.** For this lab guide, the configuration changes were made for HDFS and YARN under the configuration tab:

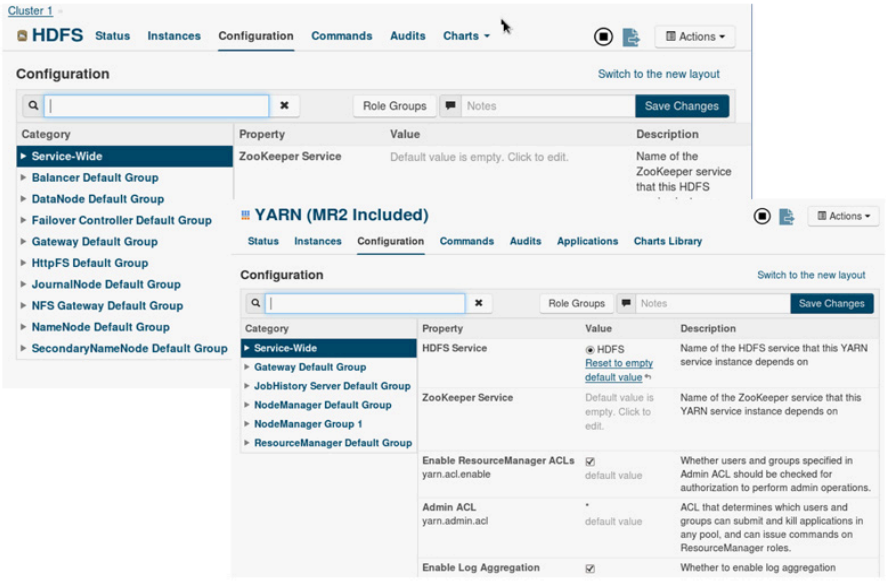

Figure 16. Configuration changes.

6. **Below are the configuration changes which were made:**

Name	Value
mapreduce.framework.name	yarn
mapreduce.jobhistory.address	monkey:10020
mapreduce.jobhistory.webapp.address	monkey:19888
yarn.app.mapreduce.am.staging-dir	/user
mapred.child.java.opts	-Xmx512m
dfs.namenode.name.dir	file:///disk1/dfs/nn,file:///disk2/dfs/nn
dfs.datanode.data.dir	file:///root/Desktop/data1/dn
yarn.resourcemanager.hostname	horse
yarn.application.classpath	Leave the value specified in the file
yarn.nodemanager.aux-services	mapreduce_shuffle
yarn.nodemanager.local-dirs	file:///root/Desktop/data1/nodemgr/local
yarn.nodemanager.log-dirs	/var/log/hadoop-yarn/containers
yarn.nodemanager.remote-app-log-dir	/var/log/hadoop-yarn/apps
n.log-aggregation-enable	true

Figure 17. Configuration parameters.

7. Deploy the configuration to the clients by going to the home page and under Cluster 1, click on "Deploy Client Configuration".

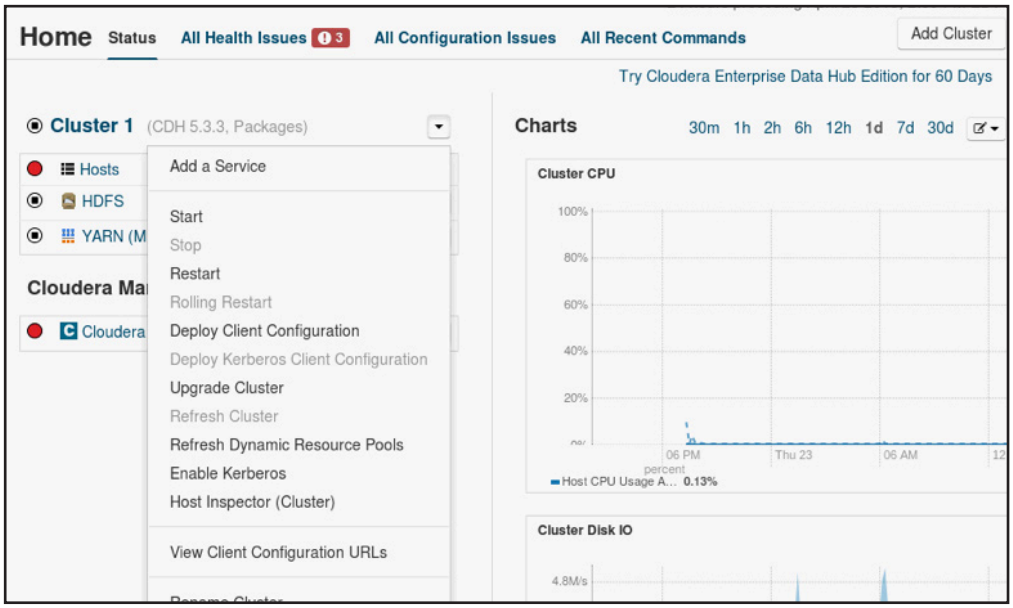

Figure 18. Deploy client configuration.

