CLEMSON UNIVERSITY

Distributed Computing using CloudLab

Linh B. Ngo
Clemson University

Introduction

- Distributed and Cluster Computing (CPSC 3620)
- Offered twice per academic year
- Average class size: 40-45 students
- Required junior-level class (typically taken at senior year)
- Contents:
  - Infrastructure/System-oriented
  - Performance/Efficiency
  - High Performance Computing
    - MPI
  - Big Data Computing
    - Hadoop MapReduce
    - Apache Spark
    - HPCCSystems

Computing Resources

- Palmetto Supercomputer
  - 2000+ nodes, open to all faculty/students
  - No administrative access
  - Cannot share nodes among students to support group assignment
  - Preemption from node owners
- CloudLab
  - Limited resources for large-scale study
  - Administrative access
  - Ease of collaboration
  - No preemption

Computing Resources

- Combine both local computing resources and CloudLab

- Learning outcomes through CloudLab
  - Administrative skills for distributed systems
  - In-depth understanding of distributed systems
- Learning outcomes through Palmetto
  - Basic understanding of parallel application development
  - Impacts of scaling and efficiency on larger systems

Tutorial

- Set up environments for distributed computing on CloudLab
- MPI
  - Two-node cluster
  - OpenMPI
- Hadoop
  - Three-node cluster
  - Hortonwork Distribution

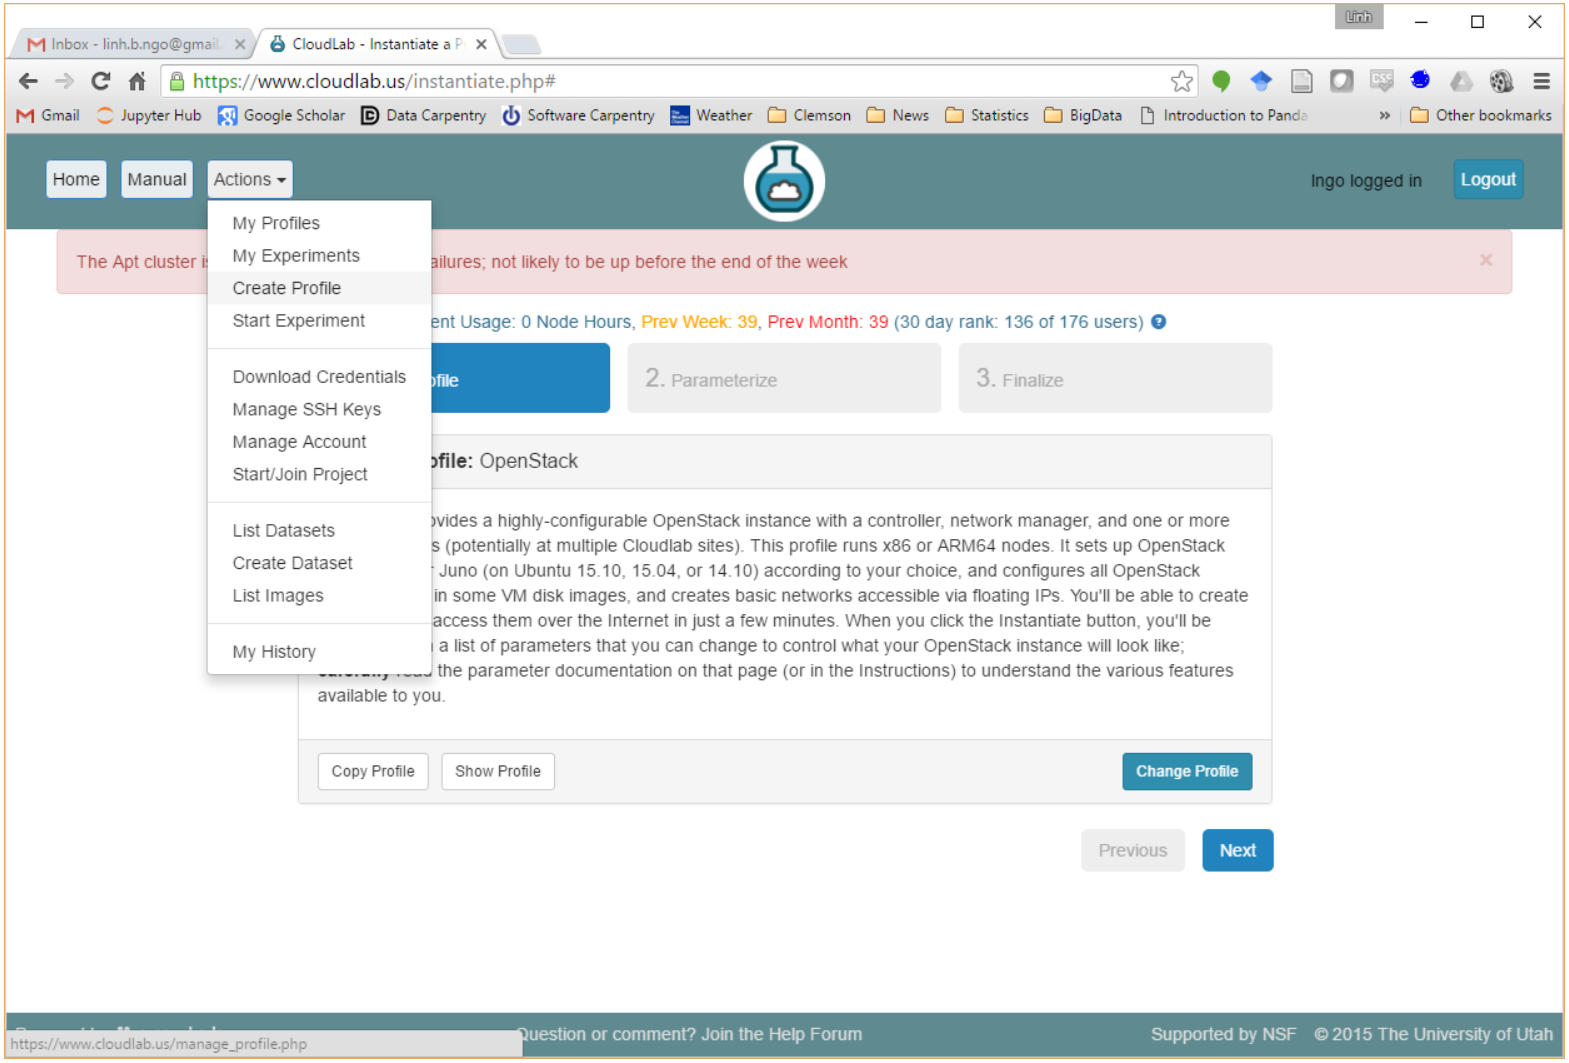
Inbox - linh.b.ngo@gmail.
CloudLab - Instantiate a P
https://www.cloudlab.us/instantiate.php#
Gmail
Jupyter Hub
Google Scholar
Data Carpentry
Software Carpentry
Weather
Clemson
News
Statistics
BigData
Introduction to Panda
Other bookmarks
Home
Manual
Actions
My Profiles
My Experiments
Create Profile
Start Experiment
Download Credentials
Manage SSH Keys
Manage Account
Start/Join Project
List Datasets
Create Dataset
List Images
My History
Ingo logged in
Logout
The Apt cluster i
ailures; not likely to be up before the end of the week
ent Usage: 0 Node Hours, Prev Week: 39, Prev Month: 39 (30 day rank: 136 of 176 users)
ofile
2. Parameterize
3. Finalize
ofile: OpenStack
ovides a highly-configurable OpenStack instance with a controller, network manager, and one or more
s (potentially at multiple Cloudlab sites). This profile runs x86 or ARM64 nodes. It sets up OpenStack
Juno (on Ubuntu 15.10, 15.04, or 14.10) according to your choice, and configures all OpenStack
in some VM disk images, and creates basic networks accessible via floating IPs. You'll be able to create
access them over the Internet in just a few minutes. When you click the Instantiate button, you'll be
a list of parameters that you can change to control what your OpenStack instance will look like;
the parameter documentation on that page (or in the Instructions) to understand the various features
available to you.
Copy Profile
Show Profile
Change Profile
Previous
Next
Question or comment? Join the Help Forum
Supported by NSF
© 2015 The University of Utah
https://www.cloudlab.us/manage_profile.php

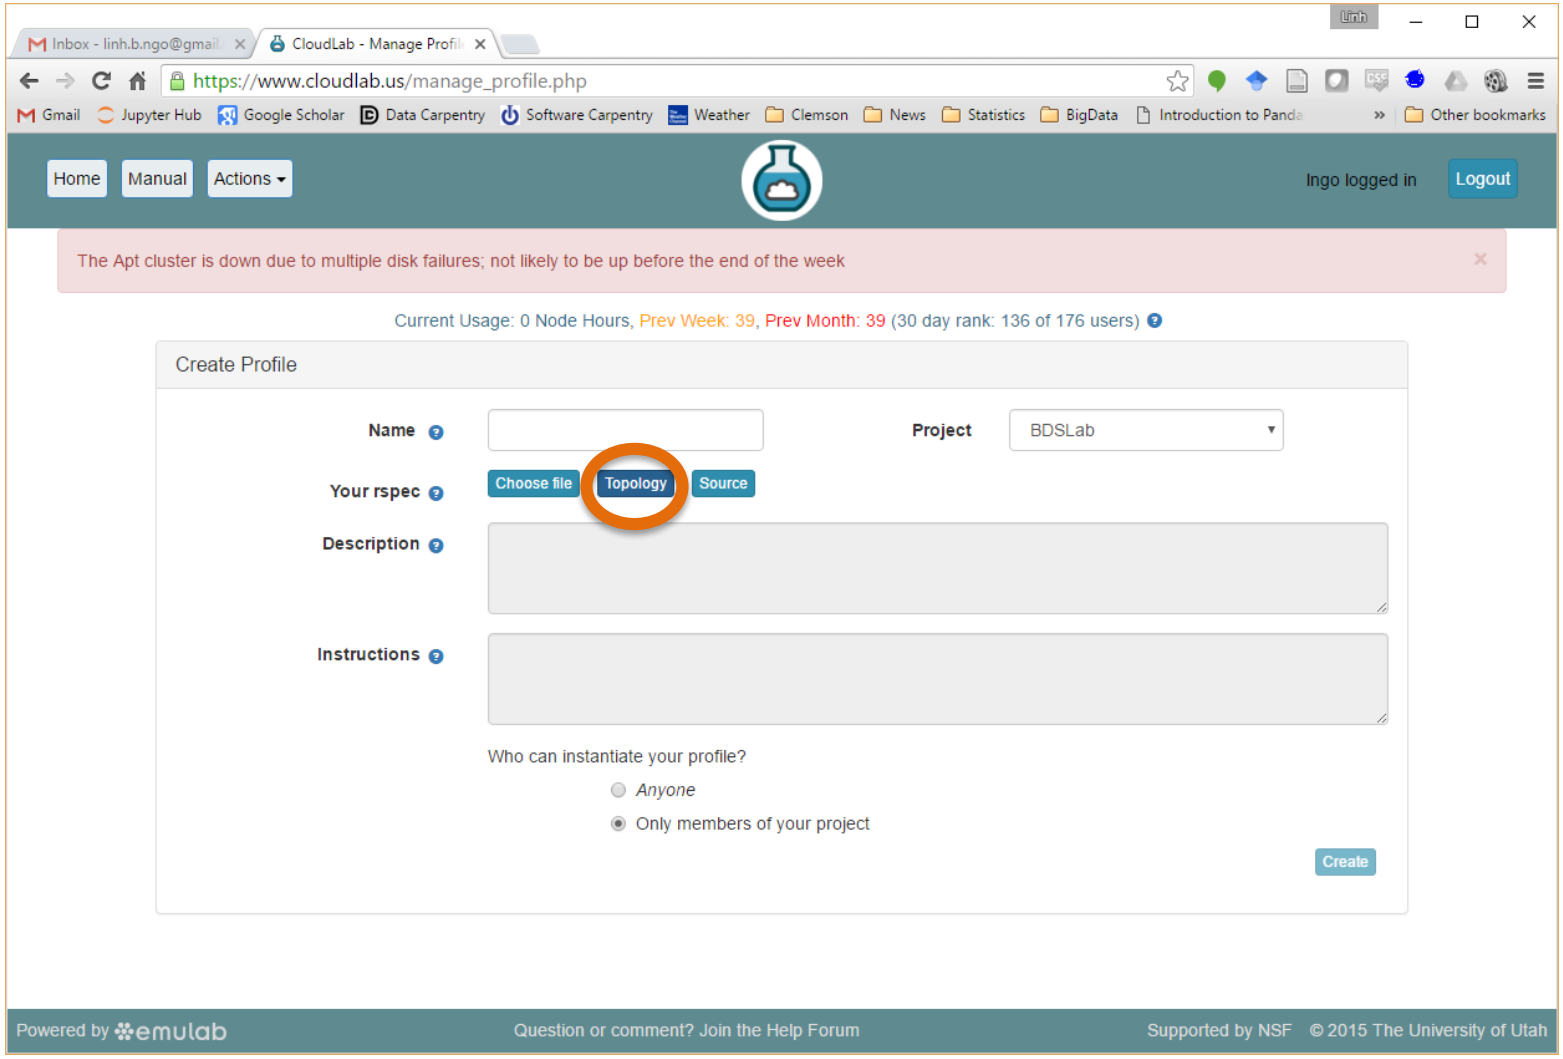
Inbox - linh.b.ngo@gmail
CloudLab - Manage Profi
https://www.cloudlab.us/manage_profile.php
Gmail
Jupyter Hub
Google Scholar
Data Carpentry
Software Carpentry
Weather
Clemson
News
Statistics
BigData
Introduction to Panda
Other bookmarks
Home
Manual
Actions
Ingo logged in
Logout
The Apt cluster is down due to multiple disk failures; not likely to be up before the end of the week
Current Usage: 0 Node Hours, Prev Week: 39, Prev Month: 39 (30 day rank: 136 of 176 users)
Create Profile
Name
Project
BDSLab
Your rspec
Choose file
Topology
Source
Description
Instructions
Who can instantiate your profile?
Anyone
Only members of your project
Create
Powered by emulab
Question or comment? Join the Help Forum
Supported by NSF
© 2015 The University of Utah

