



Build your own cluster: Session 1

Farrukh Khan

Computational Science and AI Directorate (CSAID)

Fermi National Accelerator Laboratory

May 20, 2024

Why are we here?

- In short, we are here for some hands-on exercises today to give you a system management perspective
- The Monday exercises are divided as follows:
 - Monday morning (Session 1):
 - What is system management?
 - What is the hands-on exercise environment going to be like?
 - Linux 101
 - Monday morning (Session 2):
 - What is a facility?
 - HTCondor 101

What is system management?

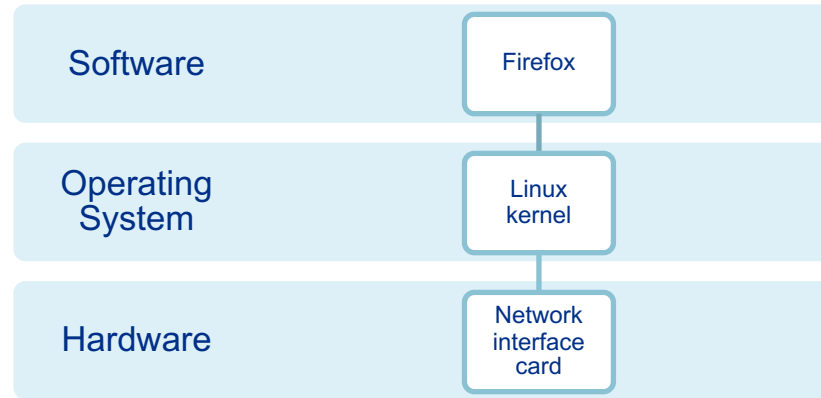
- System management is ensuring that a given system is operating efficiently and reliably. It also includes deployment of the system, monitoring the system, troubleshooting issues, providing support to system users and performing regular upkeep (upgrades, patching etc.)
- System management roughly translates to roles such as reliability engineers, system engineers and system administrators
- System managers work with a variety of system components such as:
 - Hardware
 - Software
 - Networking
 - Cybersecurity

What is system management?



What is a system?

- System is a logical entity that can refer to many things. In general, it is a collection of individual units (hardware, software, network, security) that work together to achieve a common objective
- For example, compute systems, storage systems, analysis systems etc.
- At the heart of any system is generally an operating system. An operating system is responsible for managing the underlying hardware resources based on the needs of the applications/services running on top



What is a system?

- Why is it important to know about an operating system? A system manager cannot perform their job functions without adequate mastery of the operating system!
- For computing systems at Fermilab, we use Alma Linux 9 as our operating system
- The pre-training exercises provided some key commands that you should be familiar with. We'll go over the knowledge check questions as a review on our Alma Linux testbed

How do I access my system?

- Make sure that you have Kerberos setup on your laptop and that you have a valid Kerberos token
- Please log in using your Fermilab computing/Kerberos username

```
mac-135932:~ fkhan$ kinit
fkhan@FNAL.GOV's password:
mac-135932:~ fkhan$ klist
Credentials cache: API:03A45E25-99B8-4D6D-82A2-7C202B5853E1
Principal: fkhan@FNAL.GOV

    Issued                Expires                Principal
May 16 11:05:09 2024  May 17 13:05:04 2024  krbtgt/FNAL.GOV@FNAL.GOV
mac-135932:~ fkhan$ ssh fkhan@cmswn3012.fnal.gov
Warning: Permanently added 'cmswn3012.fnal.gov' (ED25519) to the list of known hosts.

NOTICE TO USERS

This is a Federal computer (and/or it is directly connected to a
```

How do I access my system?

- A virtual environment consisting of containers has been setup on `cmswn3012.fnal.gov`
- At first login, a new container will be created for you to work in
 - This container is cordoned off to prevent cross talk with production Fermilab systems
 - This also means no Internet access
 - The container hostname is set to your Fermilab computing account
- All packages that you need for the exercise will be provided to you
 - If there is a package you absolutely must have, please talk to one of us
- The containers are persistent so you can safely log out
 - Subsequent log in will put you back inside the same container
 - If you need to 'reset' your container, please talk to one of us
- You have root access inside the container
- All student containers can reach other through a virtual network

Linux 101

- Now that you have access to your system, let's do a few exercises using Alma Linux 9

