

Xilinx ML Suite Overview

Jupyter Notebooks walkthrough

Jon Beckwith
SAE, Xilinx



Dad joke time

> What kind of snake does a math teacher own?

> A pi-thon



Before we begin, FPGAs in the cloud



Cloud Based Acceleration



Reconfigurable Acceleration in the Cloud

Faced with exponential growth in computing requirements and the inability for CPU technology to keep pace, cloud and data center architectures are moving toward accelerated computing. Accelerators compliment CPU-based architectures and deliver both performance and power efficiency.

FPGAs can deliver 10x acceleration across a broad set of applications and are reconfigurable to provide an ideal fit for the changing workloads of the modern data center.

With acceleration capabilities a full generation ahead of any other FPGA, Xilinx UltraScale™ and UltraScale+ FPGAs are empowering hardware and application developers in many of the world's largest and most innovative cloud computing services.

> <https://www.xilinx.com/products/design-tools/cloud-based-acceleration.html>

Accelerated Cloud Service Partners

Click to learn more


Alibaba Cloud aws Baidu

HUAWEI NIMBIX Tencent Cloud



Introducing Amazon EC2 F1



- > Amazon EC2 F1 is a compute instance with Xilinx FPGAs which can be programmed to create custom hardware accelerated applications
- > F1 instances are easy to program and come with everything needed to develop, simulate, debug, and compile hardware accelerators
- > Once a FPGA design is complete, it can be registered as an Amazon FPGA Image (AFI), and deployed to F1 instance in just a few clicks



