

QUESTION 1

Practice Test		
Install	32 questions	1 To 32
Configure	11 questions	33 To 43
Manage	19 questions	44 To 62
Secure	20 questions	63 To 82
Test	21 questions	83 To 103
Troubleshoot	9 questions	104 To 112

Question No :1

Create Network configuration in Linux OS instance.

Correct Answer:

Explanation/Reference:

See the explanation for Step by Step Solution and configuratio

```
root@master:~  
login as: opc  
Authenticating with public key "rsa-key-20190923"  
[opc@master ~]$ sudo su -  
[root@master ~]# ifconfig -a  
ens3: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000  
    inet 10.0.0.6 netmask 255.255.255.0 broadcast 10.0.0.255  
    inet6 fe80::200:17ff:fe02:2753 prefixlen 64 scopeid 0x20<link>  
    ether 00:00:17:02:27:53 txqueuelen 1000 (Ethernet)  
    RX packets 6331 bytes 59330324 (56.5 MiB)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 2980 bytes 392171 (382.9 KiB)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536  
    inet 127.0.0.1 netmask 255.0.0.0  
    inet6 ::1 prefixlen 128 scopeid 0x10<host>  
    loop txqueuelen 1000 (Local Loopback)  
    RX packets 6 bytes 416 (416.0 B)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 6 bytes 416 (416.0 B)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
[root@master ~]#
```

```
root@master:~  
login as: opc  
Authenticating with public key "rsa-key-20190923"  
[opc@master ~]$ sudo su -  
[root@master ~]# ifconfig -a  
ens3: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000  
    inet 10.0.0.6 netmask 255.255.255.0 broadcast 10.0.0.255  
    inet6 fe80::200:17ff:fe02:2753 prefixlen 64 scopeid 0x20<link>  
    ether 00:00:17:02:27:53 txqueuelen 1000 (Ethernet)  
    RX packets 6331 bytes 59330324 (56.5 MiB)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 2980 bytes 392171 (382.9 KiB)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536  
    inet 127.0.0.1 netmask 255.0.0.0  
    inet6 ::1 prefixlen 128 scopeid 0x10<host>  
    loop txqueuelen 1000 (Local Loopback)  
    RX packets 6 bytes 416 (416.0 B)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 6 bytes 416 (416.0 B)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
[root@master ~]#
```

```
[root@master network-scripts]# more ifcfg-ens3
# Created by cloud-init on instance boot automatically, do not edit.
#
BOOTPROTO=static
DEVICE=ens3
HWADDR=00:00:17:02:27:53
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
NM_CONTROLLED=no
IPADDR=10.0.0.6
NETMASK=255.255.255.0
[root@master network-scripts]#
```

```
[root@master network-scripts]# more ifcfg-ens3
# Created by cloud-init on instance boot automatically, do not edit.
#
BOOTPROTO=static
DEVICE=ens3
HWADDR=00:00:17:02:27:53
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
NM_CONTROLLED=no
IPADDR=10.0.0.6
NETMASK=255.255.255.0
[root@master network-scripts]#
```

```
[root@master network-scripts]# more ifcfg-ens3
# Created by cloud-init on instance boot automatically, do not edit.
#
BOOTPROTO=static
DEVICE=ens3
HWADDR=00:00:17:02:27:53
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
NM_CONTROLLED=no
IPADDR=10.0.0.6
NETMASK=255.255.255.0
[root@master network-scripts]#
```

```
[root@master network-scripts]# more ifcfg-ens3
# Created by cloud-init on instance boot automatically, do not edit.
#
BOOTPROTO=static
DEVICE=ens3
HWADDR=00:00:17:02:27:53
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
NM_CONTROLLED=no
IPADDR=10.0.0.6
NETMASK=255.255.255.0
[root@master network-scripts]#
```

QUESTION 2

Please add httpfs roles to instance node01.

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.

The screenshot shows the Cloudera Manager web interface. The browser address bar displays `master.subnet.vcn.oraclevcn.com:7180/cmf/services/5/instances`. The page title is "Cloudera Manager" and the current view is "Cluster 1" with "HDFS" selected. The navigation menu includes "Status", "Instances", "Configuration", "Commands", "File Browser", "Charts Library", "Cache Statistics", "Audits", "Web UI", and "Quick Links". A search bar is present below the navigation menu. On the left, there is a "Filters" sidebar with sections for "STATUS" (None: 1, Good Health: 14), "COMMISSION STATE", "MAINTENANCE MODE", "RACK", "ROLE GROUP", "ROLE TYPE", "STATE", and "HEALTH TESTS". The main content area shows a table of instances with columns: Role Type, State, Host, Commission State, and Role Group. The "Add Role Instances" button is circled in red. Below the table, there are buttons for "Actions for Selected", "Migrate Roles", "Add Role Instances", "Federation and High Availability", and "Role Groups".

<input type="checkbox"/>	Role Type	State	Host	Commission State	Role Group
<input type="checkbox"/>	Balancer	N/A	master.subnet.vcn.oraclevcn.com	Commissioned	Balancer Default Group
<input type="checkbox"/>	DataNode	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	DataNode Group 1
<input type="checkbox"/>	DataNode	Started	node03.subnet.vcn.oraclevcn.com	Commissioned	DataNode Default Group
<input type="checkbox"/>	DataNode	Started	node02.subnet.vcn.oraclevcn.com	Commissioned	DataNode Default Group
<input type="checkbox"/>	Failover Controller	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	Failover Controller Default Group
<input type="checkbox"/>	Failover Controller	Started	master.subnet.vcn.oraclevcn.com	Commissioned	Failover Controller Default Group

Search

Filters

- STATUS
 - None 1
 - Good Health 14
- COMMISSION STATE
- MAINTENANCE MODE
- RACK
- ROLE GROUP
- ROLE TYPE
- STATE
- HEALTH TESTS

Actions for Selected ▾ Migrate Roles Add Role Instances Federation and High Availability Role Groups

<input type="checkbox"/>	Role Type	State	Host	Commission State	Role Group
<input type="checkbox"/>	Balancer	N/A	master.subnet.vcn.oraclevcn.com	Commissioned	Balancer Default Group
<input type="checkbox"/>	DataNode	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	DataNode Group 1
<input type="checkbox"/>	DataNode	Started	node03.subnet.vcn.oraclevcn.com	Commissioned	DataNode Default Group
<input type="checkbox"/>	DataNode	Started	node02.subnet.vcn.oraclevcn.com	Commissioned	DataNode Default Group
<input type="checkbox"/>	Failover Controller	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	Failover Controller Default Group
<input type="checkbox"/>	Failover Controller	Started	master.subnet.vcn.oraclevcn.com	Commissioned	Failover Controller Default Group

Add Role Instances to HDFS

- Assign Roles
- Review Changes

Assign Roles

You can specify the role assignments for your new roles here.

You can also view the role assignments by host. [View By Host](#)

Gateway Select hosts	HttpFS × 1 New node01.subnet.vcn.oracle	SecondaryNameNode Select hosts	NFS Gateway × 4 Select hosts
JournalNode × 3 Select hosts	NameNode × 2 Select hosts	Failover Controller × 2 Select hosts	DataNode × 3 Select hosts

Add Role Instances to HDFS

- Assign Roles
- Review Changes

Assign Roles

You can specify the role assignments for your new roles here.

