


Leveraging oneAPI Containers to Deploy OpenVINO™ Notebooks for Efficient AI Research

oneAPI DevSummit for AI and HPC – December 2023

Presented by: Peter Darveau P. Eng. - Hexagon Technology Inc., Toronto Canada


Agenda

- Brief Introduction
- Brief Overview of oneAPI and OpenVINO™
- What are Containers ?
- Benefits of using containerized Jupyter notebooks
- Basic Configuration Example
- Live demo deploying an OpenVINO™ notebook container
- How containerized notebooks enhance efficiency
- Q&A



Brief Intro

Canada's Advanced Research Computing Platform



- National Host Sites
- Support Sites



oneAPI and OpenVINO™ Intro

Intel® AI Analytics Toolkit

Deep Learning

TensorFlow Optimizations from Intel®

PyTorch Optimizations from Intel®

Intel® Neural Compressor

Model Zoo for Intel® Architecture

Machine Learning

Intel® Extension for Scikit-learn

Intel® Optimization for XGBoost

Data Analytics

Intel® Distribution of Modin

OmniSci Backend

Intel-optimized Python


NumPy

SciPy


Pandas

Numba

Python




CPU

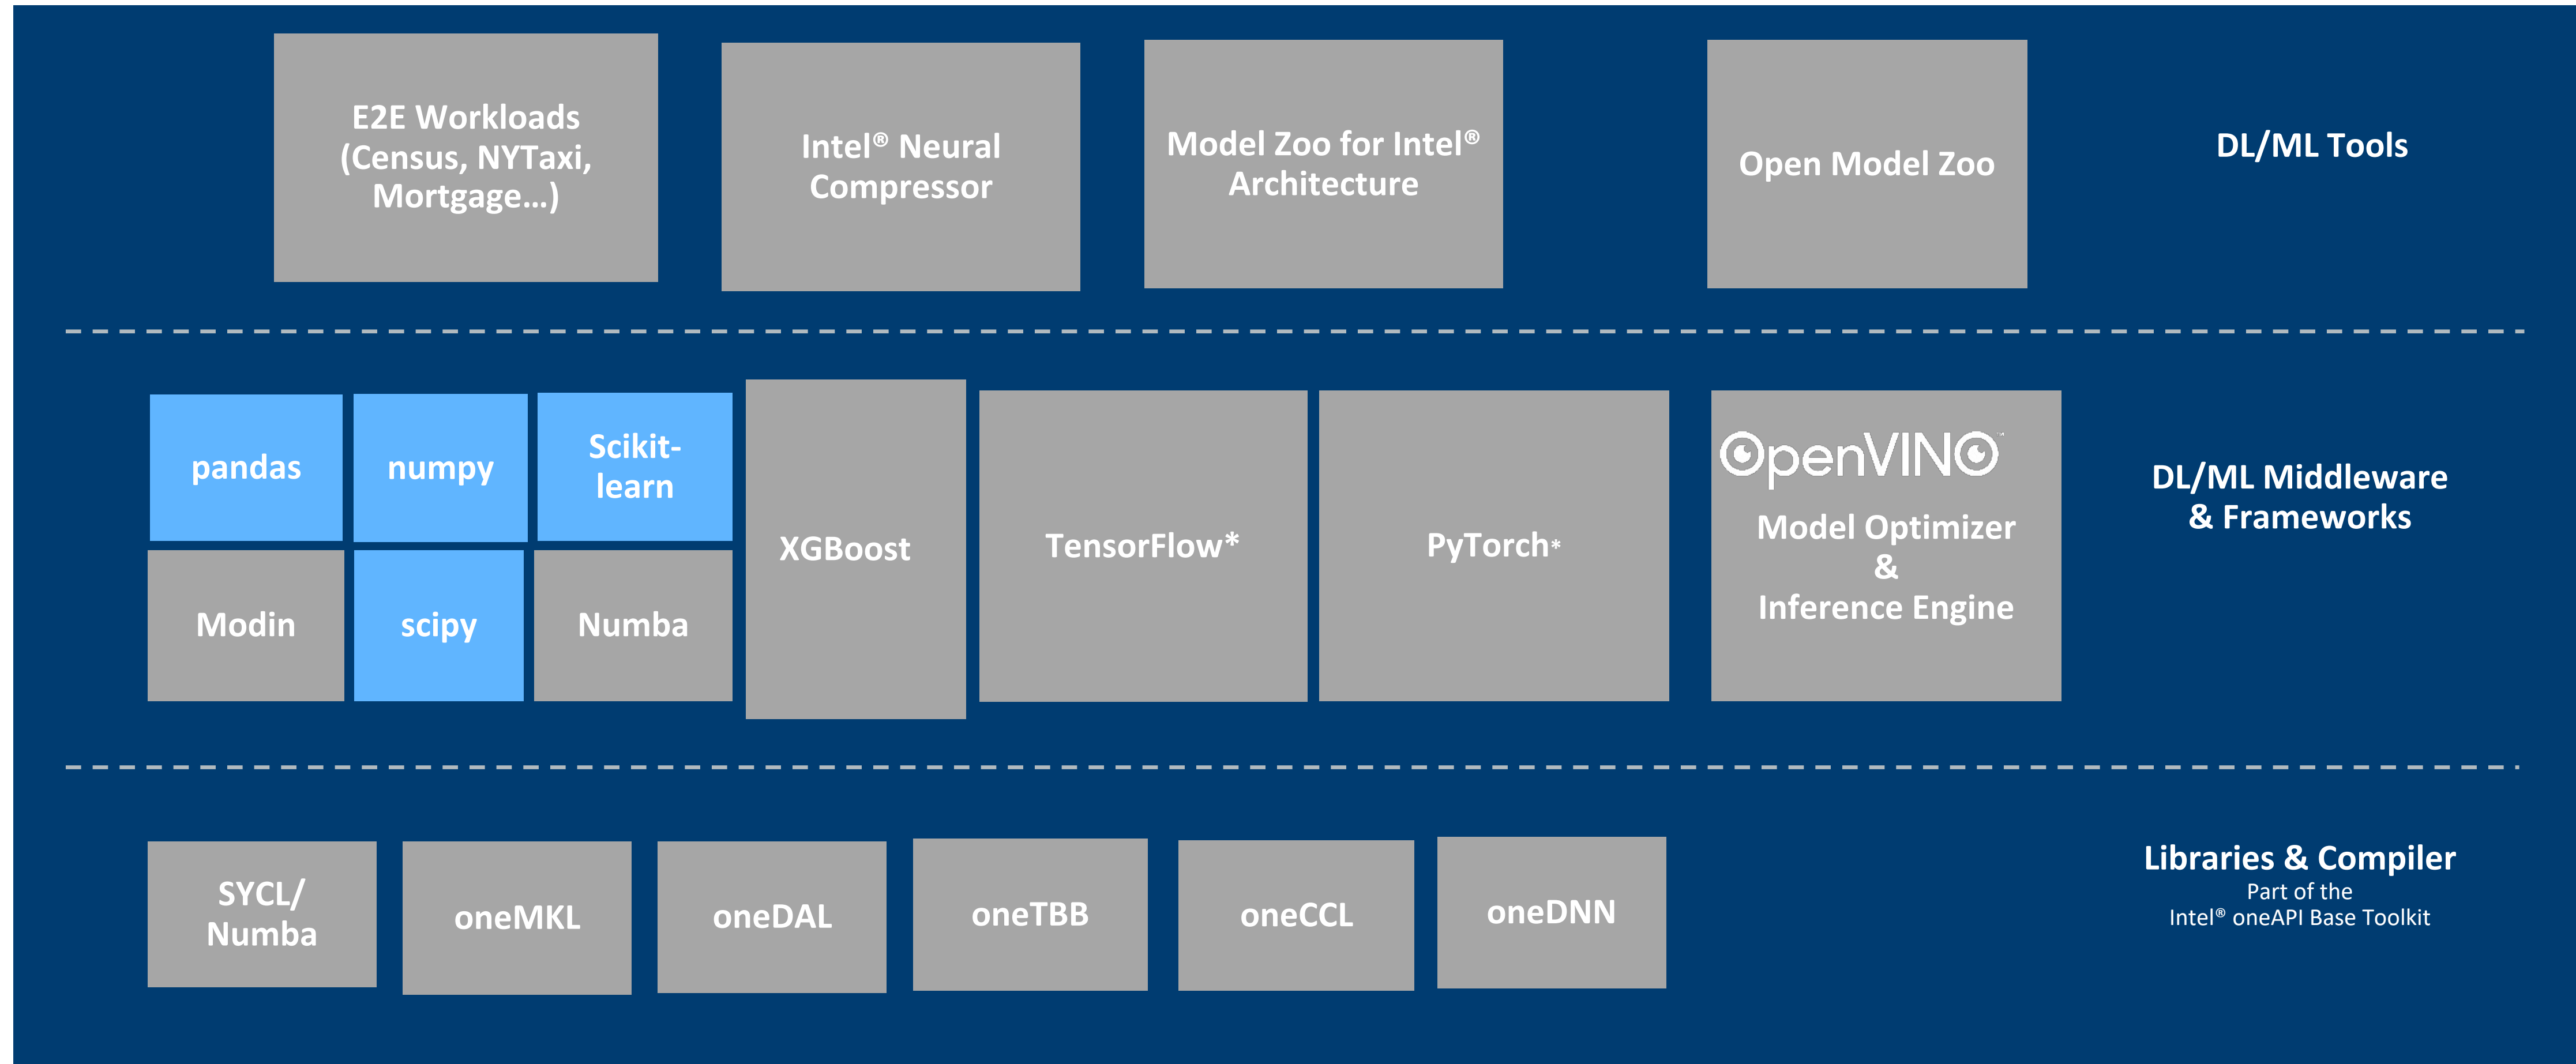


GPU

Hardware support varies by individual tool. Architecture support will be expanded over time.



