

# Homework 1

---

*EECS 495*

*Hot Topics in Distributed Systems: Data-Intensive Computing*

<http://www.eecs.northwestern.edu/~iraicu/teaching/EECS495-DIC/index.html>

**Assigned:** 01-12-2010

**Due:** 11:59PM 01-13-2010

## Overview

There is a Linux-based virtual cluster configured for you to use for this course. This cluster is actually running on a single machine, called Falkon; the cluster name is hence also Falkon. The machine is a 16-core Intel Xeon at 2.33GHz, 48GB RAM, 7TB RAID5 disk, 1Gb/s network, and Linux OpenSuse 11.2 x64. The machine (also to be considered the head-node) can be accessed via ssh at "falkon.eecs.northwestern.edu". There are 10 virtual machines (VM) configured that can be accessed from the head-node only, as they do not have a globally accessible network connectivity. From the head-node, you can ssh to any of the 10 virtual nodes with the following names: falkon01, falkon02, falkon03, falkon04, falkon05, falkon06, falkon07, falkon08, falkon09, falkon10.

Each VM already has your accounts created, and a keyless public/private key installed. Each VM is composed of a single core, 4GB RAM, 100GB virtual disk (a file on the 7TB RAID5 disk), and 1Gb/s network (shared with all other VMs). Each VM has 3 networks configured, a loopback network (127.0.0.1), a host-only network (192.168.56.1xx), and a NAT network (10.0.2.x). The loopback network can be used for local testing. The host-only network is used for communicating with the head-node and other VMs. The NAT is used to communicate with the public network (e.g. Internet). Each VM has NFS mounted in /home, from the head-node. So, essentially your /home/userid folder is shared among all nodes, head-node and all VMs. User accounts are not centralized, so if you change a password on a VM or the head-node, the operations are only local. This is why I have installed the keyless public/private keys on all VMs so you can easily connect to all VMs, and not have to worry about the passwords. The firewall has been disabled on the head-node and all VMs, to make sure it doesn't interfere with your projects.

There is also a SVN server setup to help you with code source control. Once you decide on a project, give it a name, and I'll create a repository for your project, and give you detailed instructions on how to contact the SVN repository. For example, the current project that is in the SVN server called "Falkon2" can be checked out with the following command:

```
svn co svn+ssh://falkon.eecs.northwestern.edu/home/svn/falkon2
```

**IMPORTANT:** This Falkon cluster is an experimental cluster, aimed to facilitate you with a flexible testbed for your course projects. Although the storage system is protected from hard drive failures with RAID5, the storage array can still break. There are no regular backups being conducted on this machine, and hence always make sure to backup your work to additional locations, in case of failures. Also, the machine is quite opened up, firewall disabled, and many services are enabled. I will periodically perform updates, but they will not be at regular intervals. The combination of these two things makes the machine more prone to being attacked. I would not keep any sensitive information on this machine, important passwords, important private keys, etc.

As of right now, there are the following accounts configured on the head-node and virtual machines:

```
ahug berkin cjin hgao iraicu kpzhang vrastogi ycao
```

Email me at [iraicu@eecs.northwestern.edu](mailto:iraicu@eecs.northwestern.edu) for your password.

## Assignment

Record the output of the following commands:

```
ssh falkon.eecs.northwestern.edu
pwd
cat ~/.ssh/authorized_keys
ssh falkon01
exit
ssh falkon02
exit
ssh falkon03
exit
ssh falkon04
exit
ssh falkon05
exit
ssh falkon06
exit
ssh falkon07
exit
ssh falkon08
exit
ssh falkon09
exit
ssh falkon10
exit
```

Now test the SVN server:

```
ssh falkon.eecs.northwestern.edu
svn co svn+ssh://falkon.eecs.northwestern.edu/home/svn/falkon2
cd falkon2/hw1
echo "Your Name" >> ${USER}.txt
svn add iraicu.txt
svn update
svn ci -m "HW1 for Your Name"
```

Now test the gcc and javac compiler:

- write a simple hello world program in C/C++ and in Java
- make sure they compile and run, on the head node and your VMs
- create a directory in falkon2/hw1 with your name
- put your C/C++ and Java code and appropriate makefile (but not your binaries) in that directory (e.g. falkon/hw1/iraicu")
- add the new dir, source files, and makefiles to SVN, and commit your changes