



Autonomous Driving Solutions


Oct, 2017 | DrivePX2 & DriveWorks

Marcus Oh (moh@nvidia.com)

Sr. Solution Architect, NVIDIA

This work is licensed under a Creative Commons Attribution-Share Alike 4.0 (CC BY-SA 4.0)
GENIVI is a registered trademark of the GENIVI Alliance in the USA and other countries.
Copyright © GENIVI Alliance 2017.

ADAS & Autonomous Driving



Camera base ADAS Features

Pedestrian Detection
Blind Spot Monitoring
Lane Departure Warning
Park assist

Collision Avoidance
Traffic Sign Recognition
Adaptive Cruise Control
Driver Monitoring



Optical Flow



Histogram



Feature Detection



Level of Autonomous driving car

Different Levels In a Self Driving Car



LEVEL 0



There are no autonomous features.

LEVEL 1



These cars can handle one task at a time, like automatic braking.

LEVEL 2



These cars would have at least two automated functions.

LEVEL 3



These cars handle "dynamic driving tasks" but might still need intervention.

LEVEL 4



These cars are officially driverless in certain environments.


LEVEL 5




These cars can operate entirely on their own without any driver presence.

Deep learning in Autonomous DRIVING



What is DEEP LEARNING?



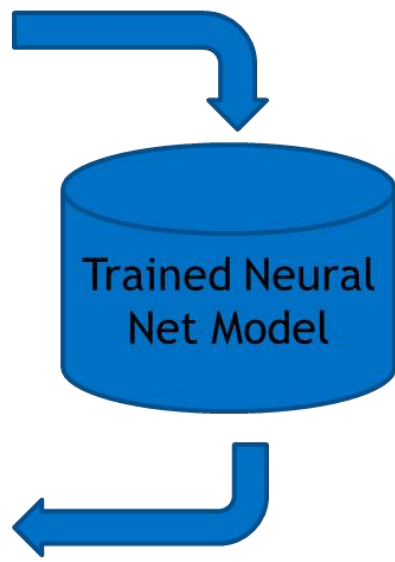

What is DEEP LEARNING?



Deep Learning Framework



Compute weight update to nudge from “turtle” towards “dog”



REINFORCEMENT learning

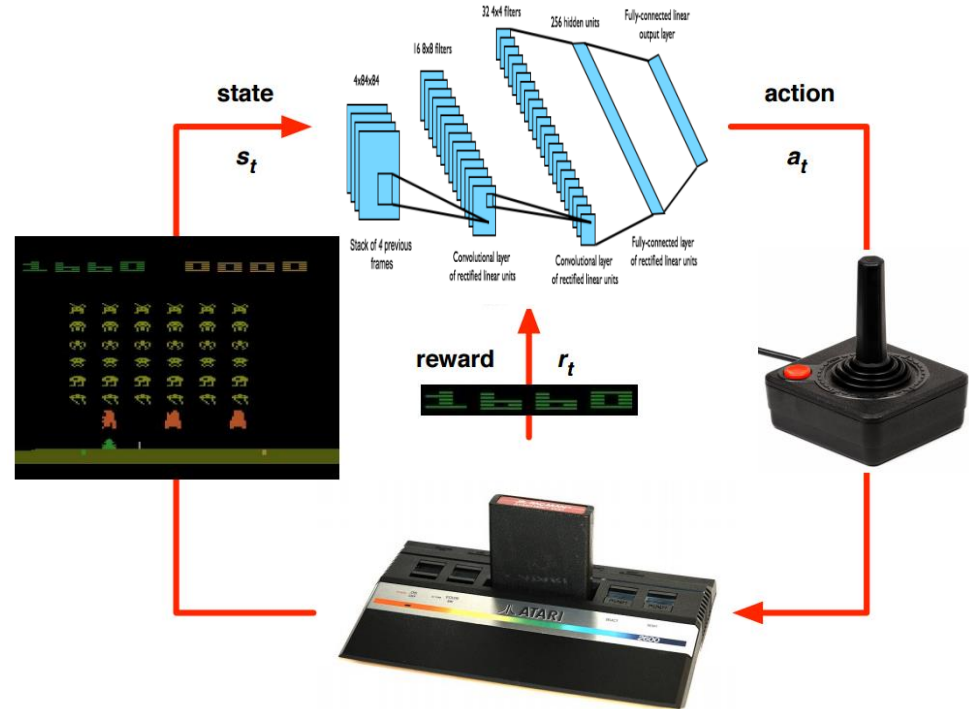
▶ A reinforcement learning agent includes:

- ▶ **state** (environment)
- ▶ **actions** (controls)
- ▶ **reward** (feedback)

▶ A **value function** predicts the future reward of performing actions in the current state

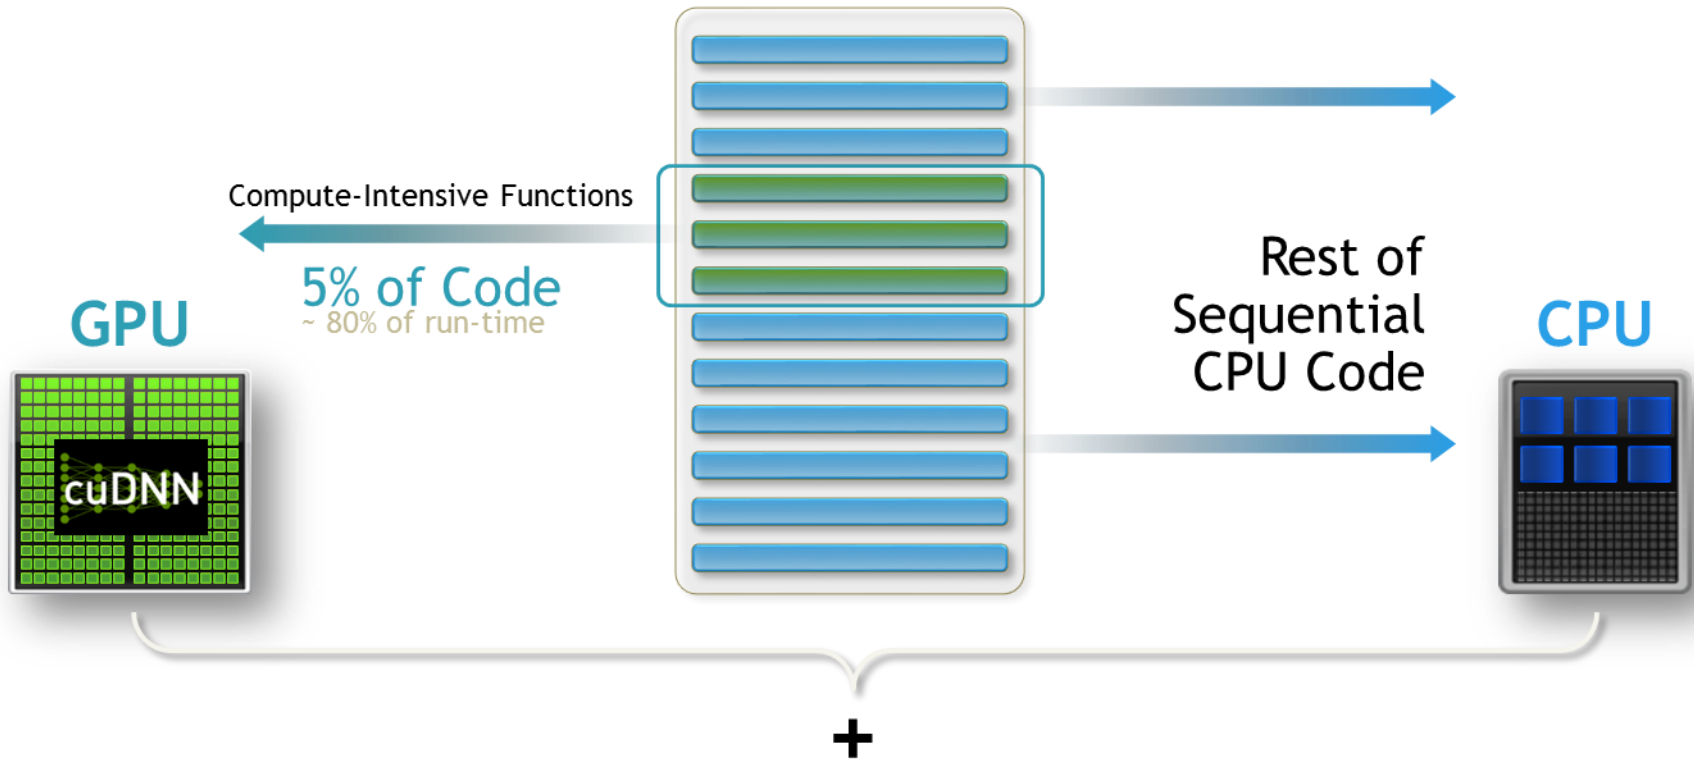
- ▶ Given the recent state, action with the maximum estimated future reward is chosen for execution