Host CPU



Xilinx FPGA

Compelling Applications for Cloud FPGA Acceleration



Genomics



Big Data Analytics



Financial Analytics



Security



Image and Video Processing



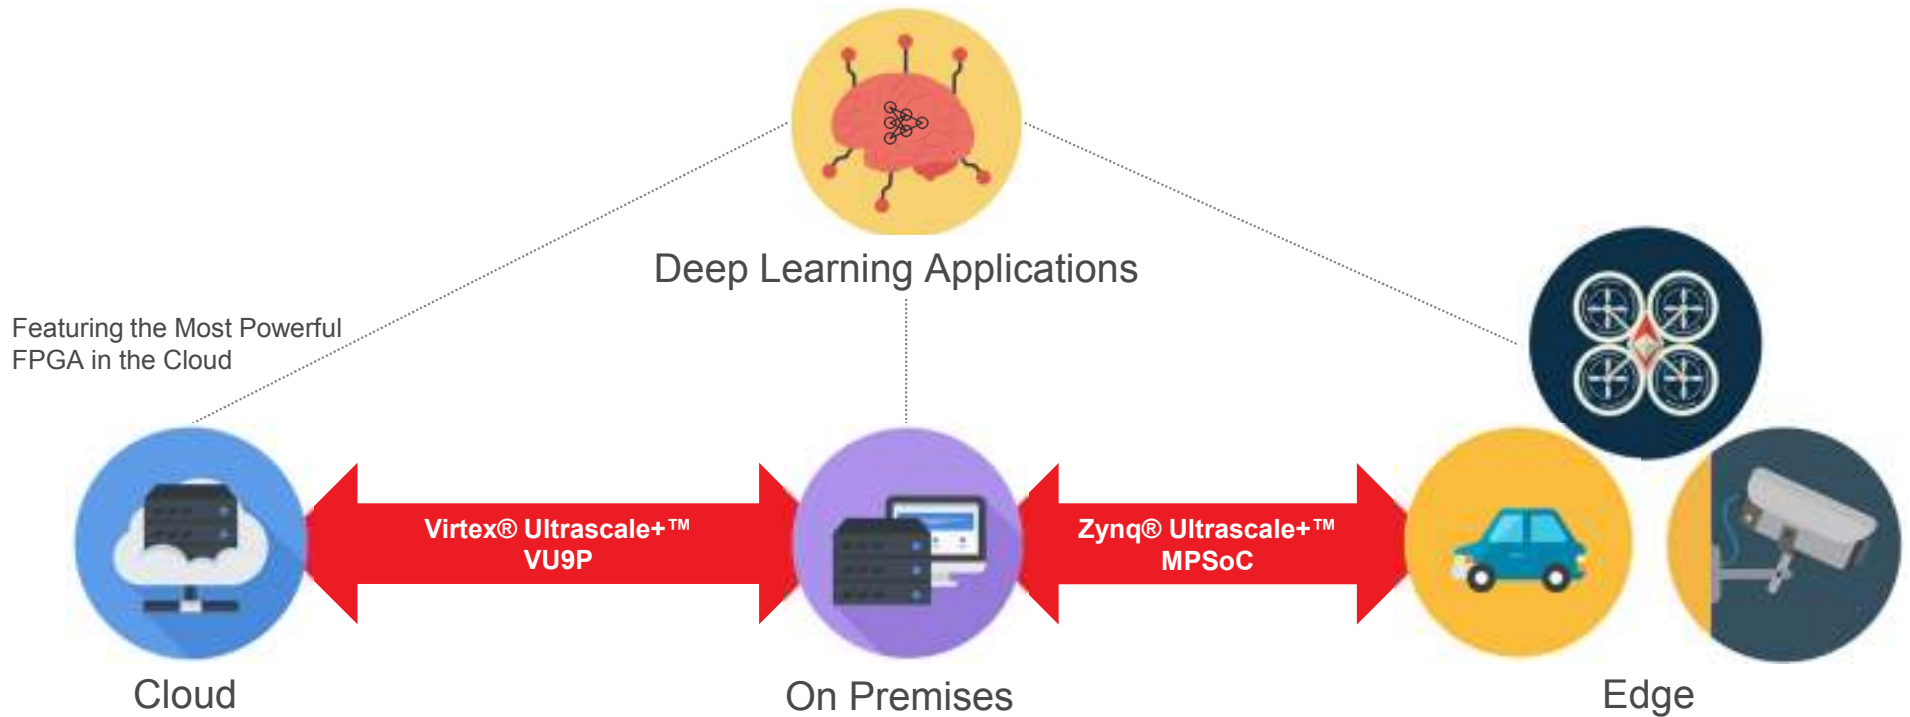
Machine Learning



ML Suite



Accelerating AI Inference into Your Cloud Applications



© Copyright 2018 Xilinx



Want to try the out Xilinx ML Suite ?

➤ <https://github.com/Xilinx/ml-suite>

The image shows a composite view of the Xilinx ML Suite. On the left is a terminal window with a dark blue header containing the name 'Jon Beckwith' and email '<jon.beckwith@xilinx.com>'. The terminal title is 'Current'. It displays the Xilinx logo and 'Xilinx ML Suite(502722)' with icons for copy, paste, and power. Below this, it lists: Command: Notebook; Status: Processing; Address: NAE-165-254-189-4.jarvice.com; Password: (click to show). On the right is a screenshot of the AWS Marketplace page for 'Xilinx ML Suite'. The page header includes 'aws marketplace' and 'XILINX'. The product title is 'Xilinx ML Suite' with a 'Continue to Subscribe' button and a 'Save to List' button. A price tag indicates 'Typical Total Price: \$1,850/hr'. The page has tabs for Overview, Pricing, Usage, Support, and Reviews. The 'Overview' section is active, showing a 'Product Overview' with text about integrating machine learning into applications and a 'Highlights' box listing supported frameworks (Caffe, MXNet, TensorFlow) and pre-trained models (GoogLeNet v1, ResNet50, Flower100, Places365, Vggv2). A link for more information is provided at the bottom.