You can also view the role assignments by host. [View By Host](#)

Gateway Select hosts	HttpFS × 1 New node01.subnet.vcn.oracle	SecondaryNameNode Select hosts	NFS Gateway × 4 Select hosts
JournalNode × 3 Select hosts	NameNode × 2 Select hosts	Failover Controller × 2 Select hosts	DataNode × 3 Select hosts

[Back](#) [Continue](#)

Cluster 1

HDFS [Actions](#)

Sep 13, 5:42 PM GMT

- Status
- Instances**
- Configuration
- Commands
- File Browser
- Charts Library
- Cache Statistics
- Audits
- Web UI
- Quick Links

Filters [Clear All](#)

- STATUS**
 - None 0
 - Good Health 1
- COMMISSION STATE**
- MAINTENANCE MODE**
- RACK**
- ROLE GROUP**

- [Actions for Selected](#)
- [Migrate Roles](#)
- [Add Role Instances](#)
- [Federation and High Availability](#)
- [Role Groups](#)

<input type="checkbox"/>	Role Type	State	Host	Commission State	Role Group
<input checked="" type="checkbox"/>	HttpFS	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	HttpFS Default Group

25 Per Page

Cloudera Manager Clusters Hosts Diagnostics Audits Charts Backup Administration Search Support admin

Cluster 1 HDFS Actions Sep 13, 5:42 PM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

Search

Filters Clear All

- STATUS
 - None 0
 - Good Health 1
- COMMISSION STATE
- MAINTENANCE MODE
- RACK
- ROLE GROUP

Actions for Selected Migrate Roles Add Role Instances Federation and High Availability Role Groups

<input type="checkbox"/>	Role Type	State	Host	Commission State	Role Group
<input checked="" type="checkbox"/>	HttpFS	Started	node01.subnet.vcn.oraclevcn.com	Commissioned	HttpFS Default Group

25 Per Page

QUESTION 3

Change hdfs thrash retention policy

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.

Cluster 1 HDFS Actions Sep 14, 12:08 PM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

trash Role Groups History and Rollback

Filters

SCOPE

HDFS (Service-Wide)	1
Balancer	0
DataNode	0
Gateway	1
HttpFS	0
JournalNode	0
NFS Gateway	0
NameNode	2
SecondaryNameNode	0
Failover Controller	0

CATEGORY

Advanced	0
----------	---

Use Trash Gateway Default Group

Filesystem Trash Interval NameNode Default Group
fs.trash.interval 1 day(s)

Filesystem Trash Checkpoint Interval NameNode Default Group
fs.trash.checkpoint.interval 1 hour(s)

Audit Event Filter HDFS (Service-Wide)
navigator.event.filter

Show All Descriptions View as JSON

Cluster 1 HDFS Actions Sep 14, 12:08 PM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

trash Role Groups History and Rollback

Filters

SCOPE

HDFS (Service-Wide)	1
Balancer	0
DataNode	0
Gateway	1
HttpFS	0
JournalNode	0
NFS Gateway	0
NameNode	2
SecondaryNameNode	0
Failover Controller	0

CATEGORY

Advanced	0
----------	---

Use Trash Gateway Default Group

Filesystem Trash Interval NameNode Default Group
fs.trash.interval 1 day(s)

Filesystem Trash Checkpoint Interval NameNode Default Group
fs.trash.checkpoint.interval 1 hour(s)

Audit Event Filter HDFS (Service-Wide)
navigator.event.filter

Show All Descriptions View as JSON

QUESTION 4

All the data stored in HDFS are from Retail Bank, which will be used to analyze the customer spending

and fraud analysis. However, you want to make sure that data stored in HDFS should not display credit card number whenever audit events and information retrieved from the Audit Server database and the Metadata Server persistent storage for Cloudera Navigator property. Queried and it must be masked.

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.

The screenshot shows the Cloudera Manager interface. The browser address bar indicates the URL: `master.subnet.vcn.oraclevcn.com:7180/cm/services/3/config#filterfreeText=PII`. The page title is "Cloudera Management Service" with a timestamp of "Sep 13, 5:39 PM GMT". The navigation menu includes "Status", "Instances", "Configuration" (selected), "Commands", "Charts Library", "Audits", and "Quick Links". A search bar contains the text "PII".

On the left, the "Filters" sidebar shows a list of services and categories:

SCOPE	Count
Cloudera Management Servi...	0
Activity Monitor	0
Host Monitor	0
Service Monitor	0
Event Server	0
Reports Manager	0
Alert Publisher	0
Navigator Audit Server	1
Navigator Metadata Server	0
Telemetry Publisher	0

CATEGORY	Count
Advanced	1
Cloudera Navigator	0
Database	0
External Authentication	0
Extractor Filter	0

The main content area displays two configuration items:

- PII Masking Regular Expression**: `navigator.pii.masking.regex`
- Navigator Audit Server Default Group**: `(4[0-9]{12})(?:[0-9]{3})?(5[1-5][0-9]{14})(3[47][0-9]{13})(3(?:0[0-5][68][0-9]{0-9})(11))(6(?:011|5[0-9]{2})[0-9]{12})((?:2131|1800|35\d{3})\d{11})`

At the bottom right, there is a "25 Per Page" dropdown menu.

The screenshot shows the Cloudera Manager interface. At the top, there's a navigation bar with 'Cloudera Manager' and various menu items like Clusters, Hosts, Diagnostics, Audits, Charts, Backup, and Administration. Below that, the 'Cloudera Management Service' is selected, with a search bar and 'Role Groups' and 'History and Rollback' links. The main content area is titled 'PII' and shows a configuration for 'PII Masking Regular'. The 'Expression' field contains a complex regular expression: `navigator.pii.masking.regex`. Below the expression, there's a list of filters categorized by 'SCOPE' and 'CATEGORY'. The 'SCOPE' list includes items like 'Cloudera Management Servi...', 'Activity Monitor', 'Host Monitor', 'Service Monitor', 'Event Server', 'Reports Manager', 'Alert Publisher', 'Navigator Audit Server' (with a count of 1), 'Navigator Metadata Server', and 'Telemetry Publisher'. The 'CATEGORY' list includes 'Advanced' (count 1), 'Cloudera Navigator', 'Database', 'External Authentication', and 'Extractor Filter'. A '25 Per Page' dropdown is visible on the right side.

QUESTION 5

Accomplish following activities using httpfs and curl command.

- List the content in /user/root directory
- Create directory in /user/root/cloud
- Get the ACL detail for /user/admin1" directory
- Get the ACL detail for /user/admin1 directory but use proxy user as cert
- Read the content of the file "user/root/hadoop_examples.txt"
- Rename file "/user/root/hadoop_examples.txt" to "/user/root/hadoop_examples_new.txt"

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.