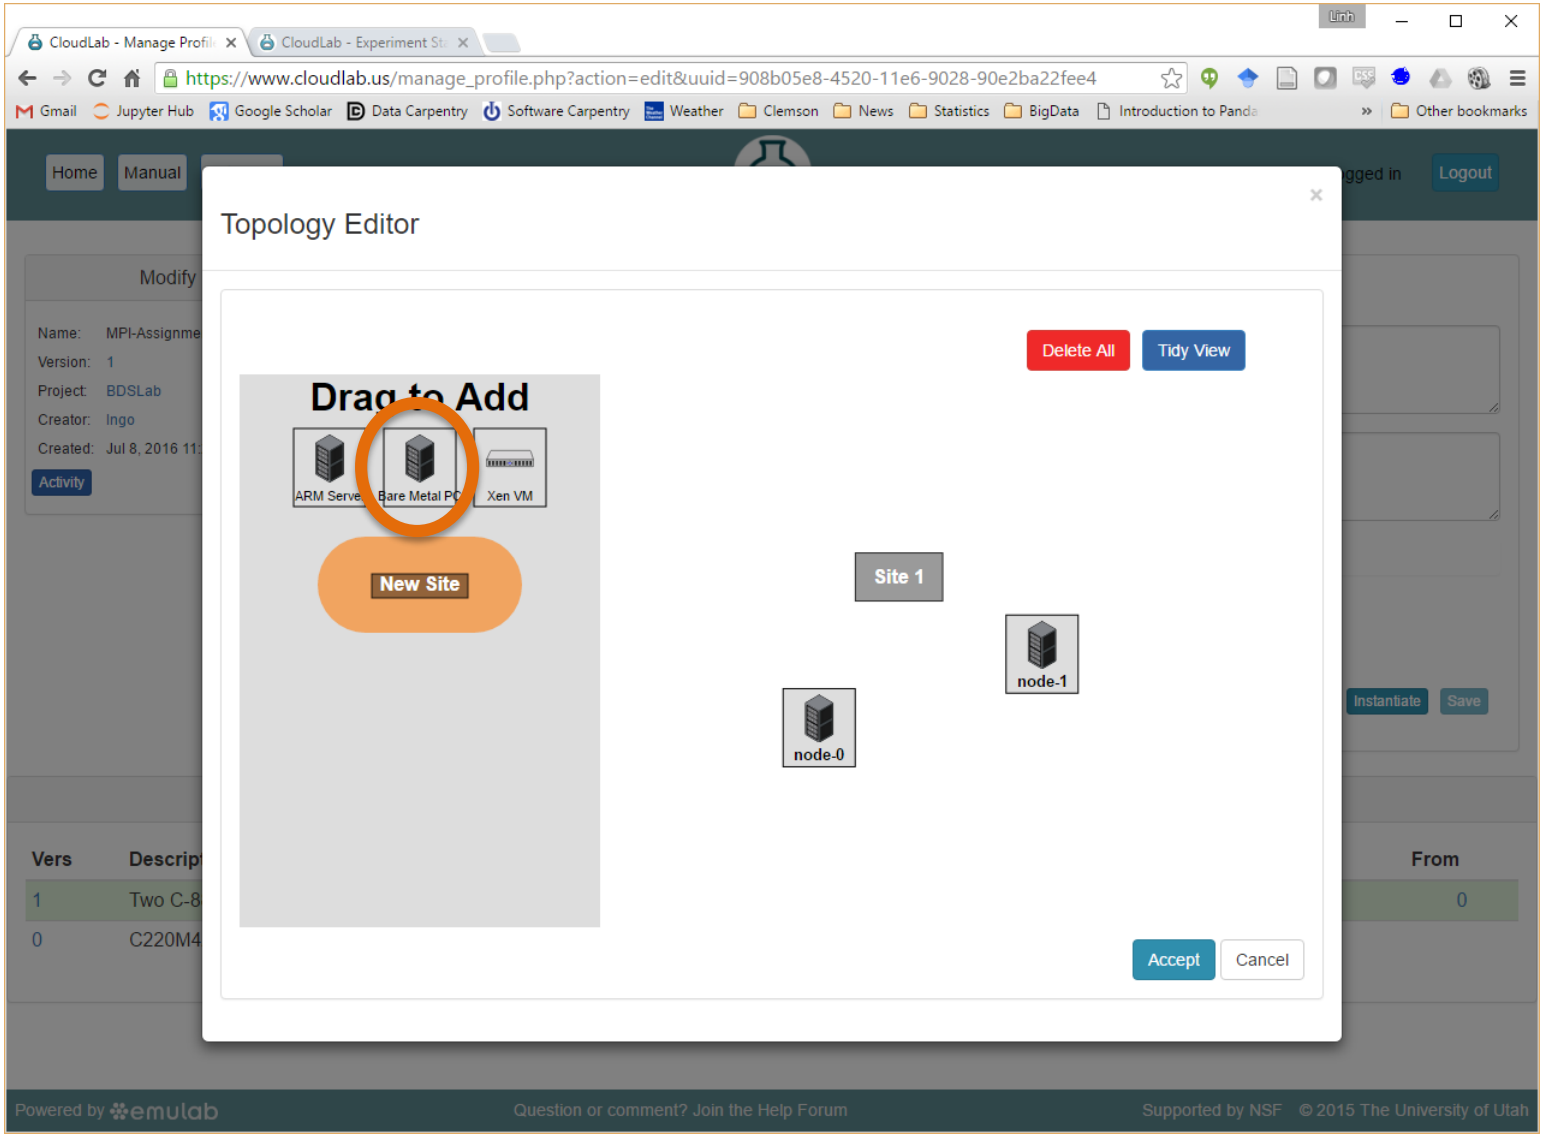
CloudLab - Manage Profil
CloudLab - Experiment St
https://www.cloudlab.us/manage_profile.php?action=edit&uuid=908b05e8-4520-11e6-9028-90e2ba22fee4
Gmail
Jupyter Hub
Google Scholar
Data Carpentry
Software Carpentry
Weather
Clemson
News
Statistics
BigData
Introduction to Panda
Other bookmarks
Home
Manual
Logout
Modify
Name: MPI-Assignme
Version: 1
Project: BDSLab
Creator: Ingo
Created: Jul 8, 2016 11
Activity
Instantiate
Save
Vers
Descrip
From
1
Two C-8
0
0
C220M4
Topology Editor
Drag to Add
ARM Serve
Bare Metal PC
Xen VM
New Site
Delete All
Tidy View
Site 1
node-1
node-0
Accept
Cancel
Powered by emulab
Question or comment? Join the Help Forum
Supported by NSF
© 2015 The University of Utah

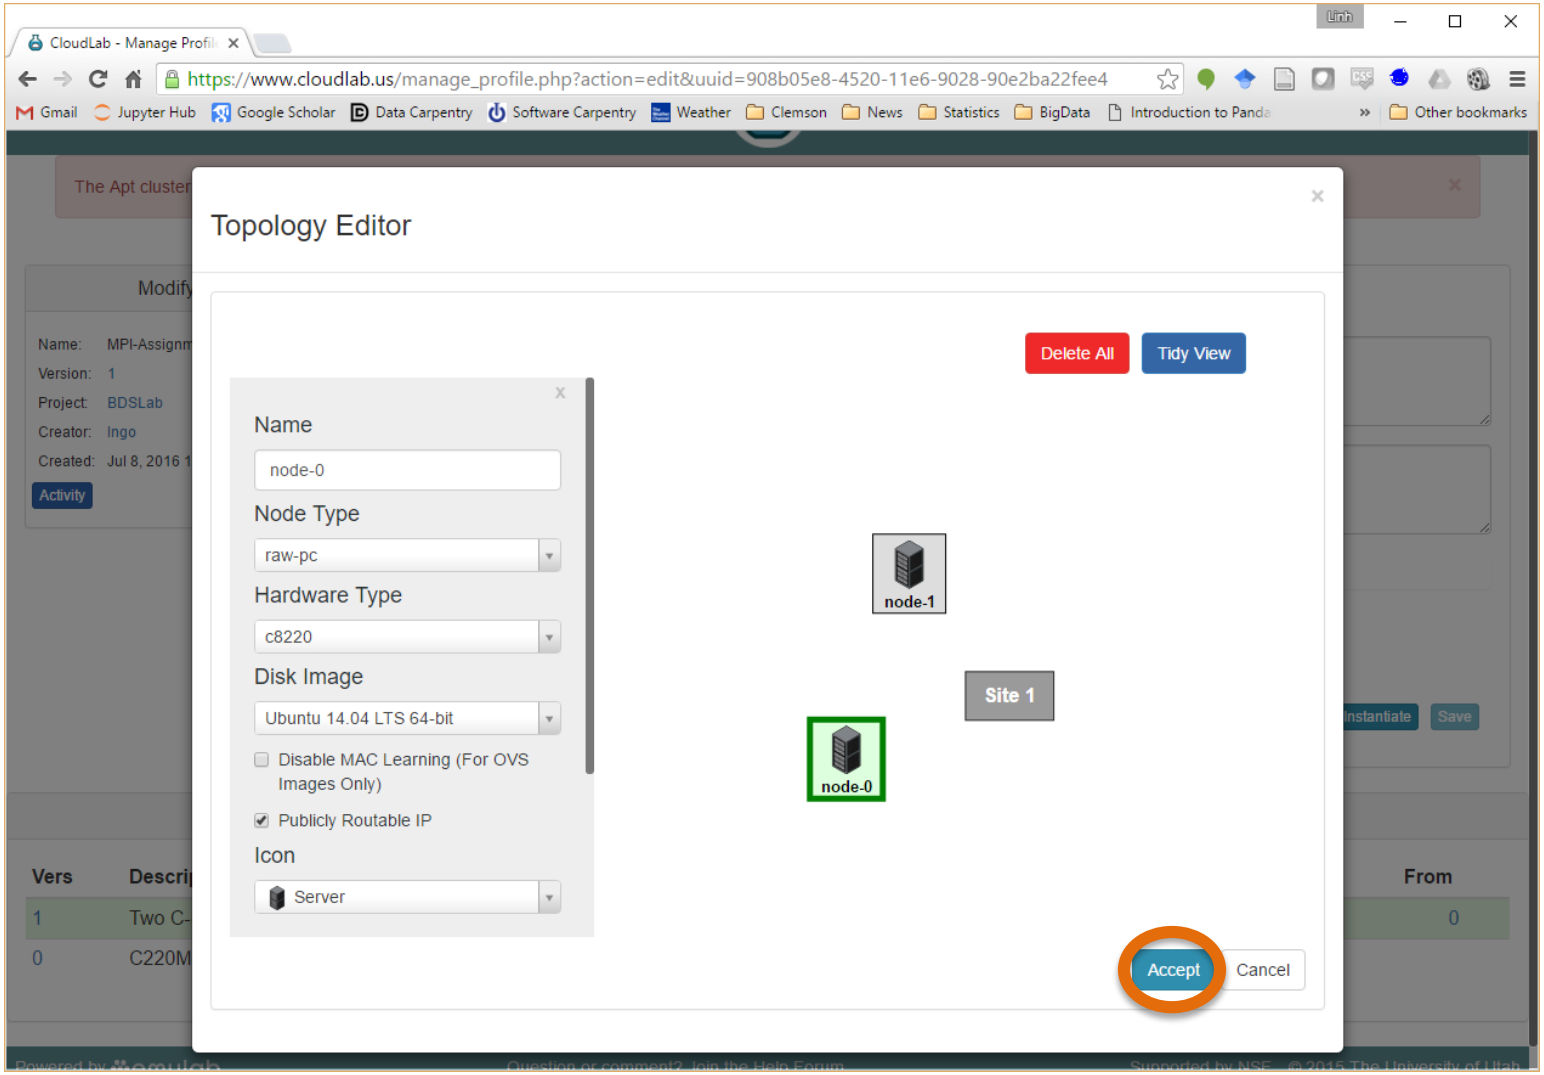

- Both nodes should have the same configuration, no network connection is needed (due to public IP)

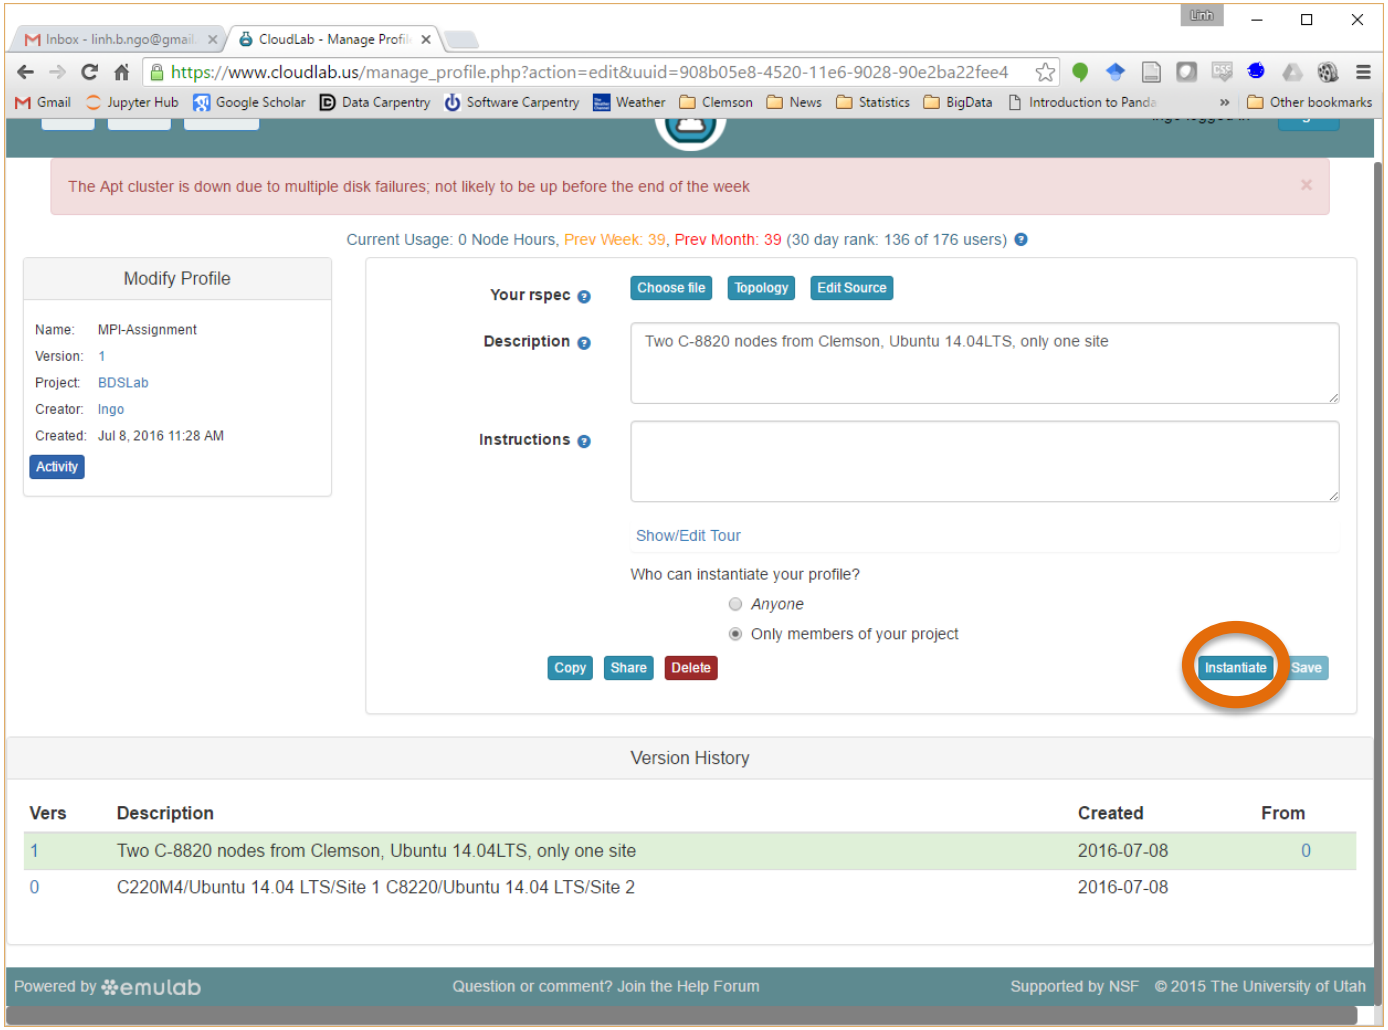

- You can save multiple versions of your topology
- Instantiate the version that you want to launch
- The launching procedure will be similar to the OpenStack tutorial

