NVIDIA

JETSON PLATFORM OVERVIEW

JETSON SUCCESS STORIES

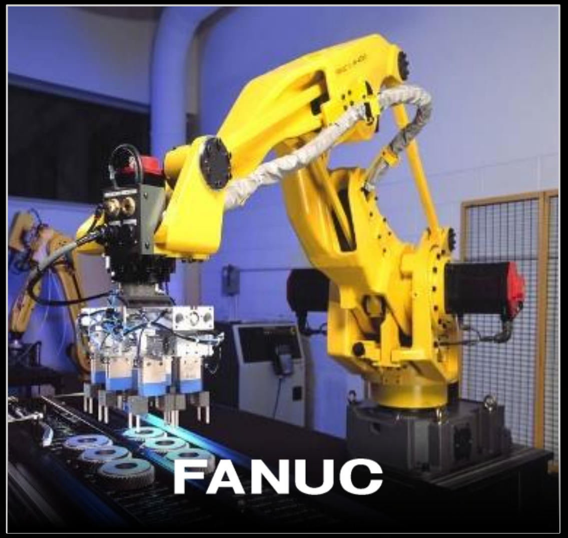

Industrial

Aerospace/Defense

Construction

Agriculture

Healthcare

Smart City

Retail

Logistics

Delivery

Inspection

Service

Collaboration

JETSON PRODUCT FAMILY OVERVIEW

THE JETSON FAMILY

From AI at the Edge to Autonomous Machines

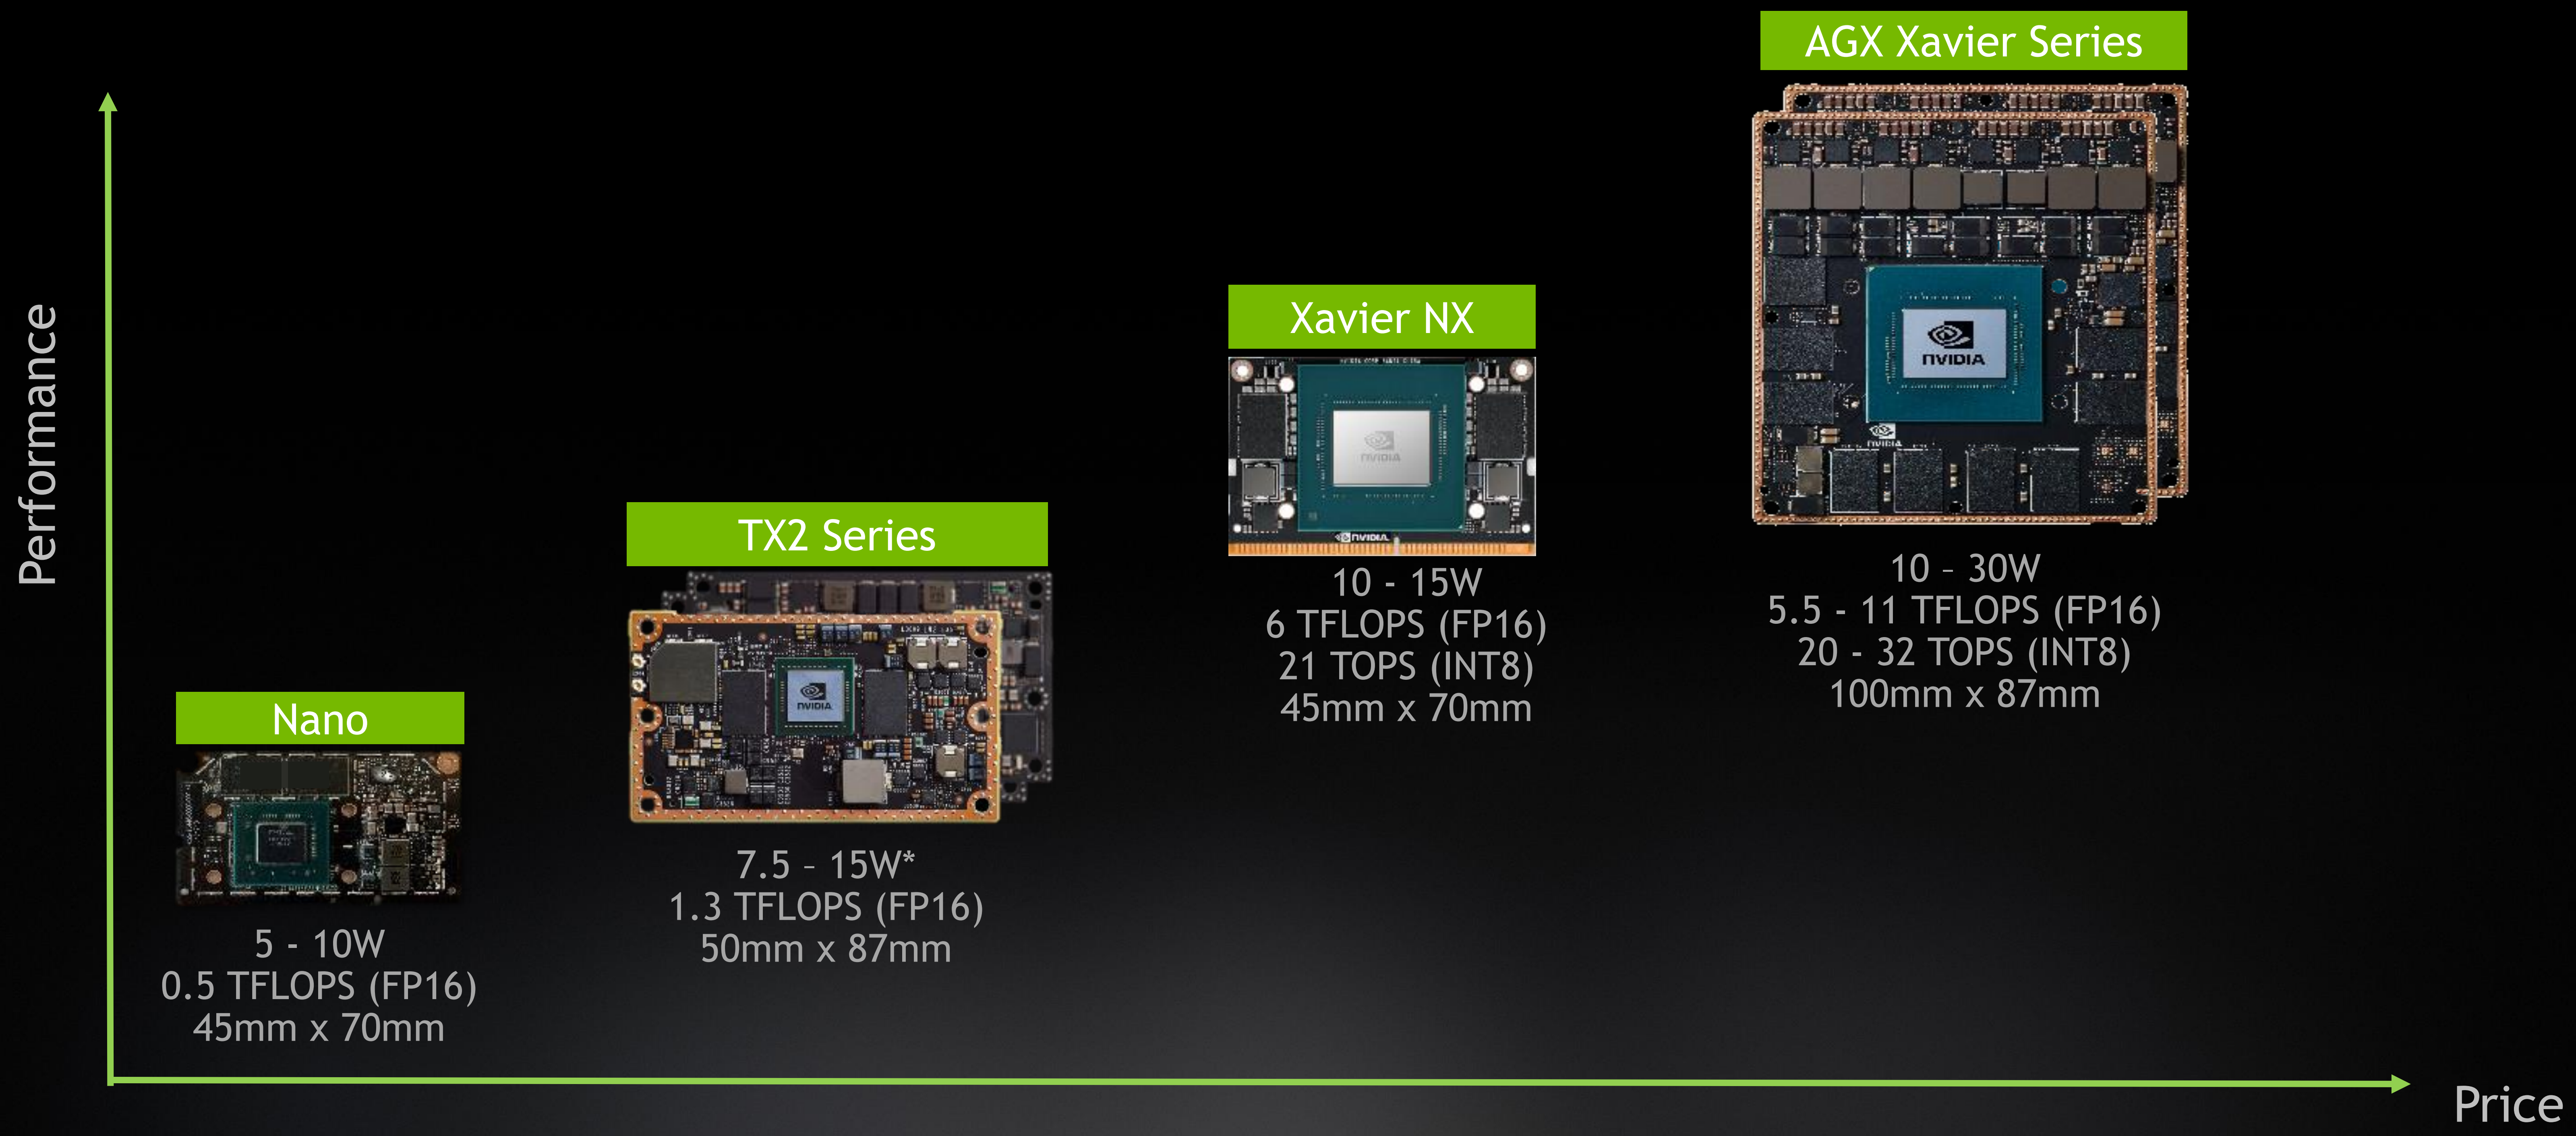

JETSON DEVELOPER KITS

For Developers, Engineers and Makers

JETSON NANO
5W | 10W
0.5 TFLOPS (FP16)
$99

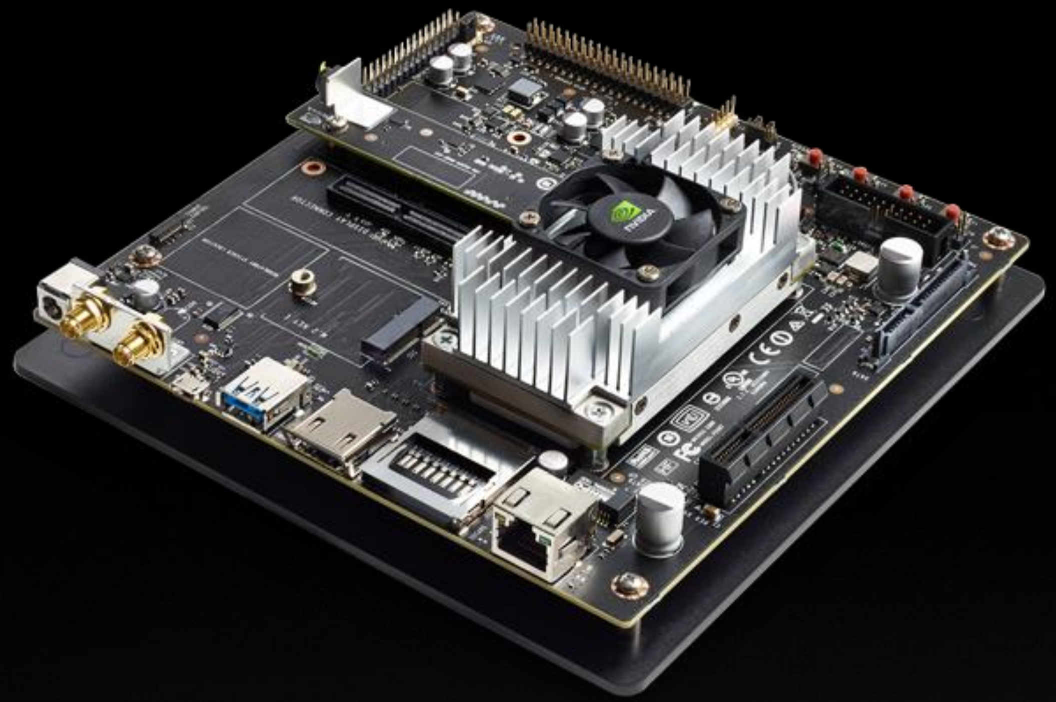

JETSON TX2
7.5W | 15W
1.3 TFLOPS (FP16)
$399 ($299 EDU)

JETSON XAVIER NX
10W | 15W
7 TFLOPS (FP16) | 21 TOPS (INT8)
$399

JETSON AGX XAVIER
10 | 15W | 30W
11 TFLOPS (FP16) | 32 TOPS (INT8)
$699

Multiple developer kits - Same software

Full specs at developer.nvidia.com/jetson

JETSON MODULES — COMMERCIAL ROADMAP

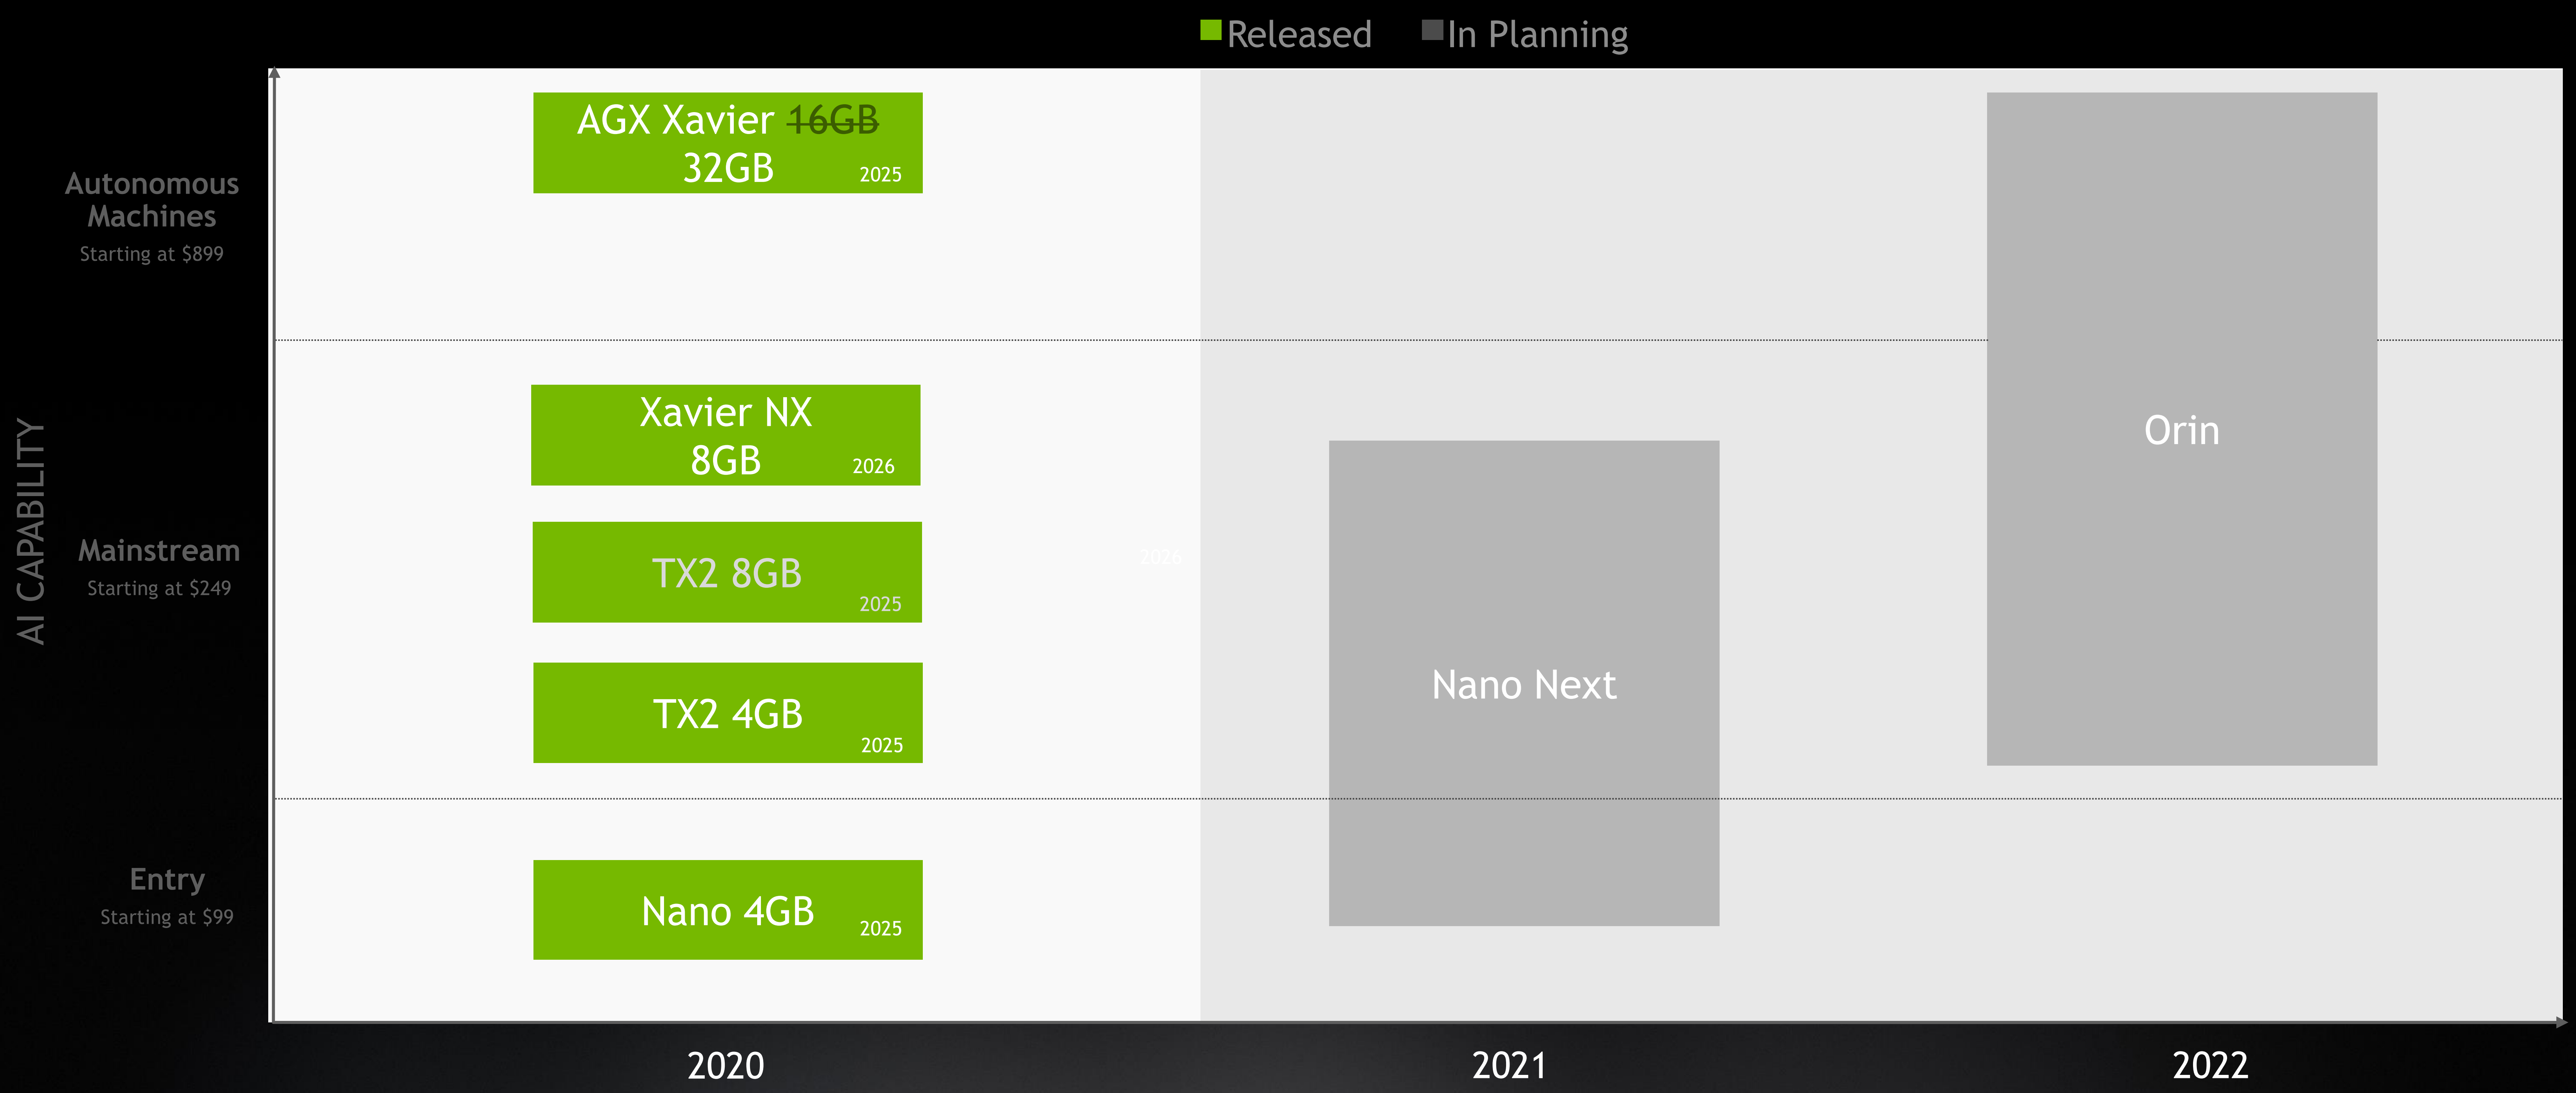

** TX2 8GB will ship till 2025 but new designs recommend transition to NX.*
The year in each box indicates supply availability at least / until. Lifecycle might be extended further, contact your partner sales manager.
Operating life of Jetson is 5 years 24x7 for commercial production modules. Products in development and planning are subject to change.

NVIDIA.

