

The background of the slide features two red fedoras resting on a white surface. A wooden clothespin is attached to the brim of the larger hat in the foreground. The text is overlaid on this scene.

SC17 Red Hat Enterprise Linux Accelerated ML Containers Advanced Development Research

Bill Gray
Red Hat Performance Engineering
bgray@redhat.com

Context and Agenda



“Advanced Development Research” means

- Some usual rules may not apply...
- This tool may or may not be real.
- Even if it is real, it may or may not work.
- Even if it is real and does work, it may or may not exist tomorrow.
- So, basically no promises, at all, about anything.

ML Container Creator:

- The MLCC tool makes it easy to create containers to use various machine learning related software frameworks and libraries
- User input is mostly just picking items from a menu
- Output is a normal Open Container Initiative dockerfile

MLCC Presumed Value:

- Quick, easy, and convenient
 - You don't need to know how to build or install the packages
- MLCC curates sets of packages that work together
 - It knows about dependencies
 - It can optimize the package “synergy”
- Separates ML environments and preserves the system OS
- MLCC can optimize installation for specific target hardware, or leave things generic to maximize container portability

Feedback desired:

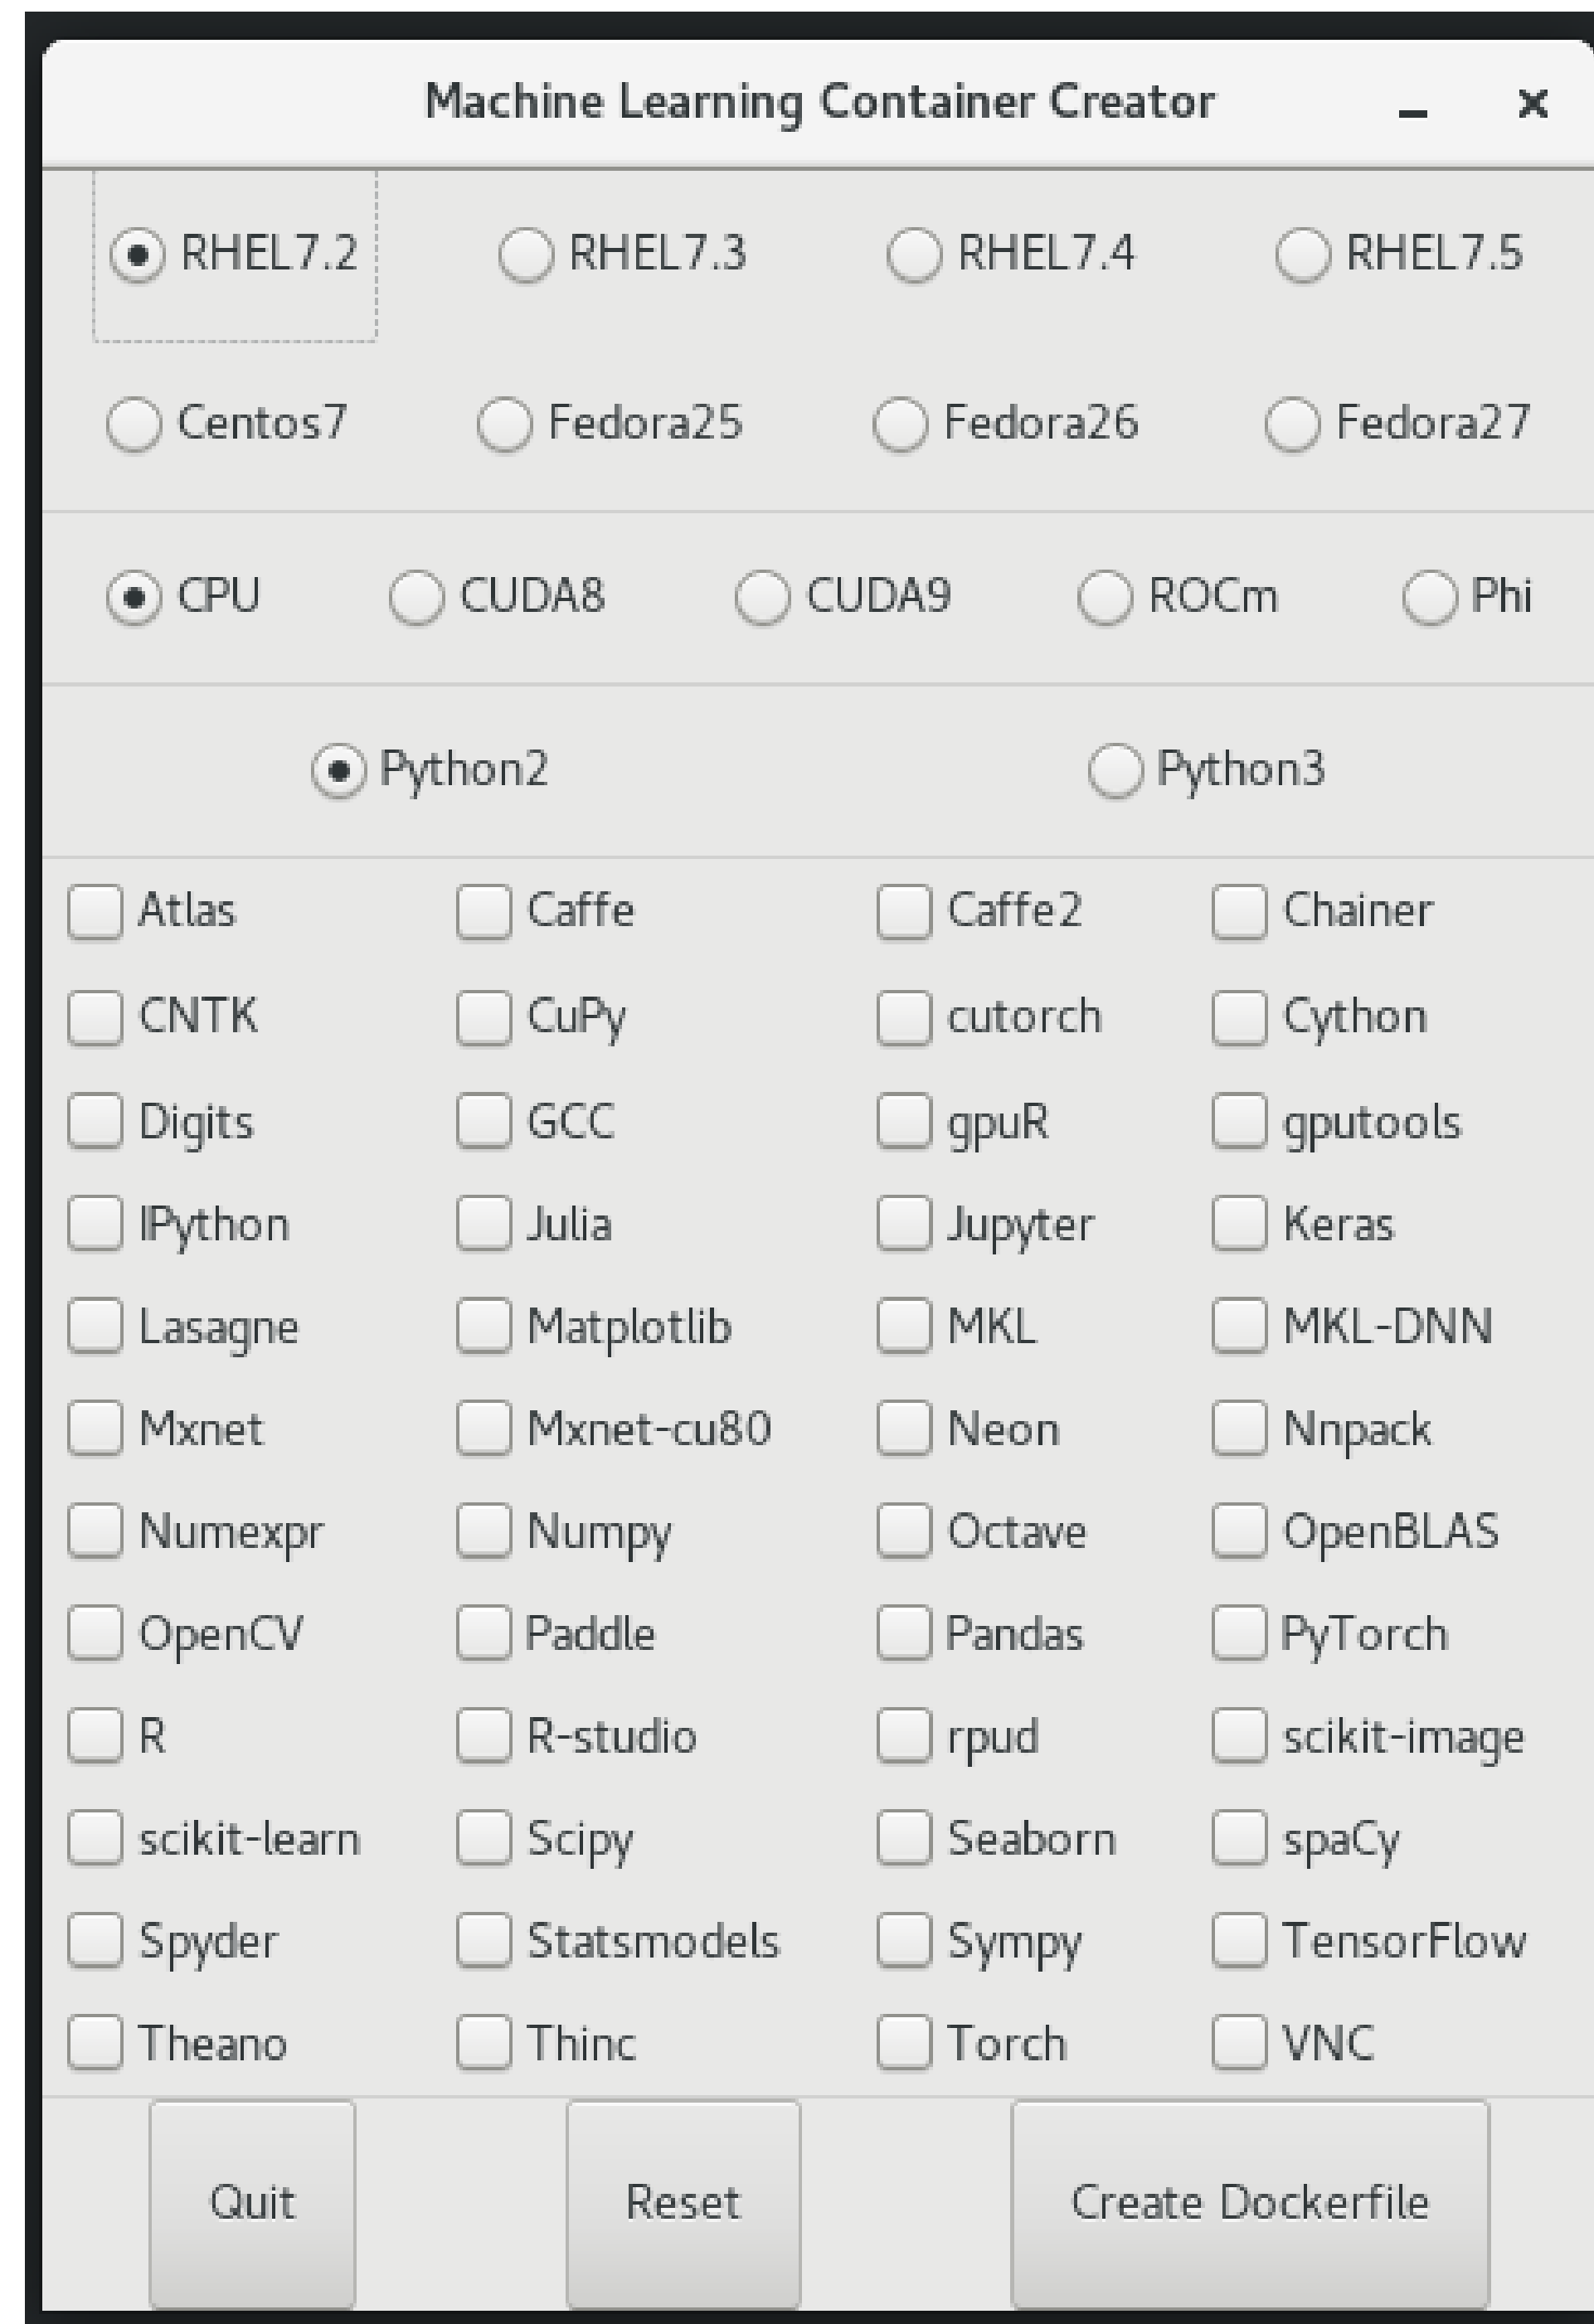
- “This is a stupid and worthless idea that will never fly because ...”
- “This is an interesting idea, and you should keep working on it because ...”
- “You should go learn about <great existing idea> because ...”
- “You should consider <great new idea> because ...”
- bgray@redhat.com