MPI on CloudLab

On each node:

```
sudo apt-get update
sudo apt-get install libibnetdisc-dev

sudo nano /etc/environment

PATH="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/
bin:/sbin:/bin:/usr/games:/usr/local/games:/home/mpi
user/.openmpi/bin"
LD_LIBRARY_PATH="/lib:/usr/lib:/usr/local/lib:/home/
mpiuser/.openmpi/lib/“

sudo adduser mpiuser
```

MPI on CloudLab

```
sudo adduser mpiuser
ssh mpiuser@localhost

wget https://www.open-mpi.org/software/ompi/v1.8/downloads/openmpi-1.8.1.tar.gz

tar xzf openmpi-1.8.1.tar.gz
cd openmpi-1.8.1

./configure --prefix="/home/mpiuser/.openmpi"
make
make install
```

MPI on CloudLab

```
ssh-keygen –t rsa

cd .ssh

cp id_rsa.pub authorized_keys

ssh-copy-id -i id_rsa.pub <hostname of the other node>

nano nodelist

<hostname of first node>
<hostname of second node>
…
```

Example program

```
#include <stdio.h>
#include <unistd.h>
#include <sys/utsname.h>
#include <mpi.h>

int main(int argc, char *argv[]){
  int rank, size;
  MPI_Status status;
  MPI_Init(&argc,&argv);
  MPI_Comm_size(MPI_COMM_WORLD,&size);
  MPI_Comm_rank(MPI_COMM_WORLD,&rank);

  struct utsname uts;
  uname(&uts);
  printf("%d at %s\n",rank,uts.nodename);
  MPI_Finalize();
  return 0;
}
```

MPI on CloudLab

```
mpicc gethostname.c -o gethostname

scp gethostname mpiuser@<the other node>:/home/mpiuser
scp nodelist mpiuser@<the other node>:/home/mpiuser

mpirun -np 2 -machinefile nodelist ./gethostname

mpirun -np 2 -machinefile nodelist --map-by node
./gethostname
```

Assignment Ideas

- Develop a work queue using various allocation strategies:
  - Normal
  - Cyclic
  - Dynamic

- Setup MPI cluster with nodes on separate sites, reduce network connection, and evaluate performance on different allocation strategies

Hadoop on CloudLab

- Enterprise Hadoop
- Hortonworks
- http://hortonworks.com/hdp/downloads/

Hadoop on CloudLab

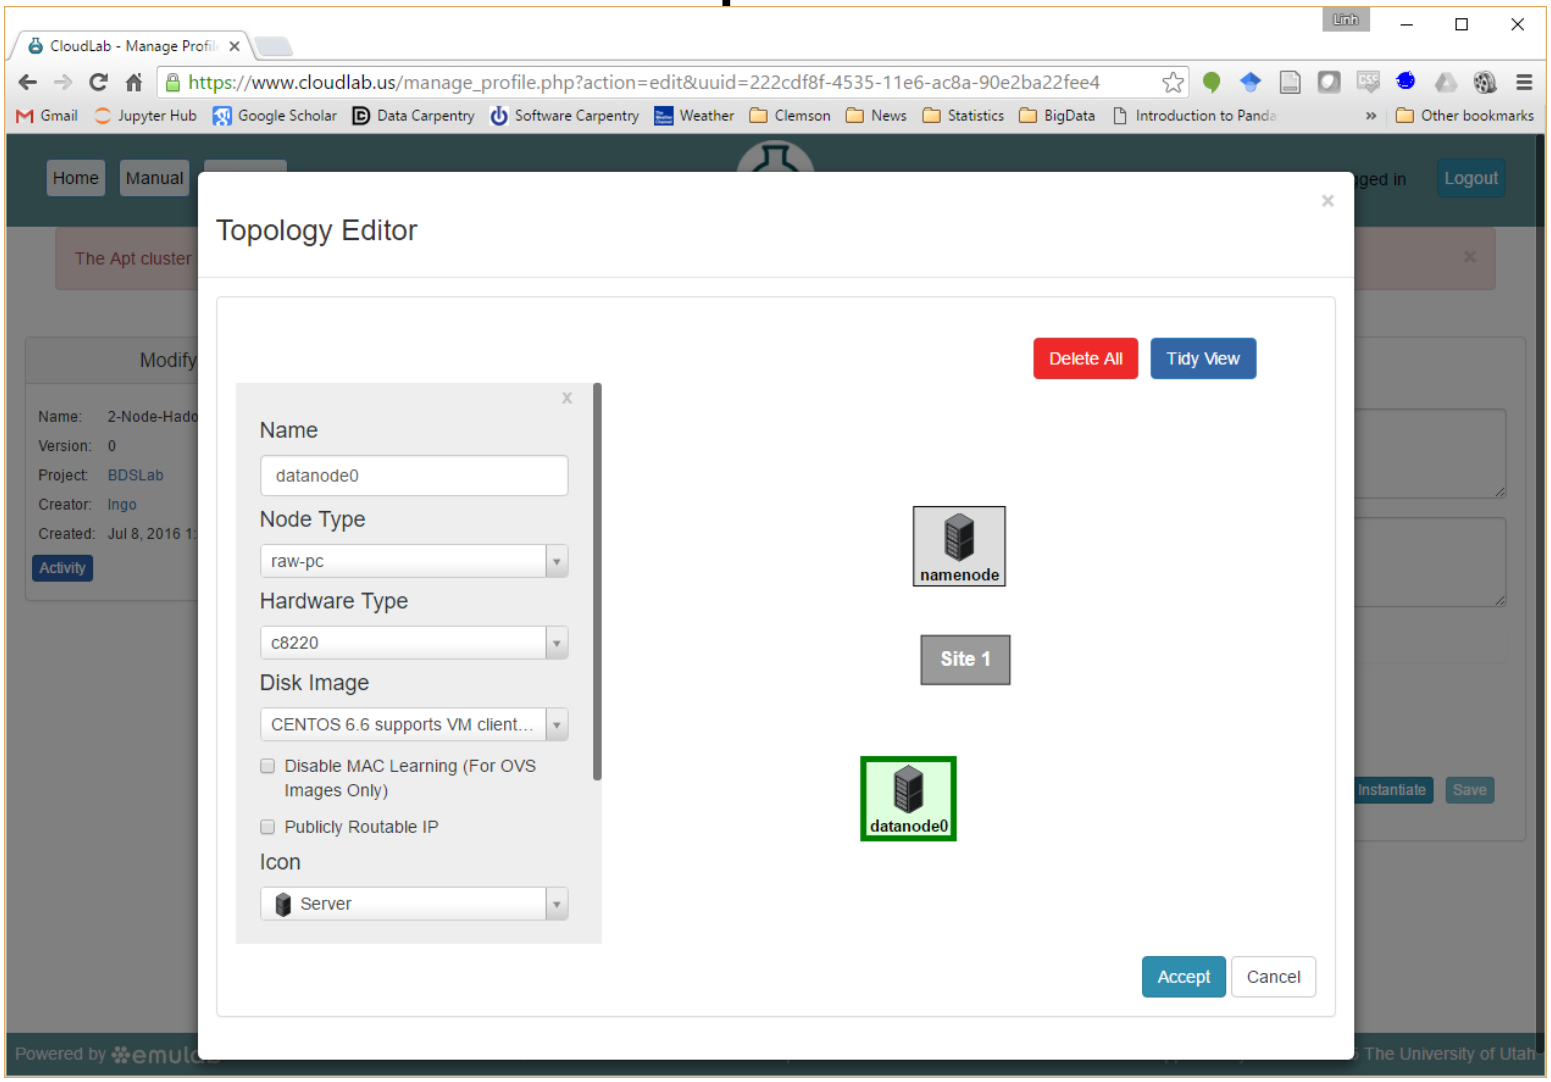

- Both nodes should have the same configuration on bare metal PC, no network connection is needed (due to public IP)

On each node

- SSH onto the node from Palmetto
- Change to root:

```
sudo su -
```

- Execute the following commands:

```
chkconfig --list ntpd
chkconfig ntpd on
service ntpd start
chkconfig iptables off
/etc/init.d/iptables stop
setenforce 0
```

On each node

- Setup Ambari download server

```
wget -nv http://public-repo-1.hortonworks.com/ambari/centos6/2.x/updates/2.1.2/ambari.repo -O /etc/yum.repos.d/ambari.repo
```

- On namenode

```
yum –y install ambari-server
yum –y install ambari-agent
```

- On datanode

```
yum –y install ambari-agent
```

On namenode

- Set up ambari server:

```
ambari-server setup
```

- Select default for all questions
- Select 1 for JDK version
- When all done, start ambari server

```
ambari-server start
```

On each node

- Using vim to edit `/etc/ambari-agent/conf/ambari-agent.ini`
- Change:

```
hostname=<hostname of namenode as shown in list view of CloudLab>
```

- Start Ambari Agent

```
ambari-agent start
```

Ambari Server (admin/admin)

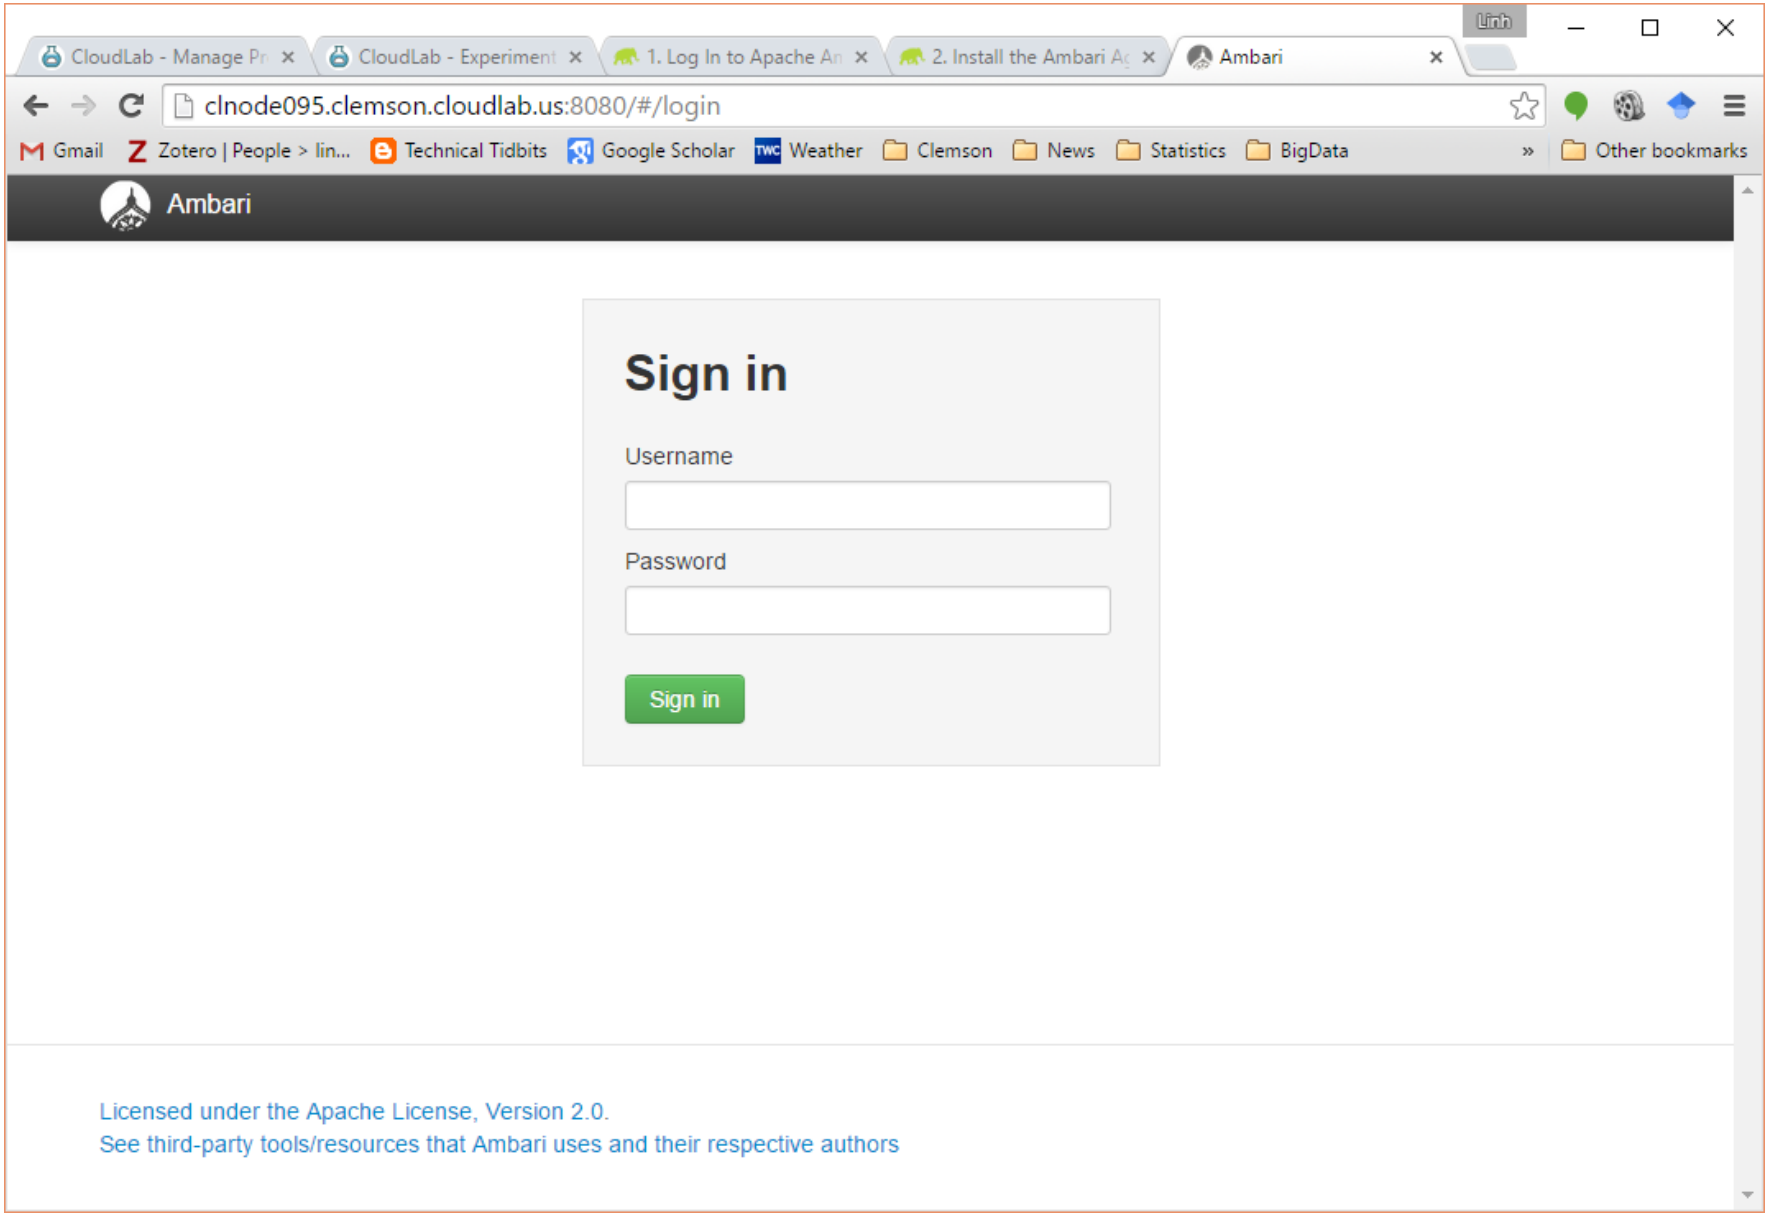

Ambari Server (admin/admin)