JETSON MODULES — INDUSTRIAL ROADMAP

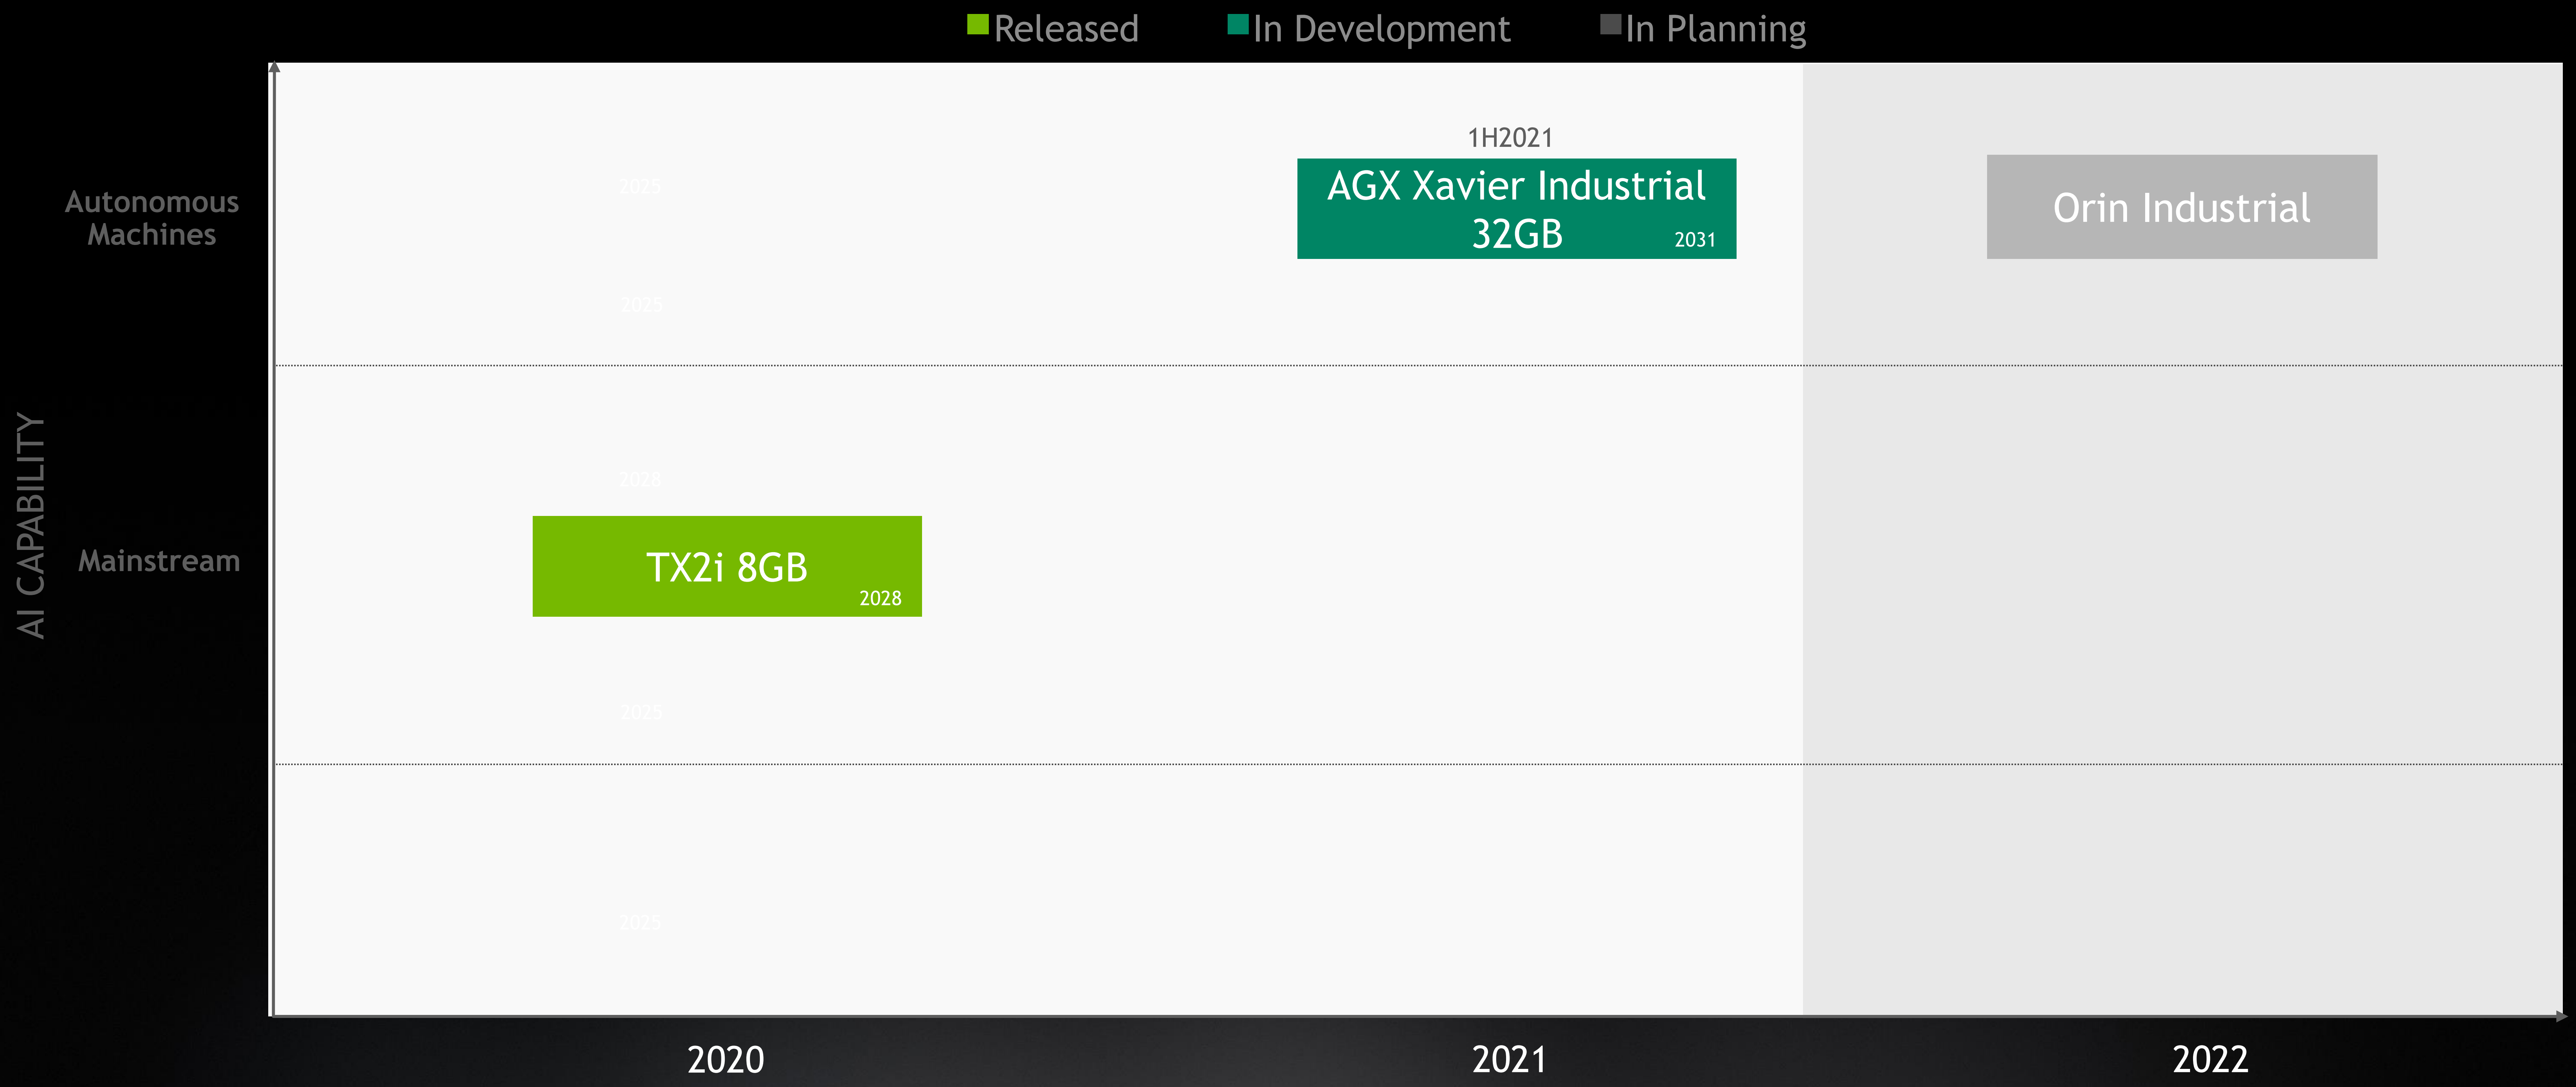

The year in each box indicates supply availability at least / until. Lifecycle might be extended further, contact your partner sales manager.
Operating life of Jetson industrial version is 10 years 24x7.
Products in development and planning are subject to change.

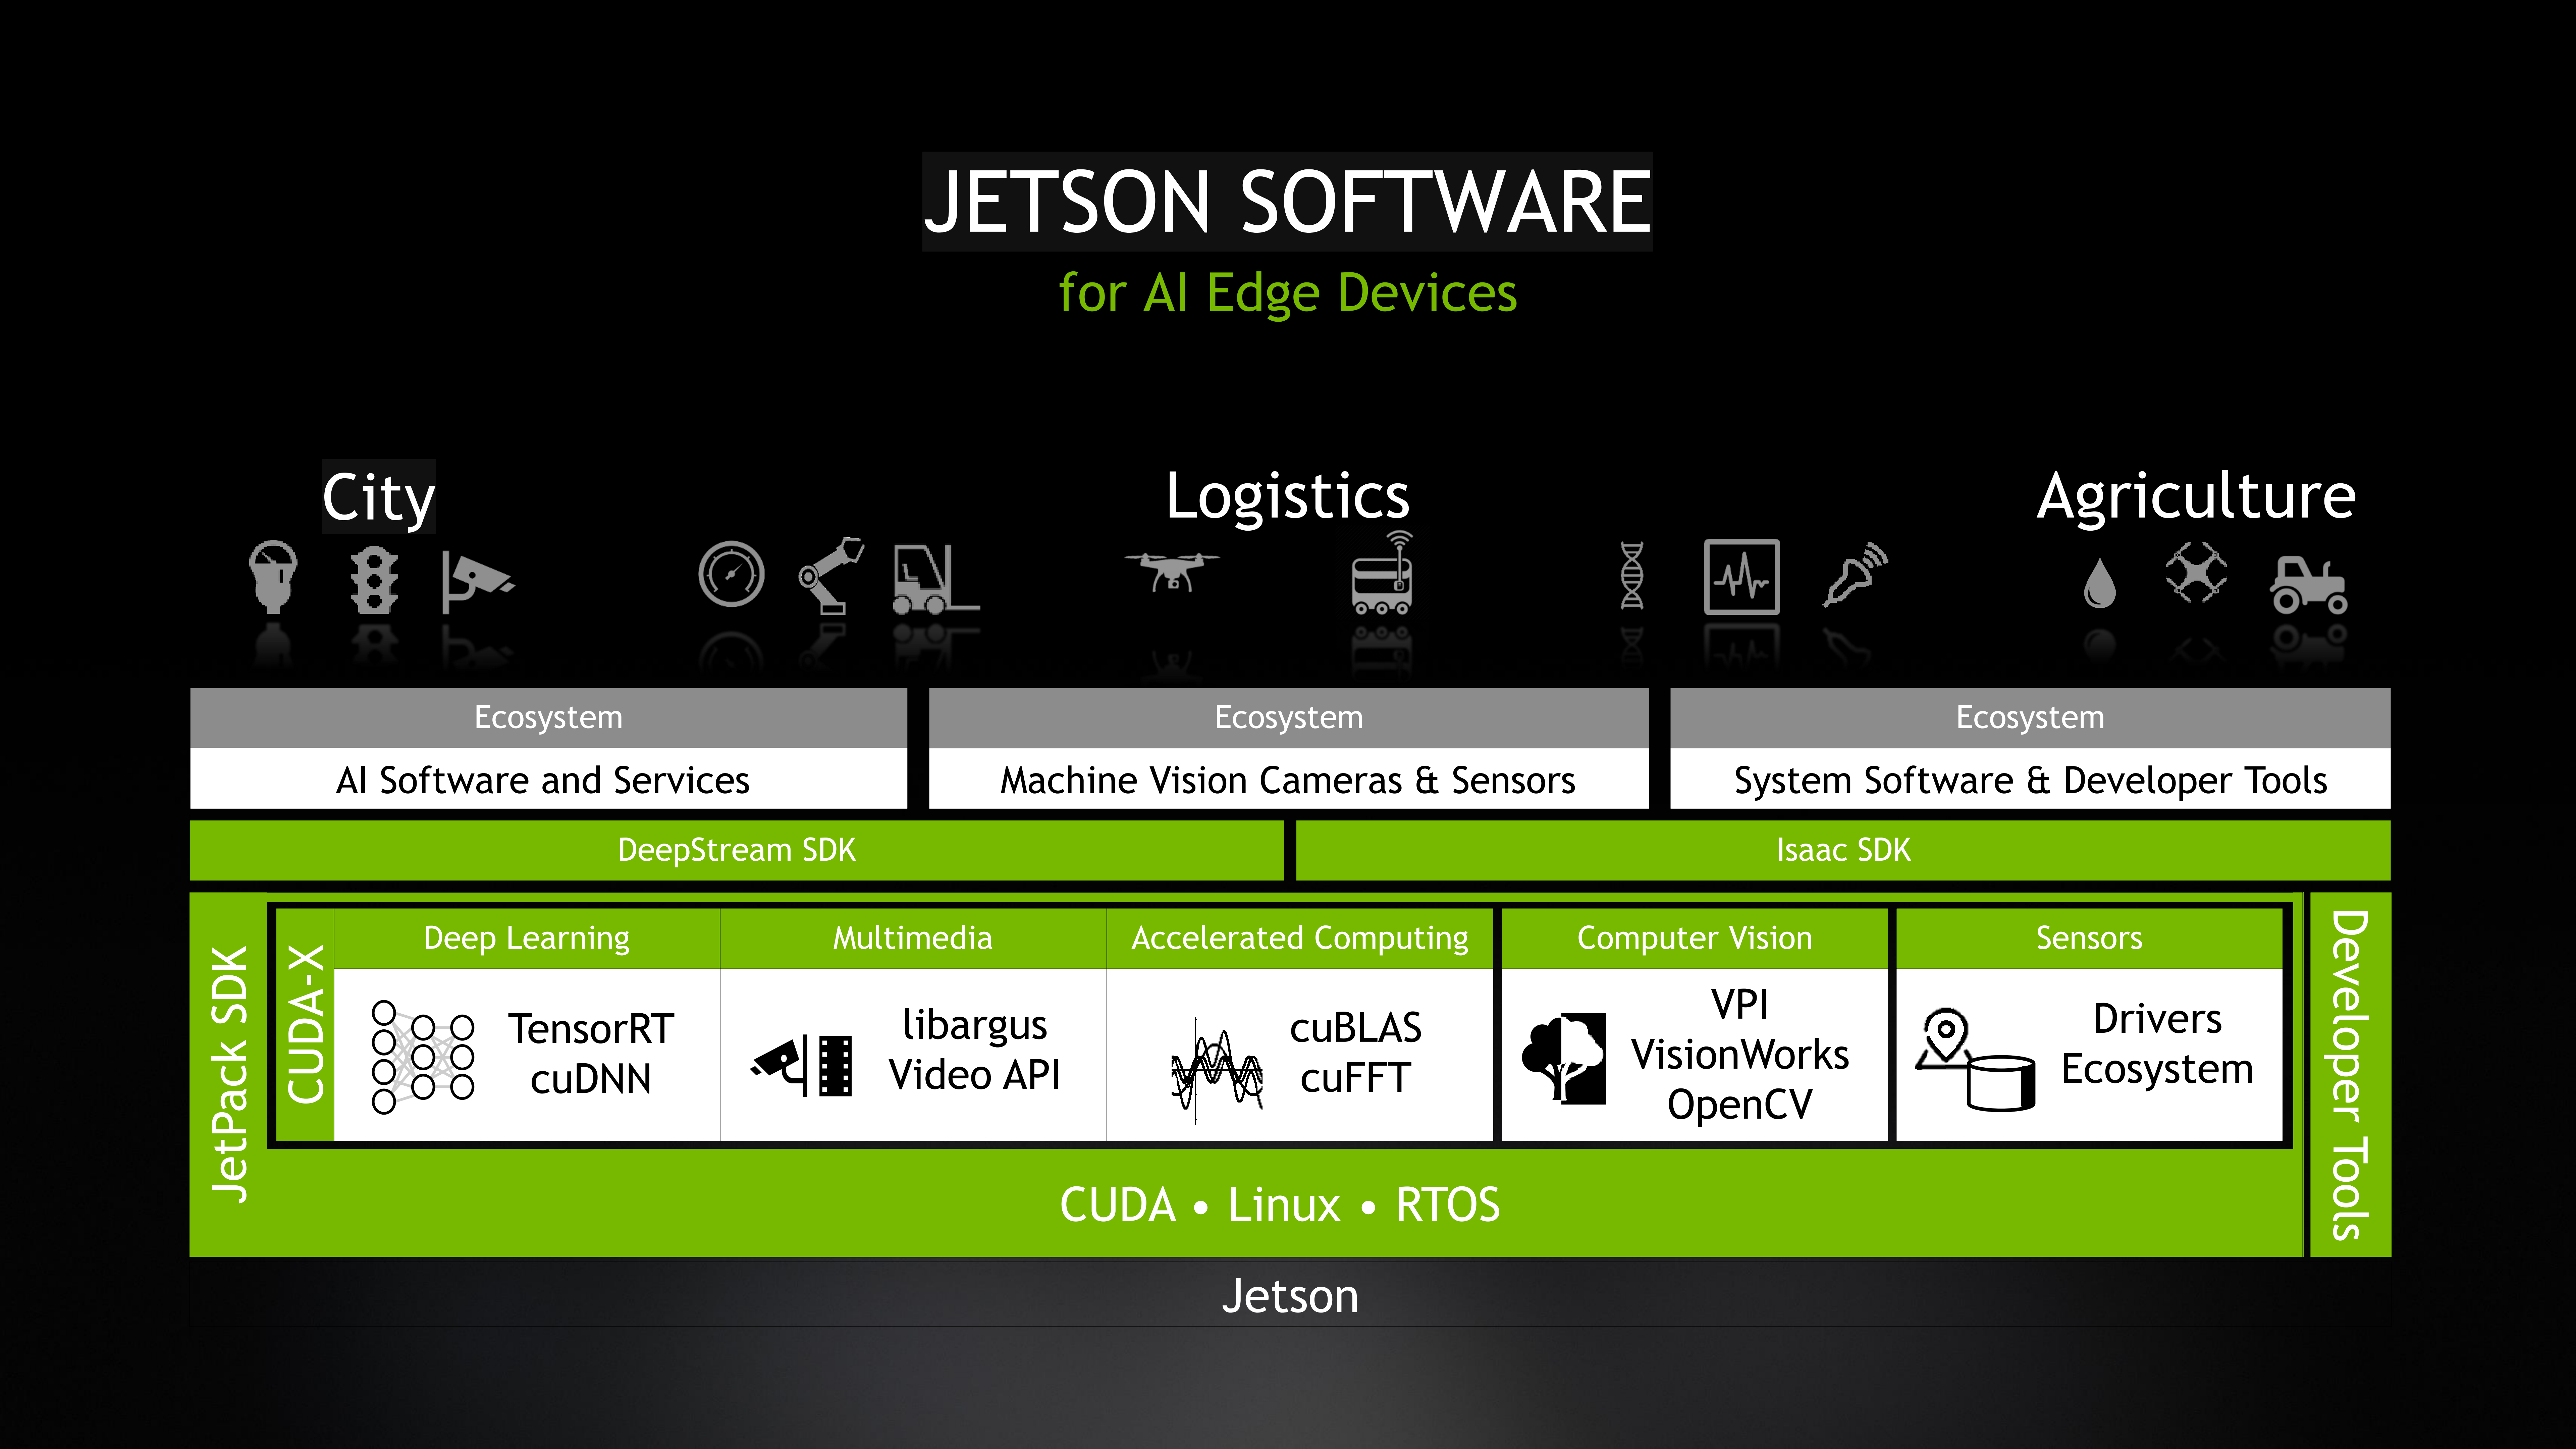
JETSON SOFTWARE
for AI Edge Devices
City
Logistics
Agriculture
Ecosystem
AI Software and Services
Ecosystem
Machine Vision Cameras & Sensors
Ecosystem
System Software & Developer Tools
DeepStream SDK
Isaac SDK
JetPack SDK
CUDA-X
Deep Learning
TensorRT
cuDNN
Multimedia
libargus
Video API
Accelerated Computing
cuBLAS
cuFFT
Computer Vision
VPI
VisionWorks
OpenCV
Sensors
Drivers
Ecosystem
CUDA • Linux • RTOS
Developer Tools
Jetson

CONTINUOUS SOFTWARE INVESTMENT

Scalable Software Defined Platform

Jetpack 3.x
+Jetson TX2i

Docker support
TensorRT adds ONNX model support

DeepStream 1.5

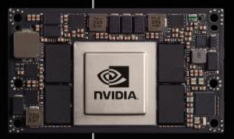

Jetpack 4 DP
+Jetson AGX Xavier

TensorRT supports DLAs and GPU concurrently
cuDNN improved perf

DeepStream 3.0 DP

JAN 2019

Jetpack 4.2.x
+Jetson Nano
+Jetson TX2 4GB

CSI virtual channels
TensorRT Python API
Jetson.GPIO Python library

DeepStream 4.0

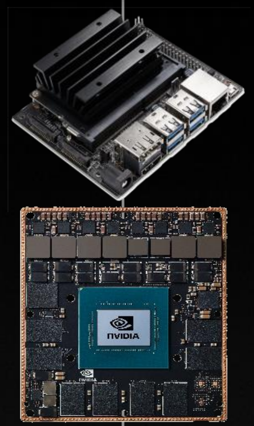

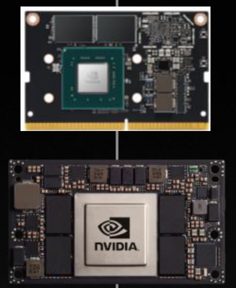

Jetpack 4.3
+Jetson AGX Xavier [32GB]

Debian package based OTA
Full support for DLA INT8 on Jetson AGX Xavier

DeepStream 4.02

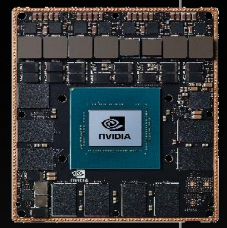

JAN 2020

Jetpack 3.x
Sustaining Release

Jetpack 4.4
+Jetson Xavier NX

Cloud-native support
Reduced memory usage for cuDNN
PowerEstimator web app

DeepStream 5.0

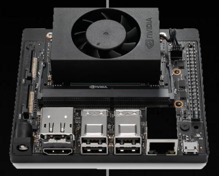

Jetpack 3.x
Sustaining Release

Jetpack 4.5

Enhance Secure Boot
Disk encryption support on selected products
Flash over NFS
VPI 1.0