▶ For agents with complex state spaces, deep networks are used as **Q-value** approximator




github.com/dusty-nv/jetson-reinforcement


HOW GPU ACCELERATION WORKS



SELF-DRIVING CARS ARE AN AI CHALLENGE



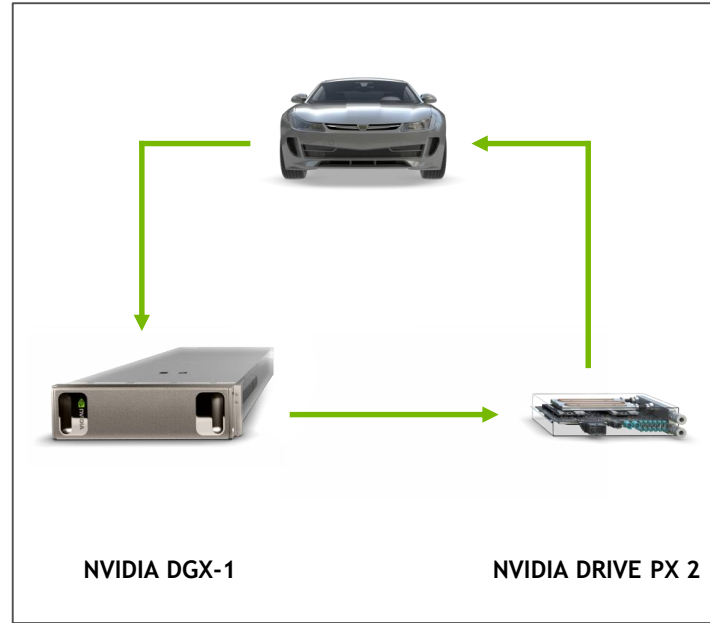
NVIDIA AI SYSTEM FOR AUTONOMOUS DRIVING


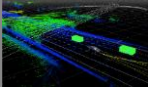
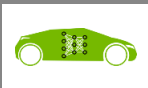


Caffe
CNTK
KALDI

TensorFlow
theano
 torch

Training on
DGX-1




MAPPING

LOCALIZATION

DRIVENET

Driving with
DriveWorks

Training Infra & Machine / DIGIT


170X SPEED-UP OVER COTS SERVER

MICROSOFT COGNITIVE TOOLKIT SUPERCHARGED ON NVIDIA DGX-1

170x Faster
(AlexNet images/sec)




8x Tesla P100 | 170TF FP16 | NVLink hybrid cube mesh



AlexNet training batch size 128, Dual Socket E5-2699v4, 44 cores CNTK 2.0b2 for CPU.
CNTK 2.0b3 (to be released) includes cuDNN 5.1.8, NCCL 1.6.1, NVLink enabled

DGX STACK

Fully integrated Analytics and Deep Learning platform



Instant productivity – plug-and-play, supports every AI framework and accelerated analytics software applications