8. Start the Cloudera Management Service by clicking on Start under "Cloudera Management Service Actions".

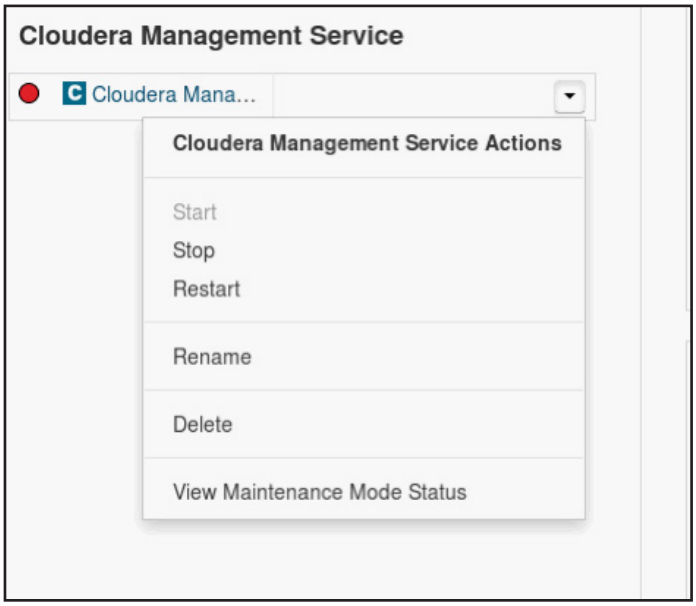

Figure 19. Start Cloudera management service.

9. Start all the services by clicking on Start under "Cluster1" on the home page.

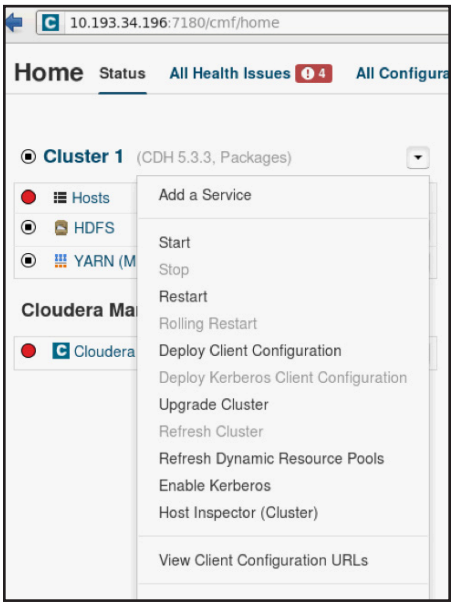

Figure 20. Start the services in the cluster.

10. The status and roles on the clients can be seen under the "Hosts" tab.

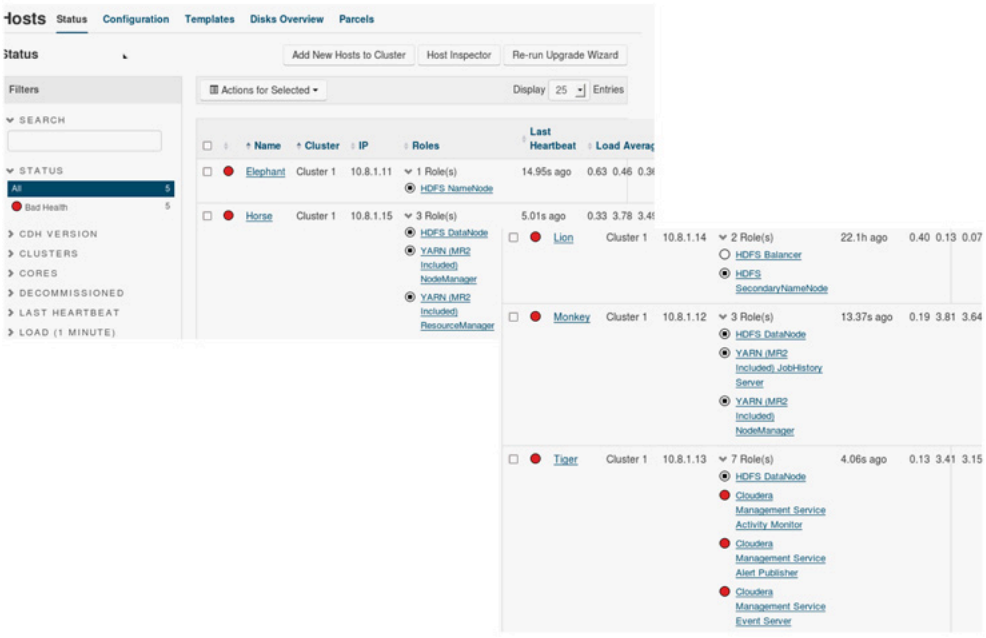

Figure 21. Service status on clients.