JAN 2021

NVIDIA

DEEPSTREAM

March 2020

INTELLIGENT VIDEO ANALYTICS (IVA) FOR EFFICIENCY AND SAFETY

Access Control

Public Transit

Industrial Inspection

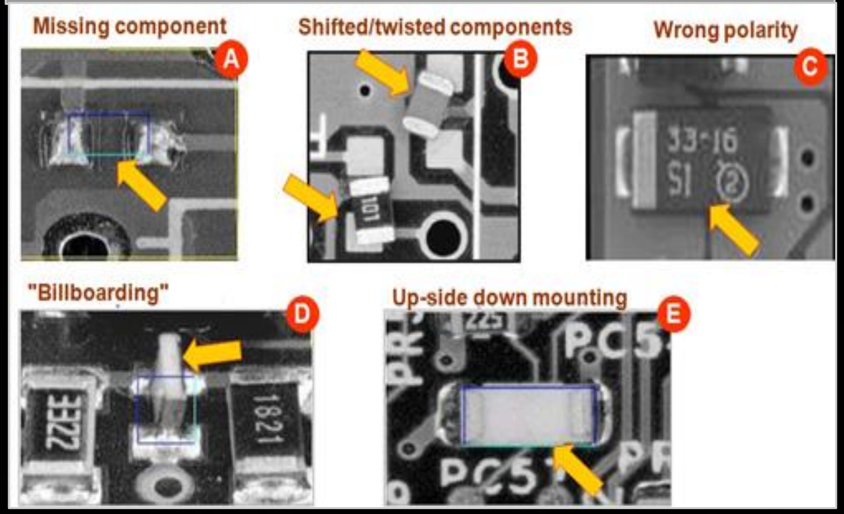

Traffic Engineering

Retail Analytics

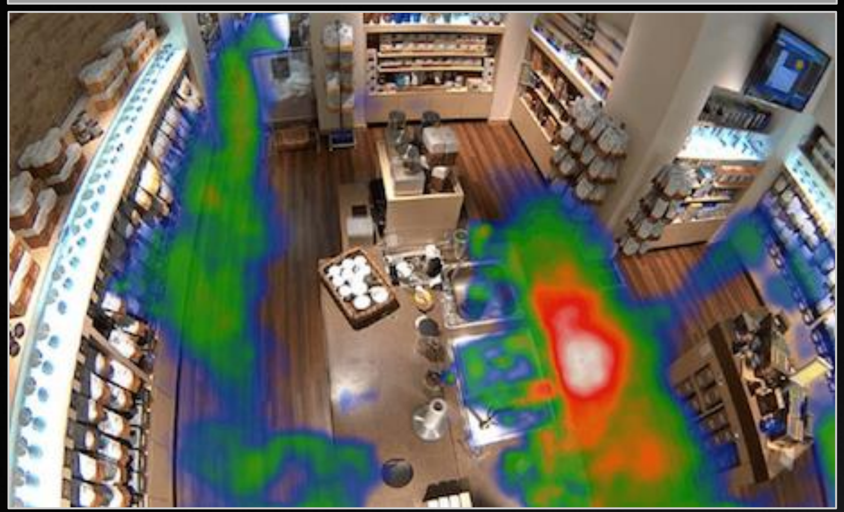

Logistics

Critical Infrastructure

Public Safety

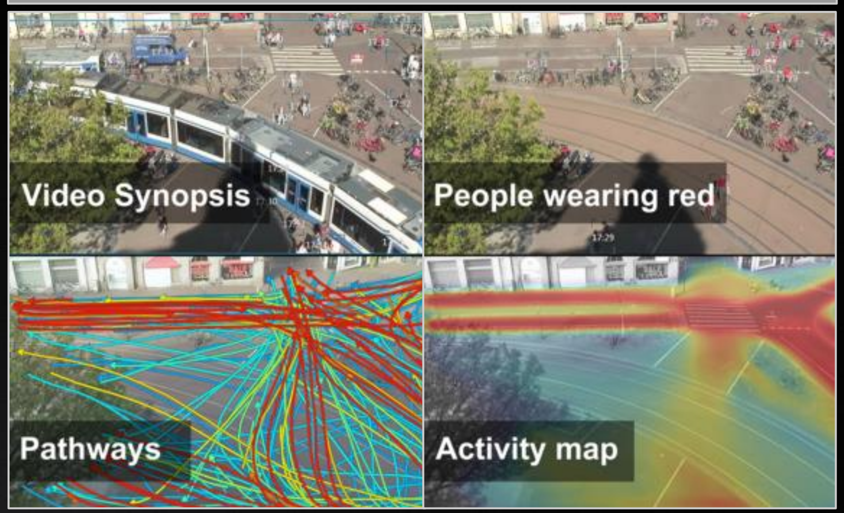

IVA APPLICATION WORKFLOW

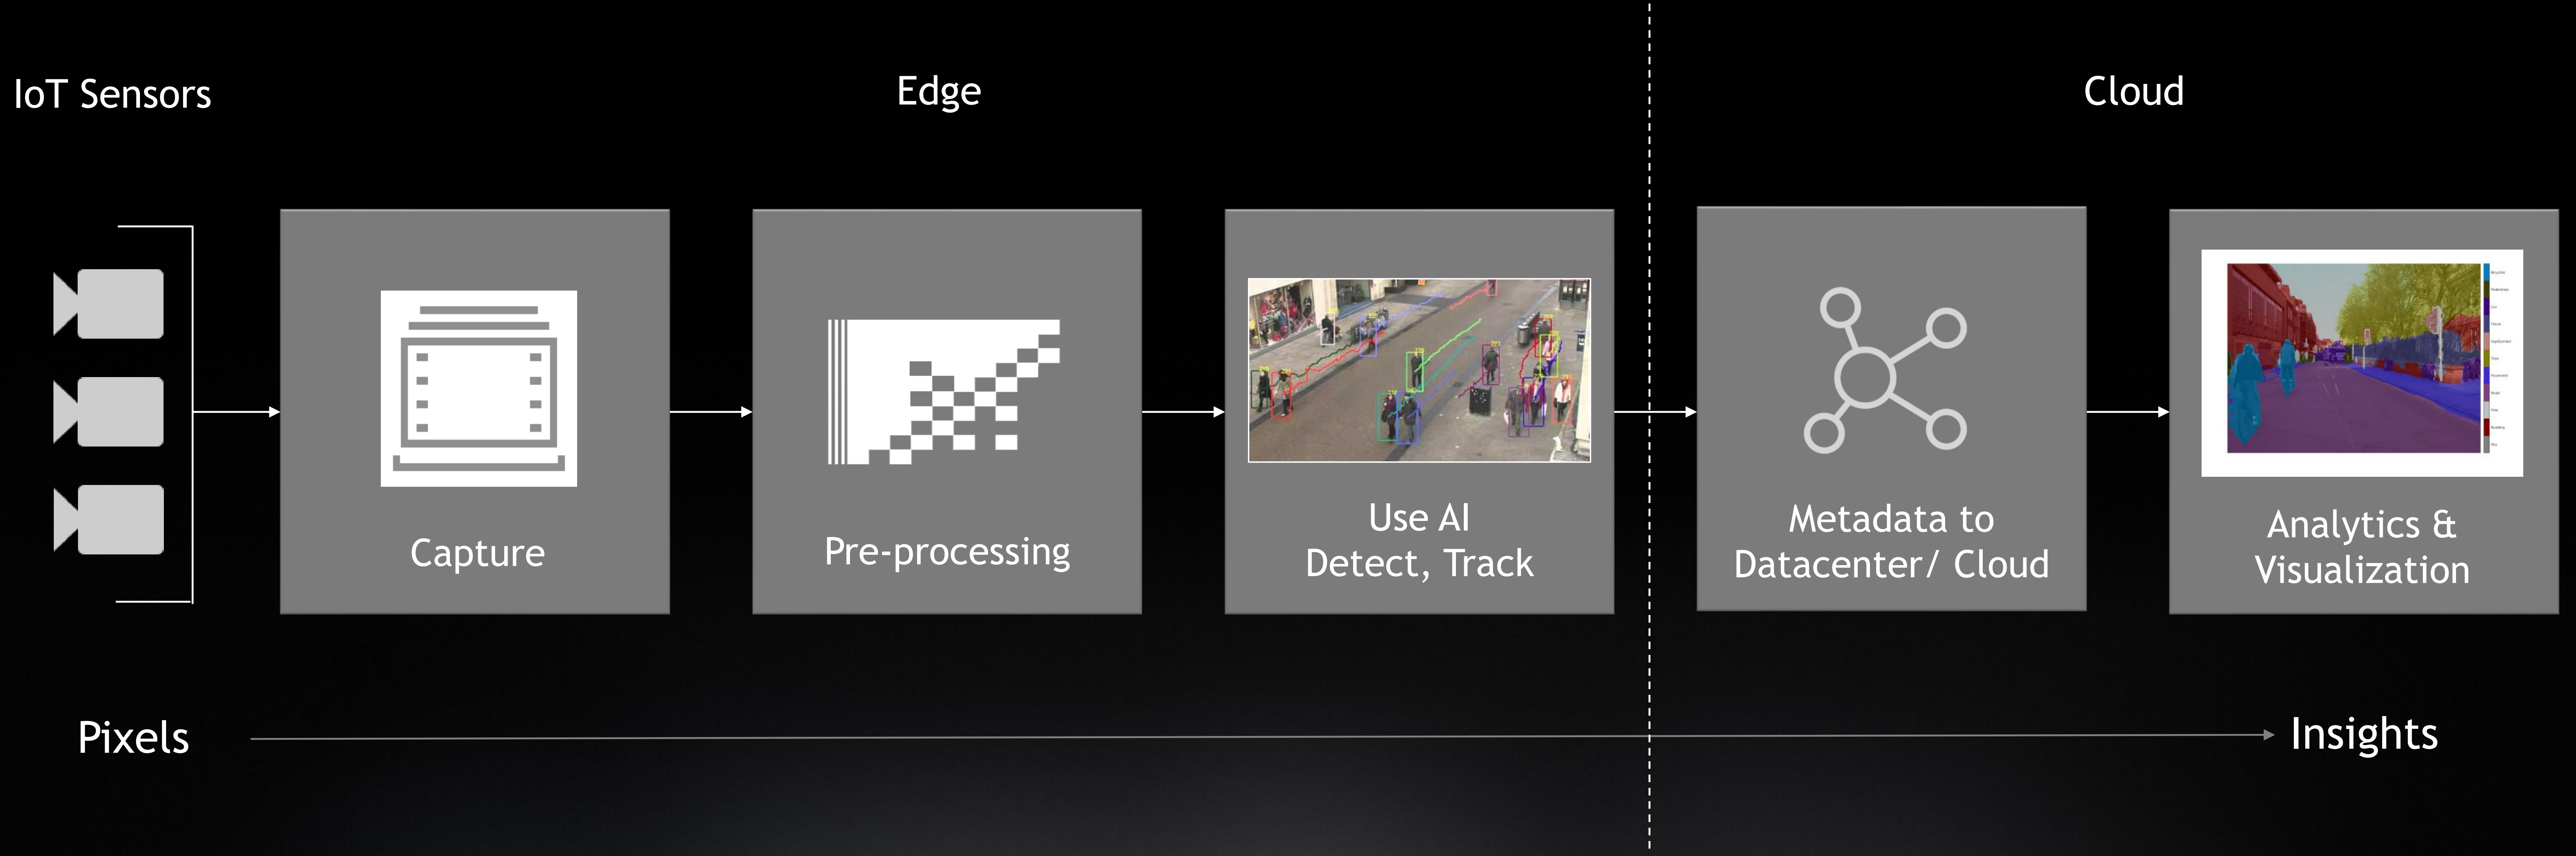

NVIDIA

PRE-TRAINED MODELS & TRANSFER LEARNING TOOLKIT

PRODUCTION READY PRE-TRAINED VISION AI MODELS

People Detection

Face Detect IR

DashCamNet

TrafficCamNet

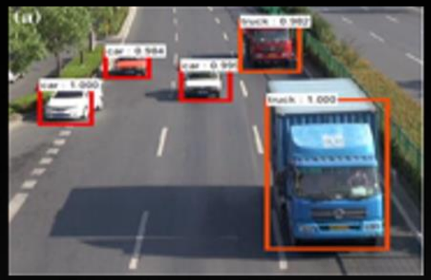

Vehicle Type net

Vehicle Make net

Free AI models for Smart City, Public Safety, Retail, Healthcare & Robotics

Over 80% accuracy

Optimized for High Throughput

Deploy with turnkey inference examples

Transfer Learn with your dataset

Check out the models on NGC

MODEL CARDS READILY AVAILABLE ON NGC

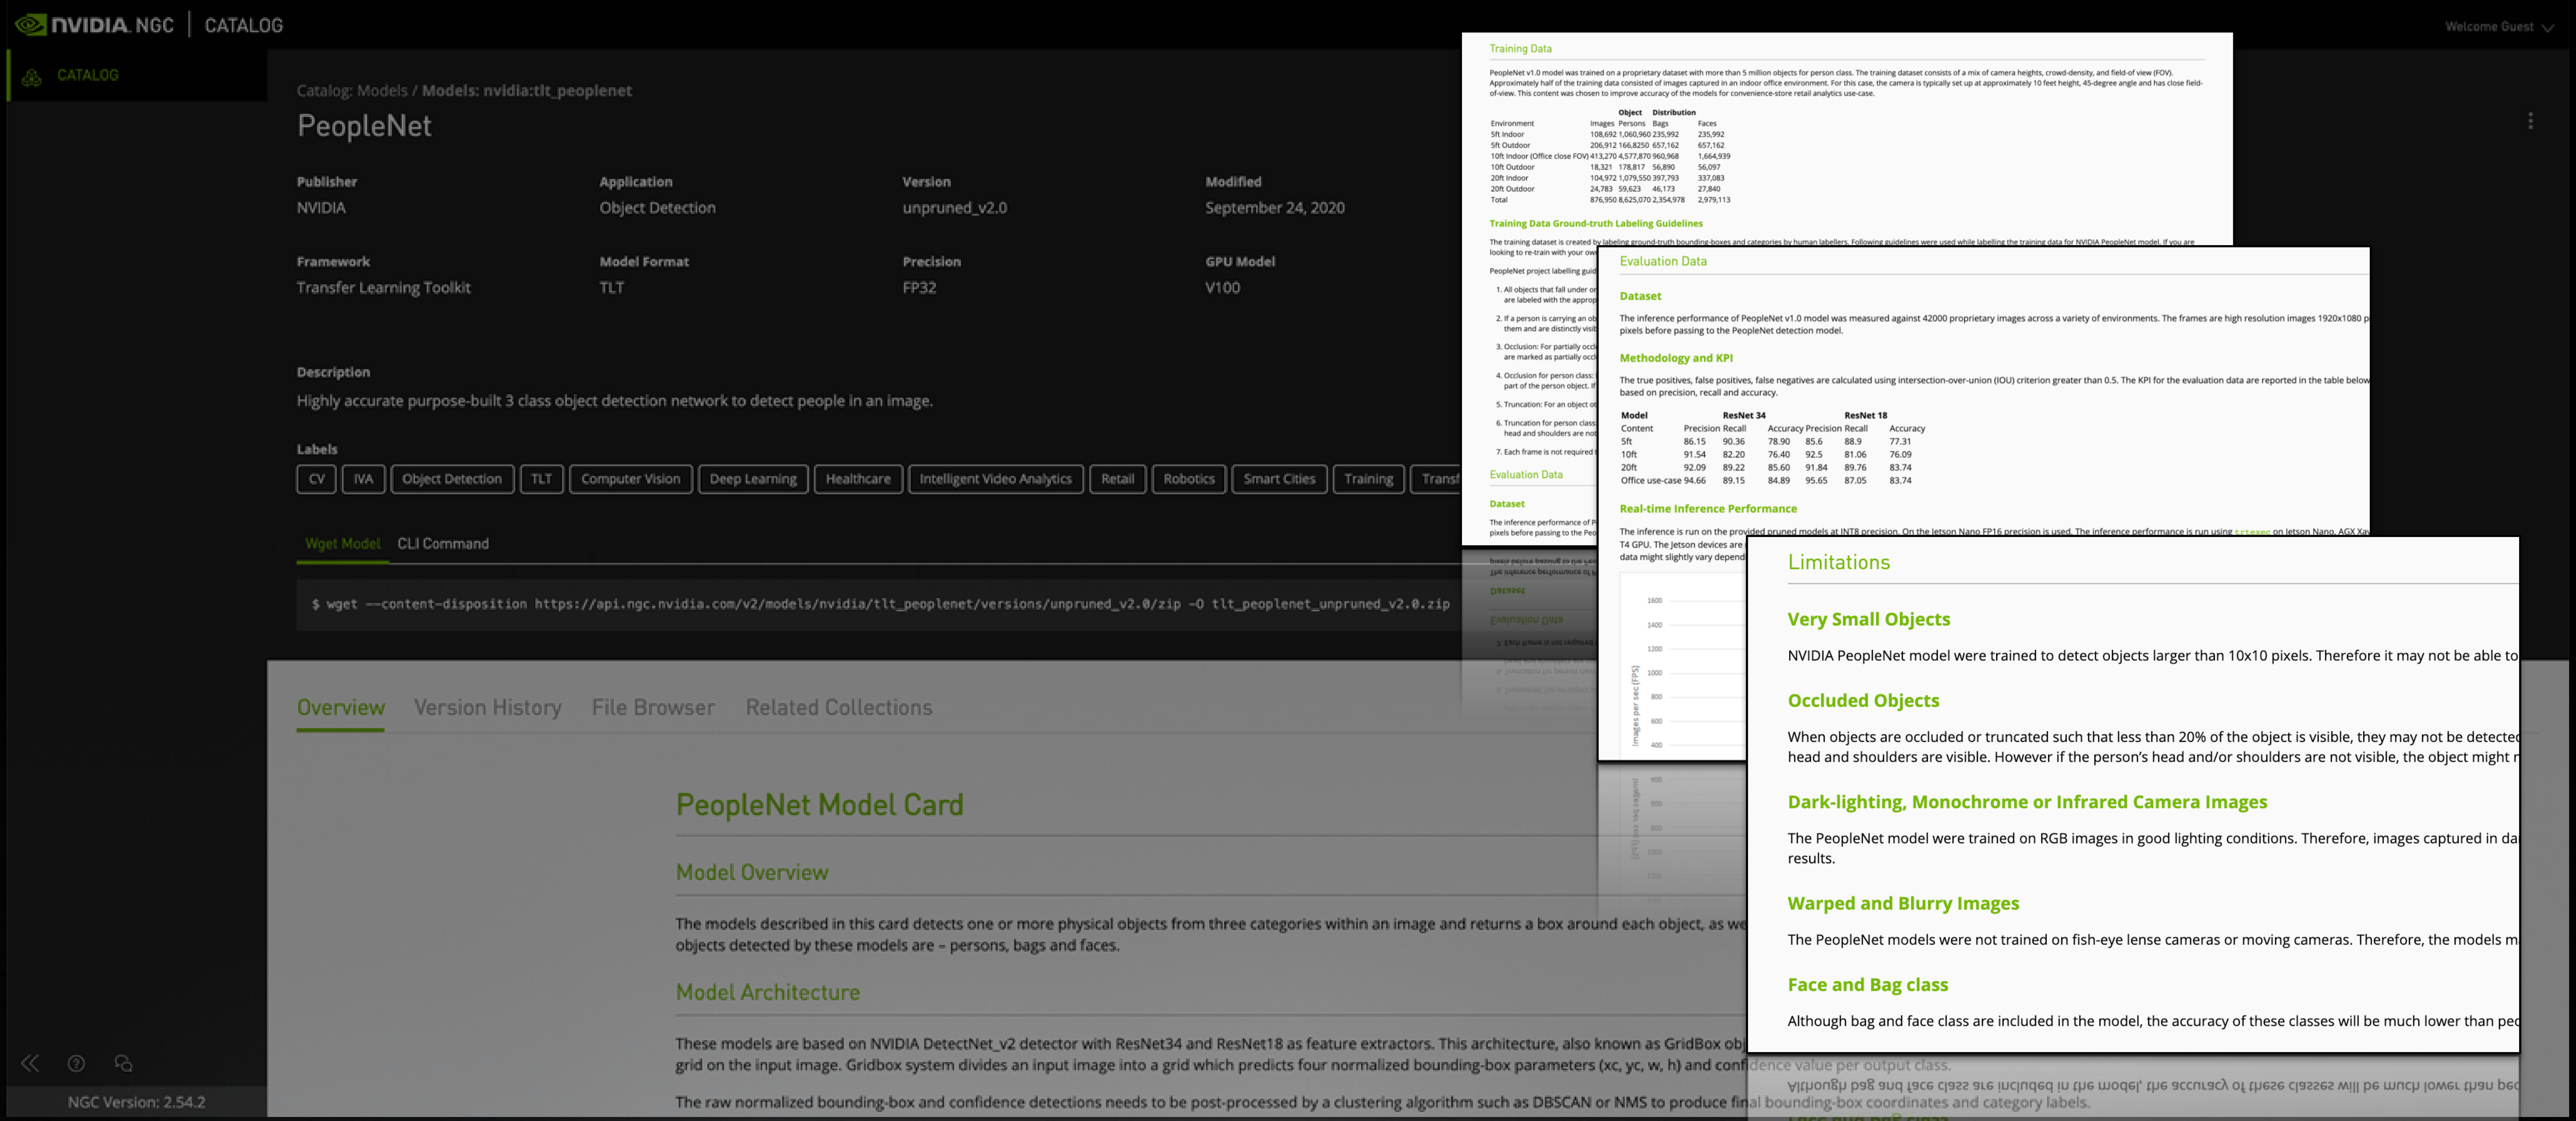

TRAINING DATASET | VALIDATION DATASET | CAMERA GUIDELINES | LIMITATIONS

NVIDIA CONFIDENTIAL. DO NOT DISTRIBUTE.

FASTEST TIME TO MARKET FOR PEOPLE DETECTION

YOLOV4 model

Lots of data required to improve accuracy
Large training time

PeopleNet model

Apply minimal fine-tuning or deploy as is

WHAT IS TRANSFER LEARNING?

"Transfer Learning is a process of transferring learned features from one model to another".

With Transfer Learning, you start with a pre-trained model and a small dataset, fine-tuning the model to learn from to your dataset. The pre-trained weights provide a better starting point to train on your dataset instead of starting from random weights.