MLCC Interfaces



GUI

- Pick an Operating System
- Pick CPU, or some acceleration technology
- Pick a python version
- Pick various packages you want to have in your container.
(Dependencies are automatically selected.)



Interactive Command Line

- Kind of klunky command line interface
- More or less equivalent functionality to GUI interface
- Good for remote terminal connections to servers

```
Selected:
 0: RHEL7.4          1: CUDA9          2: Python2        3: Keras
 4: TensorFlow

OS Choices:
 3: Centos7         12: Fedora25     13: Fedora26     14: Fedora27
41: RHEL7.2        42: RHEL7.3      43: RHEL7.5

Accelerator Choices:
 6: CPU            7: CUDA8         36: Phi          44: ROCm

Python Choices:
37: Python3

Additional Packages:
 0: Atlas          1: Caffe          2: Caffe2         4: Chainer
 5: CNTK           8: CuPy          9: cutorch       10: Cython
11: Digits        15: GCC          16: gpuR         17: gputools
18: IPython       19: Julia        20: Jupyter      21: Lasagne
22: Matplotlib   23: MKL          24: MKL-DNN      25: Mxnet
26: Mxnet-cu80   27: Neon        28: Nnpack       29: Numexpr
30: Numpy         31: Octave       32: OpenBLAS     33: OpenCV
34: Paddle        35: Pandas       38: PyTorch      39: R
40: R-studio     45: rpy2         46: scikit-image 47: scikit-learn
48: Scipy        49: Seaborn      50: spaCy        51: Spyder
52: Statsmodels  53: Sympy       54: Theano       55: Thinc
56: Torch        57: VNC

(A)dd, (R)emove, (C)reate Dockerfile, (Q)uit:
```


“Pro” Interface: when you know what you want

- mlcc -i RHEL7.3,CUDA8,Python2,Caffe
- mlcc -i RHEL7.4,CPU,Python2,Chainer
- mlcc -i RHEL7.4,CUDA8,Python2,Chainer
- mlcc -i RHEL7.4,CUDA8,Python2,Lasagne,Theano
- mlcc -i RHEL7.4,CUDA8,Python2,Tensorflow,Keras
- mlcc -i Centos7,CUDA8,Python2,Lasagne,Theano
- mlcc -i Fedora26,CPU,Python3,Pandas,scikit-learn
- mlcc -i Fedora26,CUDA8,Python2,Keras,TensorFlow
- mlcc -i Fedora26,CUDA8,Python2,Lasagne,Theano