Performance optimized across the entire stack

Always up-to-date via the cloud

Mixed framework environments – baremetal and containerized

Direct access to NVIDIA experts



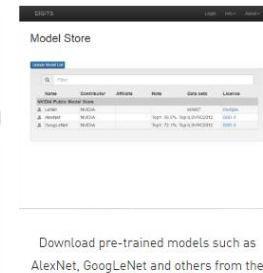
DIGIT

Deep Learning GPU Training System

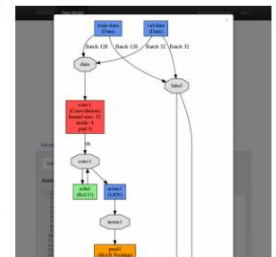
- ▶ Design, train and visualize deep neural networks for image classification, segmentation and object detection
- ▶ Download pre-trained models such as AlexNet, GoogLeNet and LeNet from the DIGITS Model Store
- ▶ Perform hyperparameter sweep of learning rate and batch size for improved model accuracy
- ▶ Schedule, monitor, and manage neural network training jobs, and analyze accuracy and loss in real time
- ▶ Import a wide variety of image formats and sources
- ▶ Scale training jobs across multiple GPUs automatically



Import data for image classification and object detection neural networks



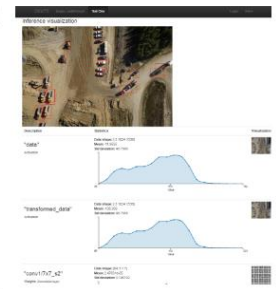
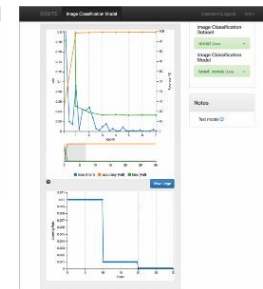
Download pre-trained models such as AlexNet, GoogLeNet and others from the DIGITS Model Store



Visualize deep neural network architectures




Schedule, monitor, and manage neural network training jobs



NVIDIA SELF-DRIVING BUILDING BLOCKS

Accelerate your development and deployment

SELF DRIVING TECHNOLOGY



SOFTWARE

PDK



cuDNN




TensorRT



VPI




DriveWorks



TOOLS

CUDA Development
Nsight Eclipse Edition



Tegra System Profiler



HARDWARE


DRIVE PX 2



DRIVE PX 2 AutoCruise



Tegra SOC, dGPU




DRIVE PX2

The image shows a close-up, angled view of the NVIDIA DRIVE PX 2 Autocruise module. It is a compact, square-shaped printed circuit board (PCB) populated with various electronic components. A prominent feature is a large, square, silver-colored integrated circuit (IC) in the center, which is the Tegra Parker SoC, bearing the NVIDIA logo. To the left of the main board, there is a separate, dark grey, rectangular component that appears to be a connector or a protective housing for the board's ports. Two bright blue connectors are visible on the bottom edge of the board. The entire assembly is set against a dark, almost black background, which makes the metallic and blue components stand out.

NVIDIA DRIVE PX 2 AUTOCRUISE

10W AI Car Computer | Passive Cooling | Automotive IO
Multiple Cameras & Sensors | DriveWorks SW/SDK
AI Highway Driving | Localization & HD Mapping
Tegra Parker SoC – 1.3 TFLOPS, 6 CPU Cores, Integrated ISP




NVIDIA DRIVE PX 2

AUTOCHAUFFEUR & FULLY AUTONOMOUS

Scalable from 1 to 4 Processors to Multiple DRIVE PX 2s
– 2x Tegra Parker SoC, 2x Pascal dGPU, 8 TFLOPS, 24 DNN TOPs
Up to 12 Cameras; plus LIDAR, Radar, Ultrasonic sensors
DriveWorks SW/SDK | AI Perception | Localization & Mapping



DriveWorks



DRIVEWORKS

Software Development Kit (SDK) for Autonomous Driving

Process sensor data through Perception, Mapping, Localization, and Path Planning steps

Provides a rich set of functionalities:

- Sensor Abstraction Layer (SAL)