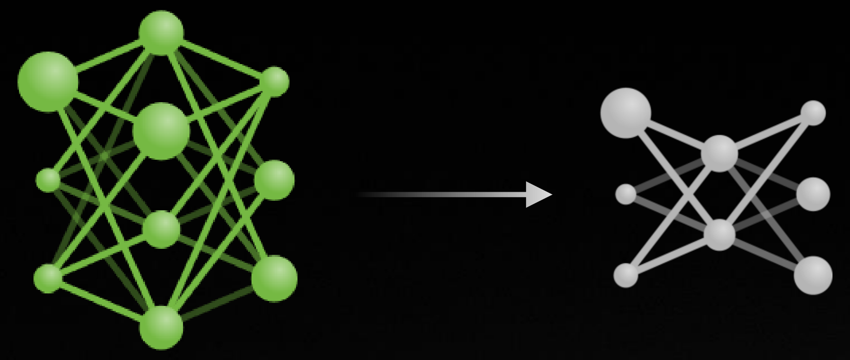

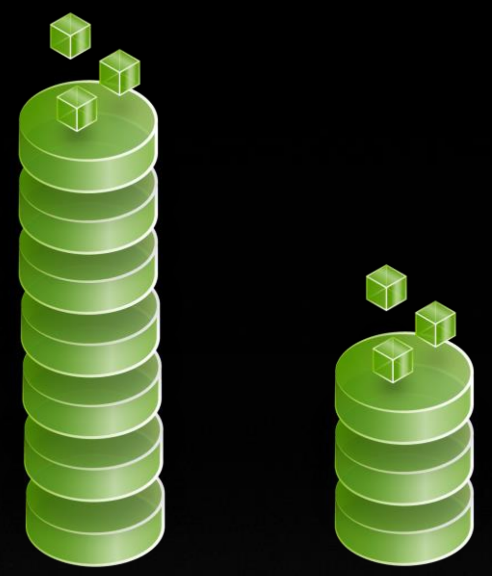

Less Data Required to Train Accurately

Reduce Training Time and Cost

NVIDIA Transfer Learning Toolkit (TLT) brings the transfer learning capability in a simplified training toolkit where you start with from a rich library of NVIDIA provided pre-trained models to build you custom AI network in 1/10th the effort and time versus starting from scratch.