- go to ssh console - run the below command

- List the content in /user/root directory

curl

"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root?op=LISTSTATUS&user.name=r
oot"

- Create directory in /user/root/cloud

curl -i -X PUT

```
"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/cloud?op=MKDIRS&user.name=root"
```

- Get the ACL detail for /user/admin1" directory

```
curl
```

```
"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root"
```

- Get the ACL detail for /user/admin1 directory but use proxy user as cert

```
curl
```

```
"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root&doas=cert"
```

- Read the content of the file "user/root/hadoop_examples.txt"

```
curl
```

```
"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.txt?op=open&user.name=root"
```

- Rename file "/user/root/hadoop_examples.txt" to "/user/root/hadoop_examples_new.txt"

```
curl -i -X PUT
```

```
"http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.txt?op=rename&destination=/user/root/hadoop_examples_new.txt&user.name=root"
```

- if you want to check results, you'll check the below command,

```
hdfs dfs -ls /user/root
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root?op=LISTSTATUS&user.name=root" >  
{  
  "FileStatuses": {  
    "FileStatus": [ {  
      "pathSuffix": ".Trash",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "700",  
      "accessTime": 0,  
      "modificationTime": 1600020000349,  
      "blockSize": 0,  
      "replication": 0,  
      "pathSuffix": ".staging",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "700",  
      "accessTime": 0,  
      "modificationTime": 1600016412129,  
      "blockSize": 0,  
      "replication": 0,  
      "pathSuffix": "hadoop_examples.txt",  
      "type": "FILE",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "644",  
      "accessTime": 1599985053186,  
      "modificationTime": 1599985053243,  
      "blockSize": 134217728,  
      "replication": 3,  
      "pathSuffix": "hadoop_root.txt",  
      "type": "FILE",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "644",  
      "accessTime": 1599929015215,  
      "modificationTime": 1599929015259,  
      "blockSize": 134217728,  
      "replication": 3,  
      "pathSuffix": "wc_out",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "755",  
      "accessTime": 0,  
      "modificationTime": 1599985153543,  
      "blockSize": 0,  
      "replication": 0,  
      "pathSuffix": "yarn_gz.txt",  
      "type": "FILE",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "644",  
      "accessTime": 1600016323726,  
      "modificationTime": 1600016323766,  
      "blockSize": 134217728,  
      "replication": 3,  
      "pathSuffix": "yarn_root.txt",  
      "type": "FILE",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "644",  
      "accessTime": 1600016038473,  
      "modificationTime": 1600016038510,  
      "blockSize": 134217728,  
      "replication": 3,  
      "pathSuffix": "yarnout",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "755",  
      "accessTime": 0,  
      "modificationTime": 1600016127680,  
      "blockSize": 0,  
      "replication": 0,  
      "pathSuffix": "yarnoutgz",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "755",  
      "accessTime": 0,  
      "modificationTime": 1600016410027,  
      "blockSize": 0,  
      "replication": 0,  
      "pathSuffix": "zx_snapshot",  
      "type": "DIRECTORY",  
      "length": 0,  
      "owner": "root",  
      "group": "hdfs",  
      "permission": "755",  
      "accessTime": 0,  
      "modificationTime": 1599931261967,  
      "blockSize": 0,  
      "replication": 0,  
      "snapshotEnabled": true } ]  
    }  
  }  
}
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root?op=LISTSTATUS&user.name=root"
{"FileStatuses":{"FileStatus":[{"pathSuffix":"Trash","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"700","accessTime":0,"modificationTime":1600020000349,"blockSize":0,"replication":0}, {"pathSuffix":"staging","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"700","accessTime":0,"modificationTime":1600016412129,"blockSize":0,"replication":0}, {"pathSuffix":"hadoop_examples.txt","type":"FILE","length":0,"owner":"root","group":"hdfs","permission":"644","accessTime":1599985053186,"modificationTime":1599985053243,"blockSize":134217728,"replication":3}, {"pathSuffix":"hadoop_root.txt","type":"FILE","length":0,"owner":"root","group":"hdfs","permission":"644","accessTime":1599929015215,"modificationTime":1599929015259,"blockSize":134217728,"replication":3}, {"pathSuffix":"wc_out","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"755","accessTime":0,"modificationTime":1599985153543,"blockSize":0,"replication":0}, {"pathSuffix":"yarn_gz.txt","type":"FILE","length":0,"owner":"root","group":"hdfs","permission":"644","accessTime":1600016323726,"modificationTime":1600016323766,"blockSize":134217728,"replication":3}, {"pathSuffix":"yarn_root.txt","type":"FILE","length":0,"owner":"root","group":"hdfs","permission":"644","accessTime":1600016038473,"modificationTime":1600016038510,"blockSize":134217728,"replication":3}, {"pathSuffix":"yarnout","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"755","accessTime":0,"modificationTime":1600016127680,"blockSize":0,"replication":0}, {"pathSuffix":"yarnoutgz","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"755","accessTime":0,"modificationTime":1600016410027,"blockSize":0,"replication":0}, {"pathSuffix":"zx_snapshot","type":"DIRECTORY","length":0,"owner":"root","group":"hdfs","permission":"755","accessTime":0,"modificationTime":1599931261967,"blockSize":0,"replication":0,"snapshotEnabled":true}}]}
[root@master ~]#
```

```
root@master:~  
[root@master ~]# curl -i -X PUT "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/cloud?op=MKDIRS&user.name=root"
HTTP/1.1 200 OK
Date: Sun, 13 Sep 2020 18:12:06 GMT
Cache-Control: no-cache
Expires: Sun, 13 Sep 2020 18:12:06 GMT
Date: Sun, 13 Sep 2020 18:12:06 GMT
Pragma: no-cache
Content-Type: application/json
X-Content-Type-Options: nosniff
X-XSS-Protection: 1; mode=block
Set-Cookie: hadoop.auth="u=root&p=root&t=simple-dt&e=1600056726495&s=jKf00fuRe80BY5mfW4zDza/lcjTCqLnQ5r7+ttwFCqQ="; Path=/; HttpOnly
Transfer-Encoding: chunked

{"boolean":true}
[root@master ~]# hdfs dfs -ls /user/root/
Found 11 items
drwx----- - root hdfs      0 2020-09-13 18:00 /user/root/.Trash
drwx----- - root hdfs      0 2020-09-13 17:00 /user/root/.staging
drwxr-xr-x - root hdfs      0 2020-09-13 18:12 /user/root/cloud
-rw-r--r--  3 root hdfs      0 2020-09-13 08:17 /user/root/hadoop_examples.txt
-rw-r--r--  3 root hdfs      0 2020-09-12 16:43 /user/root/hadoop_root.txt
drwxr-xr-x - root hdfs      0 2020-09-13 08:19 /user/root/wc_out
```

```
root@master:~  
[root@master ~]# curl -i -X PUT "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/cloud?op=MKDIRS&user.name=root"
HTTP/1.1 200 OK
Date: Sun, 13 Sep 2020 18:12:06 GMT
Cache-Control: no-cache
Expires: Sun, 13 Sep 2020 18:12:06 GMT
Date: Sun, 13 Sep 2020 18:12:06 GMT
Pragma: no-cache
Content-Type: application/json
X-Content-Type-Options: nosniff
X-XSS-Protection: 1; mode=block
Set-Cookie: hadoop.auth="u=root&p=root&t=simple-dt&e=1600056726495&s=jKf00fuRe80BY5mfW4zDza/lcjTCqLnQ5r7+ttwFCqQ="; Path=/; HttpOnly
Transfer-Encoding: chunked

{"boolean":true}
[root@master ~]# hdfs dfs -ls /user/root/
Found 11 items
drwx----- - root hdfs      0 2020-09-13 18:00 /user/root/.Trash
drwx----- - root hdfs      0 2020-09-13 17:00 /user/root/.staging
drwxr-xr-x - root hdfs      0 2020-09-13 18:12 /user/root/cloud
-rw-r--r--  3 root hdfs      0 2020-09-13 08:17 /user/root/hadoop_examples.txt
-rw-r--r--  3 root hdfs      0 2020-09-12 16:43 /user/root/hadoop_root.txt
drwxr-xr-x - root hdfs      0 2020-09-13 08:19 /user/root/wc_out
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root"  
{"AclStatus":{"owner":"admin1","group":"admingroup","stickyBit":false,"entries":[]}}  
[root@master ~]#  
[root@master ~]#  
[root@master ~]#  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root&doas=cert"  
{  
  "RemoteException": {  
    "message": "User: root is not allowed to impersonate cert",  
    "exception": "AuthorizationException",  
    "javaClassName": "org.apache.hadoop.security.authorize.AuthorizationException"  
  }  
}[root@master ~]#
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root"  
{"AclStatus":{"owner":"admin1","group":"admingroup","stickyBit":false,"entries":[]}}  
[root@master ~]#  
[root@master ~]#  
[root@master ~]#  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/admin1?op=GETACLSTATUS&user.name=root&doas=cert"  
{  
  "RemoteException": {  
    "message": "User: root is not allowed to impersonate cert",  
    "exception": "AuthorizationException",  
    "javaClassName": "org.apache.hadoop.security.authorize.AuthorizationException"  
  }  
}[root@master ~]#
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.txt?op=open&user.name=root"  
cloudera ccal31  
[root@master ~]# curl -i -X PUT "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.txt?op=rename&destination=/user/root/hadoop_examples_new.txt&user.name=root"  
HTTP/1.1 200 OK  
Date: Sun, 13 Sep 2020 18:17:40 GMT  
Cache-Control: no-cache  
Expires: Sun, 13 Sep 2020 18:17:40 GMT  
Date: Sun, 13 Sep 2020 18:17:40 GMT  
Pragma: no-cache  
X-Content-Type-Options: nosniff  
X-XSS-Protection: 1; mode=block  
Set-Cookie: hadoop.auth="u=root&p=root&t=simple-dt&e=1600057060081&s=SXLzKuAPUzs9MDRVQ/+LzZc/VzWCuQbE6HNBCATDcg4=";  
Path=/; HttpOnly  
Content-Type: application/json  
Transfer-Encoding: chunked  
  
{"boolean":true}  
[root@master ~]# hdfs dfs -ls /user/root  
Found 11 items  
drwx----- - root hdfs          0 2020-09-13 18:17 /user/root/.Trash  
drwx----- - root hdfs          0 2020-09-13 17:00 /user/root/.staging  
drwxr-xr-x - root hdfs          0 2020-09-13 18:12 /user/root/cloud  
-rw-r--r--  3 root hdfs        16 2020-09-13 18:17 /user/root/hadoop_examples_new.txt  
-rw-r--r--  3 root hdfs          0 2020-09-12 16:43 /user/root/hadoop_root.txt  
drwxr-xr-x - root hdfs          0 2020-09-13 08:19 /user/root/wc_out
```

```
root@master:~  
[root@master ~]# curl "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.txt?op=open  
&user.name=root"  
cloudera ccal31  
[root@master ~]# curl -i -X PUT "http://node01.subnet.vcn.oraclevcn.com:14000/webhdfs/v1/user/root/hadoop_examples.t  
xt?op=rename&destination=/user/root/hadoop_examples_new.txt&user.name=root"  
HTTP/1.1 200 OK  
Date: Sun, 13 Sep 2020 18:17:40 GMT  
Cache-Control: no-cache  
Expires: Sun, 13 Sep 2020 18:17:40 GMT  
Date: Sun, 13 Sep 2020 18:17:40 GMT  
Pragma: no-cache  
X-Content-Type-Options: nosniff  
X-XSS-Protection: 1; mode=block  
Set-Cookie: hadoop.auth="u=root&p=root&t=simple-dt&e=1600057060081&s=SXLzKuAPUzs9MDRvQ/+LzZc/VzWCuQbE6HNBCATDcq4=";  
Path=/; HttpOnly  
Content-Type: application/json  
Transfer-Encoding: chunked  
  
{"boolean":true}  
[root@master ~]# hdfs dfs -ls /user/root  
Found 11 items  
drwx----- - root hdfs      0 2020-09-13 18:17 /user/root/.Trash  
drwx----- - root hdfs      0 2020-09-13 17:00 /user/root/.staging  
drwxr-xr-x - root hdfs      0 2020-09-13 18:12 /user/root/cloud  
-rw-r--r--  3 root hdfs    16 2020-09-13 18:17 /user/root/hadoop_examples_new.txt  
-rw-r--r--  3 root hdfs      0 2020-09-12 16:43 /user/root/hadoop_root.txt  
drwxr-xr-x - root hdfs      0 2020-09-13 08:19 /user/root/wc_out
```

