

Deployment of Virtual Cluster on a Commercial Cloud Platform for Molecular Docking and Elasticity of the Clusters

Nara Institute of Science and Technology (NAIST), Japan
Katy Pham, 08/27/14

Overview

- Objective: upload a virtual machine that runs a protein-ligand interaction simulating program (Dock6) onto FutureGrid
- Purpose: having Dock6 onto the commercial cloud will allow tasks to be performed cheaply and efficiently
- 3 areas will be investigated:
 - Elasticity of the clusters (Katy)
 - Fault tolerance of the system (Derek)
 - Use of several virtual clusters on various commercial clouds to form a single system (Anthony)

Week 8 Progress

- Completed the yarn-pseudo distributed operation job and it worked successfully
- Set up the Hadoop cluster using Tokyo VM as the master and Fanship VM as the slave
- Edited the configuration files core-site.xml and yarn-site.xml – specified the ip addresses for the hosts
- Ran into problems while setting up the hadoop cluster: one of them was related to firewall being up (although we thought that we permanently disabled firewalls on our VMs)
- Another problem was that when turning on the services for Namenode and Datanode, only Tokyo showed up as being connected → solution: hadoop required all machines to have a list of the hostnames of all machines in the cluster inside the hostfile
- In order to facilitate the use of Hadoop with ViNe (for future usage), we decided to edit the configuration file hdfs-site.xml so that there's no need to keep the /etc/hosts file consistent on all nodes (easier because ViNe assigns an IP address for each node dynamically)
- Successful Run through of a sample yarn job on the hadoop cluster

Final Results:

- Installation of Java version 8 on the VMs in order to install Hadoop
- Successfully installed Hadoop on Barco Front End 1 VM
- Successfully tested Hadoop in standalone mode
- Successfully tested Hadoop in pseudo-distributed mode
- Successfully tested Hadoop in fully-distributed mode
- Performed cloning of Barco Front End 1 VM to make Tokyo VM and Fanship VM
- Creation of cluster made up of Tokyo VM (master) and Fanship VM (slave)
- Ran multiple sample commands in each mode to verify that Hadoop was working properly

Exploring the Culture

PRIME
2014

SHIMOJO SENSEI' S LAB TRIP

■ Amanohashidate

- Ate Japanese Seafood Donburi
- Went on a sightseeing ship
- Got to be at the beach for a little
- Went to the top of Kasamatsu park after (literally) riding the ropeway
- Went to winery to do wine tasting
- Stayed at an onsen hotel in Kinosaki (situated in city full of onsens!)
- Experienced the onsen experience! Put on yukatas, ate traditional Japanese dinners and went to onsens!
- Fireworks!

■ Izushi

- Had traditional Japanese breakfast
- Went to an onsen one more time
- Experienced making soba noodles from scratch (super interesting and fun), then got to eat our own soba!
- Walked around Izushi and did sake tasting

Exploring the Culture

PRIME 2014



Exploring the Culture

PRIME 2014



Exploring the Culture

Shimojo Sensei's Lab at Kasamatsu Park!
Thank you for inviting us, Shimojo Sensei!

PRIME 2014



Acknowledgements

- I would like to thank:
 - The graduate students, professors & assistant professors for being so friendly, welcoming and helpful
 - Dr. Haga for his constant guidance and for providing us with this incredible opportunity
 - Ichikawa-Sensei for always answering our questions and teaching us more about Japan
 - The amazing PRIME Staff (Dr. Gabriele Wienhausen – Principal Investigator, Mrs. Teri Simas, Mr. Jim Galvin) for continuous support and guidance
 - Dr. Abbie Celniker for her generous donations that are allowing for me to participate in this program
 - JASSO for the scholarship
 - NAIST for providing the accommodations
 - Karen Rodriguez and Kevin Lam for training us during Spring Quarter
 - Previous PRIME students for their excellent advice