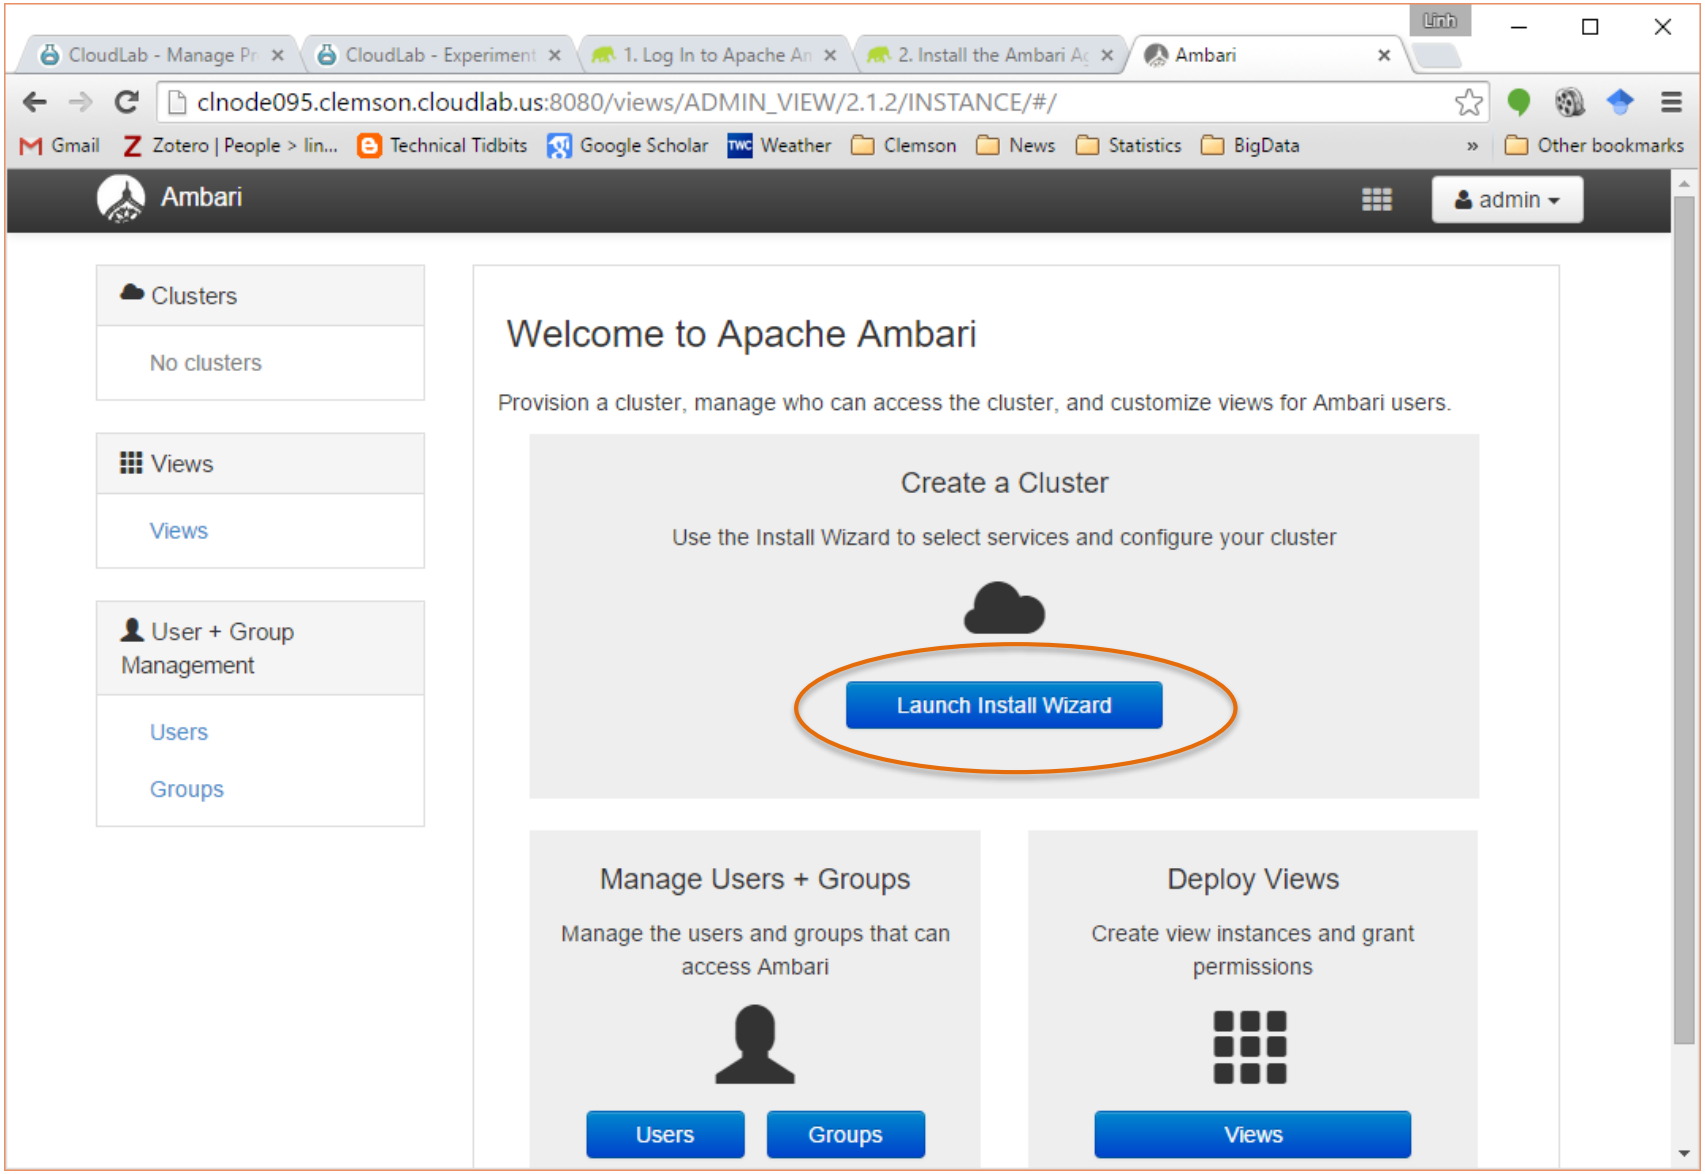

Ambari Server (admin/admin)

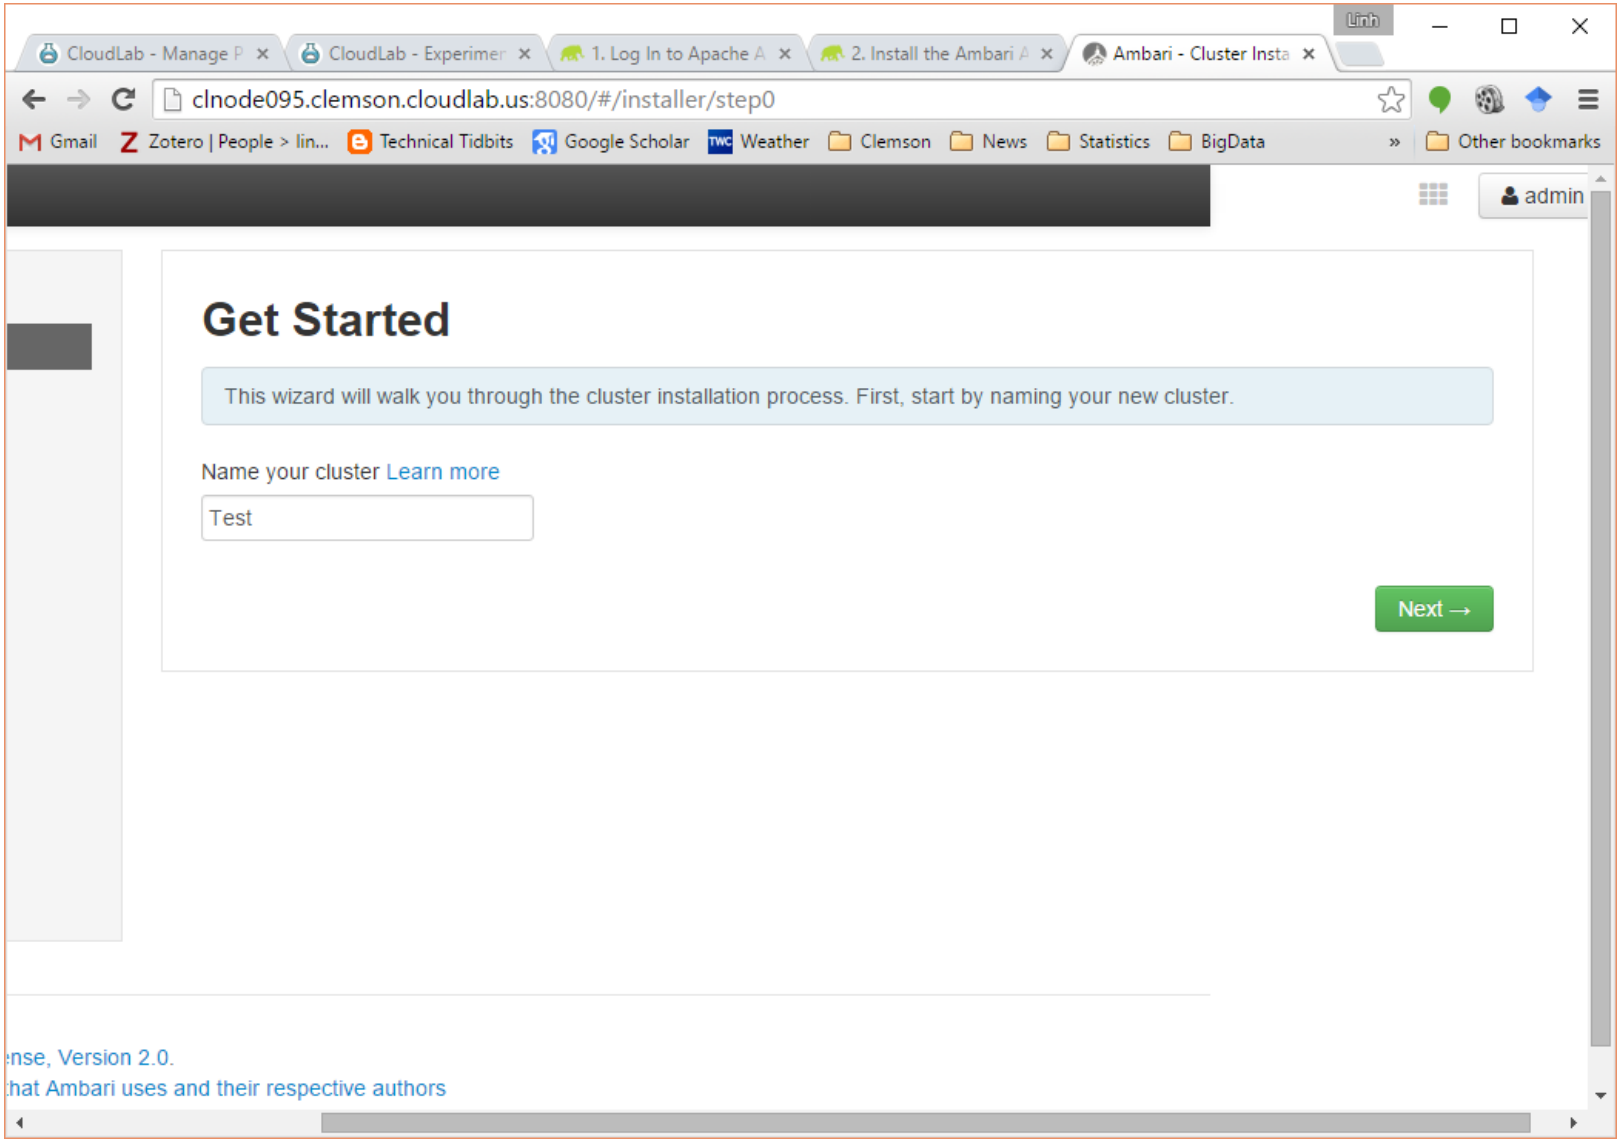

Ambari Server (admin/admin)