QUESTION 6

Configure the YARN scheduler as a Fair Scheduler and also number of threads used to handle requests through the scheduler interface should be 40.

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.

- go to CM - YARN - Configuration - Search (scheduler) - Change parameter (Scheduler Thread Count)
- Then restart service

Cluster 1

YARN (MR2 Included)

Actions

Sep 13, 7:06 PM GMT

Status Instances Configuration Commands Applications Resource Pools Charts Library Audits Web UI Quick Links

scheduler

Role Groups History and Rollback

Show All Descriptions

Filters

SCOPE

YARN (MR2 Included) (Servic...	5
Gateway	0
JobHistory Server	0
NodeManager	0
ResourceManager	23

CATEGORY

Advanced	4
Compression	0
Logs	0
Main	11
Monitoring	0
Performance	0
Ports and Addresses	1
Proxy	0
Resource Management	12

Scheduler Thread Count

yarn.resourcemanager.scheduler.client.thread-count

ResourceManager Default Group

50

Scheduler Class

yarn.resourcemanager.scheduler.class

ResourceManager Default Group

- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fair.FairScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fifo.FifoScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler

Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve)

ResourceManager Default Group

Name yarn.scheduler.capacity.root.queues

Value default

View as XML

Cluster 1

YARN (MR2 Included)

Actions

Sep 13, 7:06 PM GMT

Status Instances Configuration Commands Applications Resource Pools Charts Library Audits Web UI Quick Links

scheduler

Role Groups History and Rollback

Show All Descriptions

Filters

SCOPE

YARN (MR2 Included) (Servic...	5
Gateway	0
JobHistory Server	0
NodeManager	0
ResourceManager	23

CATEGORY

Advanced	4
Compression	0
Logs	0
Main	11
Monitoring	0
Performance	0
Ports and Addresses	1
Proxy	0
Resource Management	12