https://blogs.nvidia.com/blog/2019/02/07/what-is-transfer-learning/

MODEL PRUNING

2 Step Process

1. Reduce model size
2. Incrementally retrain model after pruning to recover accuracy

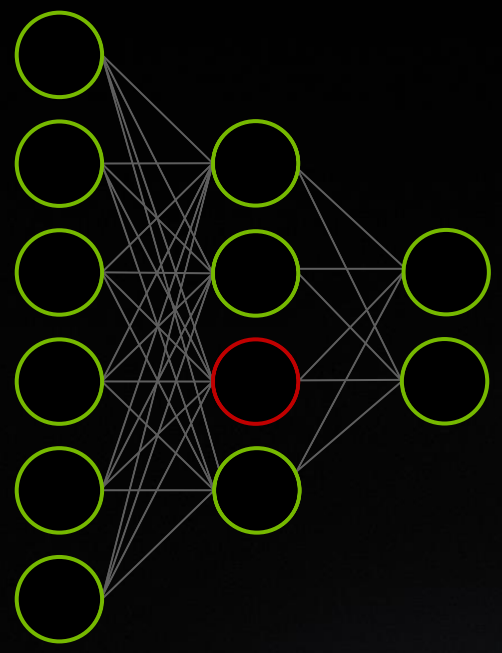

6 inputs, 6 neurons, 32 connections

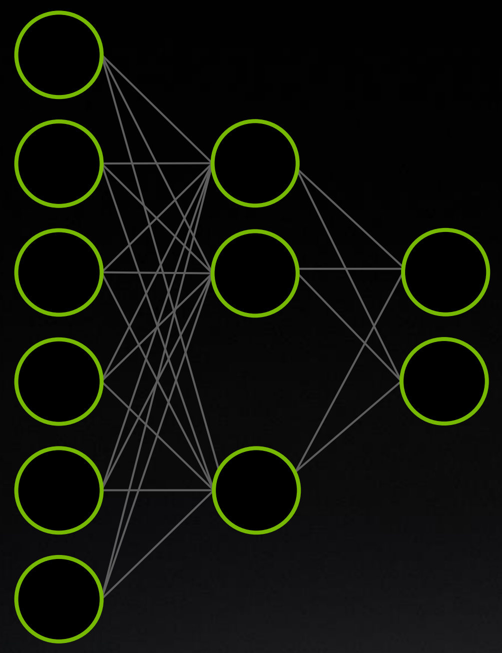

6 inputs, 5 neurons, 24 connections

Network – TrafficCamNet

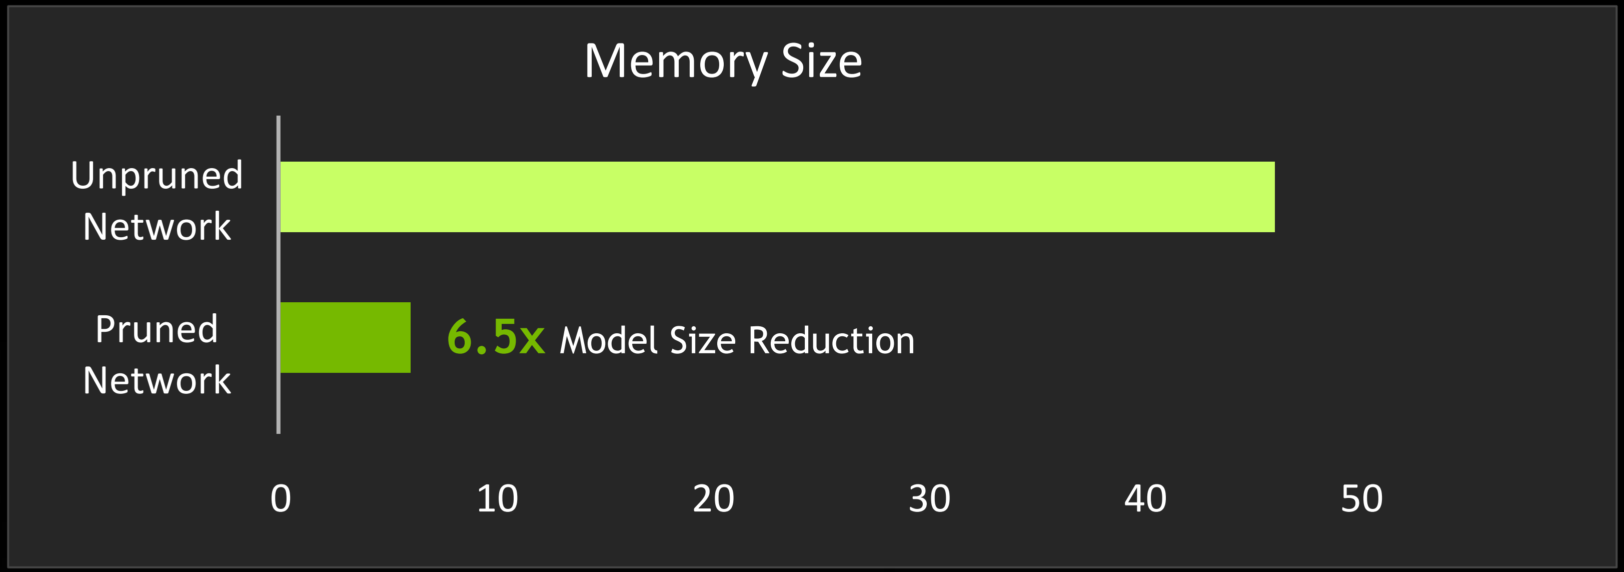

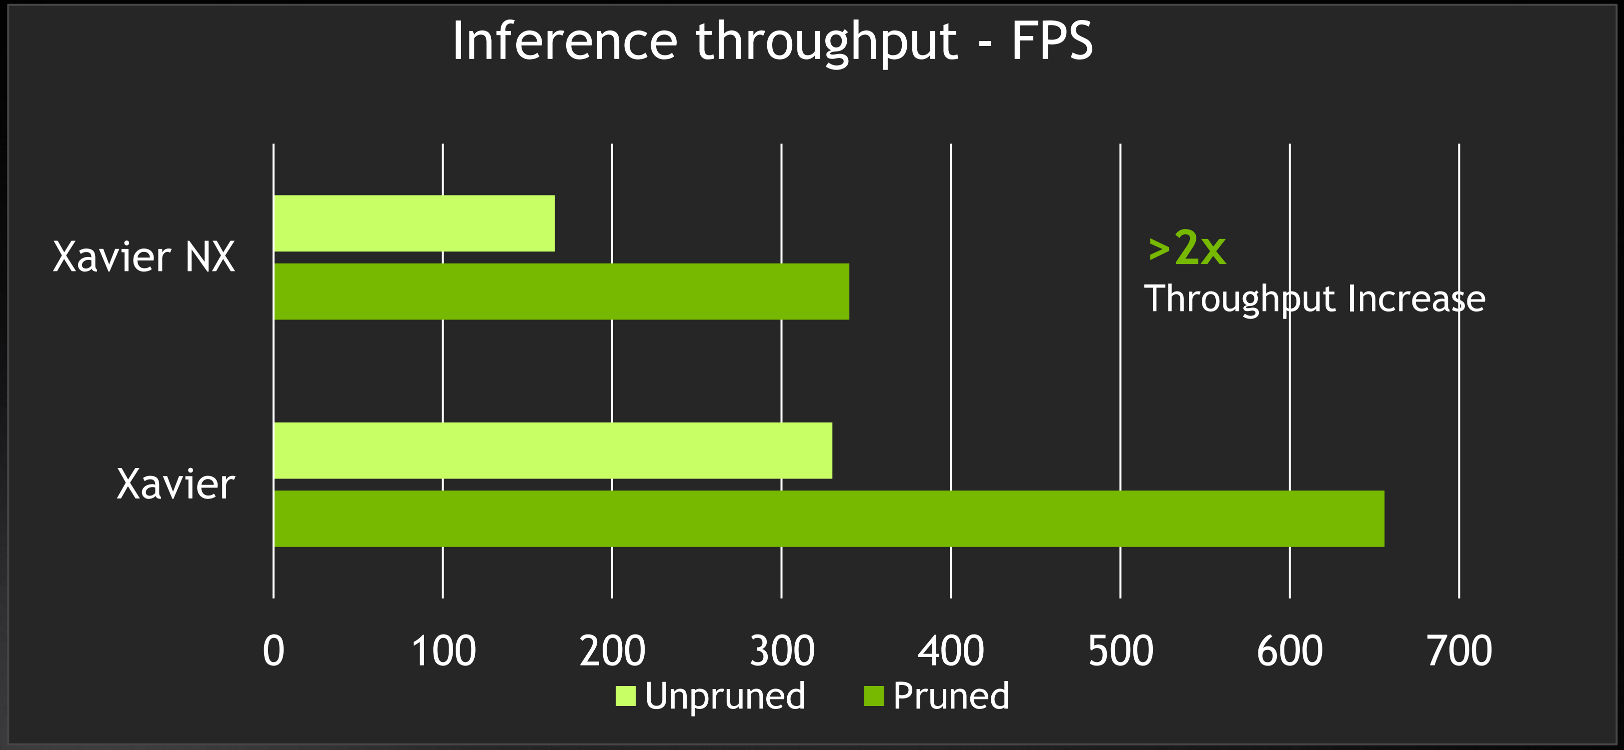

ADDING NEW CLASSES

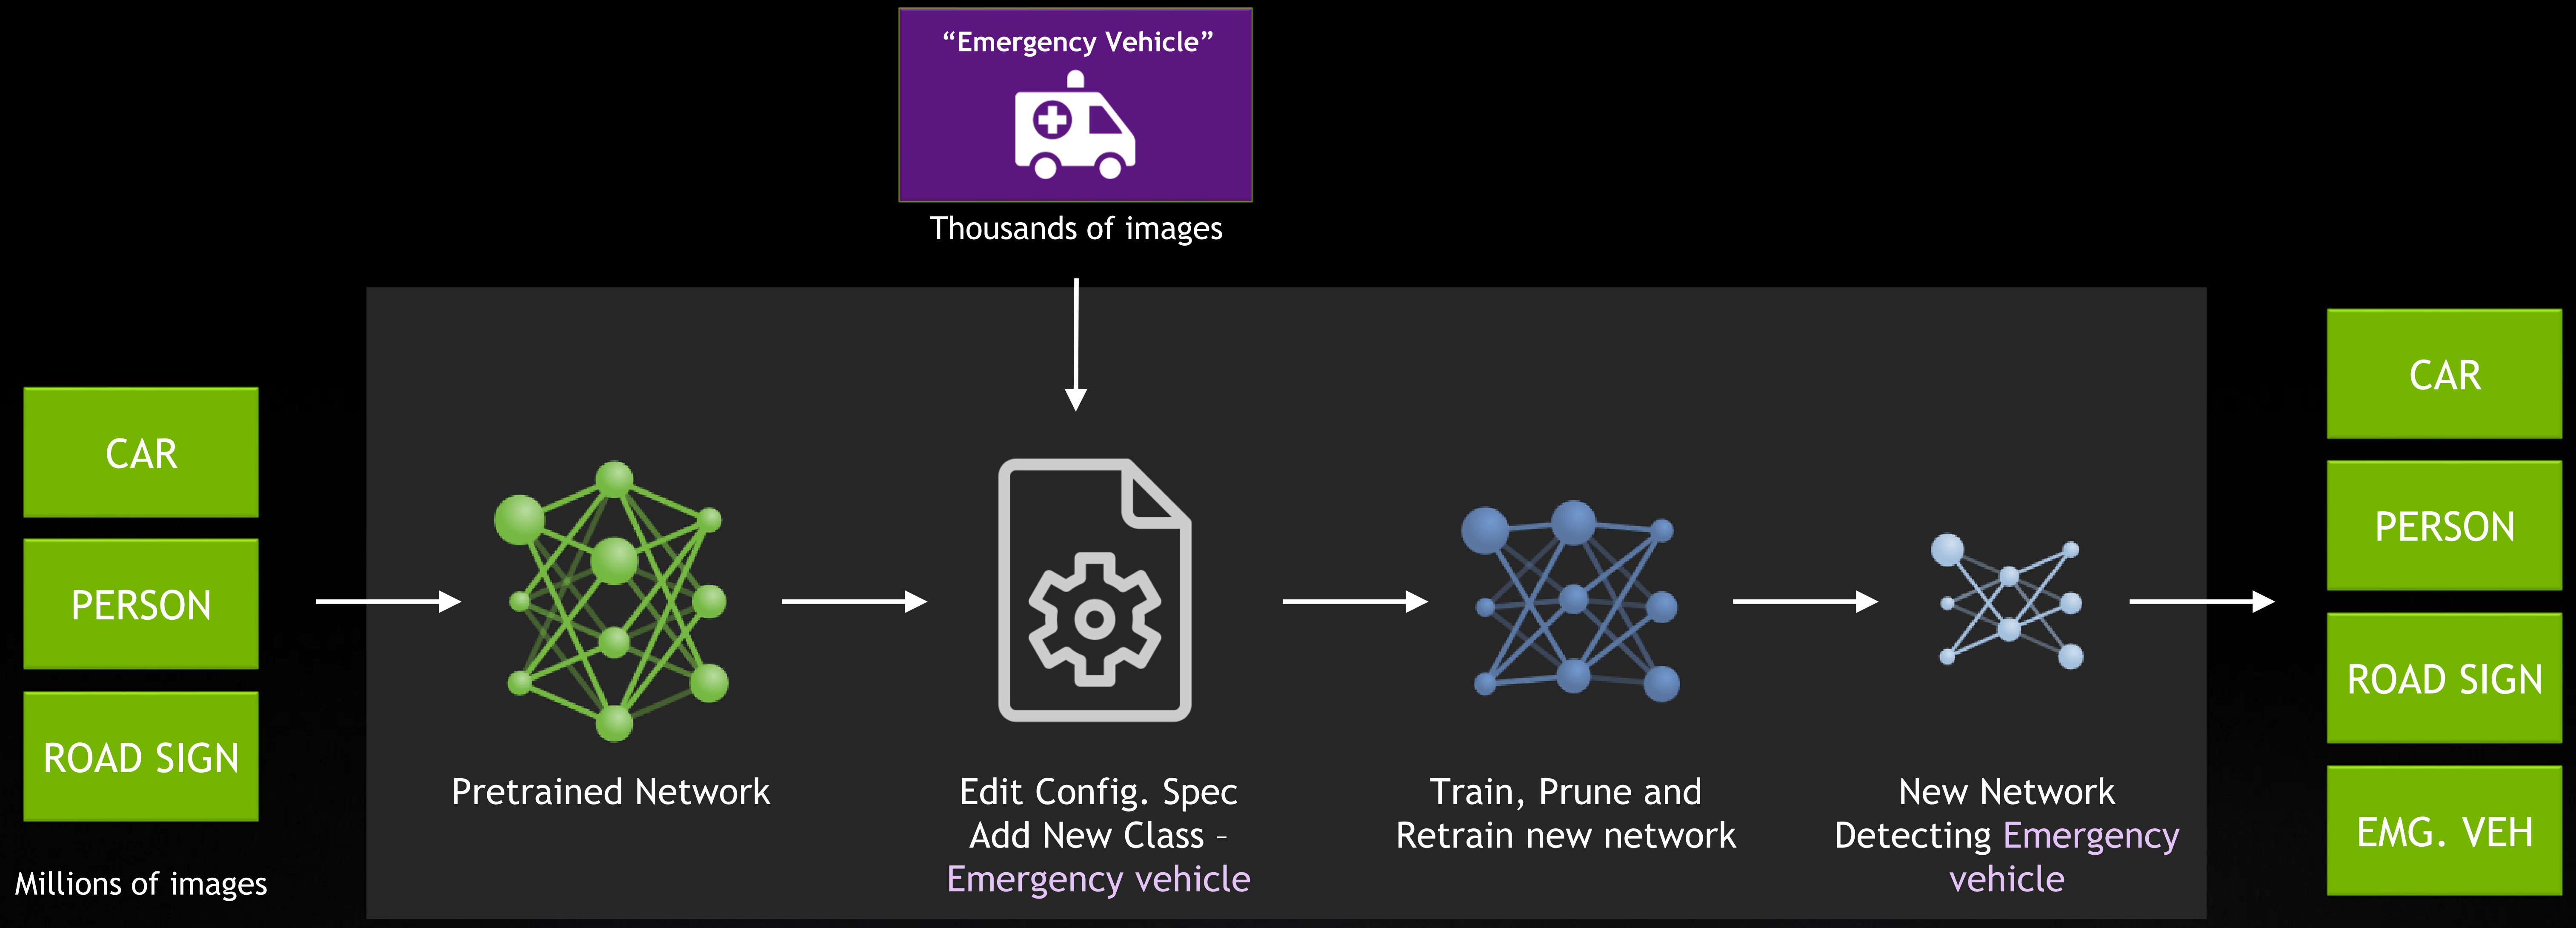

SCENE ADAPTATION

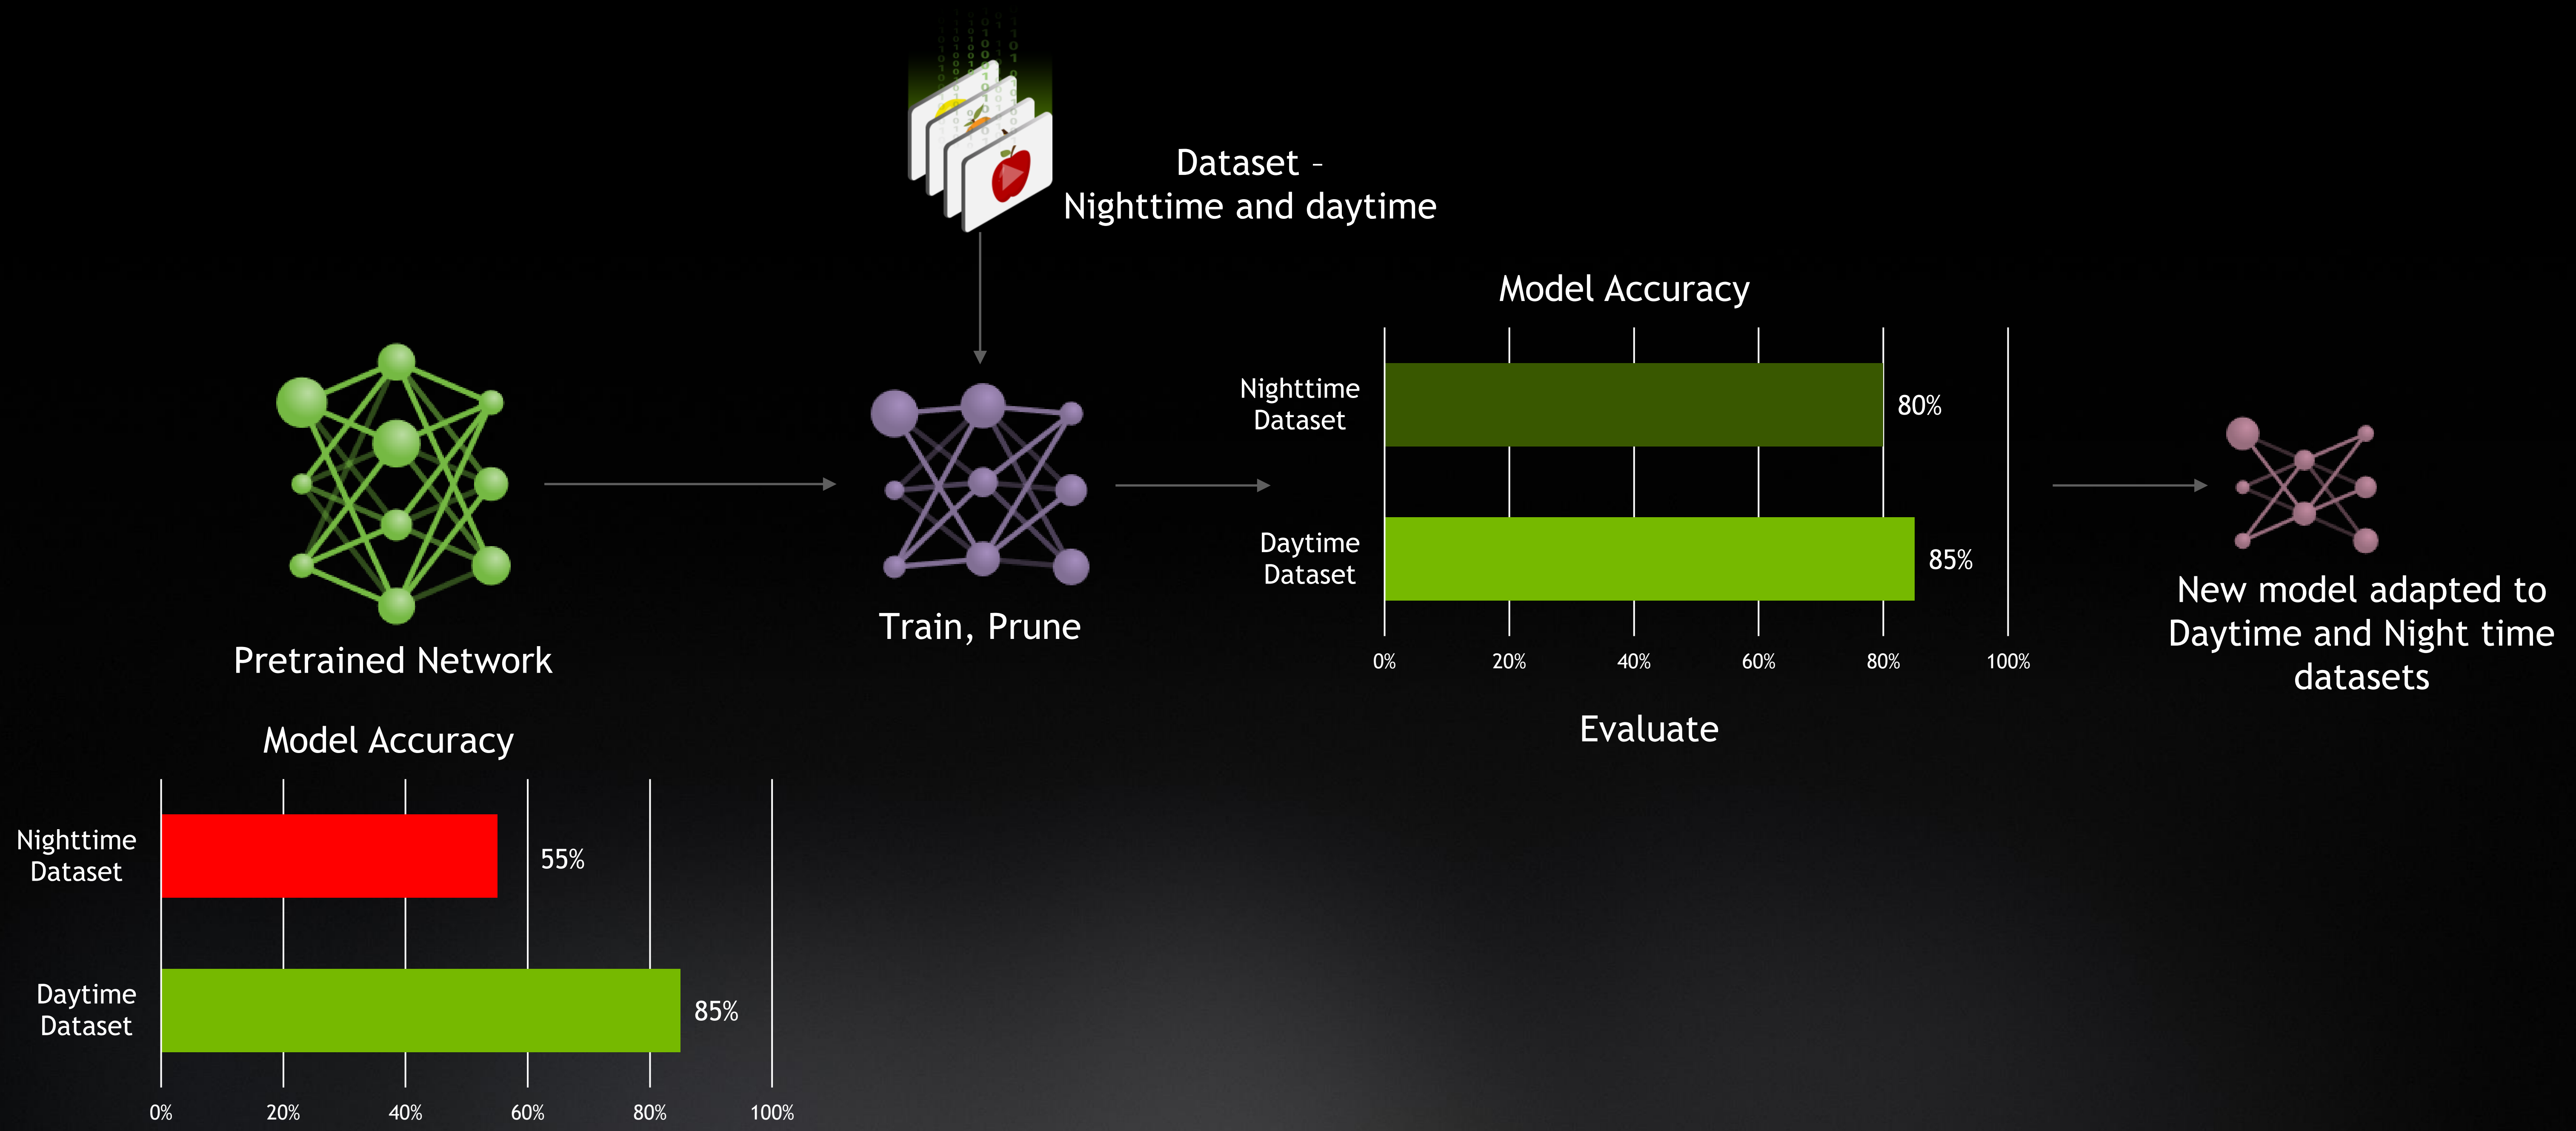

ZERO-CODING TRAINING

Train with pre-built Jupyter Notebooks

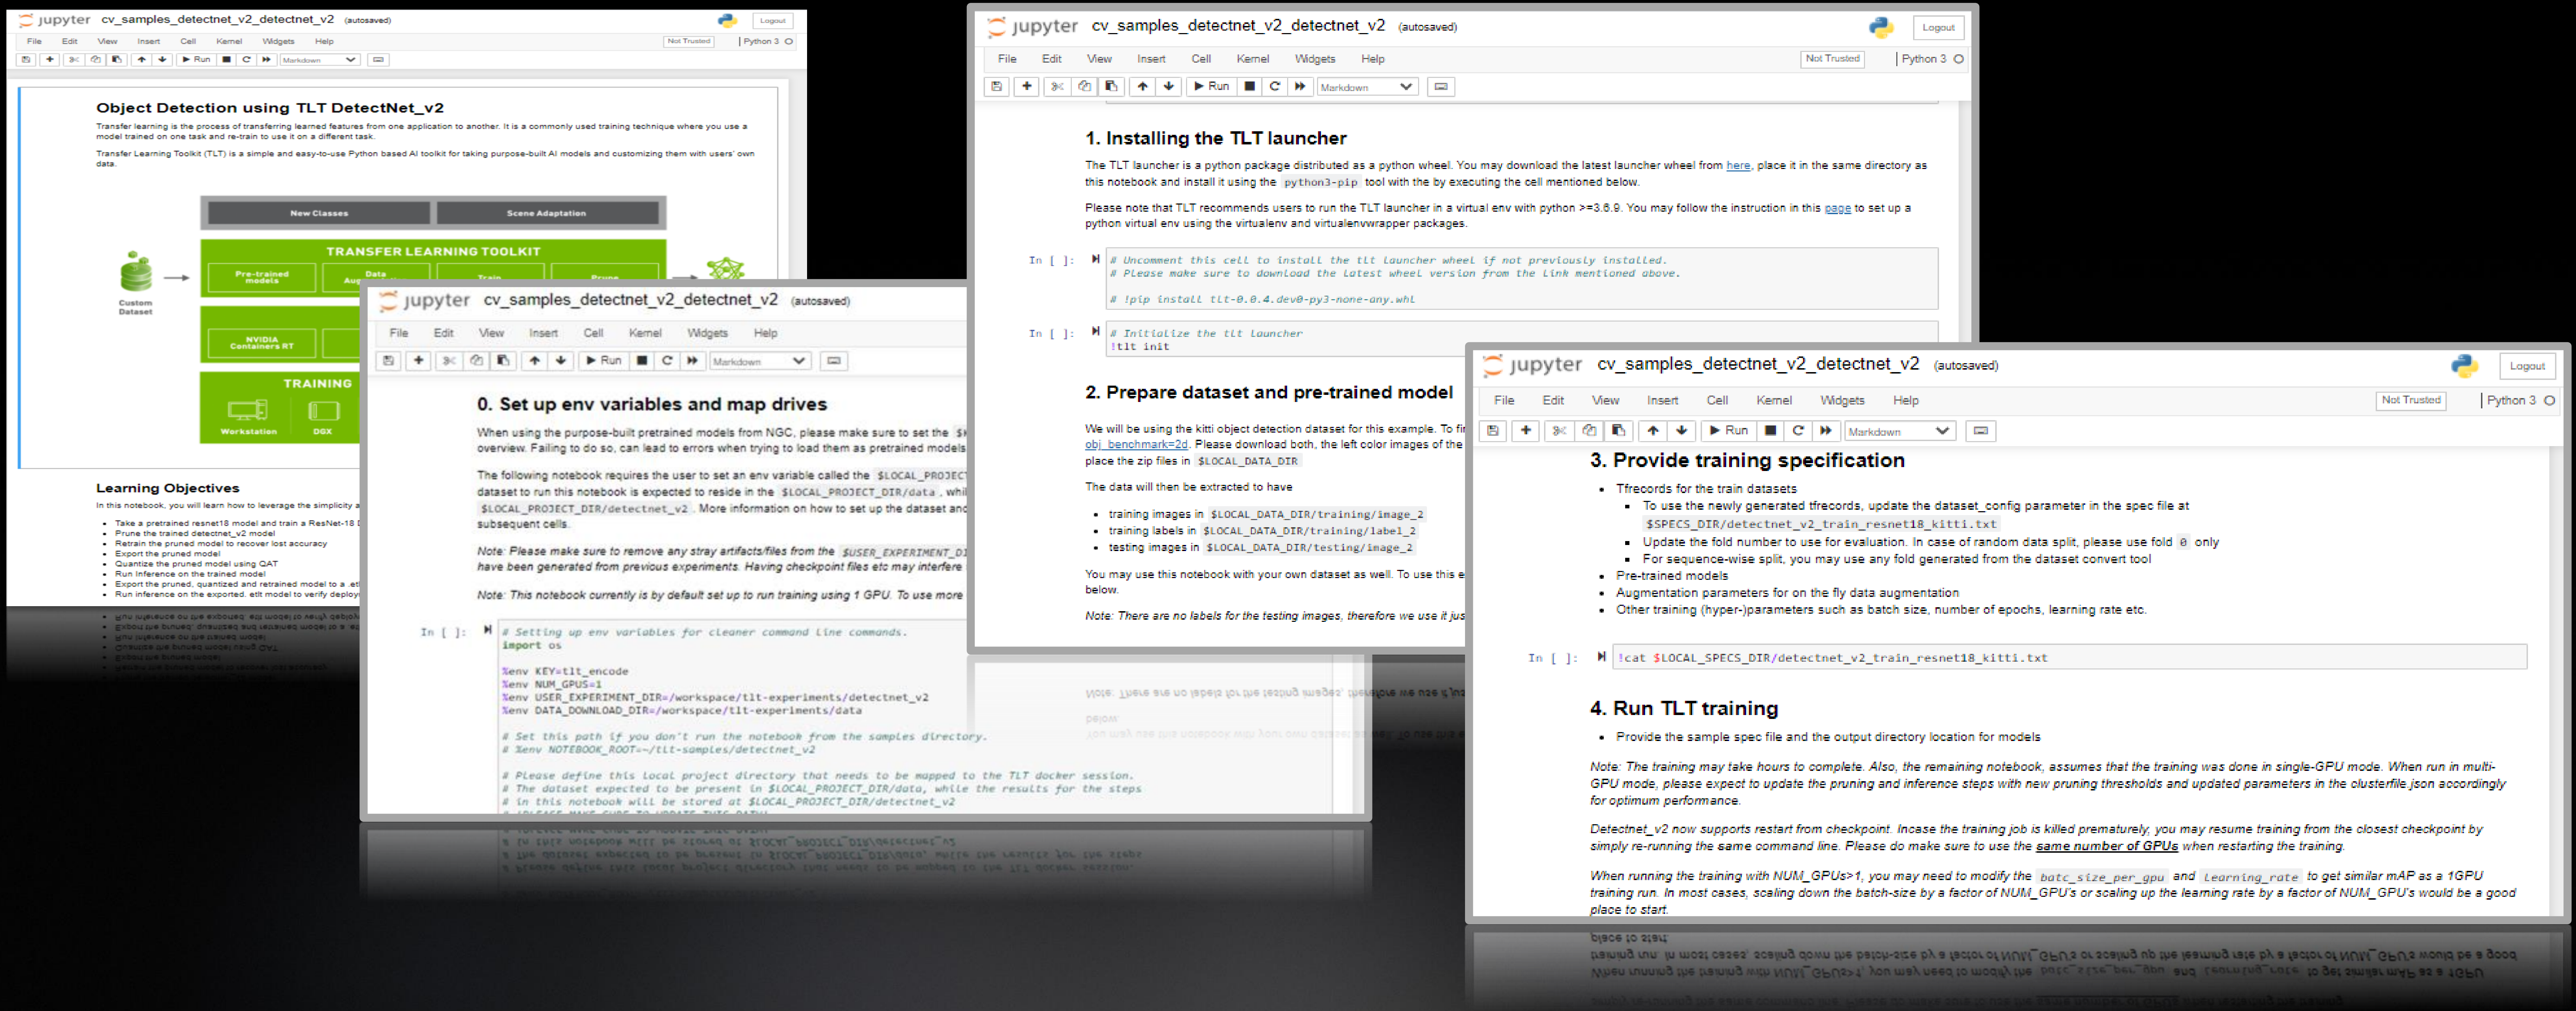

TRANSFER LEARNING TOOLKIT 2.0

	Classification	Object Detection					
Network Architecture		DetectNet_V2	FasterRCNN	SSD	YOLOV3†	RetinaNet†	DSSD†
Supported Backbones	ResNet 10/18/50	ResNet 10/18/50	ResNet 10/18/50	ResNet 10/18	ResNet 10/18/50	ResNet 10/18/50	ResNet 10/18/50
	VGG16/19	VGG16/19	VGG16/19	VGG16/19†	VGG16/19	VGG16/19	VGG16/19
	GoogLeNet	GoogLeNet	GoogLeNet	GoogLeNet†	GoogLeNet	GoogLeNet	GoogLeNet
	MobileNet V1/V2	MobileNet V1/V2	MobileNet V1/V2	MobileNet V1/V2†	MobileNet V1/V2	MobileNet V1/V2	MobileNet V1/V2
	SqueezeNet			SqueezeNet†	SqueezeNet	SqueezeNet	SqueezeNet
	ResNet 34/101†	ResNet 34/101†	ResNet 34/101†	ResNet 50/34/101†	ResNet 34/101	ResNet 34/101	ResNet 34/101
	DarkNet 19/53†		DarkNet 19/53†	DarkNet 19/53†	DarkNet 19/53	DarkNet 19/53	DarkNet 19/53

† - Available in future release

Models trained on google open images public dataset
Available to download on ngc.nvidia.com

END-TO-END DEEP LEARNING WORKFLOW

Accelerate Time to Market and Save on Compute Resources!

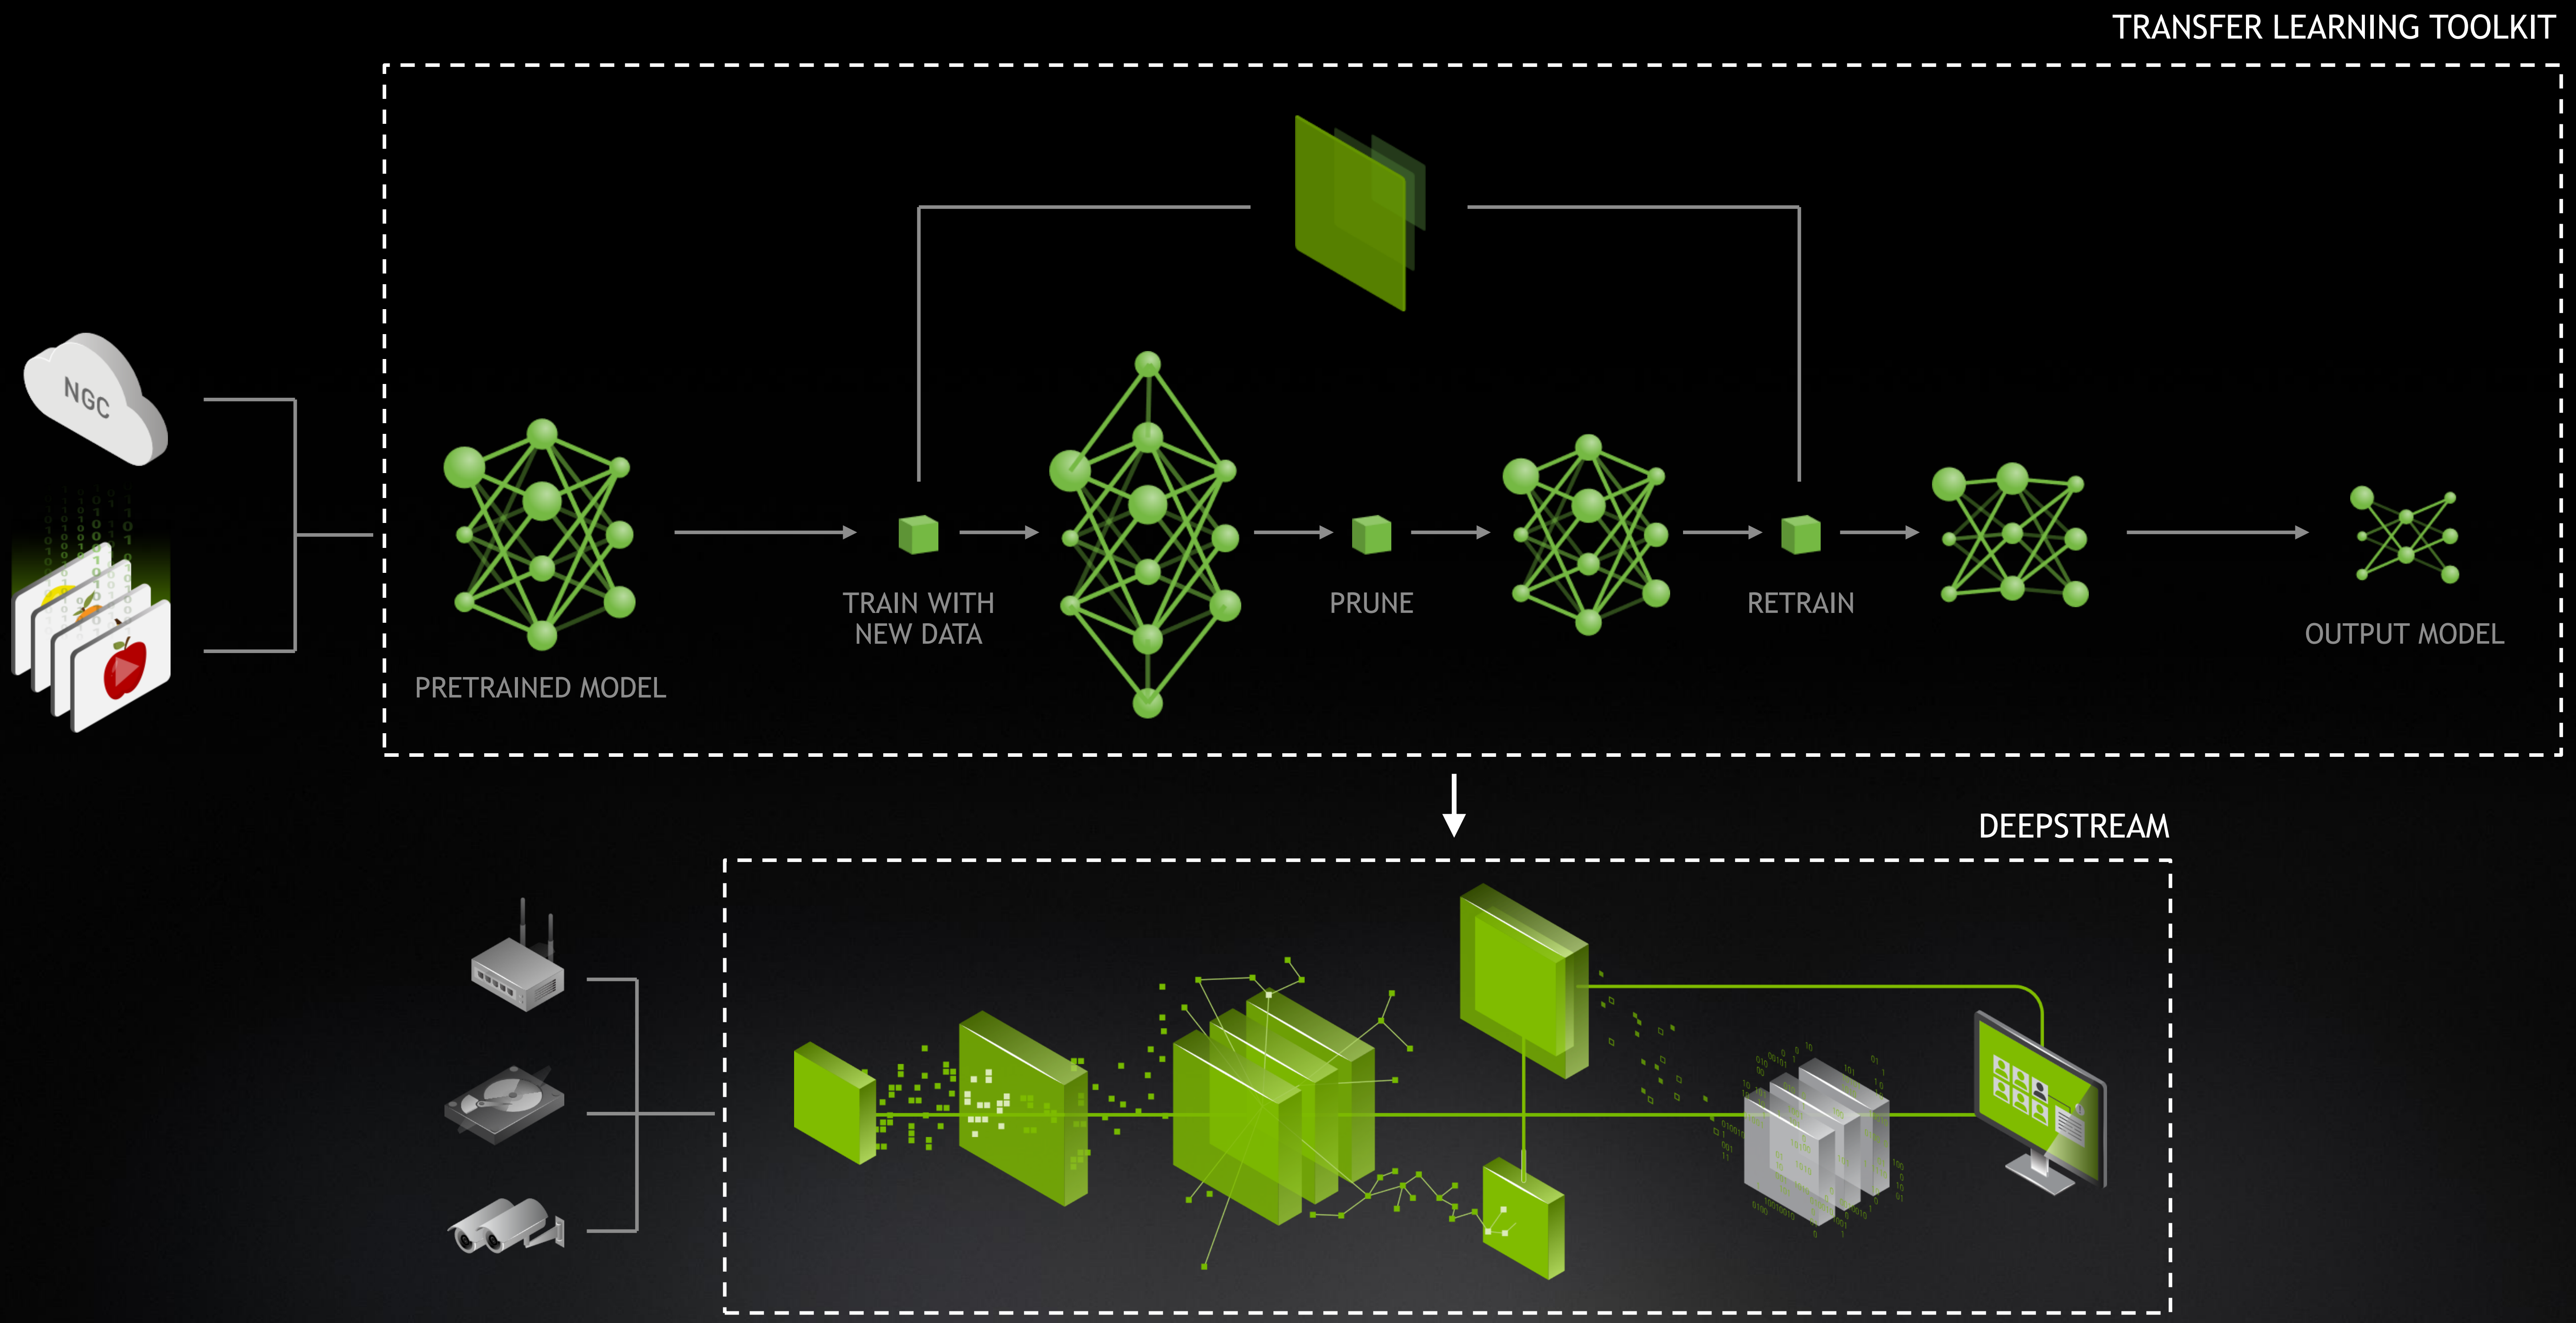

ISAAC

March 2020

NVIDIA ISAAC

- Isaac Engine
- Isaac GEMS
- Reference Designs
- Isaac Sim

ISAAC SIM

- Domain Randomization
- Scenario Management
- Sensor Models
- Robot Models

ISAAC Apps

- CARTER (Indoor Robot)
- Kaya (Getting Started)
- Leonardo (Manipulation)
- Custom Applications