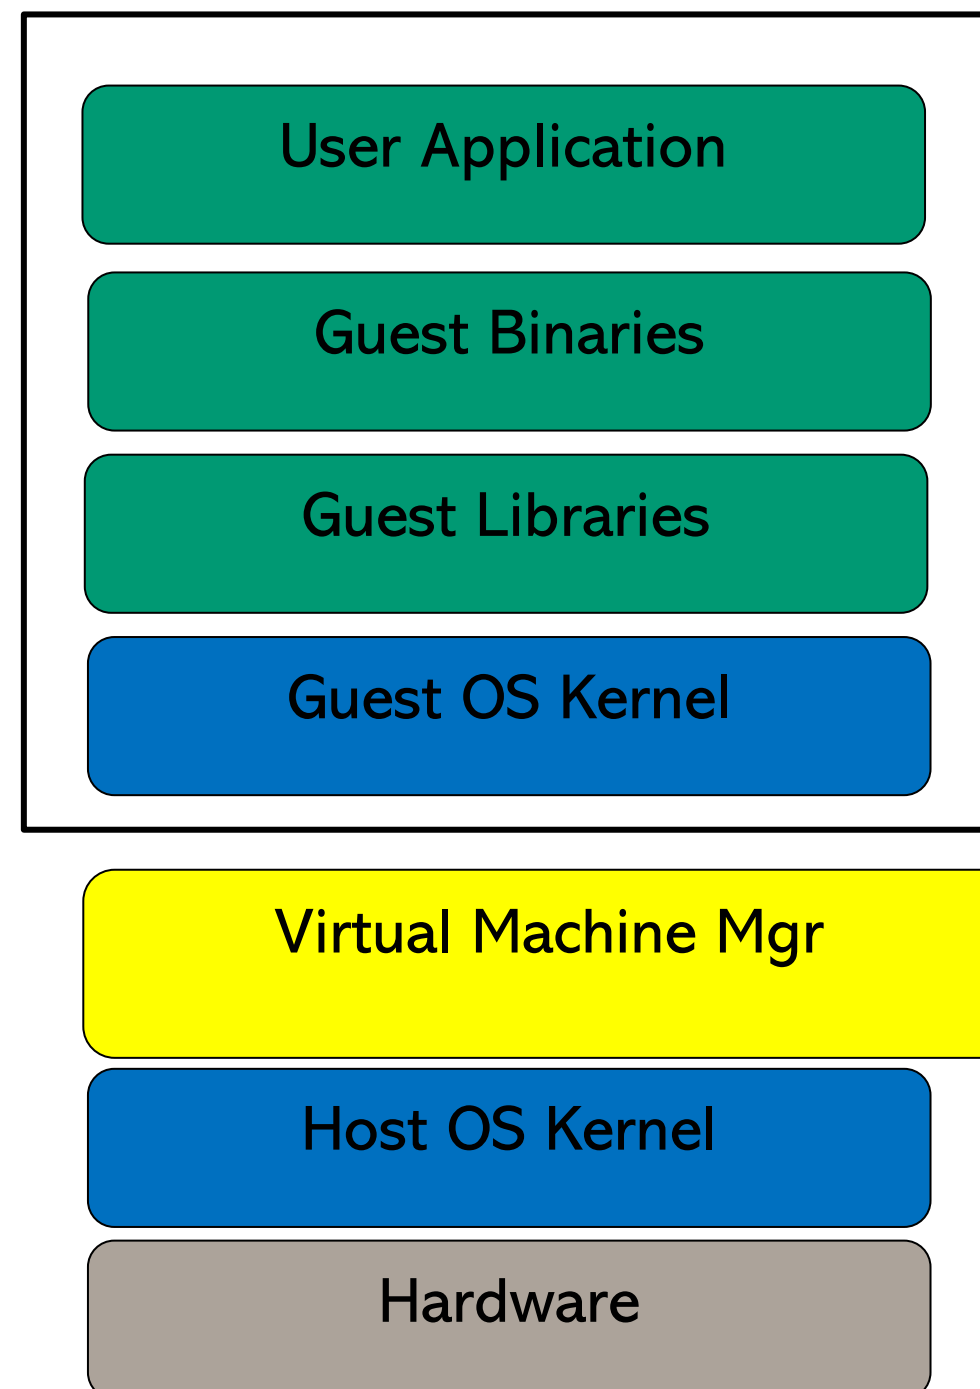
oneAPI and OpenVINO™ Intro



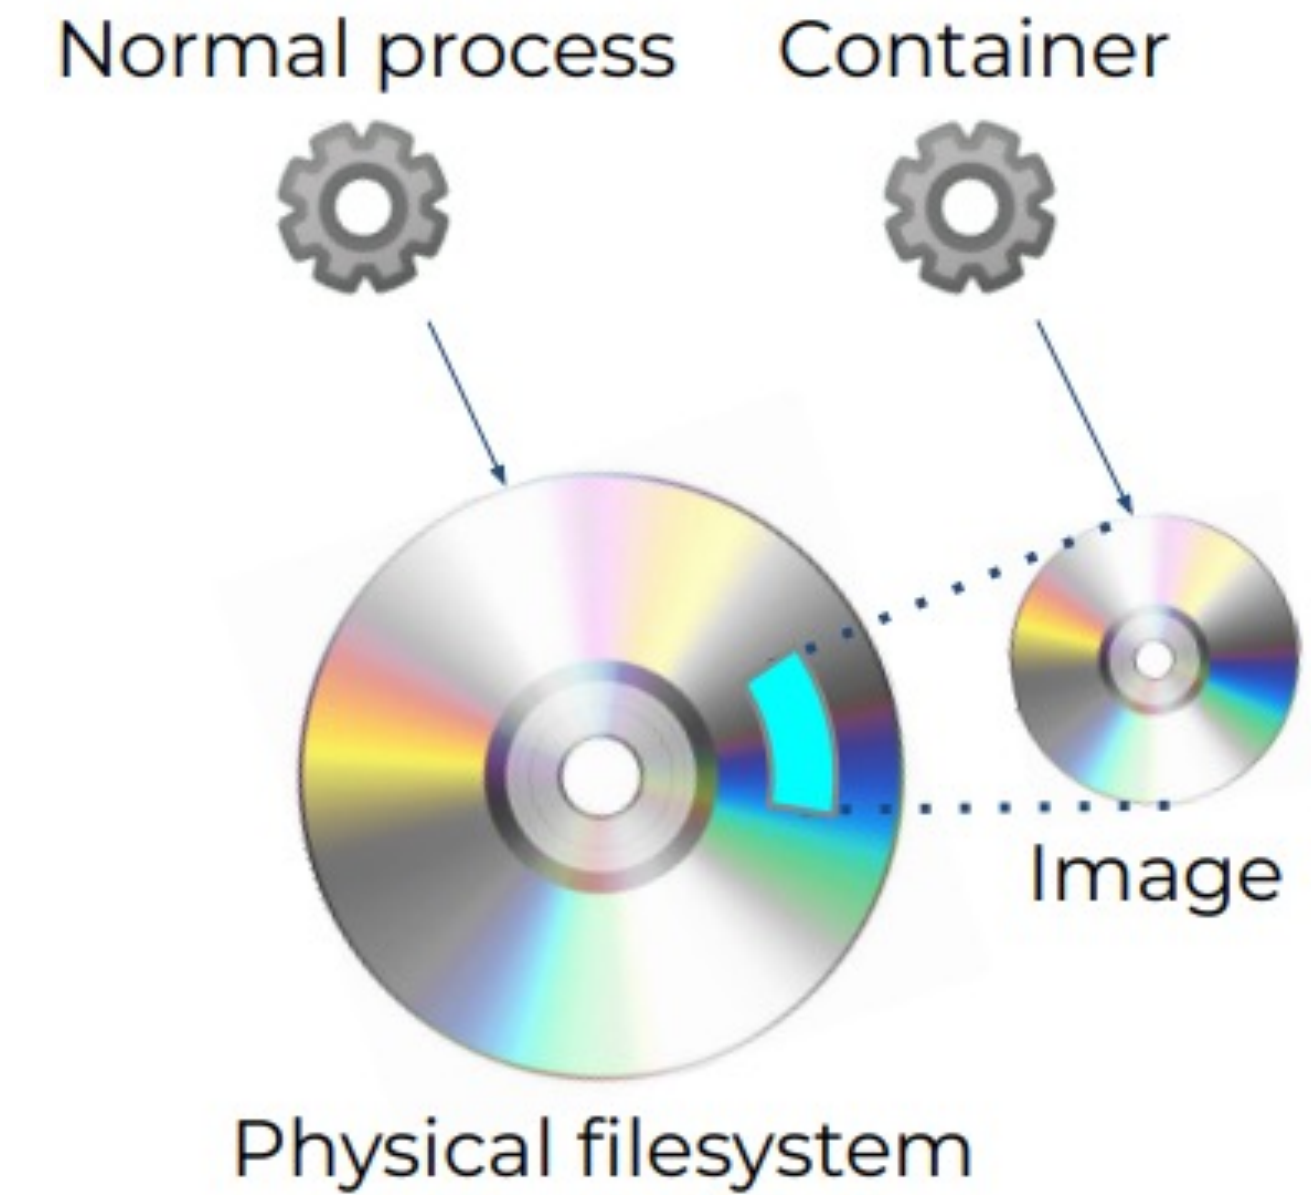