Linux 101

- Can you navigate to your home directory, create a subdirectory named `physics` and create an empty file named `computing`?

```
[root@fkhan interactive]# whoami
root
[root@fkhan interactive]# cd $HOME
[root@fkhan ~]# pwd
/root
[root@fkhan ~]# mkdir physics
[root@fkhan ~]# touch computing
[root@fkhan ~]# ls
computing  physics
[root@fkhan ~]#
```

Linux 101

- How do you tell files and directories apart on the command line? Can you remove the directory named `physics`? Can you tell what time you created the file `computing` at by using `ls`?

```
[root@fkhan ~]# ls -l
total 4
-rw-r--r-- 1 root root    0 May 14 14:38 computing
drwxr-xr-x 2 root root 4096 May 14 14:38 physics
[root@fkhan ~]# rm -rf physics
[root@fkhan ~]# date
Tue May 14 14:39:49 UTC 2024
```

Linux 101

- Create a file named `intro.txt` and add a two line introduction of yourself. Save the file and `cat` it to make sure it has what you wrote.
 - This one is highly subjective due to the choice of editor. I prefer `vi`, so the screenshots below show how to do this `vim` (`vim` is an enhanced version of `vi`)
 - Contrary to popular Internet memes, don't be alarmed! I will show you how to exit out of `vim`
 - To create a file, run: `vim intro.txt`
 - To switch to insert mode, press `i` (case sensitive)
 - Write your text
 - To switch to command mode, press `Esc`
 - In command mode, type `:wq!` (`:` command directive, `w` for write, `q` for quit, `!` for force)
 - Press enter, and you're out!

```
[root@fkhan ~]# vim intro.txt
[root@fkhan ~]# cat intro.txt
My name is Farrukh. I <3 computer science.
```

Linux 101

- Create a file named `intro.txt` and add a two line introduction of yourself. Save the file and `cat` it to make sure it has what you wrote.



Linux 101

- On your terminal, run `sleep 600 &` (& is not a typo). This will create a sleep process of 600 seconds and send it to the background. Can you use `ps` to identify the process ID (or PID) of this process? Can you use `kill` command to terminate it?

```
[root@fkhan ~]# sleep 600 &
[1] 197
[root@fkhan ~]# ps
  PID TTY          TIME CMD
  123 pts/1        00:00:00 bash
  197 pts/1        00:00:00 sleep
  198 pts/1        00:00:00 ps
[root@fkhan ~]# kill -9 197
[root@fkhan ~]#
[1]+  Killed                  sleep 600
[root@fkhan ~]# ps
  PID TTY          TIME CMD
  123 pts/1        00:00:00 bash
  199 pts/1        00:00:00 ps
[root@fkhan ~]#
```

Linux 101

- Can you install a package using `dnf`? A couple of example packages to try the command on are `iotop` (utility that lets you see how stressed your disk is) and `htop` (a utility that lets you see CPU and memory consumption)
 - This is a bit tricky on our containers because they do not have Internet access

```
[root@fkhan ~]# cd /packages/
[root@fkhan packages]# ls
condor-23.0.8-1.el9.x86_64.rpm htop-3.3.0-1.el9.x86_64.rpm python3-condor-23.0.8-1.el9.x86_64.rpm
[root@fkhan packages]# dnf install -y ./htop-3.3.0-1.el9.x86_64.rpm
Last metadata expiration check: 0:36:42 ago on Tue May 14 15:56:41 2024.
Dependencies resolved.
```

Package	Architecture	Version	Repository	Size
Installing: htop	x86_64	3.3.0-1.el9	@commandline	198 k
Transaction Summary				
Install	1 Package			

Linux 101

- Can you find three packages installed on your system using `rpm` or `dnf` command?

```
[root@fkhan packages]# dnf list installed
Installed Packages
HdrHistogram_c.x86_64
ModemManager-glib.x86_64
SDL2.x86_64
abattis-cantarell-fonts.noarch
acl.x86_64
adobe-mappings-cmap.noarch
0.11.0-6.el9 @appstream
1.20.2-1.el9 @baseos
2.26.0-1.el9 @appstream
0.301-4.el9 @appstream
2.3.1-3.el9 @baseos
20171205-12.el9 @appstream
```