ISAAC GEMS

- Multi-class Segmentation DNN
- 3D Object Pose Estimation
- Object Detection DNN
- Stereo Depth
- Stereo Visual Inertial Odometry
- Superpixels
- AprilTags
- 2D Skeleton Pose Estimation DNN
- DeepStream for Robotics
- ORB Feature Tracker
- Sensor Drivers
- Planning and Control
- Image Warping
- Navigation
- Text to Speech and Keyword Detection DNNs
- ... and more

ISAAC Engine

- Computational Graph & CUDA Messaging
- Visualization Tools
- Advanced Build System & C API

NVIDIA HW Acceleration

- TensorRT
- cuDNN
- CUDA-X
- RTX
- PhysX

NVIDIA AGX

ISAAC SIM

Simulation for Robotics

- Domain Randomization
- Scenario Management
- Import other robots

“Leonardo” Preview Release

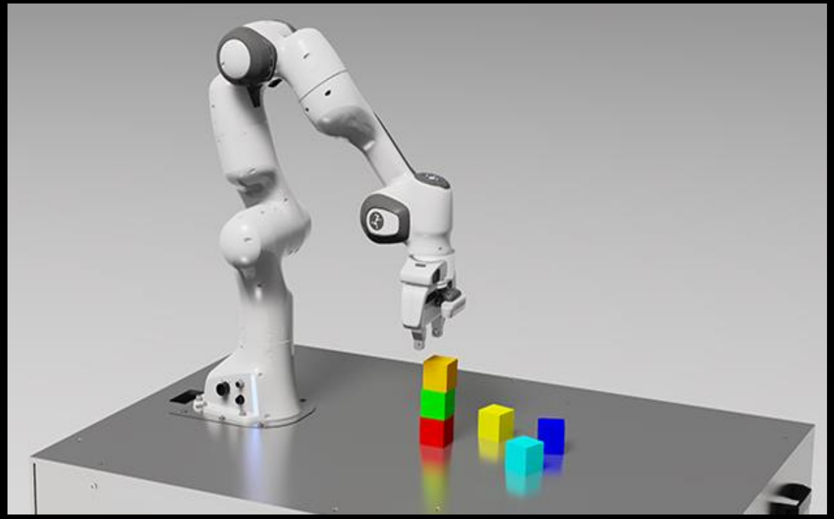

Multi Robot HIL Simulation

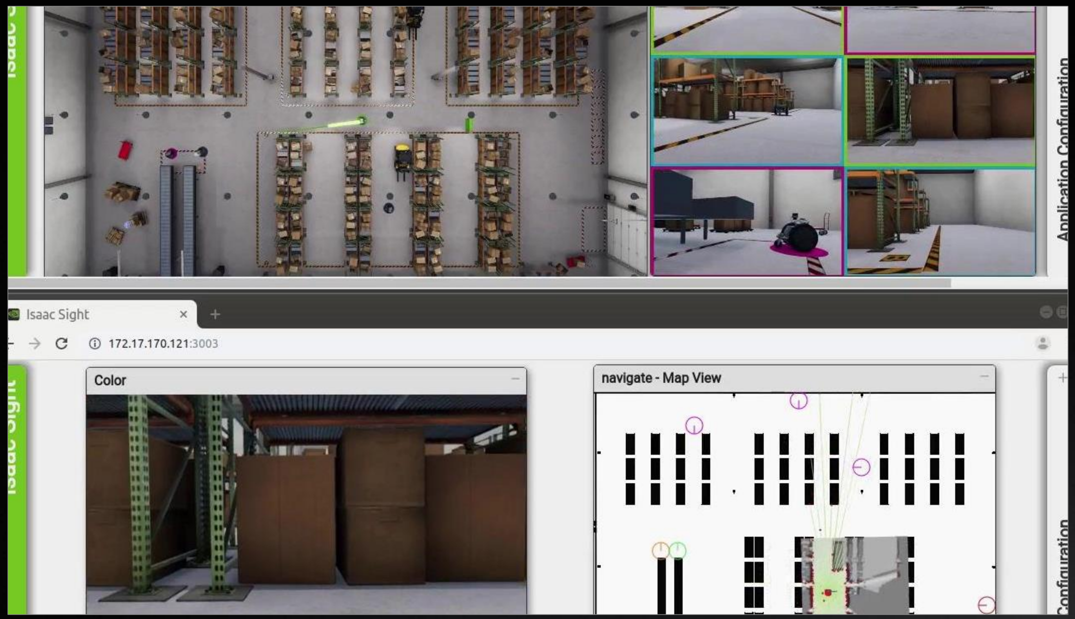

Multiple Carter robots operating simultaneously in virtual warehouse; Each operated by an independent Jetson Xavier

ML Training in Simulation

Simulated samples of a dolly (with actual CAD model) used to train object detection and pose estimation neural networks

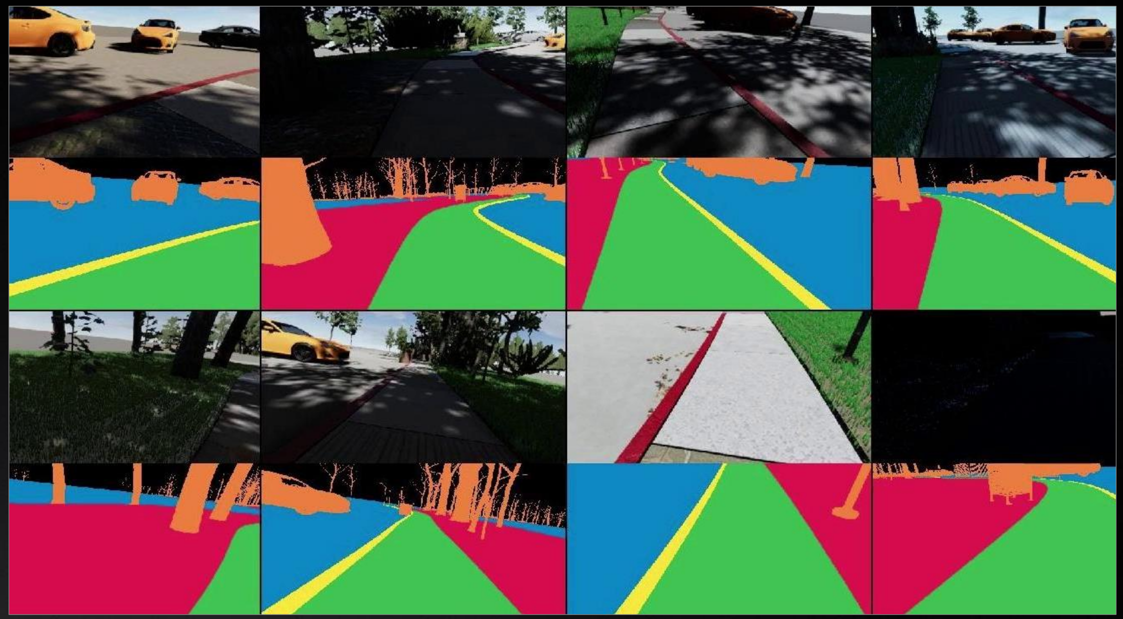

Procedurally generated simulated images used for segmentation network training


CHEEZ-IT
Domino
SUGAR
SPAM
Robotic kitchen assistant

ISAAC REFERENCE DESIGNS

Carter for Indoor Logistics; Kaya to get started; Leonardo for manipulation, more coming...

CARTER

KAYA

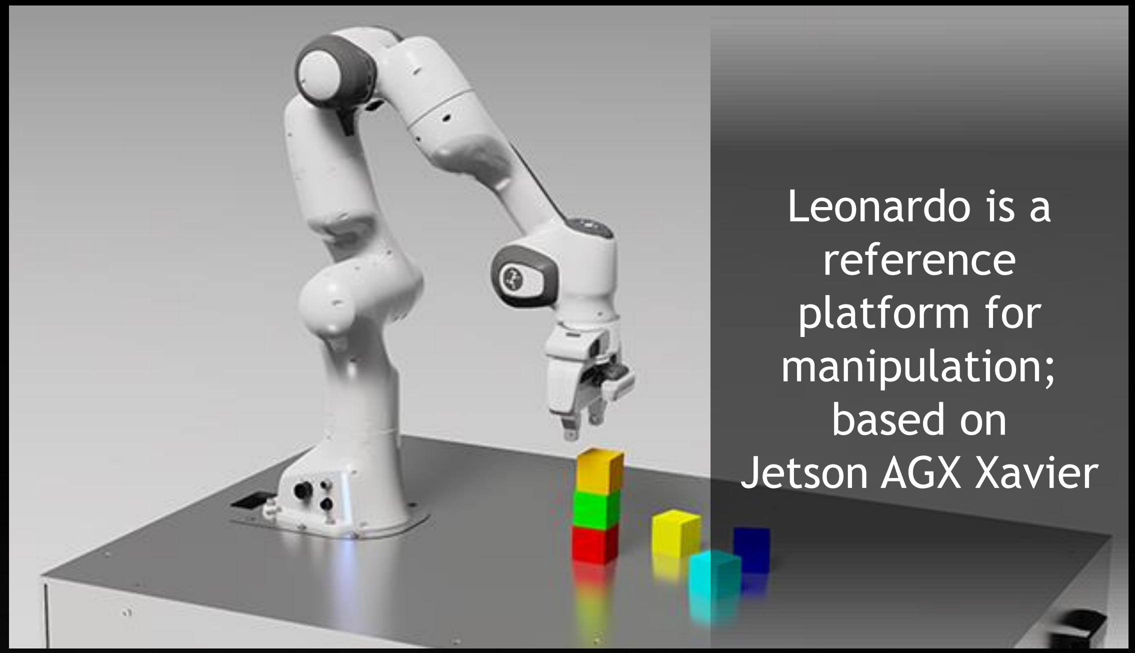

LEONARDO

ISAAC SOFTWARE

Isaac Engine

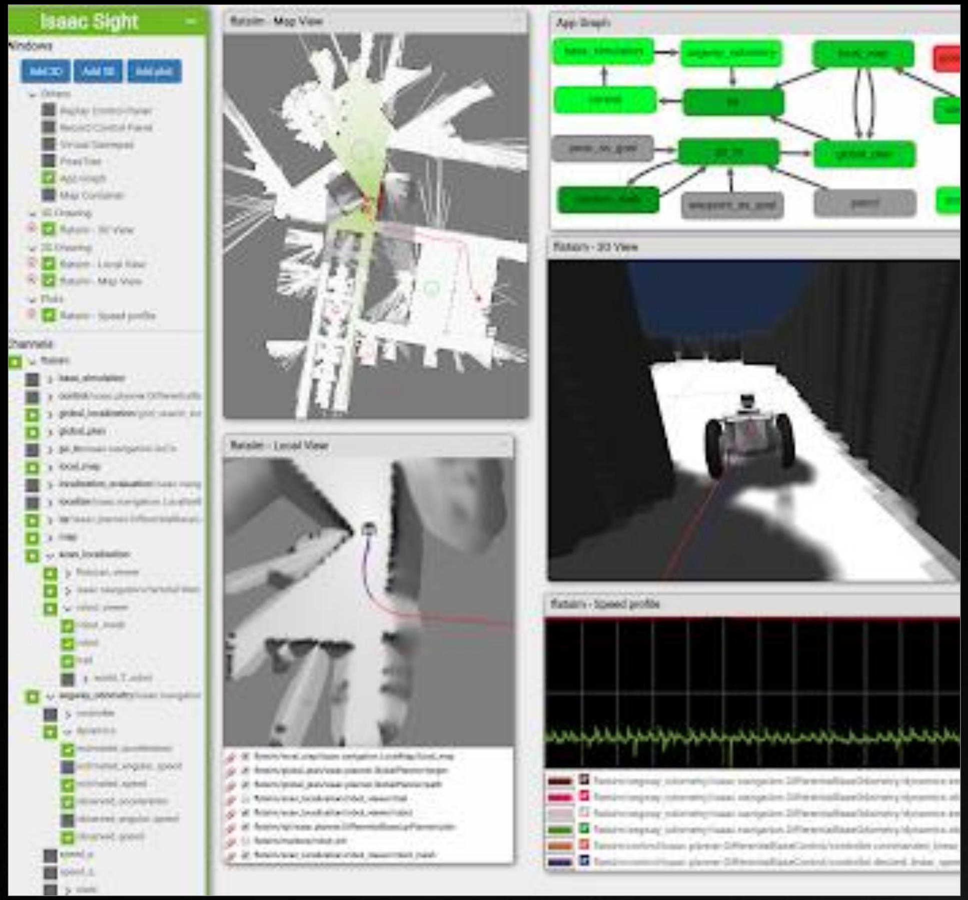

Visualization Tool

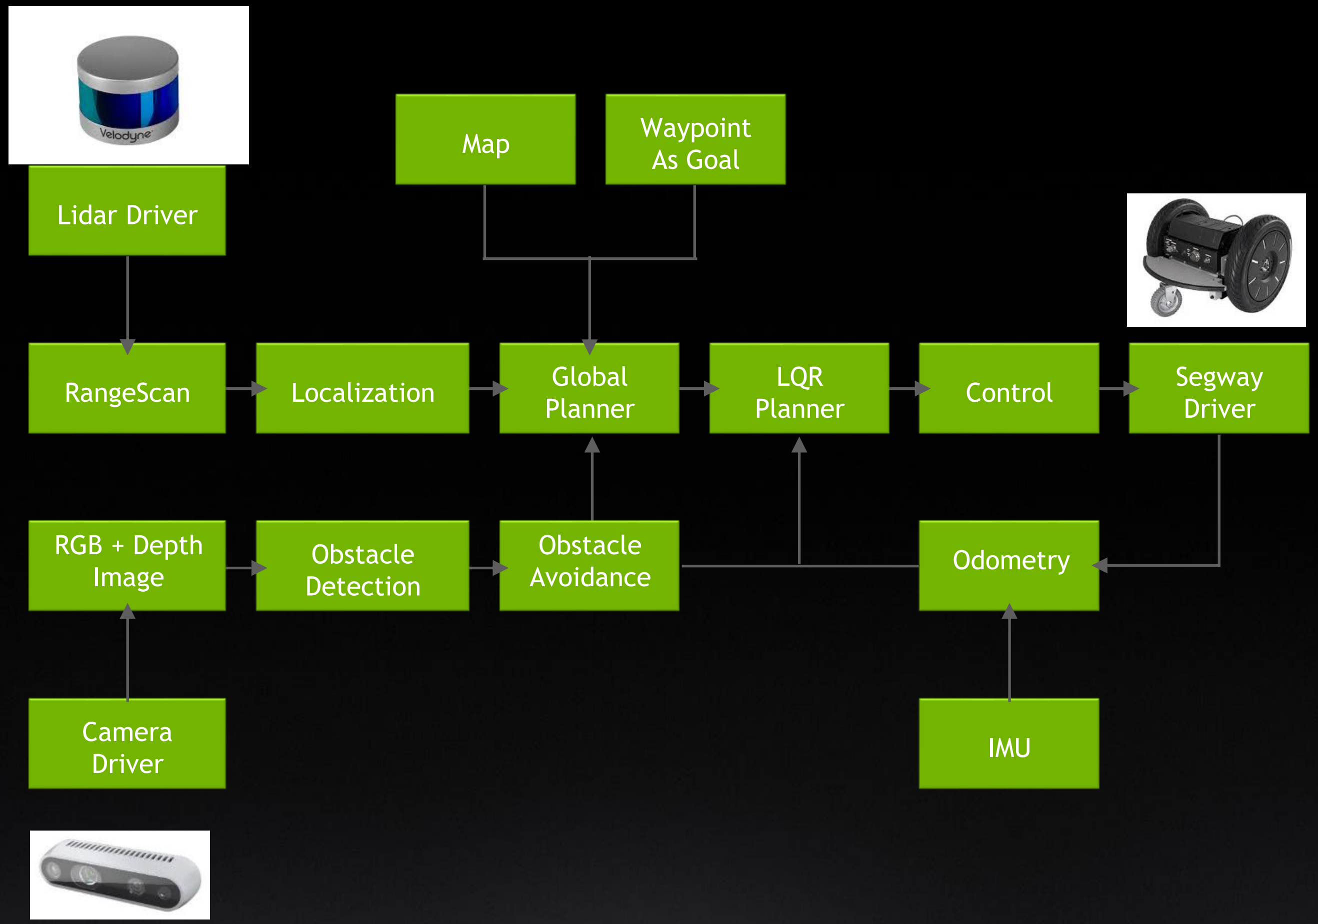

Computational Graph & CUDA Messaging

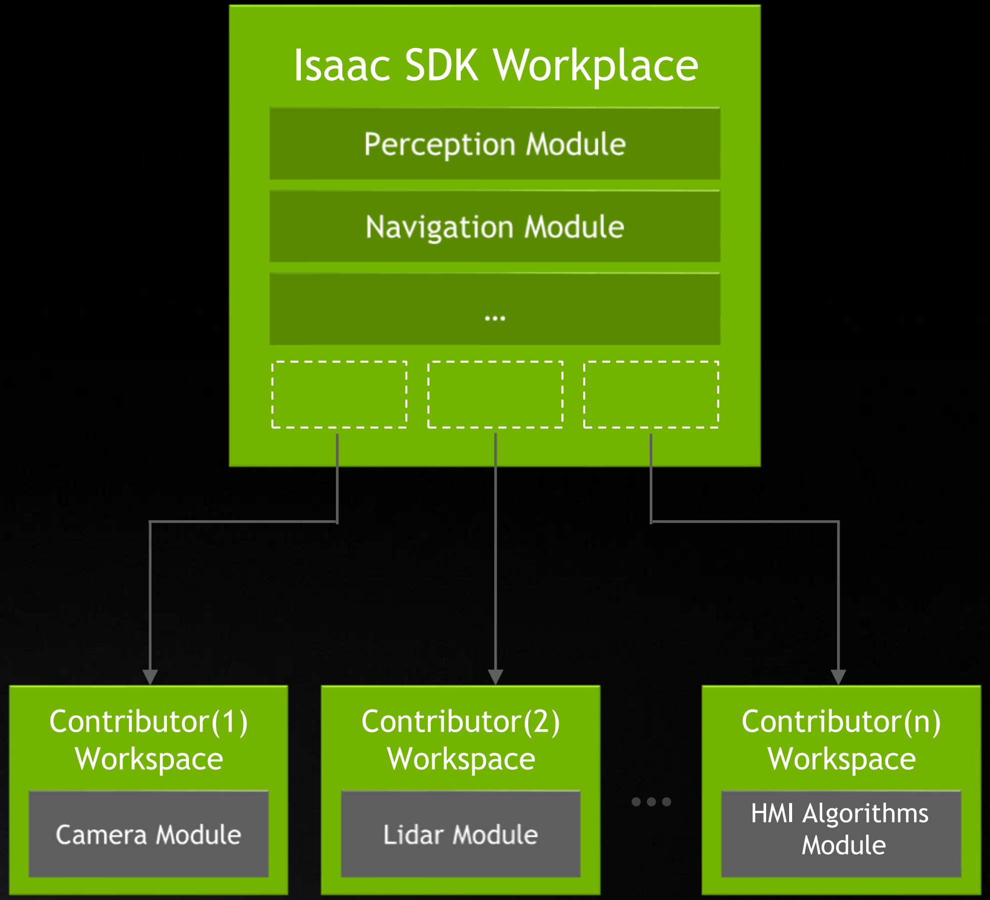

Advanced Build System & C API

ISAAC SOFTWARE

GPU Accelerated Algorithms/DNNs (GEMs)