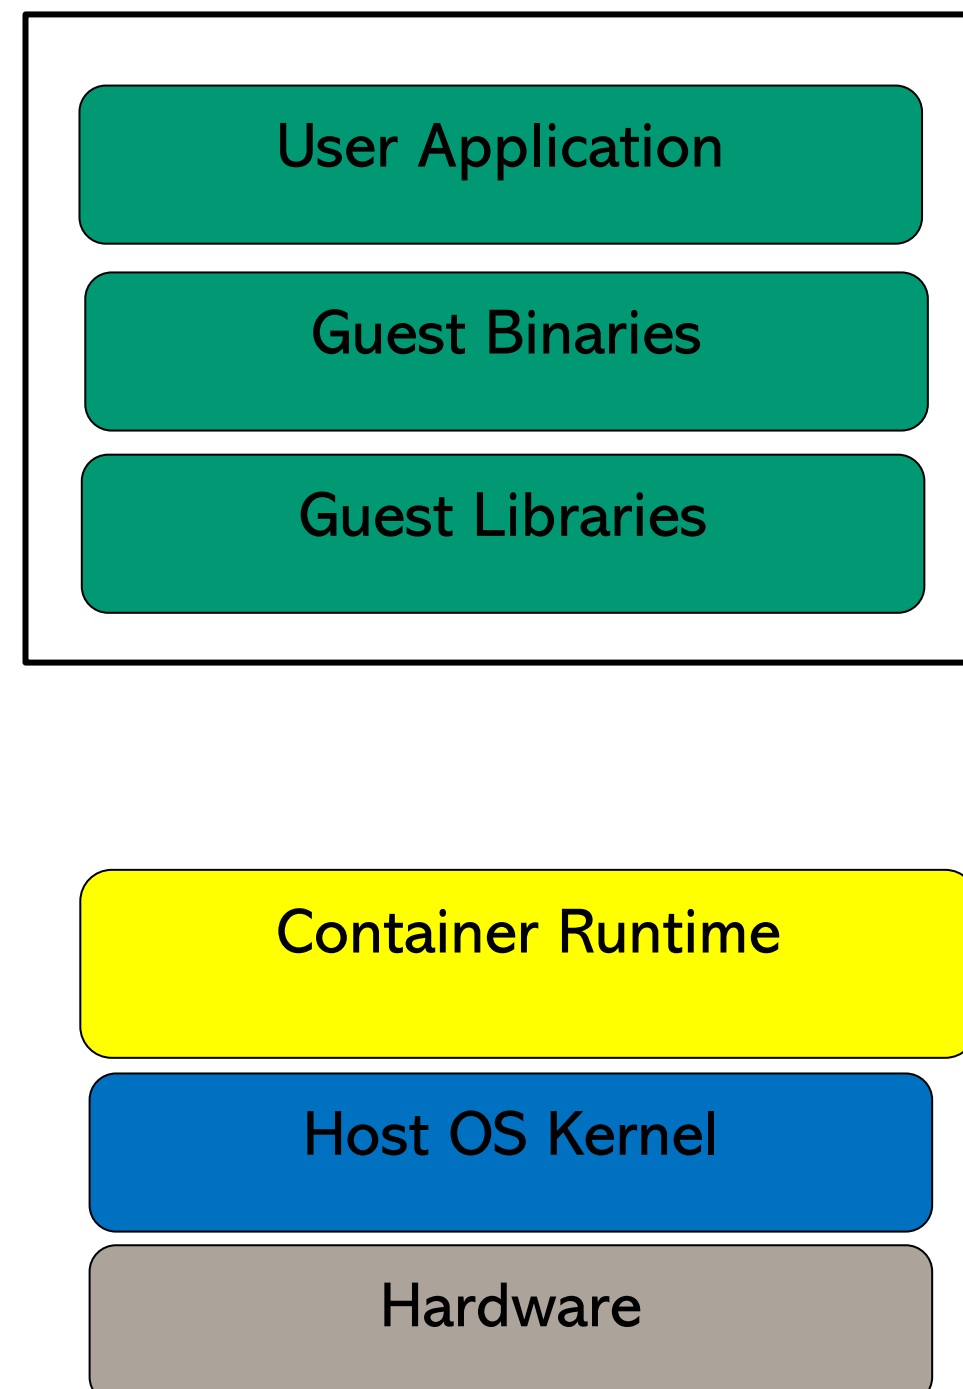
What are Containers?