Scheduler Thread Count

yarn.resourcemanager.scheduler.client.thread-count

ResourceManager Default Group

50

Scheduler Class

yarn.resourcemanager.scheduler.class

ResourceManager Default Group

- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fair.FairScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fifo.FifoScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler

Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve)

ResourceManager Default Group

Name yarn.scheduler.capacity.root.queues

Value default

View as XML

Cluster 1

YARN (MR2 Included)

Actions

Sep 13, 7:07 PM GMT

- Status
- Instances
- Configuration
- Commands
- Applications
- Resource Pools
- Charts Library
- Audits
- Web UI
- Quick Links

scheduler

Role Groups

History and Rollback

Show All Descriptions

Filters

SCOPE

YARN (MR2 Included) (Servic...	5
Gateway	0
JobHistory Server	0
NodeManager	0
ResourceManager	23

CATEGORY

Advanced	4
Compression	0
Logs	0
Main	11
Monitoring	0
Performance	0
Ports and Addresses	1
Proxy	0
Resource Management	12

Scheduler Thread Count

yarn.resourcemanager.scheduler.client.thread-count

ResourceManager Default Group

40

Scheduler Class

yarn.resourcemanager.scheduler.class

ResourceManager Default Group

- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fair.FairScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fifo.FifoScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler

Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve)

ResourceManager Default Group

Name yarn.scheduler.capacity.root.queues

Value default

View as XML

1 Edited Value Reason for change...

Save Changes

Cluster 1

YARN (MR2 Included)

Actions

Sep 13, 7:07 PM GMT

- Status
- Instances
- Configuration
- Commands
- Applications
- Resource Pools
- Charts Library
- Audits
- Web UI
- Quick Links

scheduler

Role Groups

History and Rollback

Show All Descriptions

Filters

SCOPE

YARN (MR2 Included) (Servic...	5
Gateway	0
JobHistory Server	0
NodeManager	0
ResourceManager	23

CATEGORY

Advanced	4
Compression	0
Logs	0
Main	11
Monitoring	0
Performance	0
Ports and Addresses	1
Proxy	0
Resource Management	12

Scheduler Thread Count

yarn.resourcemanager.scheduler.client.thread-count

ResourceManager Default Group

40

Scheduler Class

yarn.resourcemanager.scheduler.class

ResourceManager Default Group

- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fair.FairScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.fifo.FifoScheduler
- org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler

Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve)

ResourceManager Default Group

Name yarn.scheduler.capacity.root.queues

Value default

View as XML

1 Edited Value Reason for change...

Save Changes




QUESTION 7

Create Role Groups and assign host to it accordingly create role group with datanodes added, new role group should have default configs like disk filling policy(available space, etc)

Correct Answer: :

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration.



Cluster 1

✓ HDFS Actions

Sep 14, 8:25 AM GMT

- Status
- Instances
- Configuration**
- Commands
- File Browser
- Charts Library
- Cache Statistics
- Audits
- Web UI
- Quick Li

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

Balancer Default Group 1

DATANODE

Cluster 1

✓ HDFS Actions

Sep 14, 8:25 AM GMT

- Status
- Instances
- Configuration**
- Commands
- File Browser
- Charts Library
- Cache Statistics
- Audits
- Web UI
- Quick Li

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

Balancer Default Group 1

DATANODE

Cluster 1
HDFS

Status Instances

Role Groups

Role groups allow grouping role
For example, roles running on s
role group for easier configurat
modified and new ones can be

Create

Filters

BALANCER
Balancer Default Group 1

DATANODE
DataNode Default Group 3

FAILOVER CONTROLLER
Failover Controller Default Gr... 2

GATEWAY

Create New Role Group

Group Name

Role Type

Role Groups can only be created for non-singleton roles.

Copy from

Existing configuration values will be copied over from this group into your new group.

Cancel

Support admin

Sep 14, 8:27 AM GMT

Audits Web UI Quick Links

roles would reside in the same
created but those can be

Feedback

Cluster 1
HDFS

Status Instances

Role Groups

Role groups allow grouping role
For example, roles running on s
role group for easier configurat
modified and new ones can be

Create

Filters

BALANCER
Balancer Default Group 1

DATANODE
DataNode Default Group 3

FAILOVER CONTROLLER
Failover Controller Default Gr... 2

GATEWAY

Create New Role Group

Group Name

Role Type

Role Groups can only be created for non-singleton roles.

Copy from

Existing configuration values will be copied over from this group into your new group.

Cancel

Support admin

Sep 14, 8:27 AM GMT

Audits Web UI Quick Links

roles would reside in the same
created but those can be

Feedback

Cluster 1 HDFS Actions  Sep 14, 8:27 AM GMT

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters	
BALANCER	
Balancer Default Group	1
DATANODE	
DataNode Default Group	3
DataNode Group 1	0
FAIL OVER CONTROLLER	

Cluster 1 HDFS Actions  Sep 14, 8:27 AM GMT

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters	
BALANCER	
Balancer Default Group	1
DATANODE	
DataNode Default Group	3
DataNode Group 1	0
FAIL OVER CONTROLLER	

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

- Balancer Default Group 1

DATANODE

- DataNode Default Group 3
- DataNode Group 1 0

FAILOVER CONTROLLER

- Failover Controller Default Gr... 2

GATEWAY

- Gateway Default Group 0

Actions for Selected

Showing 1 to 3 of 3 entries (filtered from 16 total entries) < 1 >

Display 25 Entries

<input type="checkbox"/>	Name	Host	Role Group	Status
<input checked="" type="checkbox"/>	DataNode (node01)	node01.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node02)	node02.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node03)	node03.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health

Showing 1 to 3 of 3 entries (filtered from 16 total entries)

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

- Balancer Default Group 1

DATANODE

- DataNode Default Group 3
- DataNode Group 1 0

FAILOVER CONTROLLER

- Failover Controller Default Gr... 2

GATEWAY

- Gateway Default Group 0

Actions for Selected

Showing 1 to 3 of 3 entries (filtered from 16 total entries) < 1 >

Display 25 Entries

<input type="checkbox"/>	Name	Host	Role Group	Status
<input checked="" type="checkbox"/>	DataNode (node01)	node01.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node02)	node02.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node03)	node03.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health

Showing 1 to 3 of 3 entries (filtered from 16 total entries)

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

- Balancer Default Group 1

DATANODE

- DataNode Default Group 3
- DataNode Group 1 0

Change Role Group Membership

Select target role group:

You are moving the following roles from group **DataNode Default Group**.

- DataNode (node01)

Showing 1 to 3 of 3 entries

<input type="checkbox"/>	Name	Host	Role Group	Status
<input type="checkbox"/>	DataNode (node01)	node01.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node02)	node02.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health
<input type="checkbox"/>	DataNode (node03)	node03.subnet.vcn.oraclevcn.com	DataNode Default Group	Good Health

Cluster 1
 HDFS

Status Instances Configuration Commands File Browser

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

Balancer Default Group 1

DATANODE

DataNode Default Group 3

Actions for Selected

Showing 1 to 3 of 3 entries

Name Host Role Group

Any Role Type Any Host Any Status

DataNode (node01)

Change Role Group Membership

Select target role group: DataNode Group 1

You are moving the following roles from group DataNode Default Group.