Verification of Hadoop cluster

To verify the Hadoop cluster, a simple wordcount example can be run.

1. Create a sample file.

Create a sample file with a few test sentences using the VI editor on Elephant: **vi input_word_count**.

```
This is a test document to verify the wordcount output and verify the working o
 cluster
This is line two
This is a simple way to test and verify the Hadoop Cluster
```

Figure 22. Sample file.

2. Copy the file to HFDS.

For Elephant, create a directory named "in" in HDFS and copy the sample file **input_word_count** to HDFS under the "in" directory:

```
hadoop fs -mkdir /in

hadoop fs -put input_word_count /in

hadoop fs -is /in
```

The directory, file and the actual location of the file can be viewed using a web interface on the NameNode. The web address is "elephant:50070". This will also show that the replication factor of each block is 3 and it is saved at three different DataNodes (Figures 21 and 22).

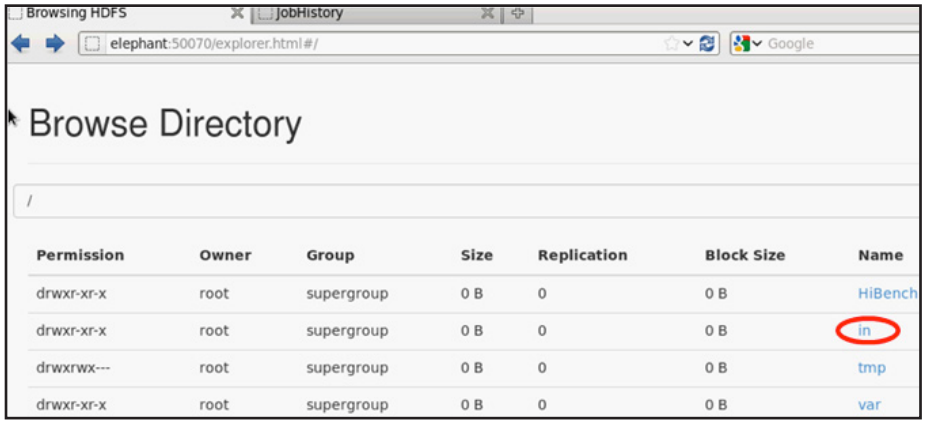

Figure 23. Verification of HFDS directory listing using web interface.

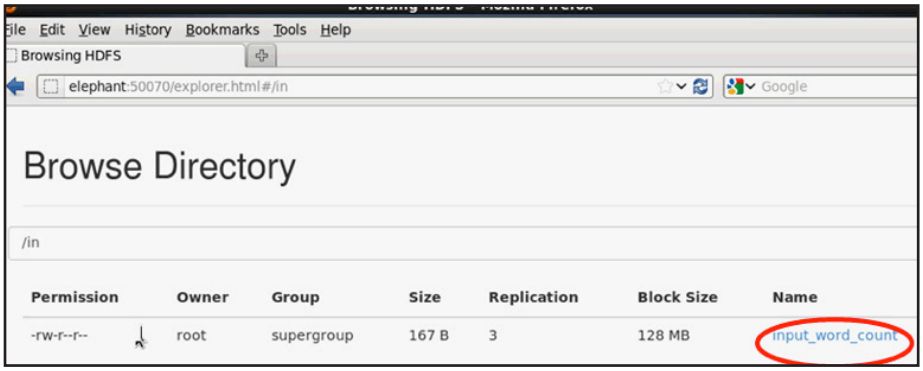

Figure 24. File size in HFDS.