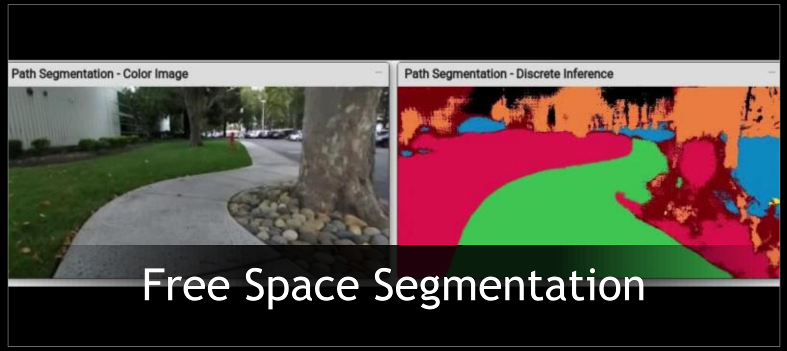

Free Space Segmentation

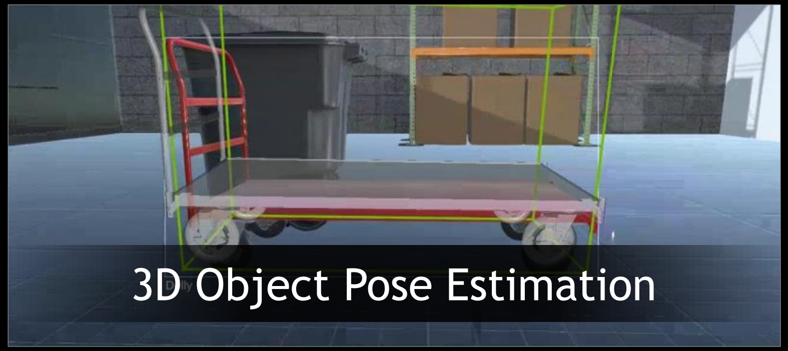
3D Object Pose Estimation


Object Detection

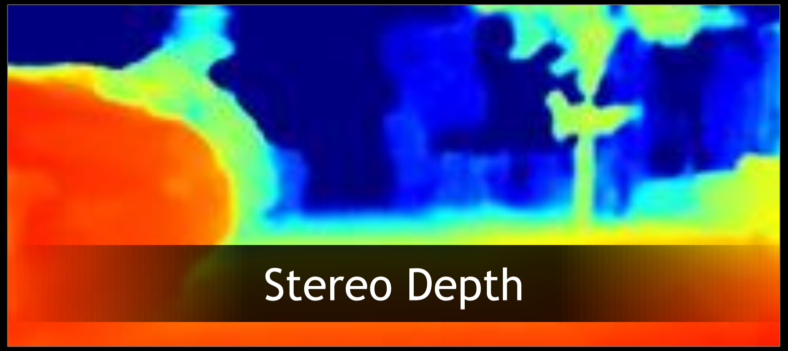
Stereo Depth


Stereo Visual Inertial Odometry

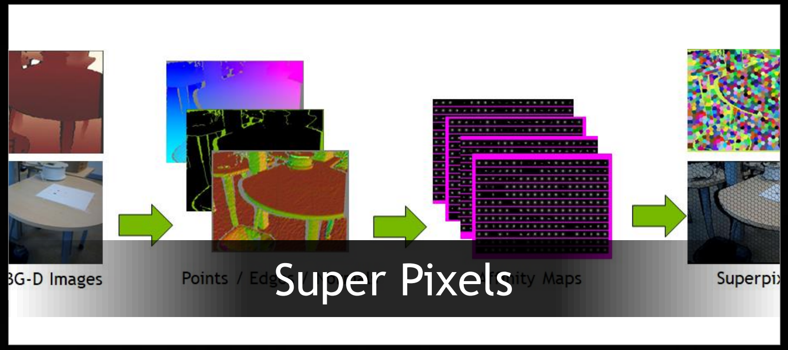
Super Pixels


April Tags


2D Skeleton Pose Estimation

DeepStream Integration

ORB Feature Tracker

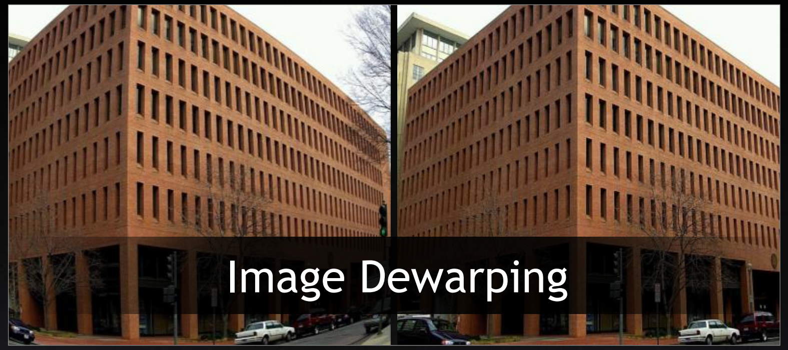
Image Dewarping


Navigation (LQR Path Planner)


Sensors

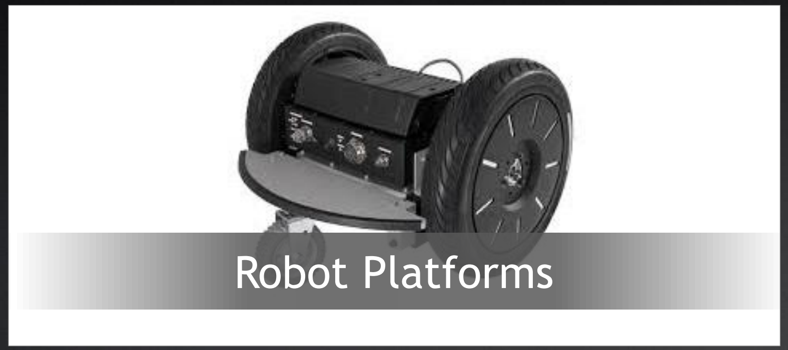
Robot Platforms

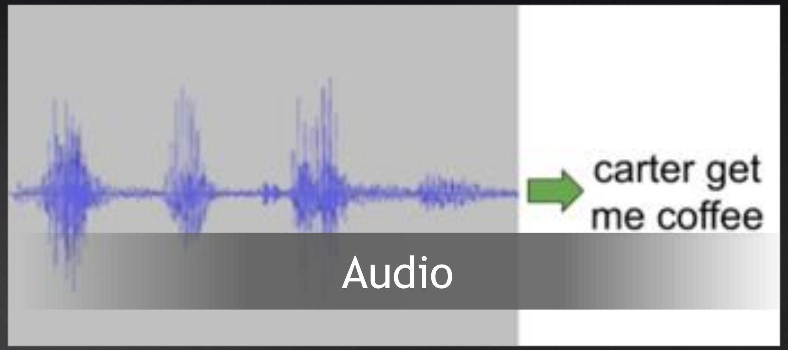

Audio

And more…

JETSON EGX SERVER

March 2020

CLOUD NATIVE SUPPORT ON JETSON

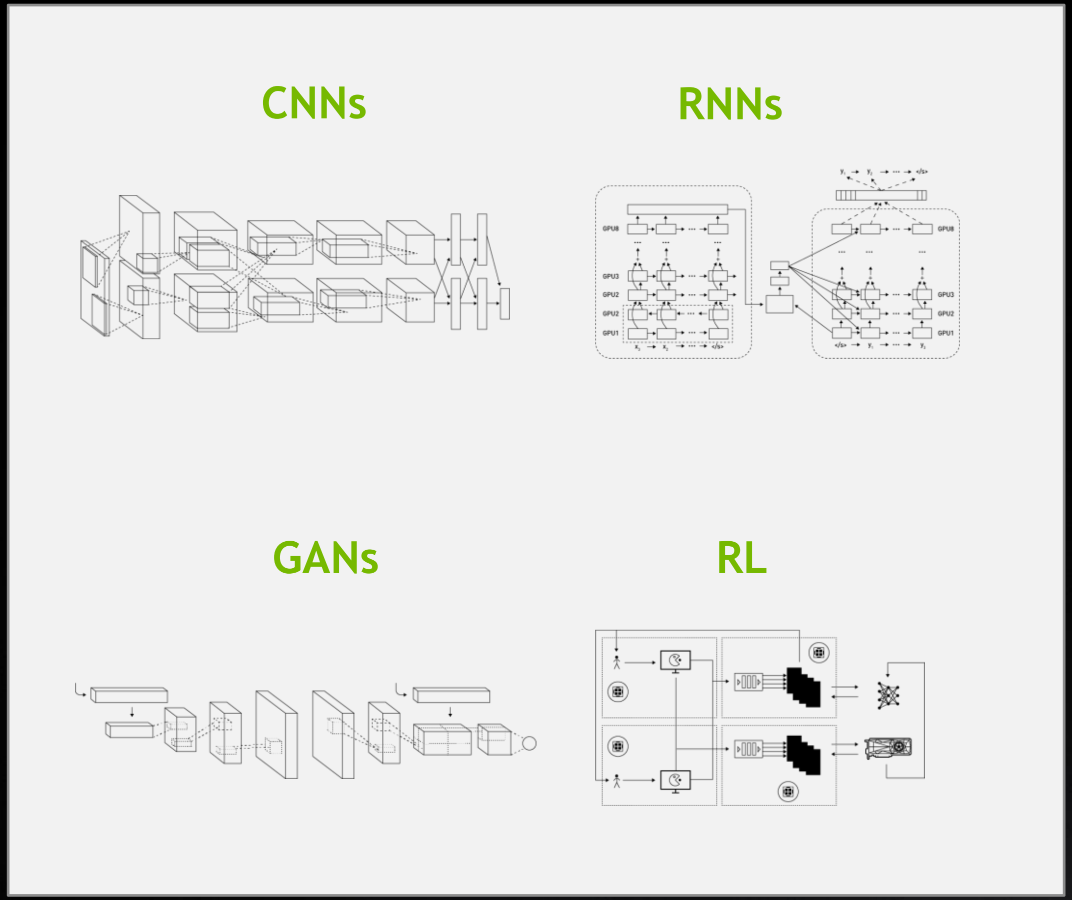

AI Inference
Rising Complexity & Diversity

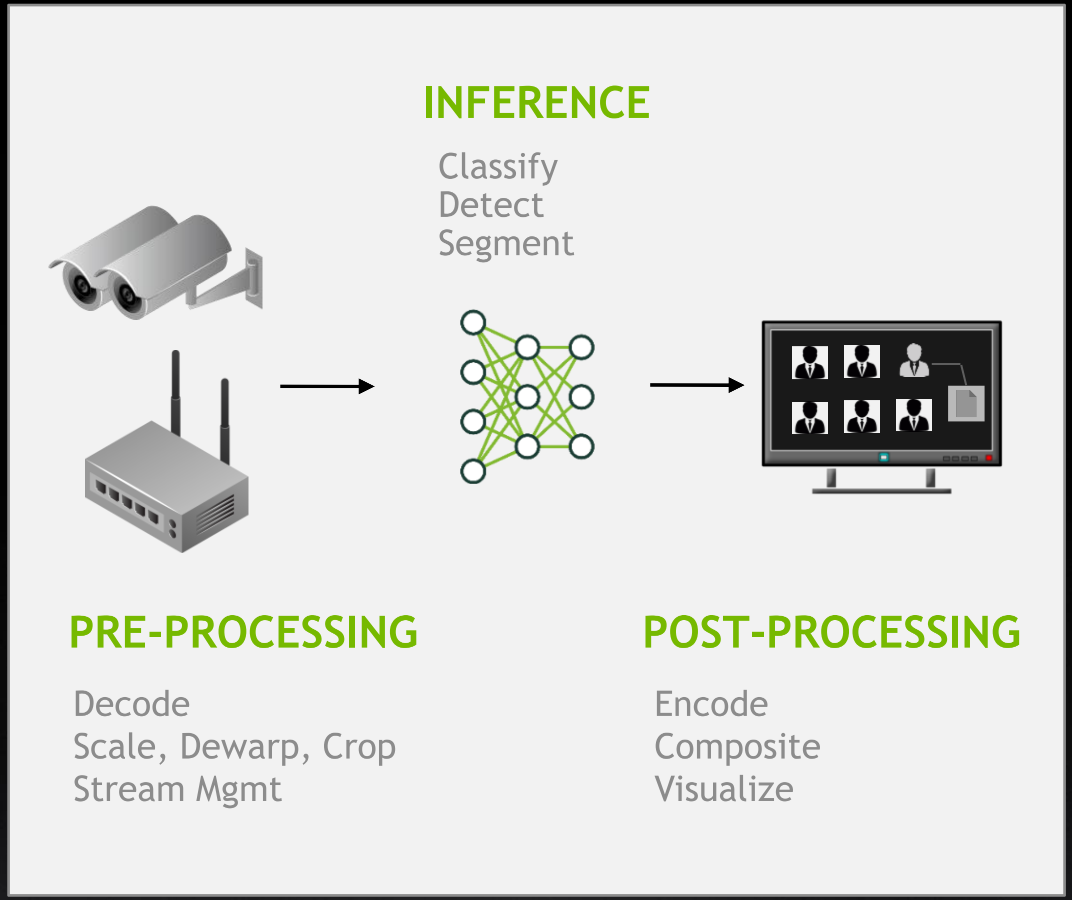

AI Applications
Accelerate the Full Stack

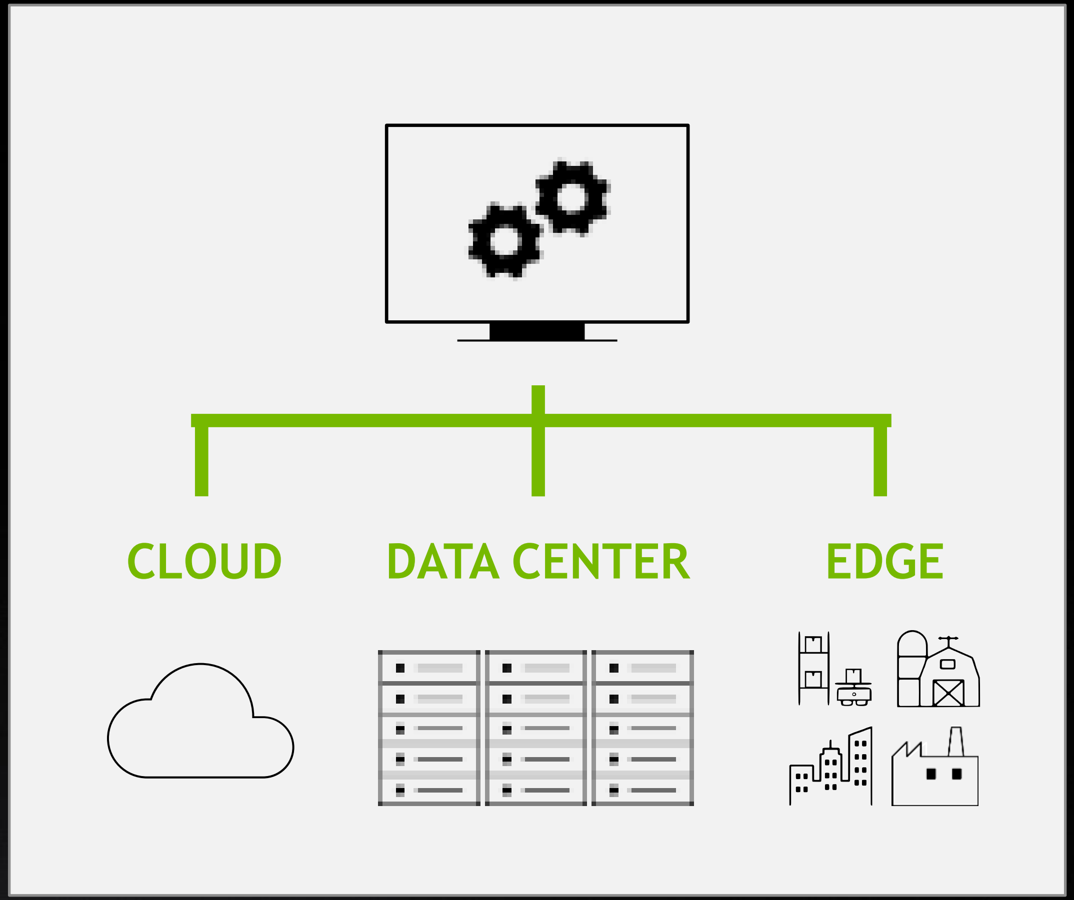

AI Product Lifecycle
Develop, Deploy and Manage at Scale

WHAT IS DOCKER ?

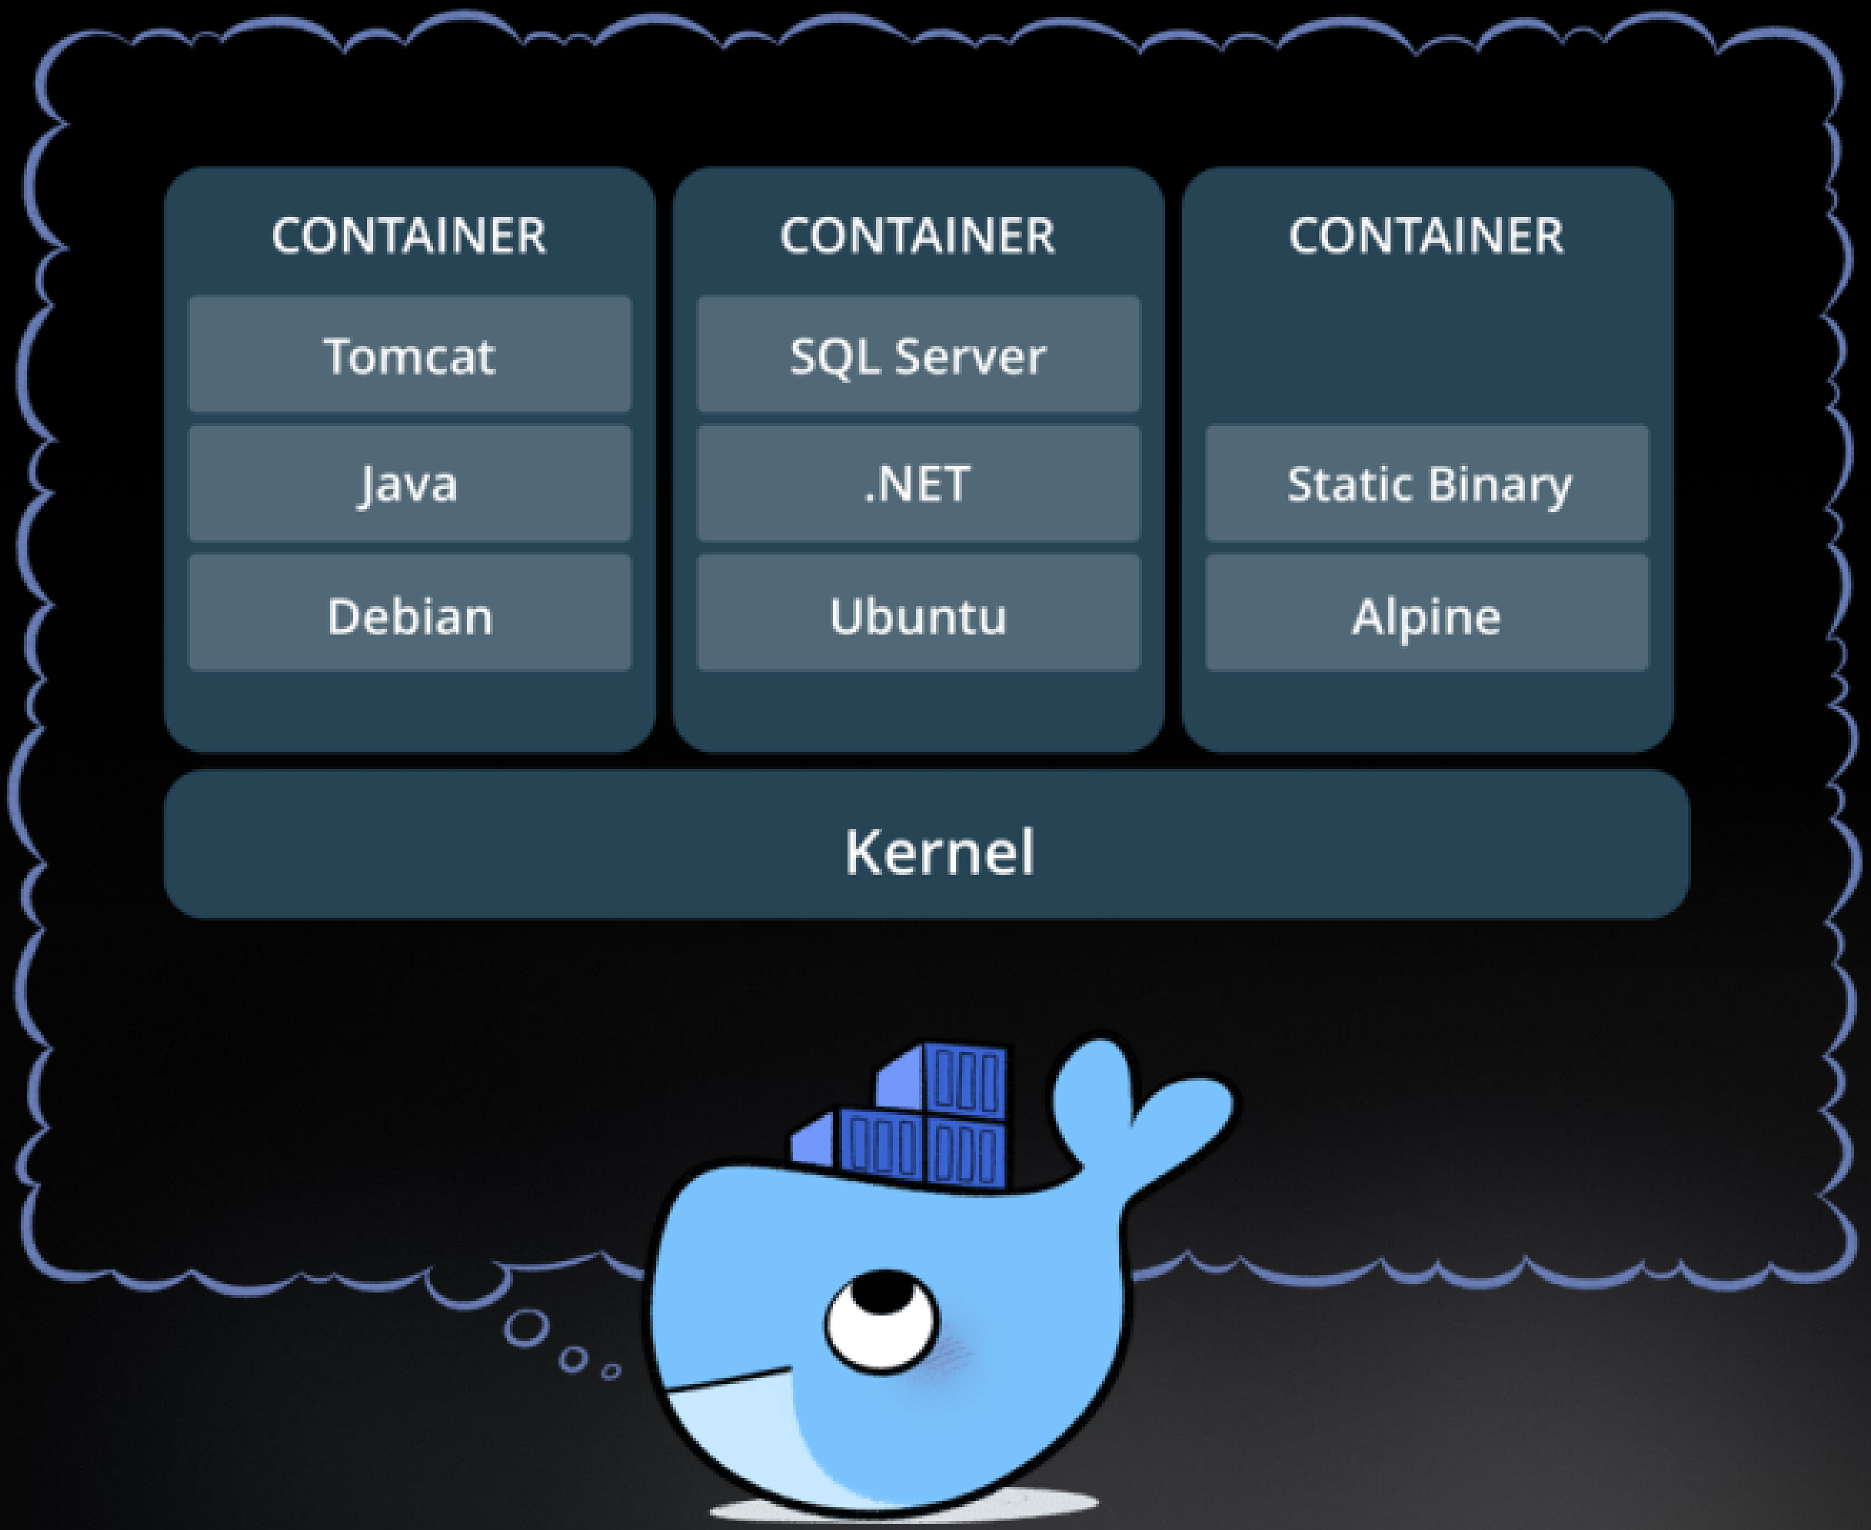

- Docker is a container runtime
- Each container can run a different function
- Each container is isolated from the others
- Containers can be started and stopped, updated, restarted independently

WHY DOCKER ON JETSON ?