- DataNode (node01)

Cancel Move

Cloudera Manager Clusters Hosts Diagnostics Audits Charts Backup Administration Search Support admin

Cluster 1
 HDFS Sep 14, 8:31 AM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

Balancer Default Group 1

DATANODE

DataNode Default Group 2

DataNode Group 1 1

FAILOVER CONTROLLER

Failover Controller Default Gr... 2

Actions for Selected Rename

Showing 1 to 1 of 1 entries (filtered from 16 total entries) < 1 >

Display 25 Entries

Name	Host	Role Group	Status
DataNode (node01)	node01.subnet.vcn.oraclevcn.com	DataNode Group 1	Good Health

Any Role Type Any Host Any Status

Cloudera Manager Clusters Hosts Diagnostics Audits Charts Backup Administration Search Support admin

Cluster 1
 HDFS Sep 14, 8:31 AM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

Role Groups

Role groups allow grouping roles of the same type that need to be configured in the same way. For example, roles running on similar hosts with similar hardware configurations and/or colocated with the same set of other roles would reside in the same role group for easier configuration management. Cloudera Manager automatically creates role groups when services are first created but those can be modified and new ones can be created as needed. [Learn more](#)

Create

Filters

BALANCER

Balancer Default Group 1

DATANODE

DataNode Default Group 2

DataNode Group 1 1

FAILOVER CONTROLLER

Failover Controller Default Gr... 2

Actions for Selected Rename

Showing 1 to 1 of 1 entries (filtered from 16 total entries) < 1 >

Display 25 Entries

Name	Host	Role Group	Status
DataNode (node01)	node01.subnet.vcn.oraclevcn.com	DataNode Group 1	Good Health

Any Role Type Any Host Any Status

Cluster 1

HDFS

Actions

Sep 14, 12:53 PM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

Space Policy

Role Groups History and Rollback

Show All Descriptions

Filters

SCOPE

HDFS (Service-Wide)	0
Balancer	0
DataNode	3
Gateway	0
HttpFS	0
JournalNode	0
NFS Gateway	0
NameNode	0
SecondaryNameNode	0
Failover Controller	0

CATEGORY

Advanced	3
Checkpointing	0
Cloudera Navigator	0
Erasure Coding	0
High Availability	0
Logs	0

DataNode Volume

Choosing Policy
dfs.datanode.fsdataset.volume.c
hoosing.policy
Edit Individual Values

DataNode Default Group ...and 1 other

Round Robin
 Available Space

Available Space Policy

Balanced Threshold
dfs.datanode.available-space-
volume-choosing-
policy.balanced-space-threshold
Edit Identical Values

DataNode Default Group

10 GiB

DataNode Group 1

10 GiB

Available Space Policy

Balanced Preference
dfs.datanode.available-space-
volume-choosing-
policy.balanced-space-
preference-fraction
Edit Identical Values

DataNode Default Group

0.75

DataNode Group 1

0.75

Save Changes

Cluster 1

HDFS

Actions

Sep 14, 12:53 PM GMT

Status Instances Configuration Commands File Browser Charts Library Cache Statistics Audits Web UI Quick Links

Space Policy

Role Groups History and Rollback

Show All Descriptions

Filters

SCOPE

HDFS (Service-Wide)	0
Balancer	0
DataNode	3
Gateway	0
HttpFS	0
JournalNode	0
NFS Gateway	0
NameNode	0
SecondaryNameNode	0
Failover Controller	0

CATEGORY

Advanced	3
Checkpointing	0
Cloudera Navigator	0
Erasure Coding	0
High Availability	0
Logs	0

DataNode Volume

Choosing Policy
dfs.datanode.fsdataset.volume.c
hoosing.policy
Edit Individual Values

DataNode Default Group ...and 1 other

Round Robin
 Available Space

Available Space Policy

Balanced Threshold
dfs.datanode.available-space-
volume-choosing-
policy.balanced-space-threshold
Edit Identical Values

DataNode Default Group

10 GiB

DataNode Group 1

10 GiB

Available Space Policy

Balanced Preference
dfs.datanode.available-space-
volume-choosing-
policy.balanced-space-
preference-fraction
Edit Identical Values

DataNode Default Group

0.75

DataNode Group 1

0.75

Save Changes

QUESTION 8

So it is always better to have local software repository created and on each node same software will be download from local repository.


Hence, please accomplish following activity.

- Create a local repository on one of the node (It's not necessary to have this node to be part of CDH cluster)
- Also setup name server, hence this node can reach internet to download required software.
- Store Cloudera Manager-6 Server in that repo. Hence, later it can be installed on any node in cluster

Correct Answer: .

Explanation/Reference:

Answer: See the explanation for Step by Step Solution and configuration



```
root@master:/etc/sysconfig/network-scripts
[root@master network-scripts]# yum install yum-utils createrepo
Loaded plugins: fastestmirror, langpacks
Determining fastest mirrors
epel/x86_64/metalink | 17 kB 00:00
* base: mirrors.advancedhosters.com
* epel: dl.fedoraproject.org
* extras: mirror.ash.fastserv.com
* updates: mirror.cs.pitt.edu
base | 3.6 kB 00:00
epel | 4.7 kB 00:00
extras | 2.9 kB 00:00
updates | 2.9 kB 00:00
(1/7): base/7/x86_64/group_gz | 153 kB 00:00
(2/7): base/7/x86_64/primary_db | 6.1 MB 00:00
(3/7): epel/x86_64/group_gz | 95 kB 00:00
(4/7): epel/x86_64/updateinfo | 1.0 MB 00:00
(5/7): extras/7/x86_64/primary_db | 206 kB 00:00
(6/7): updates/7/x86_64/primary_db | 4.5 MB 00:00
(7/7): epel/x86_64/primary_db | 6.9 MB 00:00
```

```
[root@master network-scripts]# yum install yum-utils createrepo
Loaded plugins: fastestmirror, langpacks
Determining fastest mirrors
epel/x86_64/metalink | 17 kB 00:00
* base: mirrors.advancedhosters.com
* epel: dl.fedoraproject.org
* extras: mirror.ash.fastserv.com
* updates: mirror.cs.pitt.edu
base | 3.6 kB 00:00
epel | 4.7 kB 00:00
extras | 2.9 kB 00:00
updates | 2.9 kB 00:00
(1/7): base/7/x86_64/group_gz | 153 kB 00:00
(2/7): base/7/x86_64/primary_db | 6.1 MB 00:00
(3/7): epel/x86_64/group_gz | 95 kB 00:00
(4/7): epel/x86_64/updateinfo | 1.0 MB 00:00
(5/7): extras/7/x86_64/primary_db | 206 kB 00:00
(6/7): updates/7/x86_64/primary_db | 4.5 MB 00:00
(7/7): epel/x86_64/primary_db | 6.9 MB 00:00
```

```
[root@master network-scripts]# yum install httpd -y
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
* base: mirrors.advancedhosters.com
* epel: dl.fedoraproject.org
* extras: mirror.ash.fastserv.com
* updates: mirror.cs.pitt.edu
Resolving Dependencies
--> Running transaction check
---> Package httpd.x86_64 0:2.4.6-93.el7.centos will be installed
--> Processing Dependency: httpd-tools = 2.4.6-93.el7.centos for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: /etc/mime.types for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: libaprutil-1.so.0()(64bit) for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: libapr-1.so.0()(64bit) for package: httpd-2.4.6-93.el7.centos.x86_64
--> Running transaction check
---> Package apr.x86_64 0:1.4.8-5.el7 will be installed
---> Package apr-util.x86_64 0:1.5.2-6.el7 will be installed
---> Package httpd-tools.x86_64 0:2.4.6-93.el7.centos will be installed
---> Package mailcap.noarch 0:2.1.41-2.el7 will be installed
--> Finished Dependency Resolution
```

```
[root@master network-scripts]# yum install httpd -y
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirrors.advancedhosters.com
 * epel: dl.fedoraproject.org
 * extras: mirror.ash.fastserv.com
 * updates: mirror.cs.pitt.edu
Resolving Dependencies
--> Running transaction check
---> Package httpd.x86_64 0:2.4.6-93.el7.centos will be installed
--> Processing Dependency: httpd-tools = 2.4.6-93.el7.centos for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: /etc/mime.types for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: libaprutil-1.so.0()(64bit) for package: httpd-2.4.6-93.el7.centos.x86_64
--> Processing Dependency: libapr-1.so.0()(64bit) for package: httpd-2.4.6-93.el7.centos.x86_64
--> Running transaction check
---> Package apr.x86_64 0:1.4.8-5.el7 will be installed
---> Package apr-util.x86_64 0:1.5.2-6.el7 will be installed
---> Package httpd-tools.x86_64 0:2.4.6-93.el7.centos will be installed
---> Package mailcap.noarch 0:2.1.41-2.el7 will be installed
--> Finished Dependency Resolution
```

```
[root@master network-scripts]# service httpd status
Redirecting to /bin/systemctl status httpd.service
● httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled)
   Active: inactive (dead)
     Docs: man:httpd(8)
          man:apachectl(8)
[root@master network-scripts]# service httpd start
Redirecting to /bin/systemctl start httpd.service
[root@master network-scripts]# service httpd status
Redirecting to /bin/systemctl status httpd.service
● httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor preset: disabled)
   Active: active (running) since Sat 2020-09-12 11:57:37 GMT; 2s ago
     Docs: man:httpd(8)
          man:apachectl(8)
 Main PID: 11044 (httpd)
  Status: "Processing requests..."
  Memory: 3.0M
  CGroup: /system.slice/httpd.service
          └─11044 /usr/sbin/httpd -DFOREGROUND
          └─11045 /usr/sbin/httpd -DFOREGROUND
```

```
[root@master network-scripts]# service httpd status
Redirecting to /bin/systemctl status httpd.service
● httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor prese
t: disabled)
   Active: inactive (dead)
     Docs: man:httpd(8)
          man:apachectl(8)
[root@master network-scripts]# service httpd start
Redirecting to /bin/systemctl start httpd.service
[root@master network-scripts]# service httpd status
Redirecting to /bin/systemctl status httpd.service
● httpd.service - The Apache HTTP Server
   Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled; vendor prese
t: disabled)
   Active: active (running) since Sat 2020-09-12 11:57:37 GMT; 2s ago
     Docs: man:httpd(8)
          man:apachectl(8)
 Main PID: 11044 (httpd)
  Status: "Processing requests..."
  Memory: 3.0M
 CGroup: /system.slice/httpd.service
         └─11044 /usr/sbin/httpd -DFOREGROUND
         └─11045 /usr/sbin/httpd -DFOREGROUND
```

```
[root@master ~]# systemctl enable httpd
Created symlink from /etc/systemd/system/multi-user.target.wants/httpd.service t
o /usr/lib/systemd/system/httpd.service.
[root@master ~]# █
```

```
root@master:~  
[root@master ~]# systemctl enable httpd  
Created symlink from /etc/systemd/system/multi-user.target.wants/httpd.service to /usr/lib/systemd/system/httpd.service.  
[root@master ~]#
```

```
root@master:/var/www/html/cm6  
[root@master html]# mkdir cm6  
[root@master html]# cd cm6/  
[root@master cm6]# wget https://archive.cloudera.com/cm6/6.1.1/repo-as-tarball/cm6.1.1-redhat7.tar.gz  
--2020-09-12 12:05:41-- https://archive.cloudera.com/cm6/6.1.1/repo-as-tarball/cm6.1.1-redhat7.tar.gz  
Resolving archive.cloudera.com (archive.cloudera.com)... 199.232.64.167  
Connecting to archive.cloudera.com (archive.cloudera.com)|199.232.64.167|:443... connected.  
HTTP request sent, awaiting response... 200 OK  
Length: 1358893518 (1.3G) [application/x-tar]  
Saving to: 'cm6.1.1-redhat7.tar.gz'  
  
0% [ ] 1,313,089 1.51MB/s
```

root@master:/var/www/html/cm6

```
[root@master html]# mkdir cm6
[root@master html]# cd cm6/
[root@master cm6]# wget https://archive.cloudera.com/cm6/6.1.1/repo-as-tarball/cm6.1.1-redhat7.tar.gz
--2020-09-12 12:05:41-- https://archive.cloudera.com/cm6/6.1.1/repo-as-tarball/cm6.1.1-redhat7.tar.gz
Resolving archive.cloudera.com (archive.cloudera.com)... 199.232.64.167
Connecting to archive.cloudera.com (archive.cloudera.com)|199.232.64.167|:443...
connected.
HTTP request sent, awaiting response... 200 OK
Length: 1358893518 (1.3G) [application/x-tar]
Saving to: `cm6.1.1-redhat7.tar.gz'

0% [          ] 1,313,089  1.51MB/s
```

root@master:/var/www/html/cm6

```
[root@master cm6]# tar xvfz cm6.1.1-redhat7.tar.gz
cm6.1.1/
cm6.1.1/RPMS/
cm6.1.1/RPMS/x86_64/
cm6.1.1/RPMS/x86_64/cloudera-manager-server-6.1.1-853290.el7.x86_64.rpm
cm6.1.1/RPMS/x86_64/cloudera-manager-server-db-2-6.1.1-853290.el7.x86_64.rpm
cm6.1.1/RPMS/x86_64/oracle-j2sdk1.8-1.8.0+update181-1.x86_64.rpm
cm6.1.1/RPMS/x86_64/cloudera-manager-daemons-6.1.1-853290.el7.x86_64.rpm
```

root@master:var/www/html/cm6

```
[root@master cm6]# tar xvfz cm6.1.1-redhat7.tar.gz
cm6.1.1/
cm6.1.1/RPMS/
cm6.1.1/RPMS/x86_64/
cm6.1.1/RPMS/x86_64/cloudera-manager-server-6.1.1-853290.el7.x86_64.rpm
cm6.1.1/RPMS/x86_64/cloudera-manager-server-db-2-6.1.1-853290.el7.x86_64.rpm
cm6.1.1/RPMS/x86_64/oracle-j2sdk1.8-1.8.0+update181-1.x86_64.rpm
cm6.1.1/RPMS/x86_64/cloudera-manager-daemons-6.1.1-853290.el7.x86_64.rpm
```

root@master:var/www/html

```
[root@master html]# chmod -R ugo+rX cm6
[root@master html]#
```

```
[root@master html]# chmod -R ugo+rX cm6  
[root@master html]# █
```