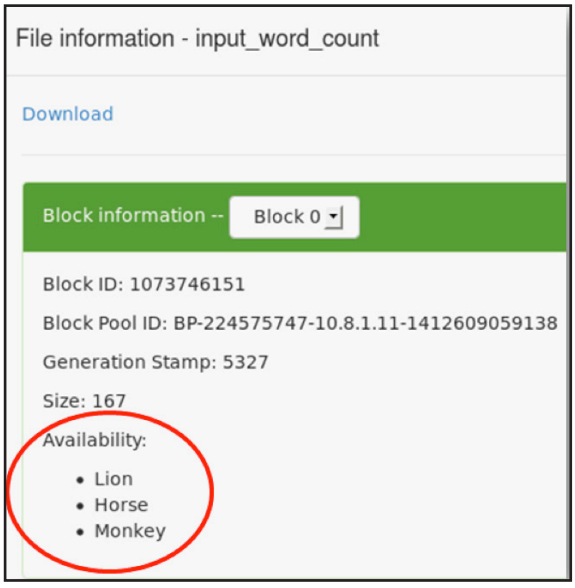

Figure 25. Replication factor and file location information on the HFDS.

3. Run the inbuilt *.jar application.

Run the command listed below to start the wordcount application.

```
/usr/lib/hadoop/bin/hadoop-mapreduce/hadoop-mapreduce-examples.jar wordcount /in/input_word_count /out
```

The output will be displayed on the terminal. The result location can also be viewed using the web interface.

```
[root@Elephant hadoop-2.4.1]# hadoop fs -tail /out/part-r-00000
14/10/16 15:04:12 WARN util.NativeCodeLoader: Unable to load nat
Cluster 1
Hadoop  1
This    3
a       2
and     2
cluster 1
document        1
is      3
line    1
of      1
output  1
simple  1
test    2
the     3
to      2
two     1
verify  3
way     1
wordcount       1
working 1
```

Figure 26. Sample output.

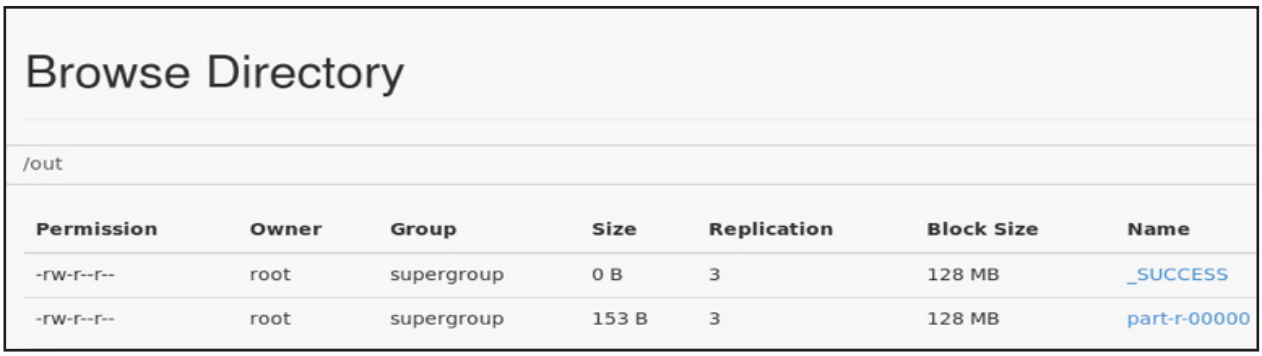

Figure 27. Sample web interface output in HFDS.

Conclusion

A complete overview of hardware and software components required for successfully deploying and evaluating Cloudera CDH Hadoop with Emulex OCe14000 Network Adapters was presented in this lab guide. More details about the Hadoop architecture and Cloudera can be obtained from the Apache Hadoop official website: http://hadoop.apache.org/ and http://www.cloudera.com/content/cloudera/en/home.html respectively. The solution described in this lab guide is scalable for a larger number of DataNodes and racks to suit the environment and needs of your organization. Emulex OCe14000 Network Adapters can be used in the cluster to move the data efficiently across the nodes.

References

Solution Implementer's Lab:
http://www.implementerslab.com/

Cloudera :
http://www.cloudera.com/content/cloudera/en/home.html

Emulex Ethernet networking and storage connectivity products:
http://www.emulex.com/products/ethernet-networking-storage-connectivity/

Apache Hadoop overview:
http://hadoop.apache.org/

How to set up a multimode Hadoop cluster:
http://wiki.apache.org/hadoop/#Setting_up_a_Hadoop_Cluster

Cloudera resources:
http://www.cloudera.com/content/cloudera/en/resources.html

Wikipedia:
http://en.wikipedia.org/wiki/Apache_Hadoop

DataNode:
http://wiki.apache.org/hadoop/DataNode

NameNode:
http://wiki.apache.org/hadoop/NameNode

EMULEX
An Avago Technologies Company

For product information and a complete list of distributors, please visit our website at www.emulex.com

ELX15-2579· 5/15