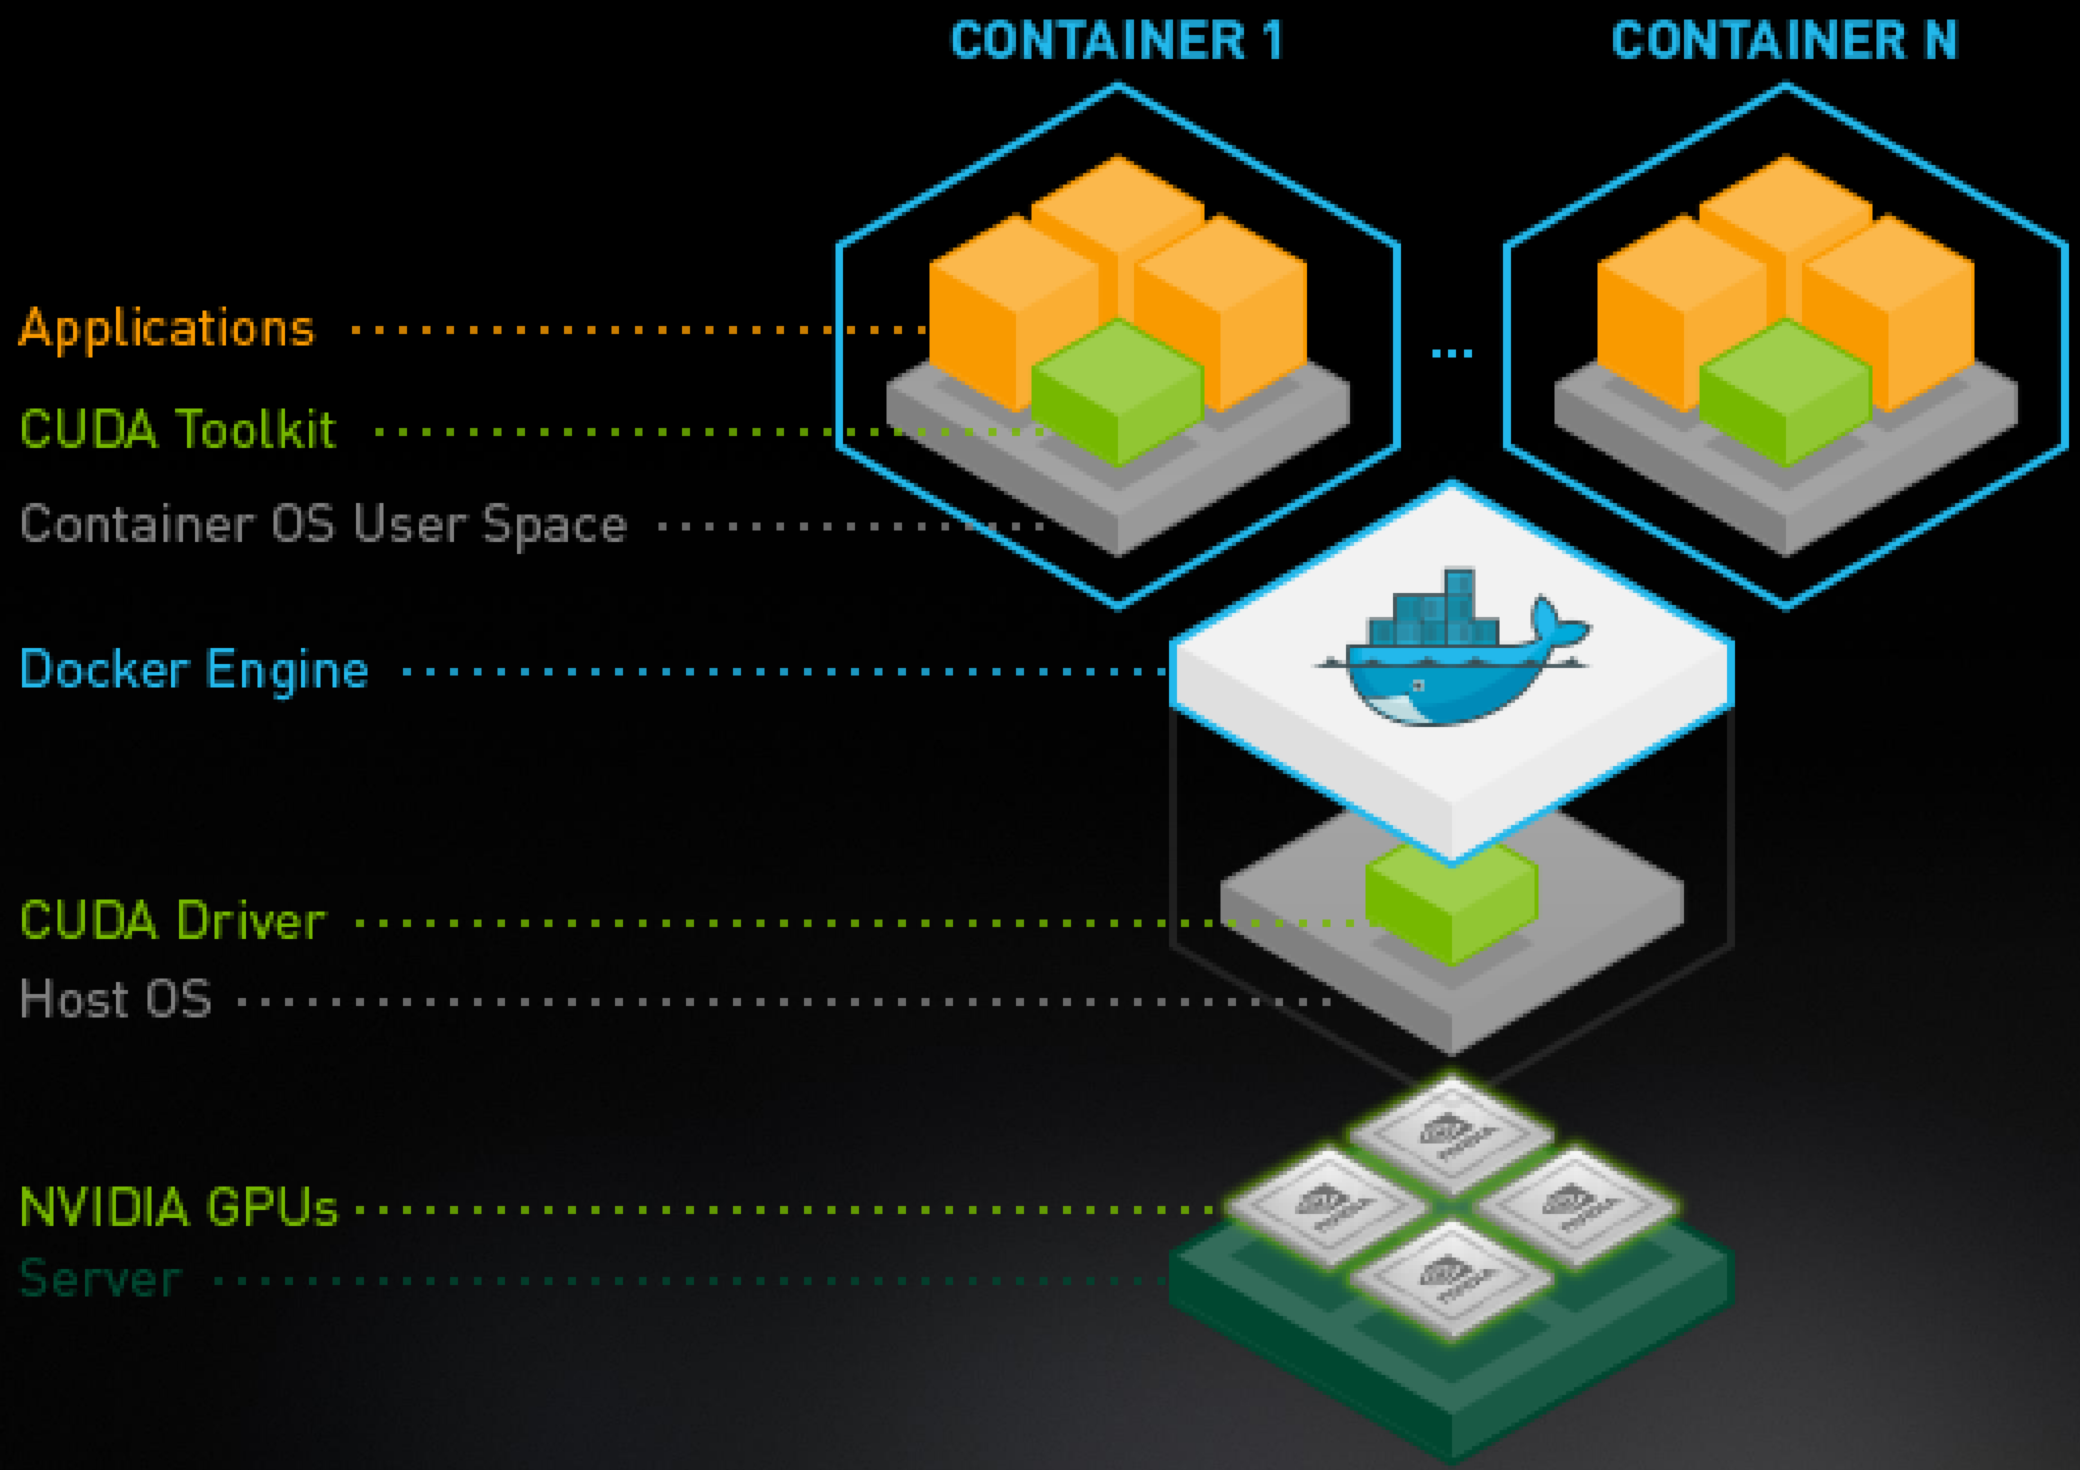

- Full support for CUDA acceleration
- Full support for Camera interface, multimedia acceleration, DLA
- Containers can be managed remotely through an Orchestrator (Kubernetes)
- Allow for easy deployment, maintenance and update

BENEFIT OF ORCHESTRATOR

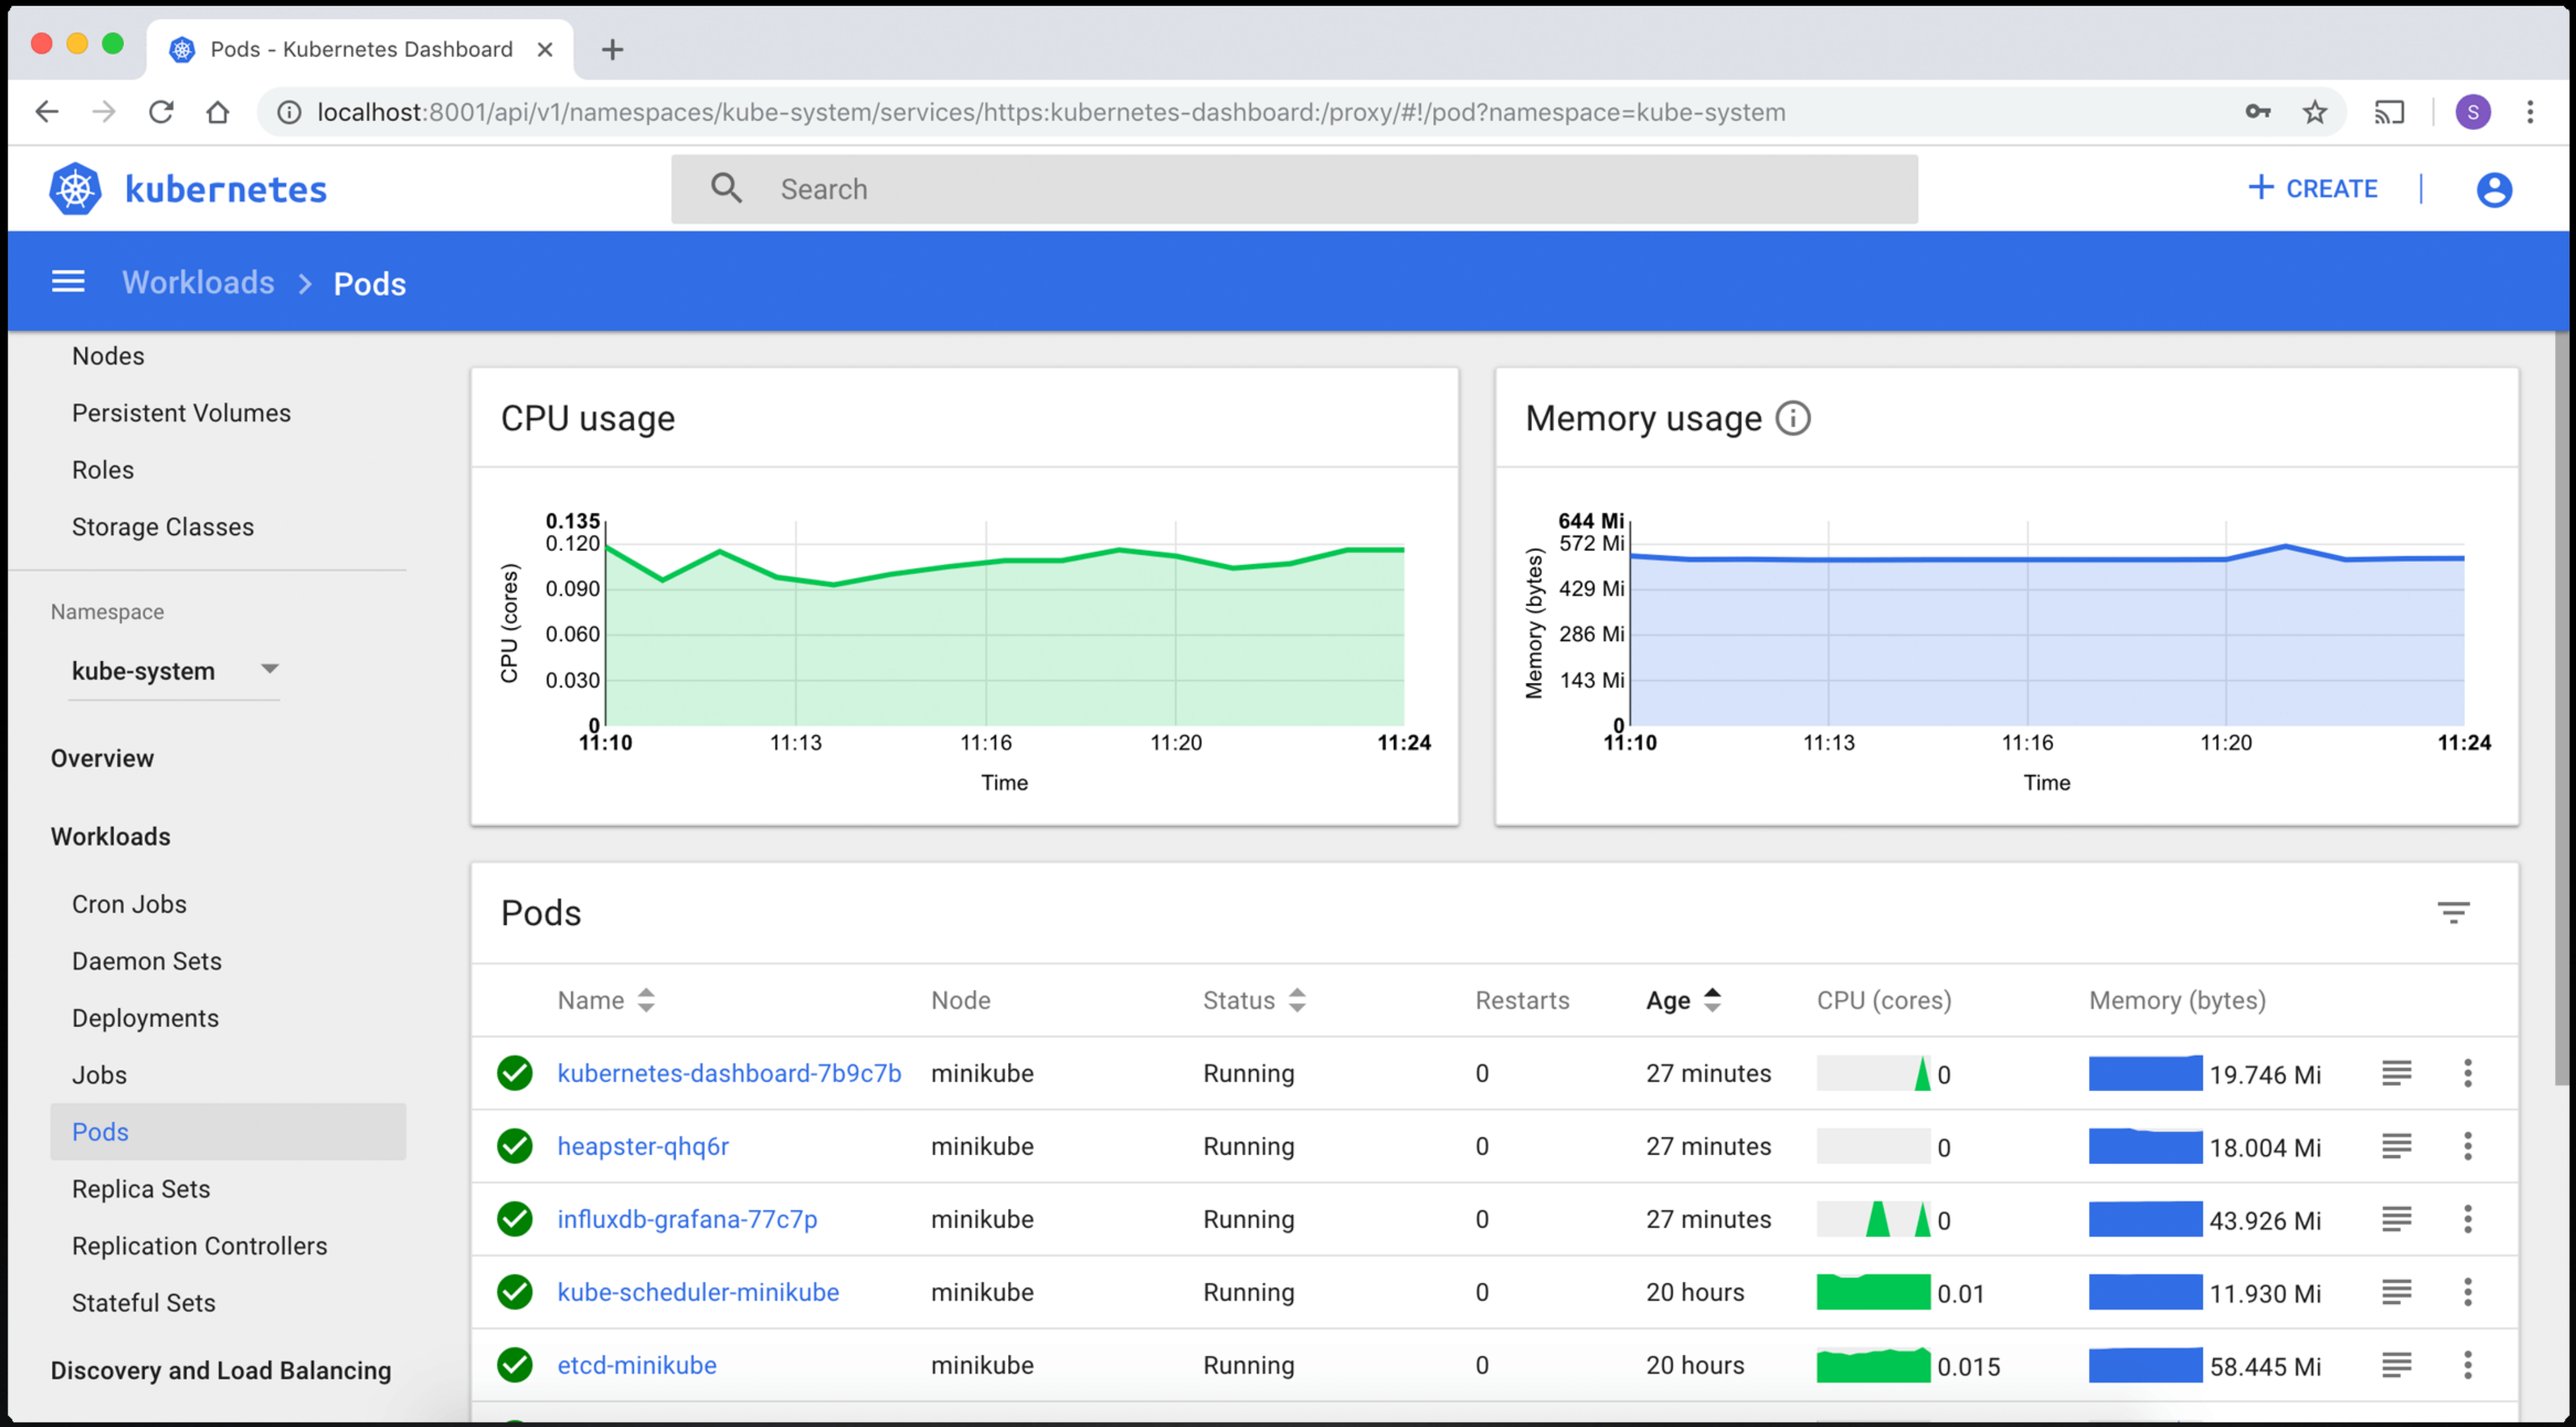

RESOURCES

JETSON DEVELOPER KIT

Start developing now

Starting at $299

developer.nvidia.com/buy-jetson

TWO DAYS TO A DEMO

Create your first demo today

developer.nvidia.com/embedded/twodaystoademo

DEEP LEARNING INSTITUTE

Training • Labs Nanodegrees

nvidia.com/DLI

GTC

Largest event for GPU developers

https://www.nvidia.com/en-us/gtc/

NVIDIA Jetson Nano Developer Kit

Bringing the Power of Modern AI to Millions of Devices

Buy Now

Meet Jetson, the Platform for AI at the Edge.

NVIDIA Jetson is the world's leading embedded AI computing platform. Its high-performance, low-power computing for deep learning and computer vision makes it possible to build software-defined autonomous machines.

The Jetson platform includes small form-factor Jetson modules with GPU-accelerated parallel processing, the JetPack SDK with developer tools and comprehensive libraries for building AI applications, along with an ecosystem of partners with services and products that accelerate development.

View Platform > View Roadmap >

BUY | DOWNLOAD | TUTORIALS | RESOURCES | FAQ | FORUM | ECOSYSTEM

NVIDIA