Xilinx ML Suite - AWS Marketplace



> ML Suite

>> Supported Frameworks:

- Caffe
- MxNet
- Tensorflow
- Keras
- Python Support
- Darknet

>> Jupyter Notebooks available:

- Image Classification with Caffe
- Using the xFDNN Compiler w/ a Caffe Model
- Using the xFDNN Quantizer w/ a Caffe Model

>> Pre-trained Models

- Caffe 8/16-bit
 - GoogLeNet v1
 - ResNet50
 - Flowers102
 - Places365
- Python 8/16-bit
 - Yolov2
- MxNet 8/16-bit
 - GoogLeNet v1

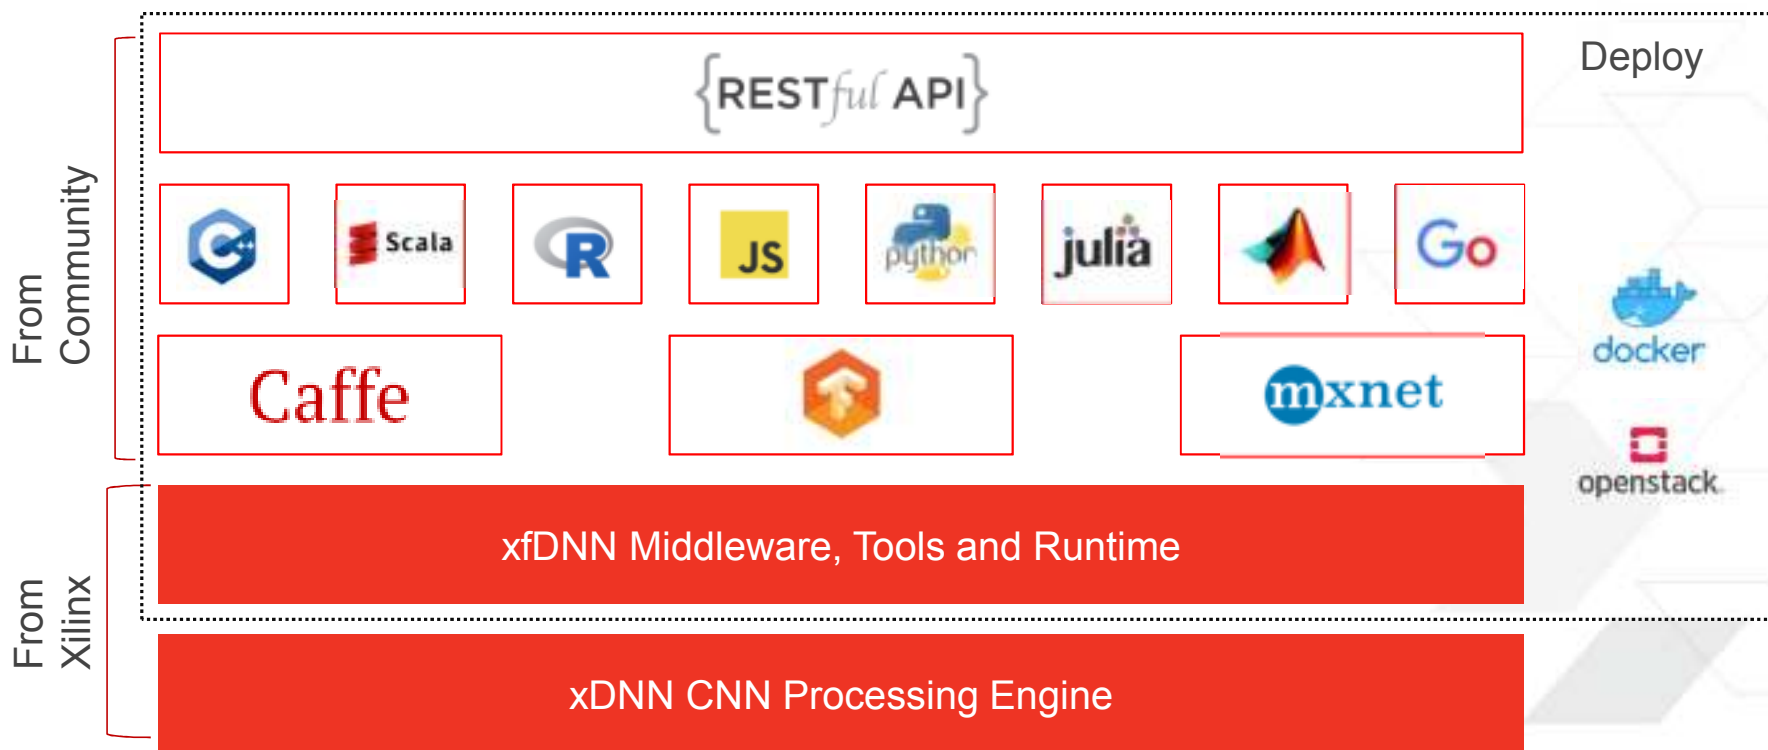
>> xFDNN Tools

- Compiler
- Quantizer

The screenshot shows the AWS Marketplace listing for Xilinx ML Suite. The header includes the AWS Marketplace logo and navigation options. The main content area features the Xilinx logo and the product name 'Xilinx ML Suite'. A 'Continue to Subscribe' button is prominent, along with a 'Save to List' button. The price is listed as '\$1.850/hr'. Below the pricing, there is a 'Product Overview' section with text describing the suite's capabilities and a 'Highlights' box listing supported frameworks and pre-trained models.

https://aws.amazon.com/marketplace/pp/B077FM2JNS?gid=1544477354556&sr=0-2&ref_=srh_res_product_title