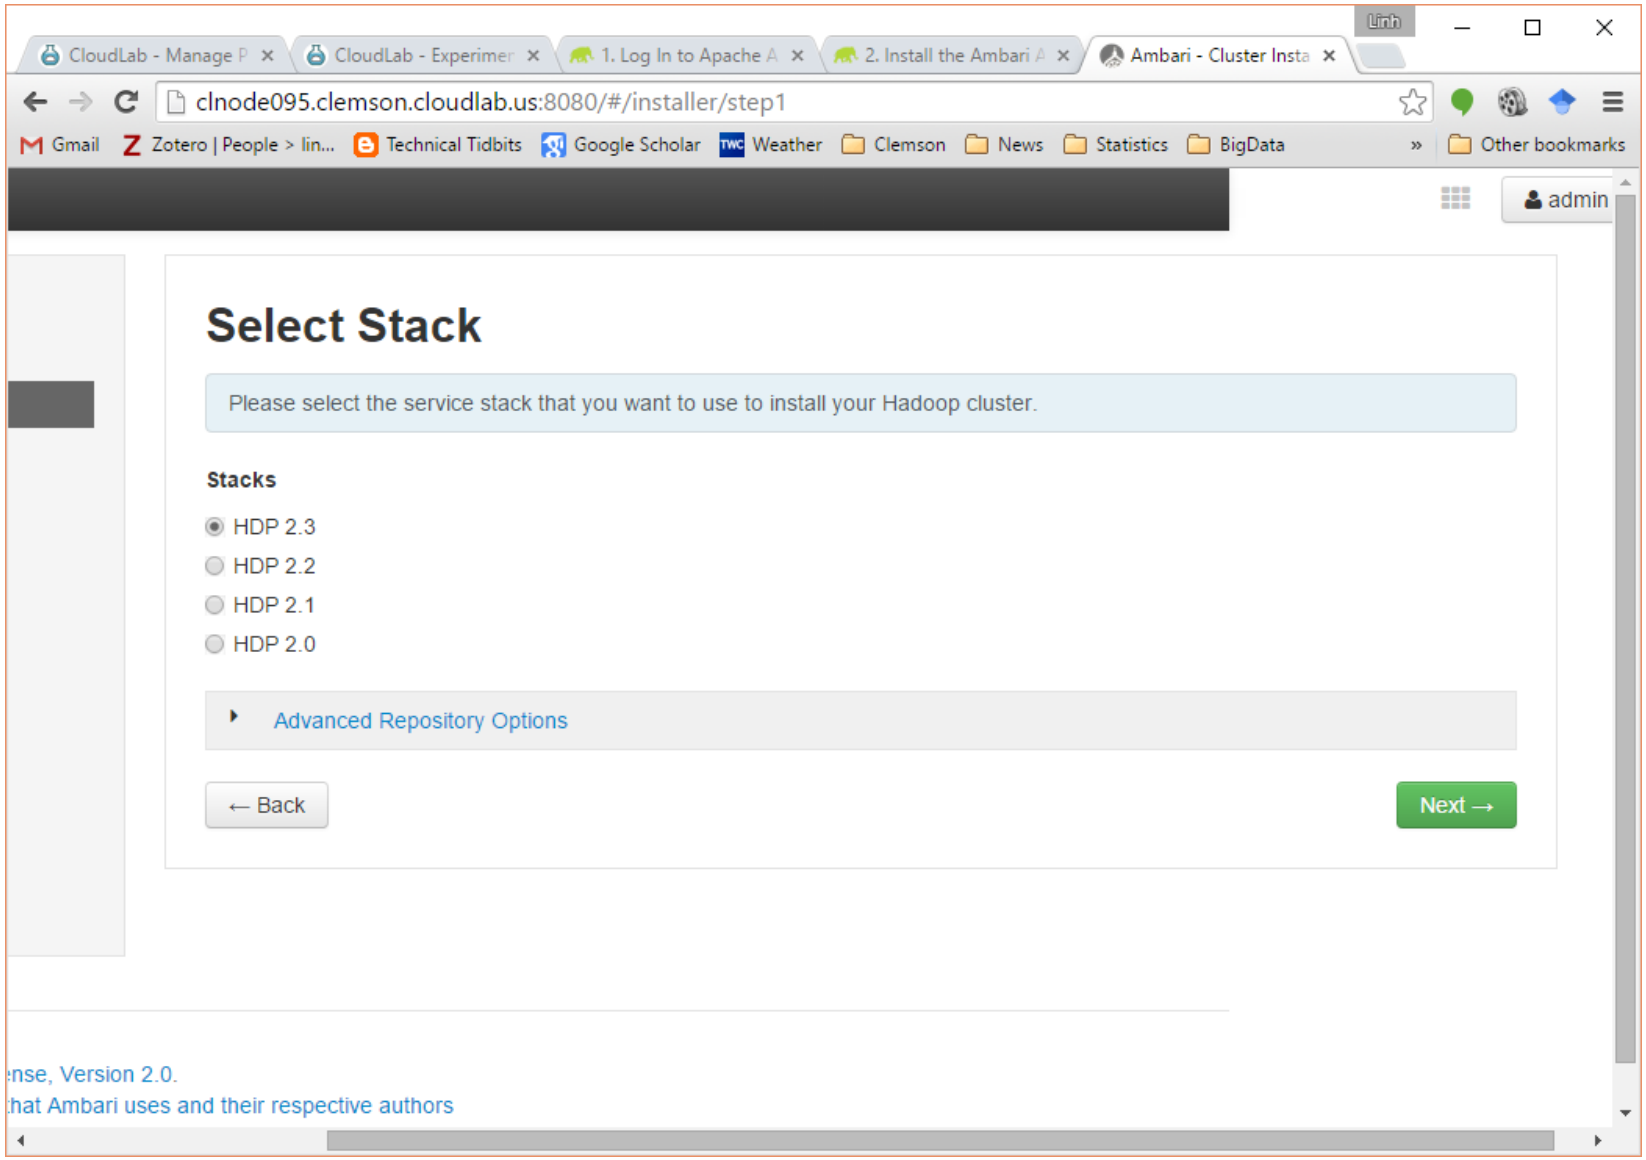

Ambari Server (admin/admin)

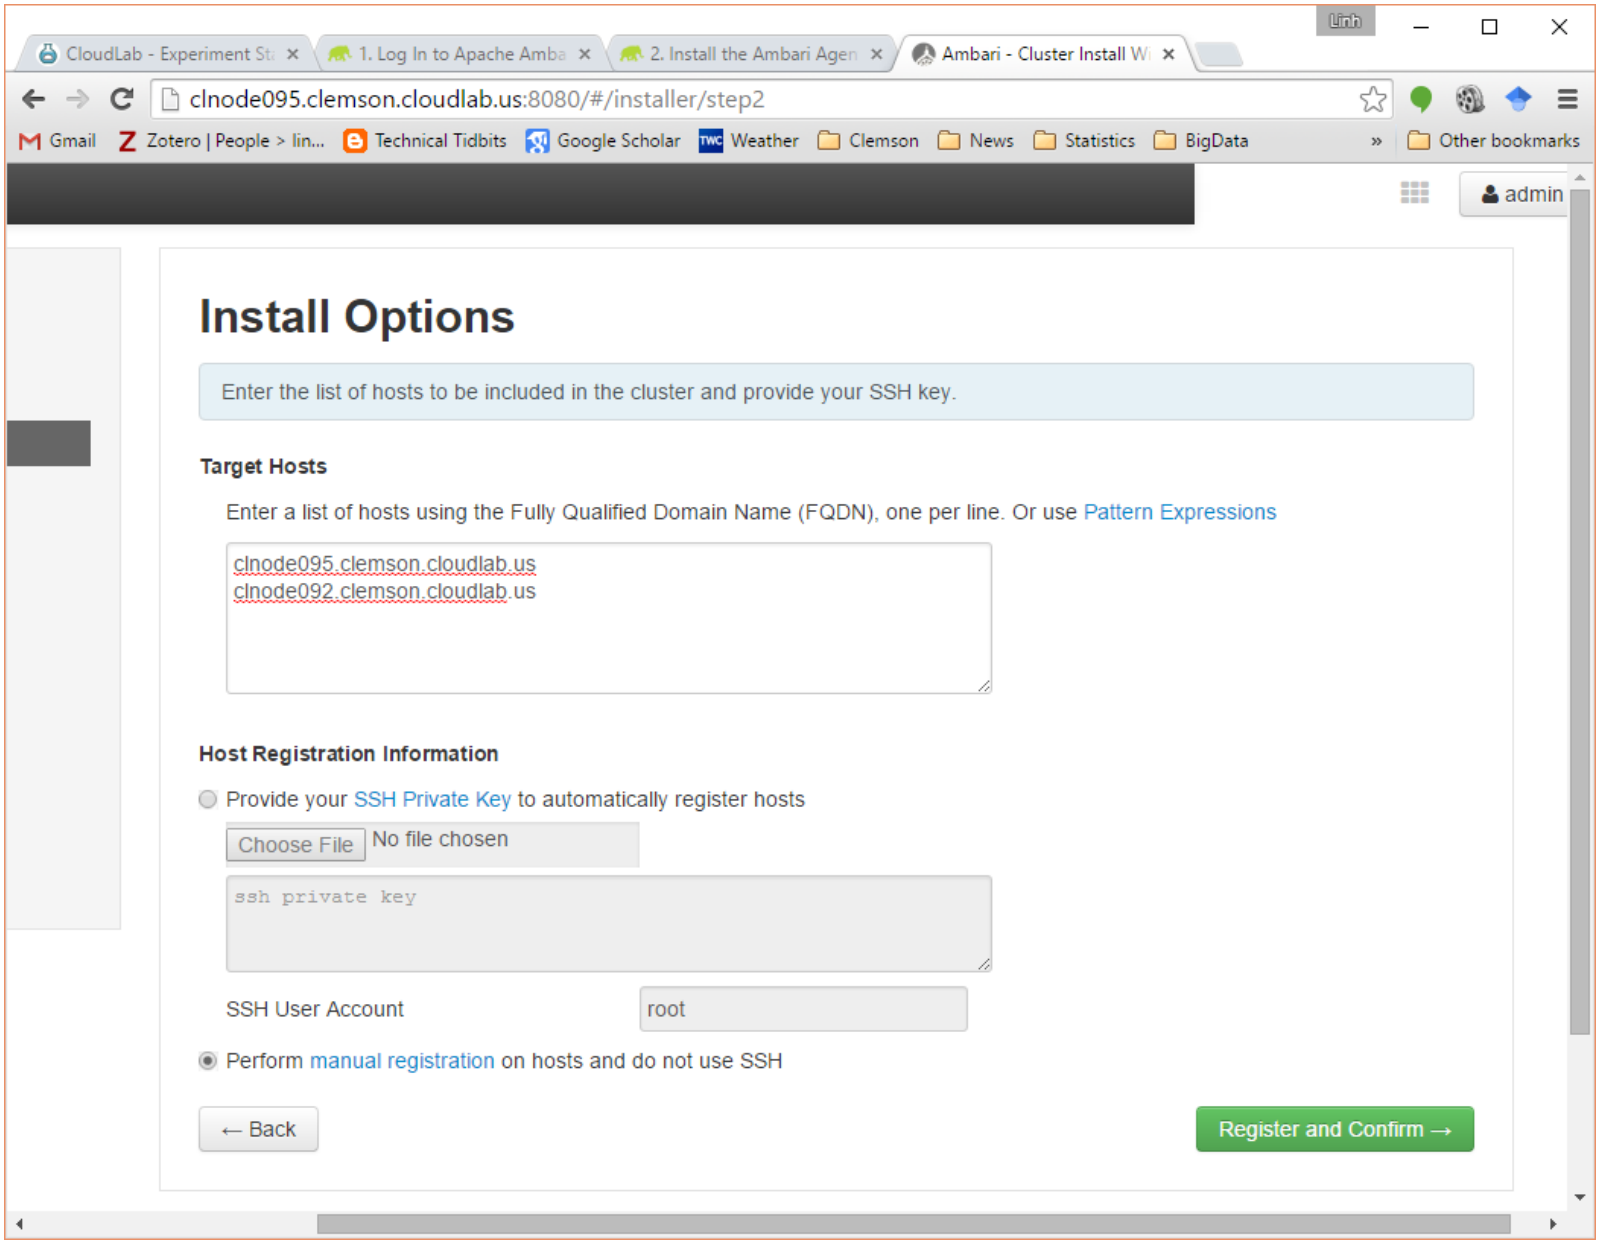

Assuming you had ambari agents up and running …

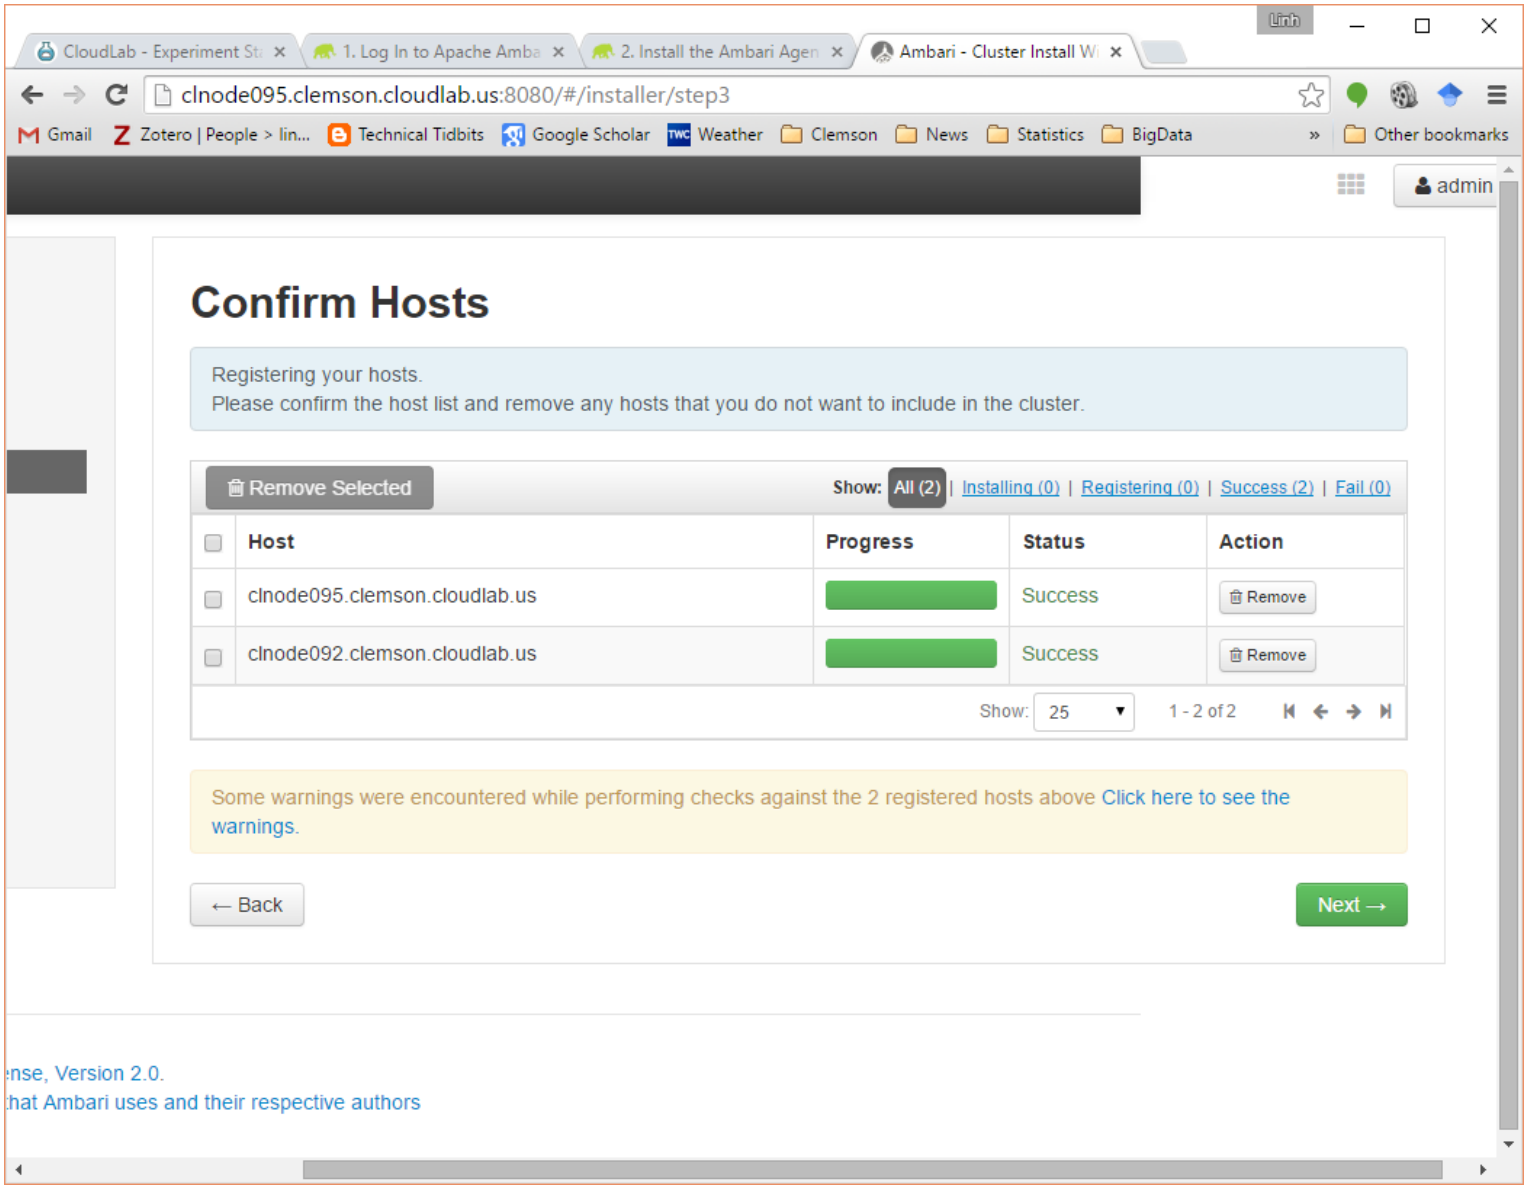

Ambari Server (admin/admin)