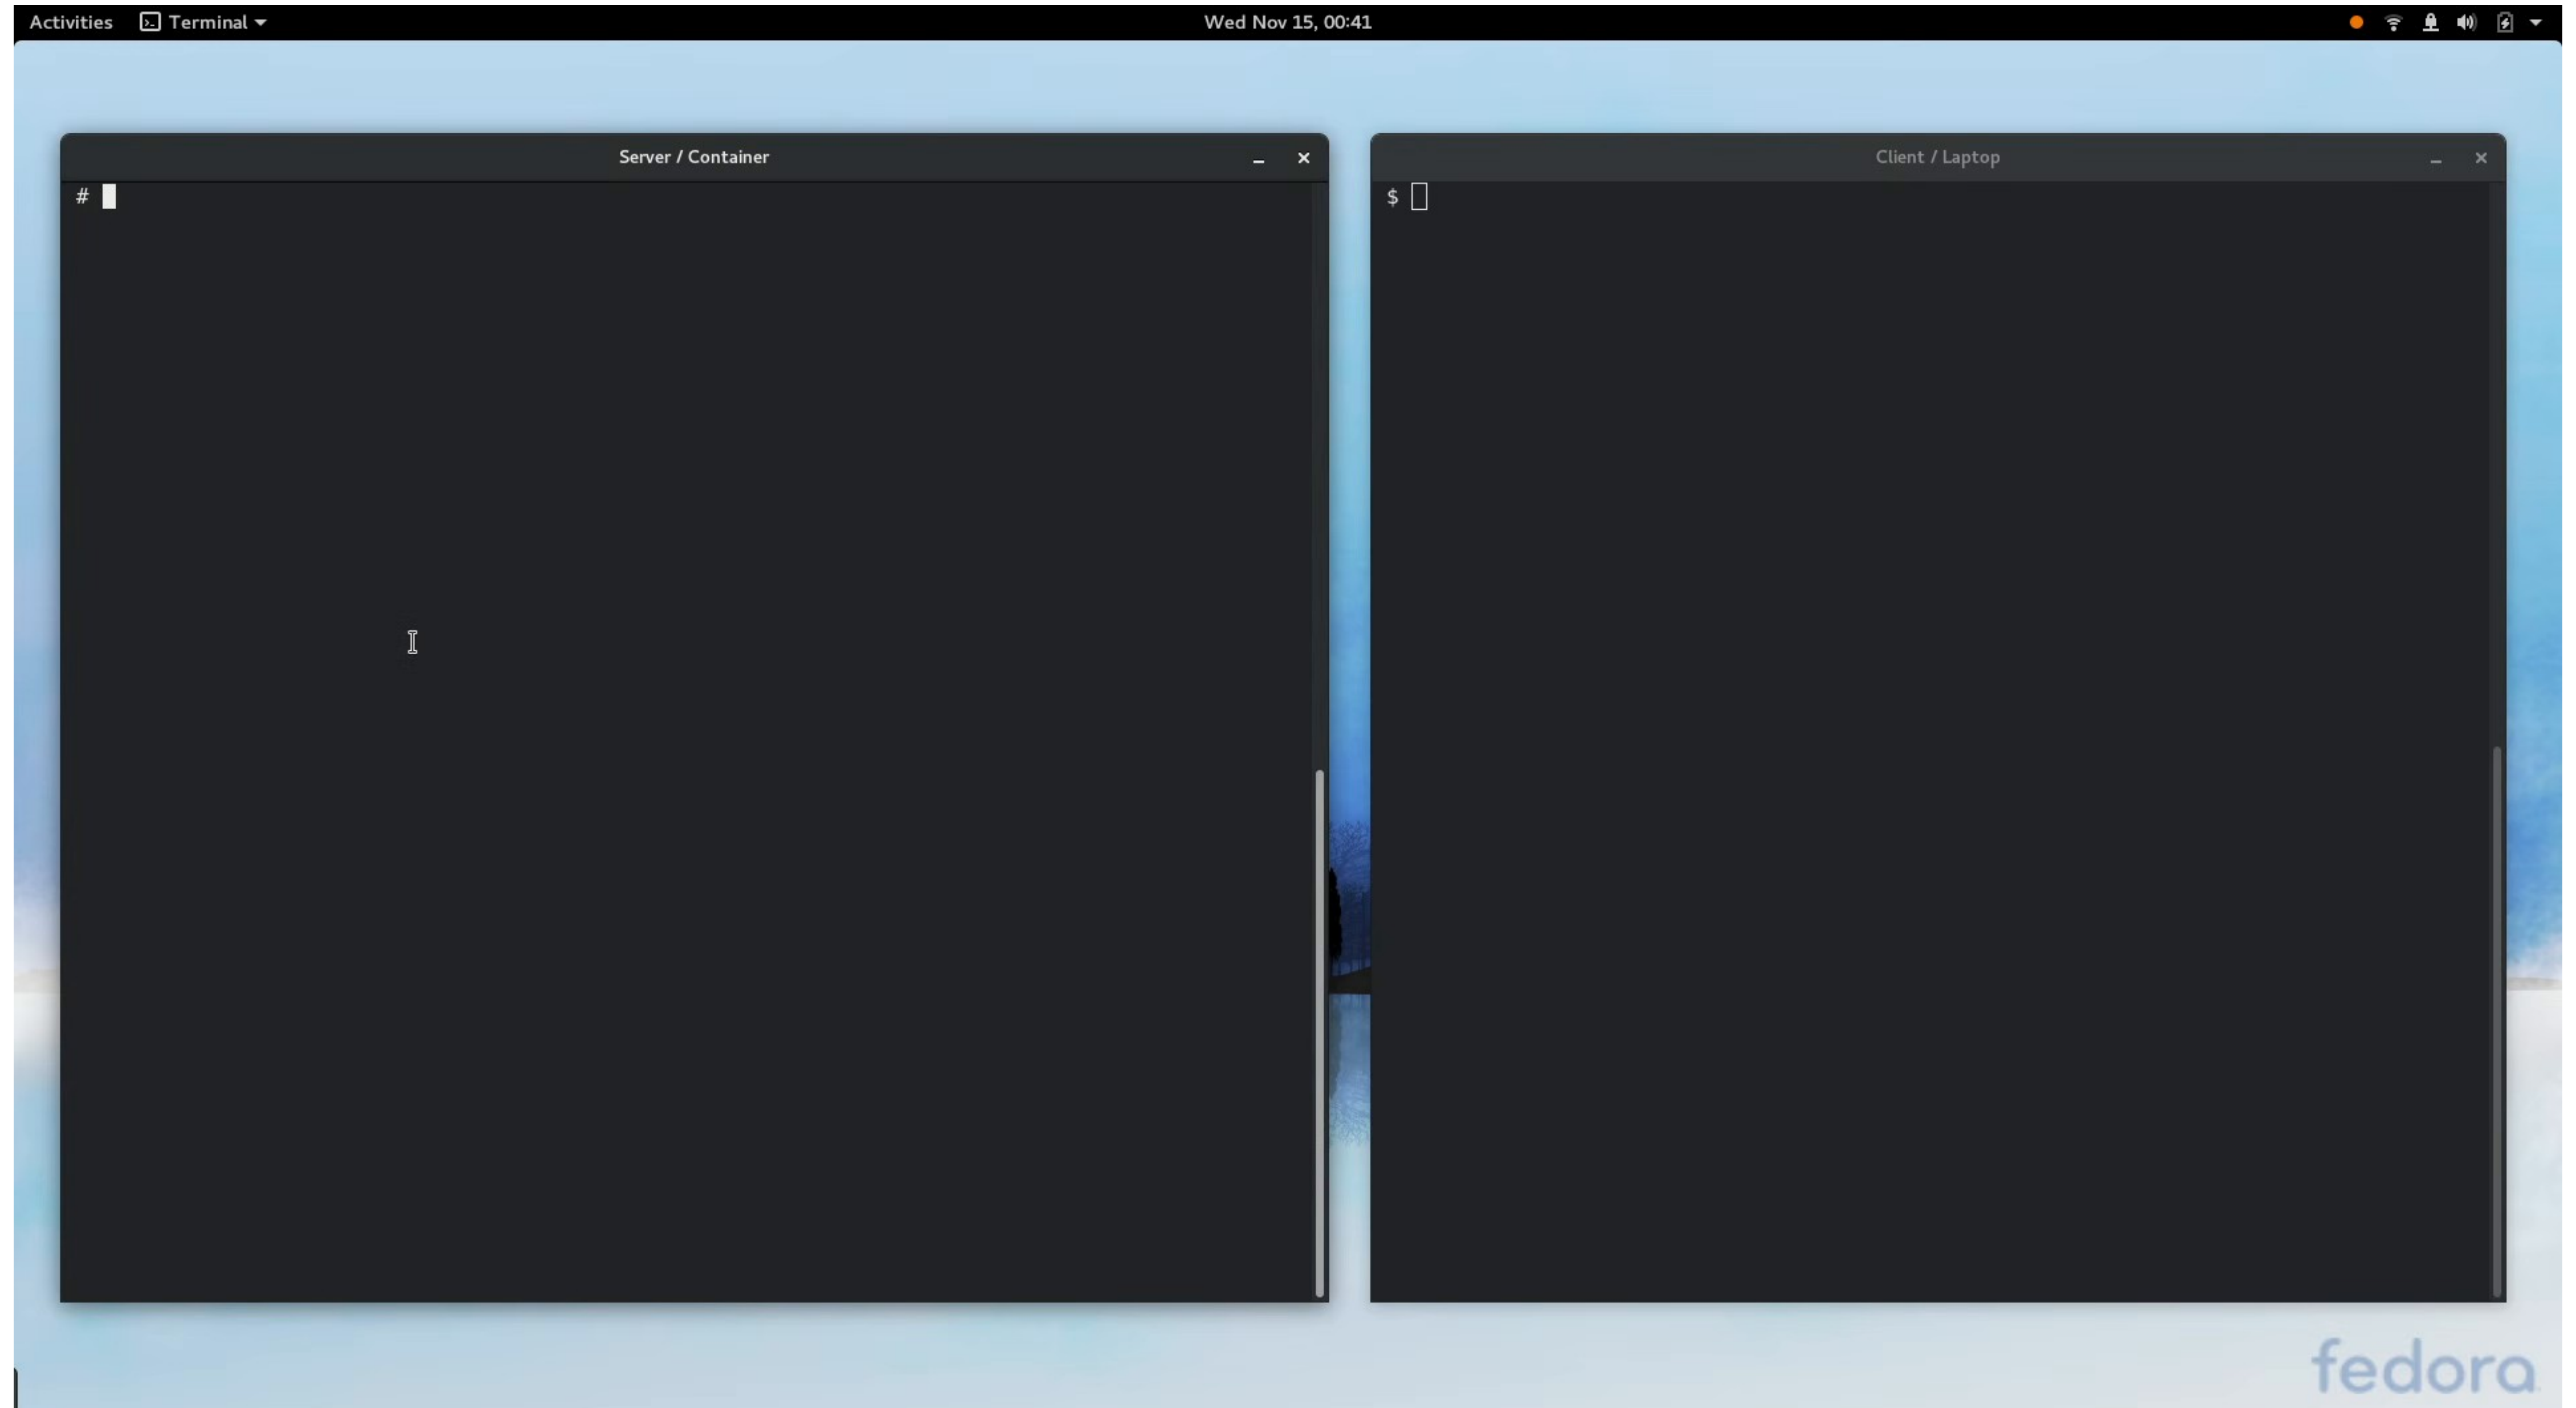
Example Videos



Cuda, TensorFlow, Keras running mnist test



R, R-studio, CUDA, rpud lib, display via vnc



Questions?