Seamless Deployment with Open Source Software

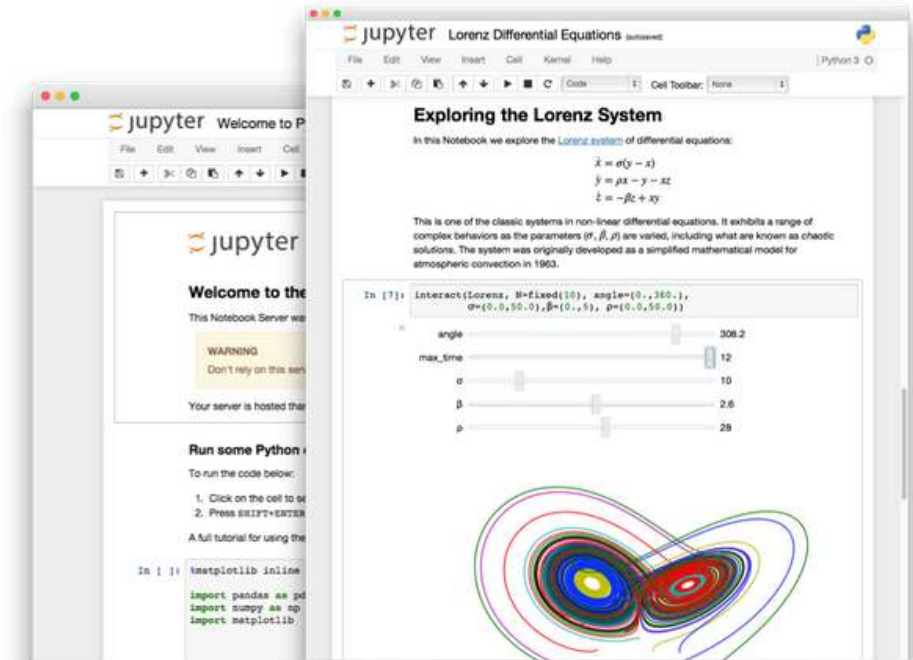


*TensorFlow Q4 2017

Jupyter Notebooks Walkthrough



Jupyter Notebooks



The Jupyter Notebook

The Jupyter Notebook is an open-source web application that allows you to create and share documents that contain live code, equations, visualizations and narrative text. Uses include: data cleaning and transformation, numerical simulation, statistical modeling, data visualization, machine learning, and much more.

Try it in your browser

Install the Notebook

➤ ML Suite AMIs launch Jupyter Notebooks at boot


➤ http://<aws_ip_addr>:8888



Demo

xfDNN Inference Toolbox

Graph Compiler




- Python tools to quickly compile networks from common Frameworks – Caffe, MxNet and Tensorflow

Network Optimization



- Automatic network optimizations for lower latency by fusing layers and buffering on-chip memory

xfDNN Quantizer



- Quickly reduce precision of trained models for deployment
- Maintains 32bit accuracy at 8 bit within 2%

Adaptable.
Intelligent.

 XILINX.

© Copyright 2018 Xilinx