- A container is a process that has its own view of local resources.
- In the picture, the container sees the image of the physical filesystem

Virtual Machine



Container



Benefits of Containerized Notebooks

- Shareability
 - Share to a public repo
 - Use images shared by others
- Portability
 - Package dependencies and environments
- Reproducibility
 - Largely unaffected by changes to the cluster environments
- Efficiency
 - Optimized for performance out of the box
- Scalability
 - Easily launch on different compute resources
- Security
 - Isolated environments to secure development

Basic Configuration and Build Example

```
%post -c /bin/bash
  export TZ=America/Toronto
  export DEBIAN_FRONTEND=noninteractive
  #Installing all dependencies
  apt-get update && apt-get -y upgrade
  apt-get -y install python3-venv build-essential python3-dev intel-opencl-icd git-all libgl1

#Clone Repo
python3 -m venv openvino_env
source openvino_env/bin/activate
git clone --depth=1 https://github.com/openvinotoolkit/openvino_notebooks.git
cd openvino_notebooks

#Install packages
python -m pip install --upgrade --no-input pip
pip install --no-input wheel setuptools
pip install tensorflow>=2.5 --no-cache-dir
grep -v "tensorflow" requirements.txt > tmpfile && mv tmpfile requirements.txt
pip install --no-input -r requirements.txt

%startscript

%runscript
#!/bin/bash
source /openvino_env/bin/activate
jupyter notebook --allow-root $@
```

Per Intel's recommended set up



%runscript allows you to tunnel into container



Live Demo – run Jupyter

```
[I 14:08:32.499 NotebookApp] Serving notebooks from local directory: /root
[I 14:08:32.499 NotebookApp] Jupyter Notebook 6.5.4 is running at:
[I 14:08:32.499 NotebookApp] http://etasc-dev:8888/?token=43b5d92d63f966b188f073ff818d1cb9e8a54d703edead08
[I 14:08:32.499 NotebookApp] or http://127.0.0.1:8888/?token=43b5d92d63f966b188f073ff818d1cb9e8a54d703edead08
[I 14:08:32.499 NotebookApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
[W 14:08:32.675 NotebookApp] No web browser found: could not locate runnable browser.
[C 14:08:32.675 NotebookApp]

To access the notebook, open this file in a browser:
    file:///root/.local/share/jupyter/runtime/nbserver-516238-open.html
Or copy and paste one of these URLs:
    http://etasc-dev:8888/?token=43b5d92d63f966b188f073ff818d1cb9e8a54d703edead08
    or http://127.0.0.1:8888/?token=43b5d92d63f966b188f073ff818d1cb9e8a54d703edead08
```


Live Demo – Tunneling into container

The screenshot shows a web browser window with the address bar displaying `localhost:8888/tree/openvino_notebooks/notebooks`. The page header features the Jupyter logo and the text "jupyter", along with "Quit" and "Logout" buttons. Below the header, there are tabs for "Files", "Running", and "Clusters". A message says "Select items to perform actions on them." with "Upload", "New", and refresh icons to the right. The main content is a file browser for the `/ openvino_notebooks / notebooks` directory. It shows a list of folders with columns for Name, Last Modified, and File size.

<input type="checkbox"/> 0	Name	Last Modified	File size
<input type="checkbox"/>	..	seconds ago	
<input type="checkbox"/>	001-hello-world	7 months ago	
<input type="checkbox"/>	002-openvino-api	7 months ago	
<input type="checkbox"/>	003-hello-segmentation	7 months ago	
<input type="checkbox"/>	004-hello-detection	7 months ago	
<input type="checkbox"/>	101-tensorflow-to-openvino	5 months ago	
<input type="checkbox"/>	102-pytorch-onnx-to-openvino	5 months ago	
<input type="checkbox"/>	103-paddle-to-openvino	7 months ago	
<input type="checkbox"/>	104-model-tools	7 months ago	
<input type="checkbox"/>	105-language-quantization	7 months ago	


Enhanced Efficiency ... What we observed

- Faster experimentation by eliminating environment setup
- Cross-platform deployability
- Leverage Intel Optimization out of the box
- More efficient workflows for researchers



Key Takeaways

- Native and robust support for x86 architecture (CPU) only
- Performance improvements from OpenVINO™ were also observed
- Intel GPU and enclaves (SGX) support are not native to any known Apptainer images and require custom install - To be tested



Q&A

Contact:

E-mail for Peter @ University of Ottawa Research IT
pdarveau@uottawa.ca

[My LinkedIn Profile](#)

Read my ML model papers at uOttawa Research
[Recherche uO Research: Peter Darveau](#)



Peter Darveau P. Eng., CED, 6sigma
GB, MBA

| RESEARCH IT & ENGINEERING
PROFESSIONAL | MACHINE LEARNING | INT...