- HDFS
- YARN+MapReduce2
- Tez
- ZooKeeper
- Ambari Metrics

Ambari Server (admin/admin)

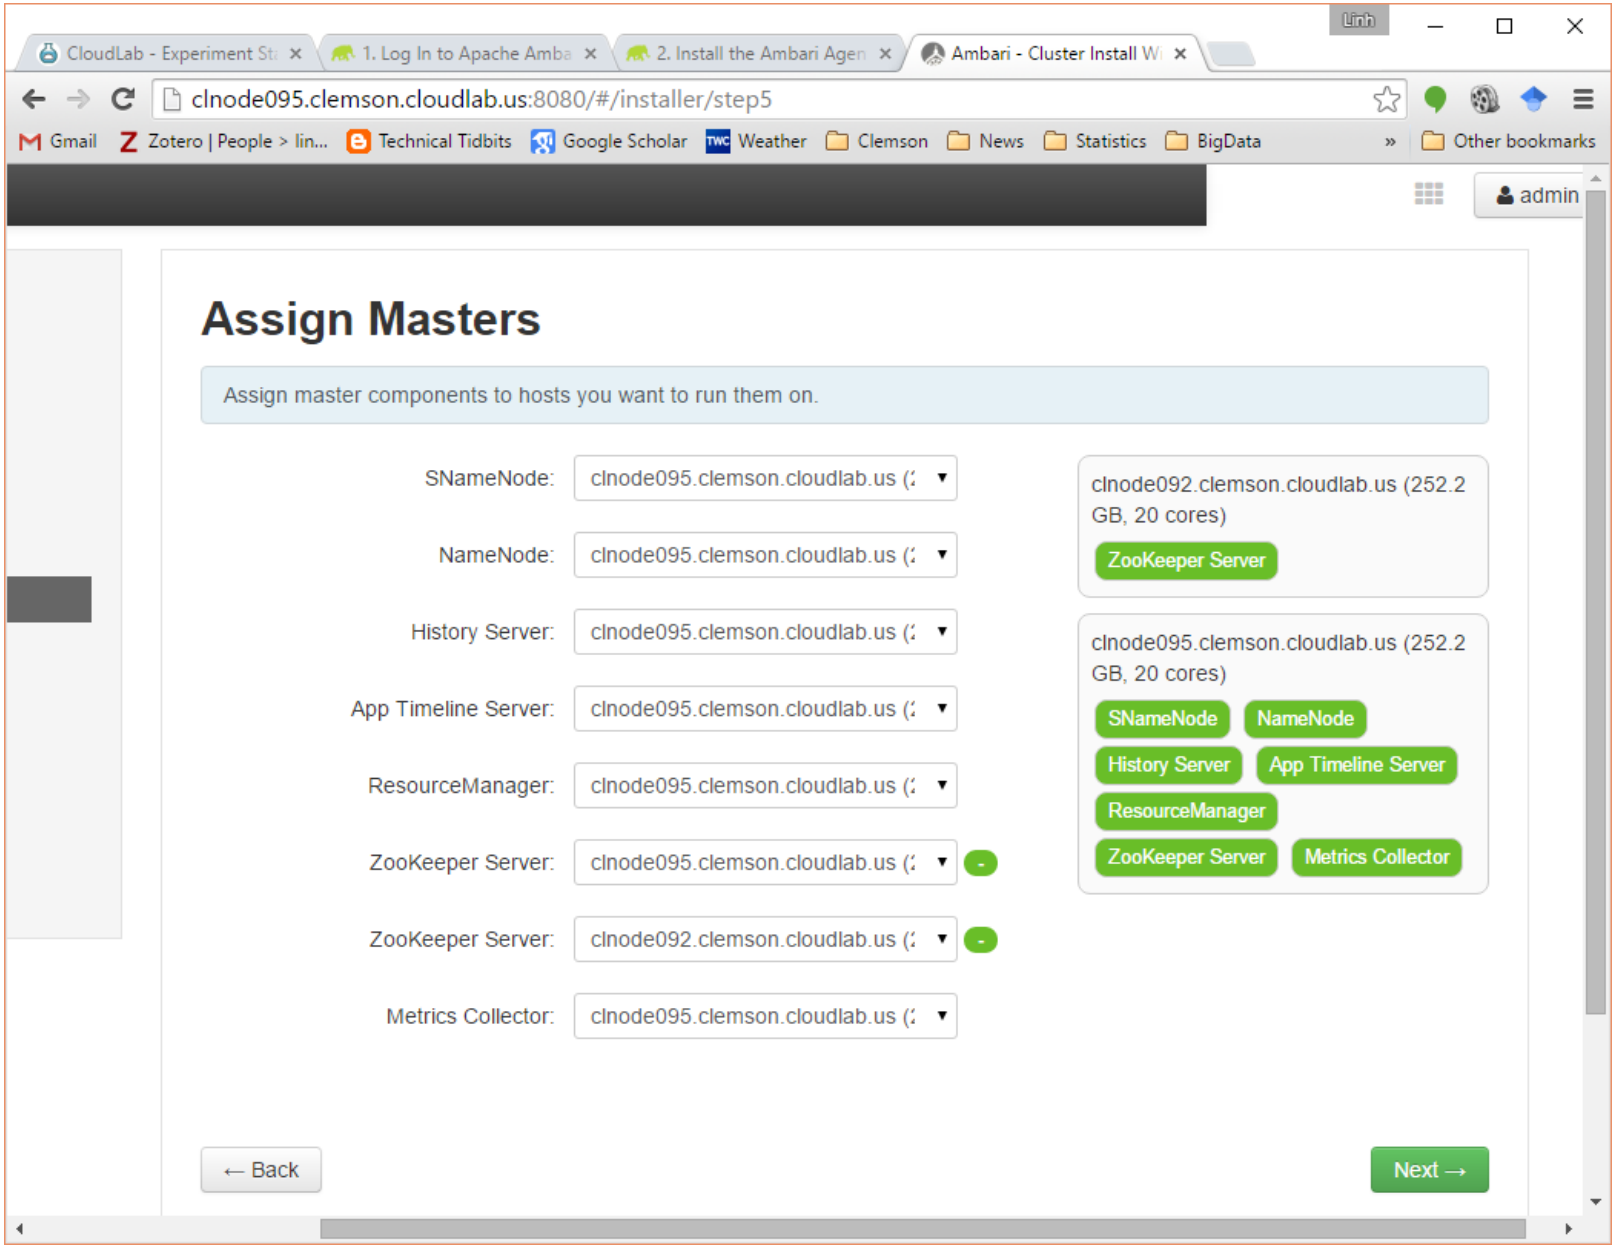

Ambari Server (admin/admin)