Index of /cm6

<u>Name</u>	<u>Last modified</u>	<u>Size</u>	<u>Description</u>
 Parent Directory		-	
 RPM-GPG-KEY-cloudera	2019-02-19 09:57	1.7K	
 RPMS/	2019-05-29 10:16	-	
 allkeys.asc	2019-02-19 09:57	14K	
 cloudera-manager.repo	2019-02-19 09:57	232	
 repodata/	2019-05-29 10:16	-	

Index of /cm6

<u>Name</u>	<u>Last modified</u>	<u>Size</u>	<u>Description</u>
 Parent Directory		-	
 RPM-GPG-KEY-cloudera	2019-02-19 09:57	1.7K	
 RPMS/	2019-05-29 10:16	-	
 allkeys.asc	2019-02-19 09:57	14K	
 cloudera-manager.repo	2019-02-19 09:57	232	
 repodata/	2019-05-29 10:16	-	

```
root@master:/etc/yum.repos.d
-rw-r--r--. 1 root root 649 Apr 7 22:01 CentOS-Debuginfo.repo
-rw-r--r--. 1 root root 314 Apr 7 22:01 CentOS-fasttrack.repo
-rw-r--r--. 1 root root 630 Apr 7 22:01 CentOS-Media.repo
-rw-r--r--. 1 root root 1331 Apr 7 22:01 CentOS-Sources.repo
-rw-r--r--. 1 root root 7577 Apr 7 22:01 CentOS-Vault.repo
-rw-r--r--. 1 root root 616 Apr 7 22:01 CentOS-x86_64-kernel.repo
-rw-r--r--. 1 root root 951 Oct 2 2017 epel.repo
-rw-r--r--. 1 root root 1050 Oct 2 2017 epel-testing.repo
[root@master yum.repos.d]# wget https://archive.cloudera.com/cm6/6.1.1/redhat7/yum/cloudera-manager.repo
--2020-09-12 12:25:14-- https://archive.cloudera.com/cm6/6.1.1/redhat7/yum/cloudera-manager.repo
Resolving archive.cloudera.com (archive.cloudera.com)... 151.101.200.167
Connecting to archive.cloudera.com (archive.cloudera.com)|151.101.200.167|:443..
. connected.
HTTP request sent, awaiting response... 200 OK
Length: 232 [binary/octet-stream]
Saving to: 'cloudera-manager.repo'

100%[=====>] 232 --.-K/s in 0s

2020-09-12 12:25:15 (16.0 MB/s) - 'cloudera-manager.repo' saved [232/232]

[root@master yum.repos.d]#
```

root@master:/etc/yum.repos.d

```
-rw-r--r--. 1 root root 649 Apr 7 22:01 CentOS-Debuginfo.repo
-rw-r--r--. 1 root root 314 Apr 7 22:01 CentOS-fasttrack.repo
-rw-r--r--. 1 root root 630 Apr 7 22:01 CentOS-Media.repo
-rw-r--r--. 1 root root 1331 Apr 7 22:01 CentOS-Sources.repo
-rw-r--r--. 1 root root 7577 Apr 7 22:01 CentOS-Vault.repo
-rw-r--r--. 1 root root 616 Apr 7 22:01 CentOS-x86_64-kernel.repo
-rw-r--r--. 1 root root 951 Oct 2 2017 epel.repo
-rw-r--r--. 1 root root 1050 Oct 2 2017 epel-testing.repo
[root@master yum.repos.d]# wget https://archive.cloudera.com/cm6/6.1.1/redhat7/yum/cloudera-manager.repo
--2020-09-12 12:25:14-- https://archive.cloudera.com/cm6/6.1.1/redhat7/yum/cloudera-manager.repo
Resolving archive.cloudera.com (archive.cloudera.com)... 151.101.200.167
Connecting to archive.cloudera.com (archive.cloudera.com)|151.101.200.167|:443..
. connected.
HTTP request sent, awaiting response... 200 OK
Length: 232 [binary/octet-stream]
Saving to: 'cloudera-manager.repo'

100%[=====>] 232 --.-K/s in 0s

2020-09-12 12:25:15 (16.0 MB/s) - 'cloudera-manager.repo' saved [232/232]

[root@master yum.repos.d]#
```

root@master:/etc/yum.repos.d

```
[root@master yum.repos.d]# vi cloudera-manager.repo
[cloudera-manager]
name=Cloudera Manager 6.1.1
baseurl=http://master.subnet.vcn.oraclevcn.com/cm6/
gpgkey=http://master.subnet.vcn.oraclevcn.com/cm6/RPM-GPG-KEY-cloudera
gpgcheck=1
enabled=1
autorefresh=0
type=rpm-md
```

root@master:/etc/yum.repos.d

```
[root@master yum.repos.d]# vi cloudera-manager.repo
[cloudera-manager]
name=Cloudera Manager 6.1.1
baseurl=http://master.subnet.vcn.oraclevcn.com/cm6/
gpgkey=http://master.subnet.vcn.oraclevcn.com/cm6/RPM-GPG-KEY-cloudera
gpgcheck=1
enabled=1
autorefresh=0
type=rpm-md
```

root@master:/etc/yum.repos.d

```
[root@master yum.repos.d]# yum repolist
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirrors.advancedhosters.com
 * epel: dl.fedoraproject.org
 * extras: mirror.ash.fastserv.com
 * updates: mirror.cs.pitt.edu
cloudera-manager                | 2.9 kB      00:00
cloudera-manager/primary_db    | 8.5 kB      00:00
repo id                          repo name                      status
base/7/x86_64                   CentOS-7 - Base                10,070
cloudera-manager              Cloudera Manager 6.1.1        6
epel/x86_64                     Extra Packages for Enterprise Linux 7 - x86_64 13,445
extras/7/x86_64                 CentOS-7 - Extras              413
updates/7/x86_64                CentOS-7 - Updates             1,127
repolist: 25,061
[root@master yum.repos.d]#
```

```
[root@master yum.repos.d]# yum repolist
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirrors.advancedhosters.com
 * epel: dl.fedoraproject.org
 * extras: mirror.ash.fastserv.com
 * updates: mirror.cs.pitt.edu
cloudera-manager                | 2.9 kB      00:00
cloudera-manager/primary_db    | 8.5 kB      00:00
repo id                          repo name                                status
base/7/x86_64                   CentOS-7 - Base                          10,070
cloudera-manager                Cloudera Manager 6.1.1                   6
epel/x86_64                      Extra Packages for Enterprise Linux 7 - x86_64 13,445
extras/7/x86_64                 CentOS-7 - Extras                         413
updates/7/x86_64                CentOS-7 - Updates                        1,127
repolist: 25,061
[root@master yum.repos.d]#
```

Solution provided above

https://docs.cloudera.com/documentation/enterprise/6/6.1/topics/cm_ig_create_local_parcel_repo.html