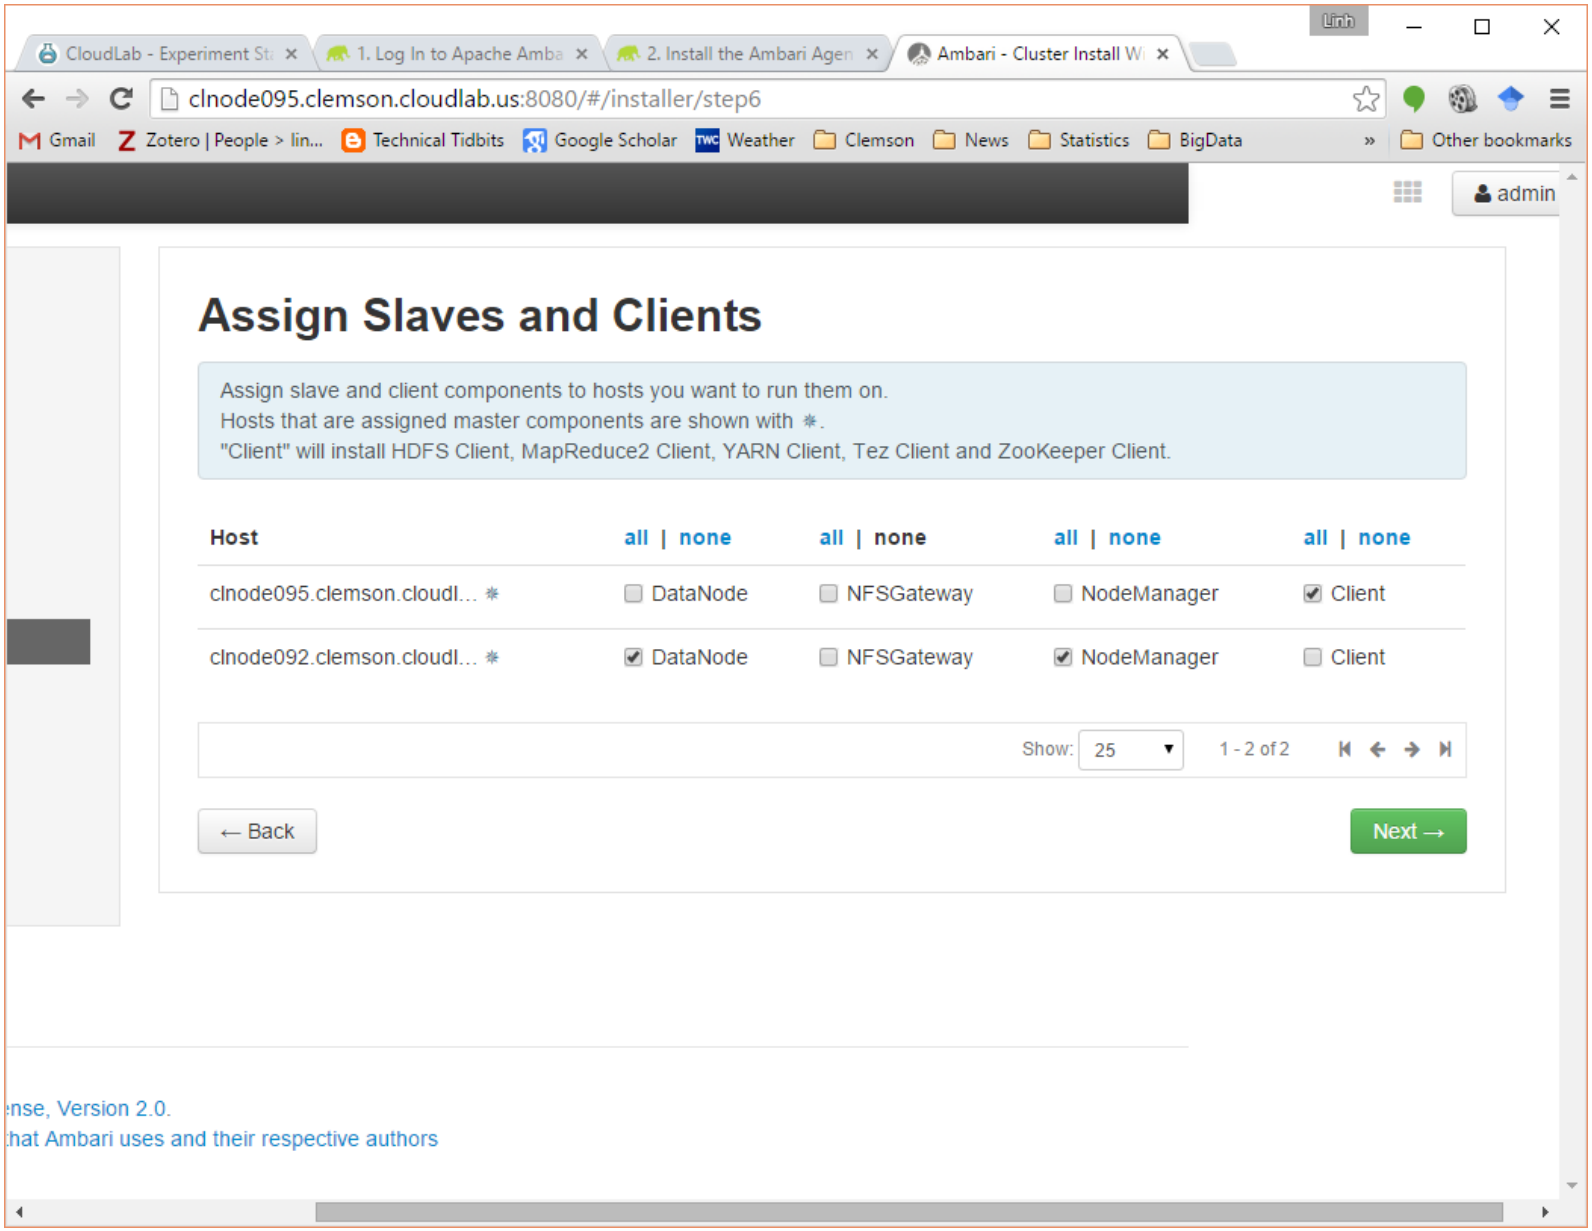

Edit configuration as you see fit

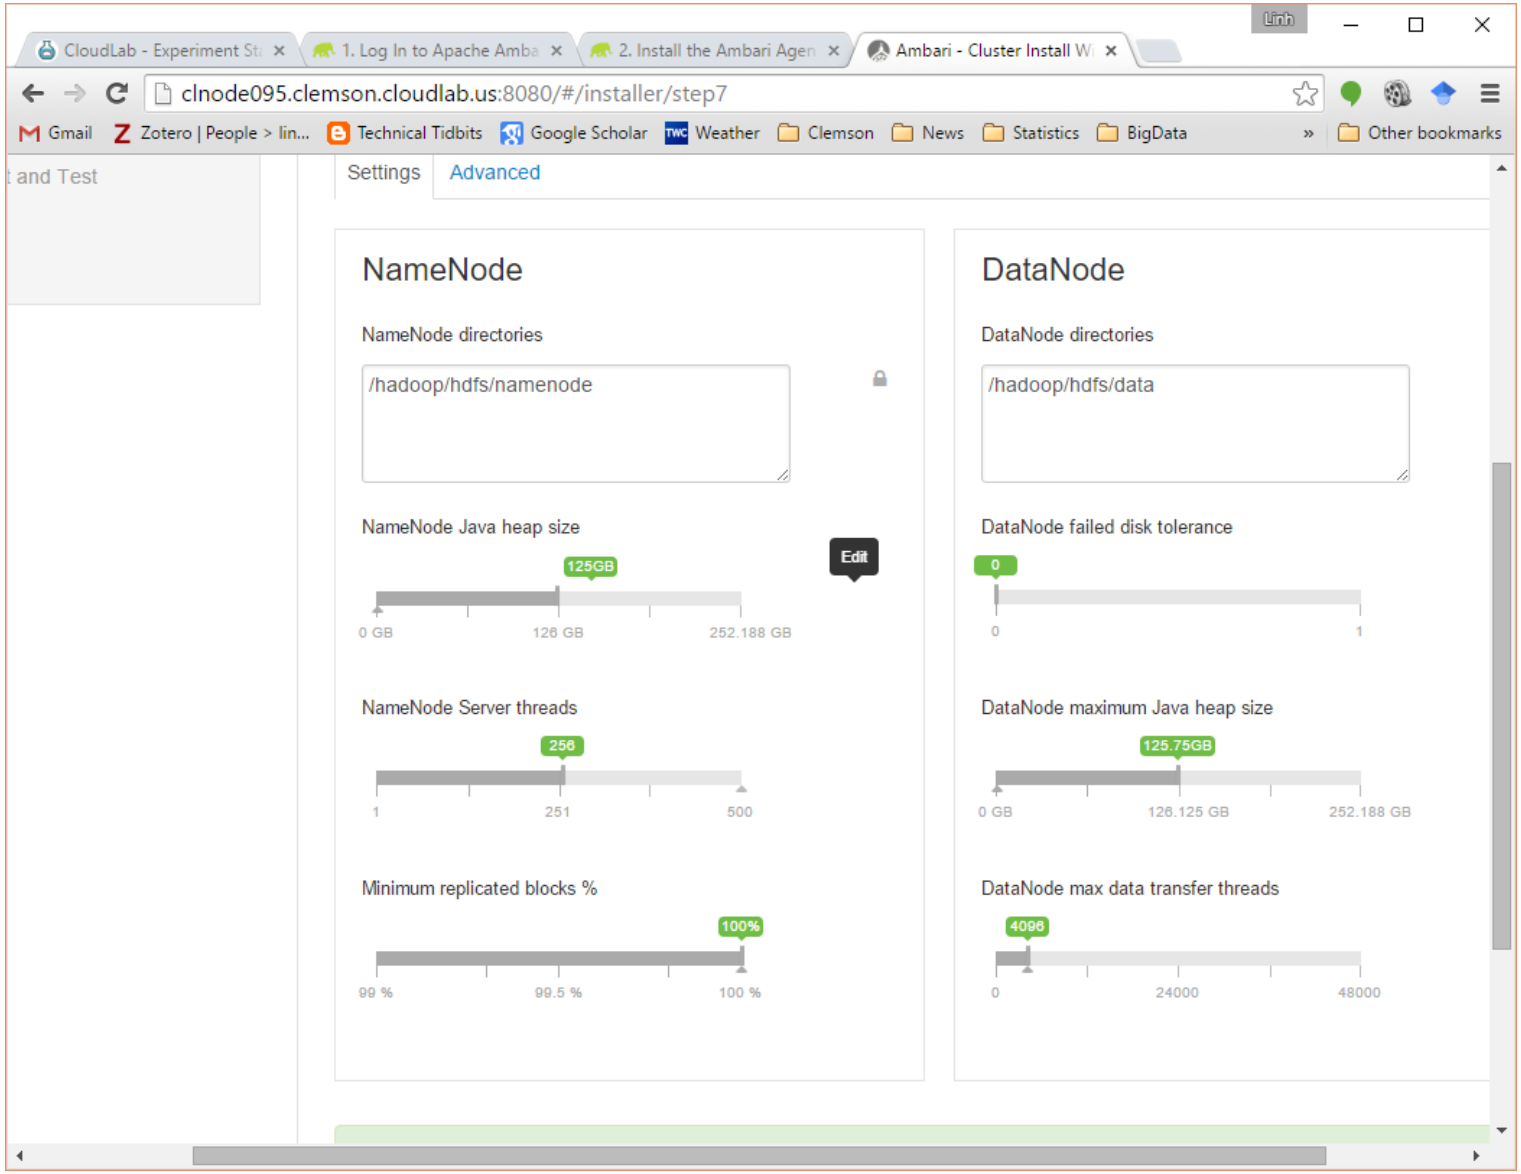

Deploy …

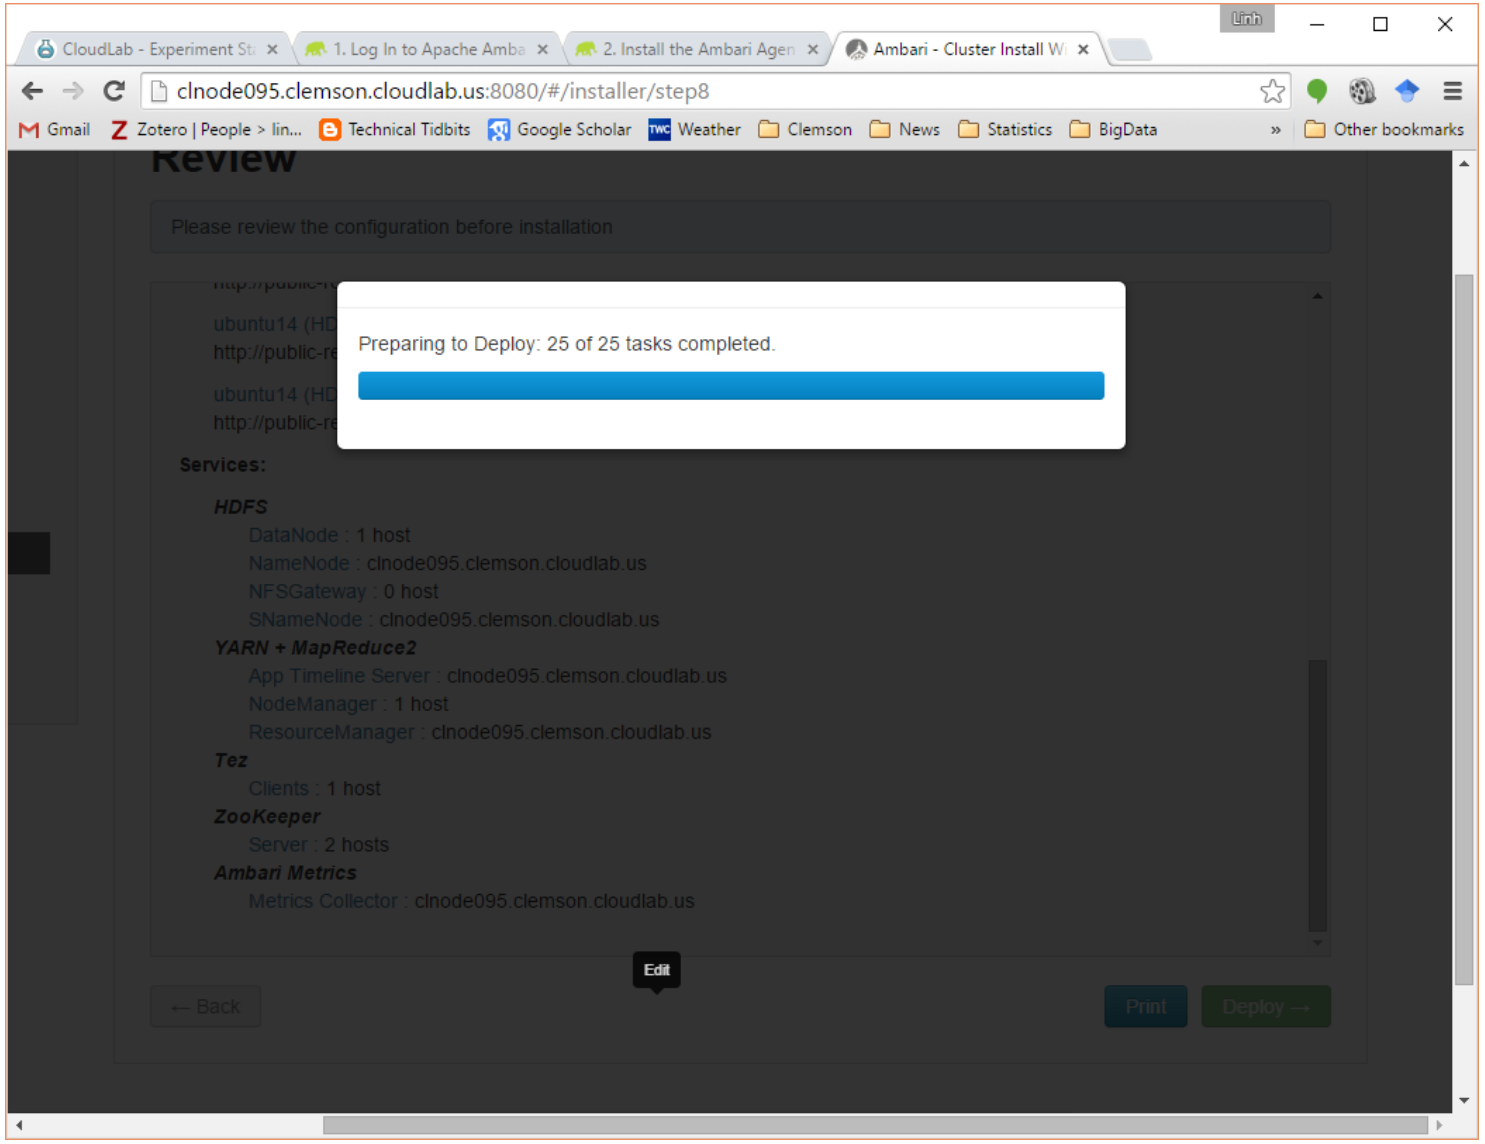
CloudLab - Experiment St
1. Log In to Apache Amba
2. Install the Ambari Agen
Ambari - Cluster Install Wi
Linh
clnode095.clemson.cloudlab.us:8080/#/installer/step8
Gmail
Zotero | People > lin...
Technical Tidbits
Google Scholar
Weather
Clemson
News
Statistics
BigData
Other bookmarks
Review
Please review the configuration before installation
Preparing to Deploy: 25 of 25 tasks completed.
Services:
HDFS
DataNode : 1 host
NameNode : clnode095.clemson.cloudlab.us
NFSGateway : 0 host
SNameNode : clnode095.clemson.cloudlab.us
YARN + MapReduce2
App Timeline Server : clnode095.clemson.cloudlab.us
NodeManager : 1 host
ResourceManager : clnode095.clemson.cloudlab.us
Tez
Clients : 1 host
ZooKeeper
Server : 2 hosts
Ambari Metrics
Metrics Collector : clnode095.clemson.cloudlab.us
Edit
← Back
Print
Deploy →

Warning due to lack of space and failed checks (ignore)

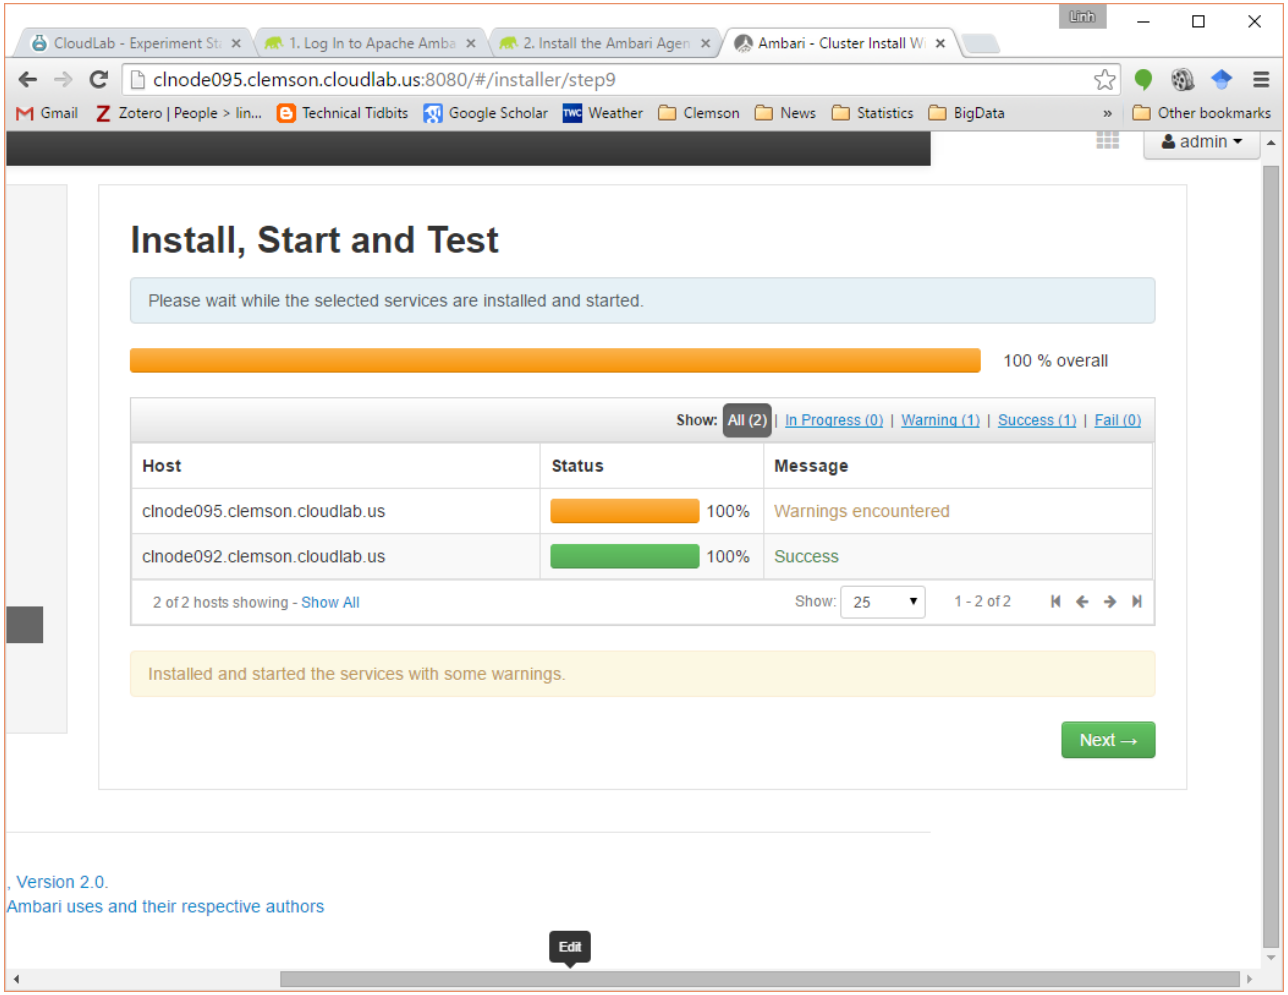

HDFS

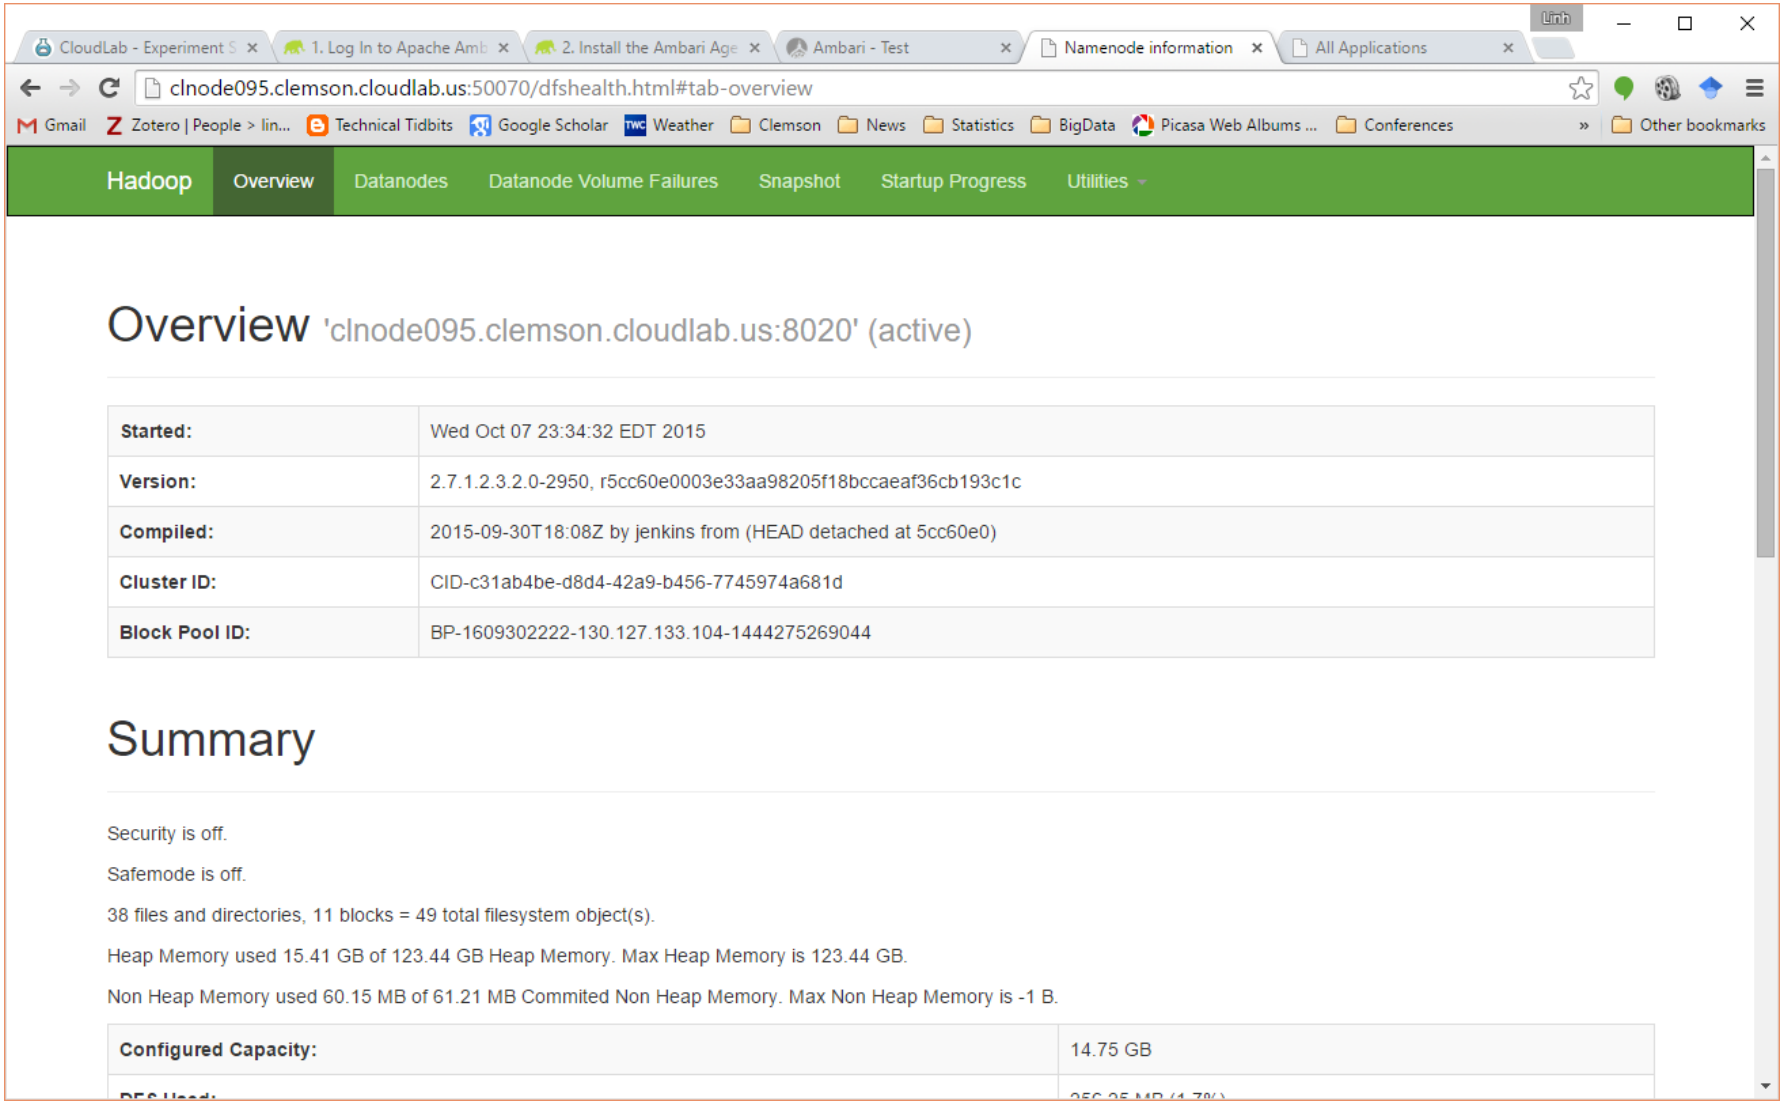
CloudLab - Experiment S
1. Log In to Apache Amb
2. Install the Ambari Age
Ambari - Test
Namenode information
All Applications
Linh
clnode095.clemson.cloudlab.us:50070/dfshealth.html#tab-overview
Gmail
Zotero | People > lin...
Technical Tidbits
Google Scholar
Weather
Clemson
News
Statistics
BigData
Picasa Web Albums ...
Conferences
Other bookmarks
Hadoop
Overview
Datanodes
Datanode Volume Failures
Snapshot
Startup Progress
Utilities
Overview 'clnode095.clemson.cloudlab.us:8020' (active)
Started: Wed Oct 07 23:34:32 EDT 2015
Version: 2.7.1.2.3.2.0-2950, r5cc60e0003e33aa98205f18bccaeaf36cb193c1c
Compiled: 2015-09-30T18:08Z by jenkins from (HEAD detached at 5cc60e0)
Cluster ID: CID-c31ab4be-d8d4-42a9-b456-7745974a681d
Block Pool ID: BP-1609302222-130.127.133.104-1444275269044
Summary
Security is off.
Safemode is off.
38 files and directories, 11 blocks = 49 total filesystem object(s).
Heap Memory used 15.41 GB of 123.44 GB Heap Memory. Max Heap Memory is 123.44 GB.
Non Heap Memory used 60.15 MB of 61.21 MB Commited Non Heap Memory. Max Non Heap Memory is -1 B.
Configured Capacity: 14.75 GB

YARN

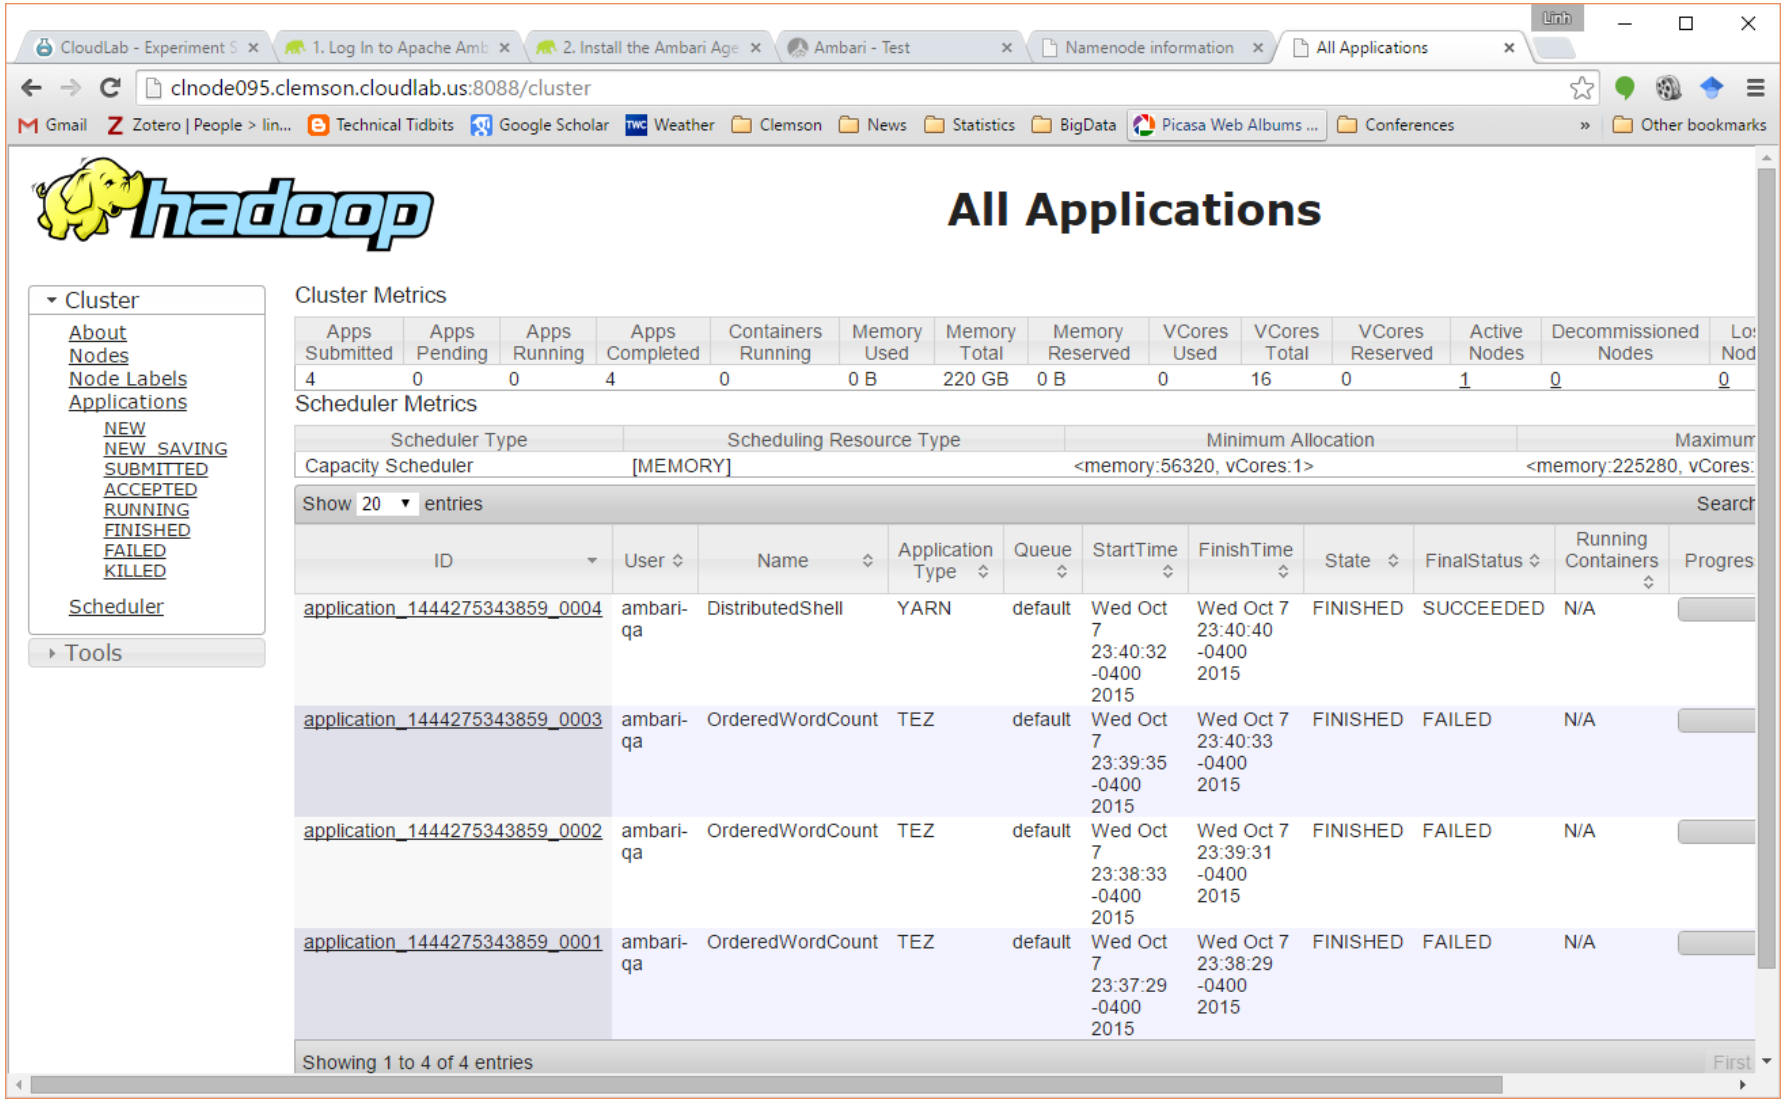
CloudLab - Experiment S
1. Log In to Apache Amb
2. Install the Ambari Age
Ambari - Test
Namenode information
All Applications
clnode095.clemson.cloudlab.us:8088/cluster
Gmail
Zotero | People > lin...
Technical Tidbits
Google Scholar
Weather
Clemson
News
Statistics
BigData
Picasa Web Albums ...
Conferences
Other bookmarks
hadoop
All Applications
Cluster
About
Nodes
Node Labels
Applications
NEW
NEW_SAVING
SUBMITTED
ACCEPTED
RUNNING
FINISHED
FAILED
KILLED
Scheduler
Tools
Cluster Metrics
Apps Submitted 4
Apps Pending 0
Apps Running 0
Apps Completed 4
Containers Running 0
Memory Used 0 B
Memory Total 220 GB
Memory Reserved 0 B
VCores Used 0
VCores Total 16
VCores Reserved 0
Active Nodes 1
Decommissioned Nodes 0
Scheduler Metrics
Scheduler Type: Capacity Scheduler
Scheduling Resource Type: [MEMORY]
Minimum Allocation: <memory:56320, vCores:1>
Maximum: <memory:225280, vCores:
Show 20 entries
Search
ID | User | Name | Application Type | Queue | StartTime | FinishTime | State | FinalStatus | Running Containers | Progres
application_1444275343859_0004 | ambari-qa | DistributedShell | YARN | default | Wed Oct 7 23:40:32 -0400 2015 | Wed Oct 7 23:40:40 -0400 2015 | FINISHED | SUCCEEDED | N/A
application_1444275343859_0003 | ambari-qa | OrderedWordCount | TEZ | default | Wed Oct 7 23:39:35 -0400 2015 | Wed Oct 7 23:40:33 -0400 2015 | FINISHED | FAILED | N/A
application_1444275343859_0002 | ambari-qa | OrderedWordCount | TEZ | default | Wed Oct 7 23:38:33 -0400 2015 | Wed Oct 7 23:39:31 -0400 2015 | FINISHED | FAILED | N/A
application_1444275343859_0001 | ambari-qa | OrderedWordCount | TEZ | default | Wed Oct 7 23:37:29 -0400 2015 | Wed Oct 7 23:38:29 -0400 2015 | FINISHED | FAILED | N/A
Showing 1 to 4 of 4 entries
First

Tutorial

```
sudo su hdfs
hdfs dfs -mkdir /user/<username>
hdfs dfs -chown <username>:<username> /user/<username>
exit to <username>
hdfs dfs -ls /user/
git clone https://github.com/clemsoncoe/Introduction-to-Hadoop-data.git
cd Introduction-to-Hadoop-data
hdfs dfs -put gutenberg-shakespeare.txt /user/<username>/
yarn jar /usr/hdp/current/hadoop-mapreduce-client/hadoop-mapreduce-examples-2.7.1.2.3.6.0-3796.jar wordcount gutenberg-Shakespeare.txt output/
hdfs dfs -ls output
hdfs dfs -cat output/part-r-00000
```

Assignment Ideas

- Deploy a Hadoop cluster and upload a large data set (Airline on-time performance data: http://stat-computing.org/dataexpo/2009/the-data.html)
- Examine and investigate performance of Hadoop MapReduce as data nodes are killed/added to the cluster
- Examine performance of Hadoop MapReduce as data nodes are located on different sites