- Algorithm Modules, DNNs

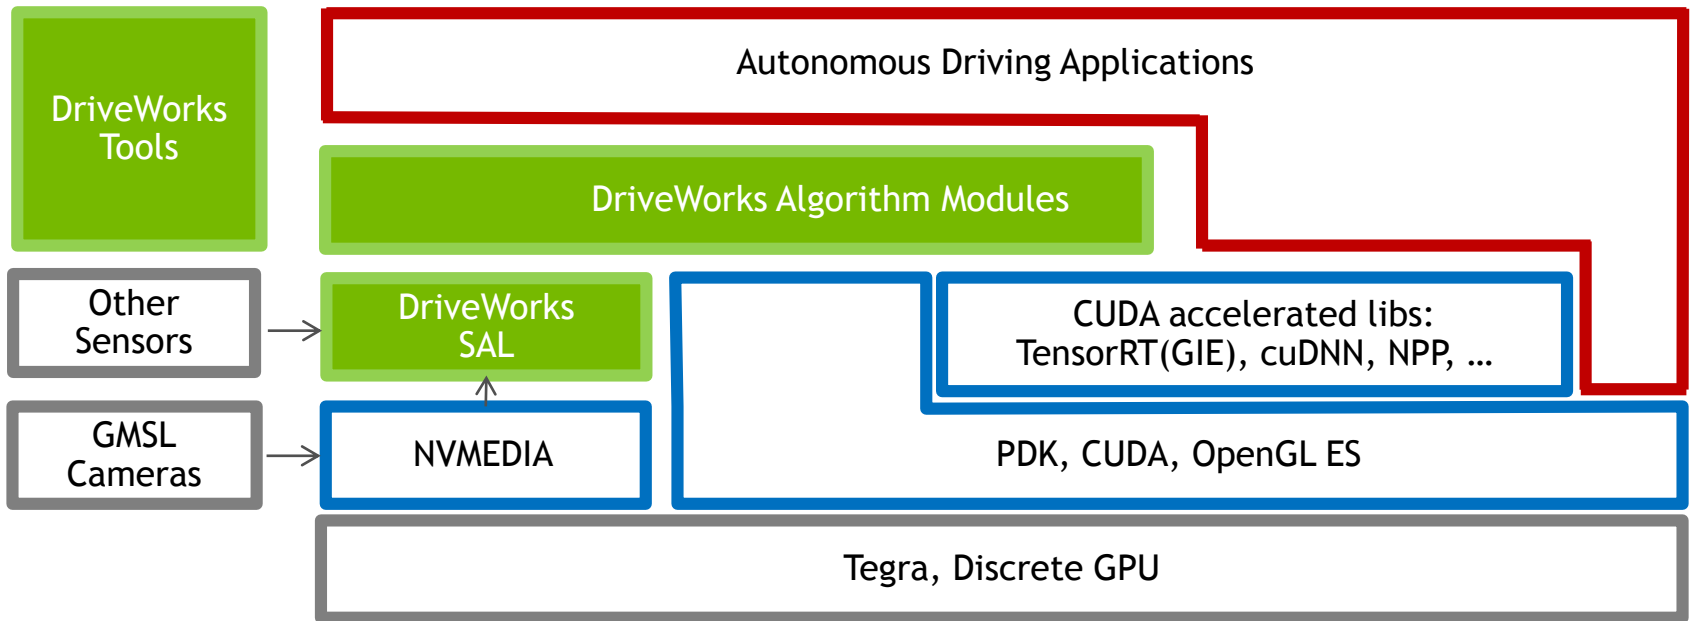
- Applications

- Tools for sensor setup and management

Flexible, modular and optimized for GPU

Runs on top of OS, CUDA/CuDNN, TensorRT, VPI

SOFTWARE STACK WITH DRIVEWORKS



HW

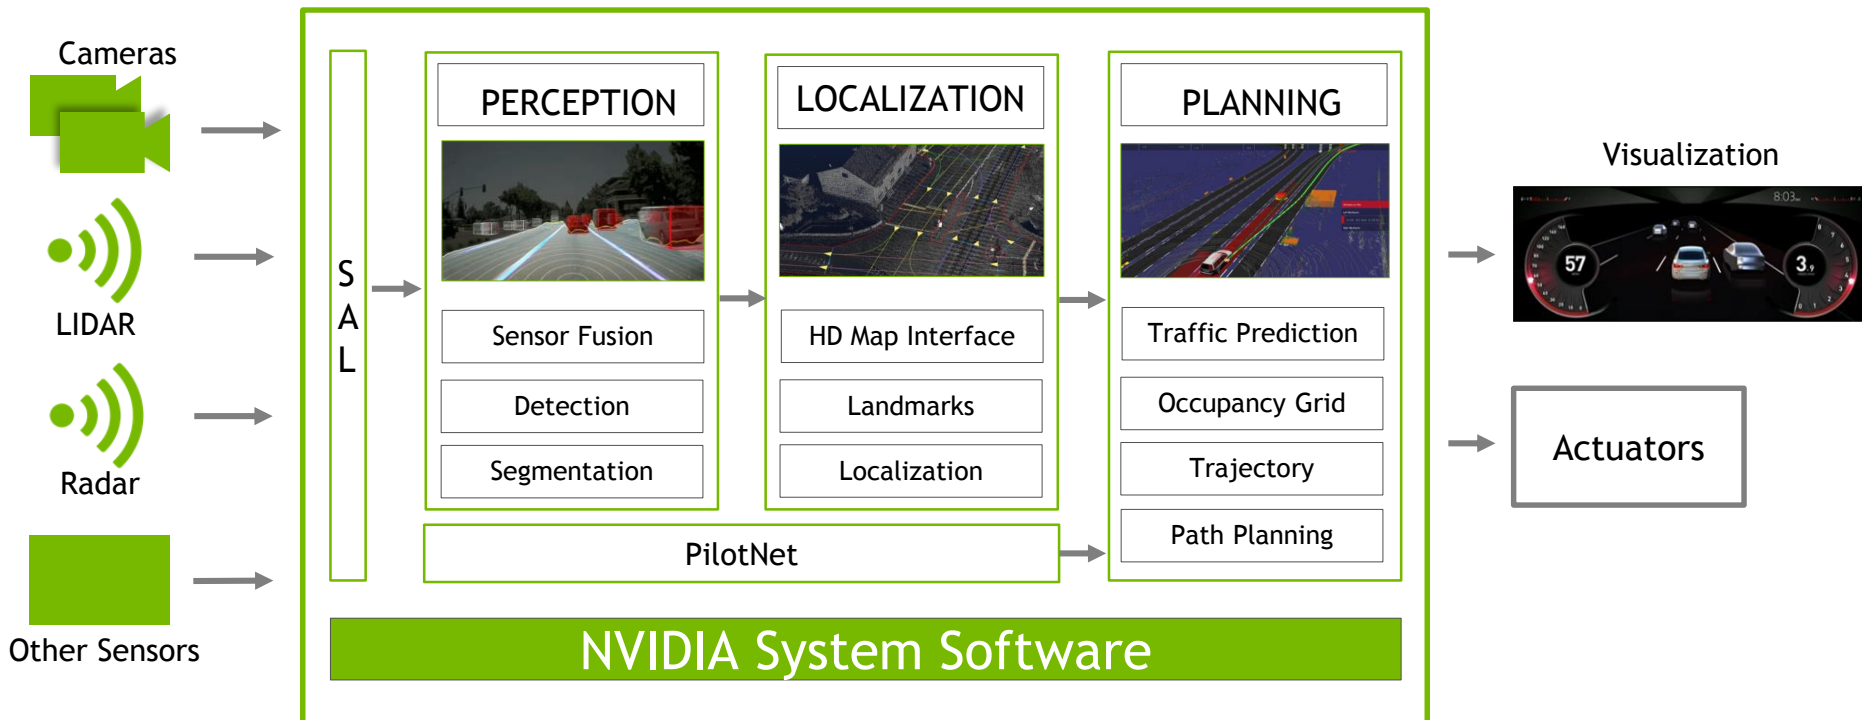
Linux PDK

DriveWorks

OEM/Sample Applications

DRIVEWORKS PROCESSING PIPELINE

End-to-End processing for Autonomous Driving



DEEP NEURAL NETWORKS IN DRIVEWORKS

DriveNet



Multi-class detection: Cars, Pedestrian, Bicycles
Upcoming: Lanes, Traffic signs

LaneNet



Lane detection

OpenRoadNet



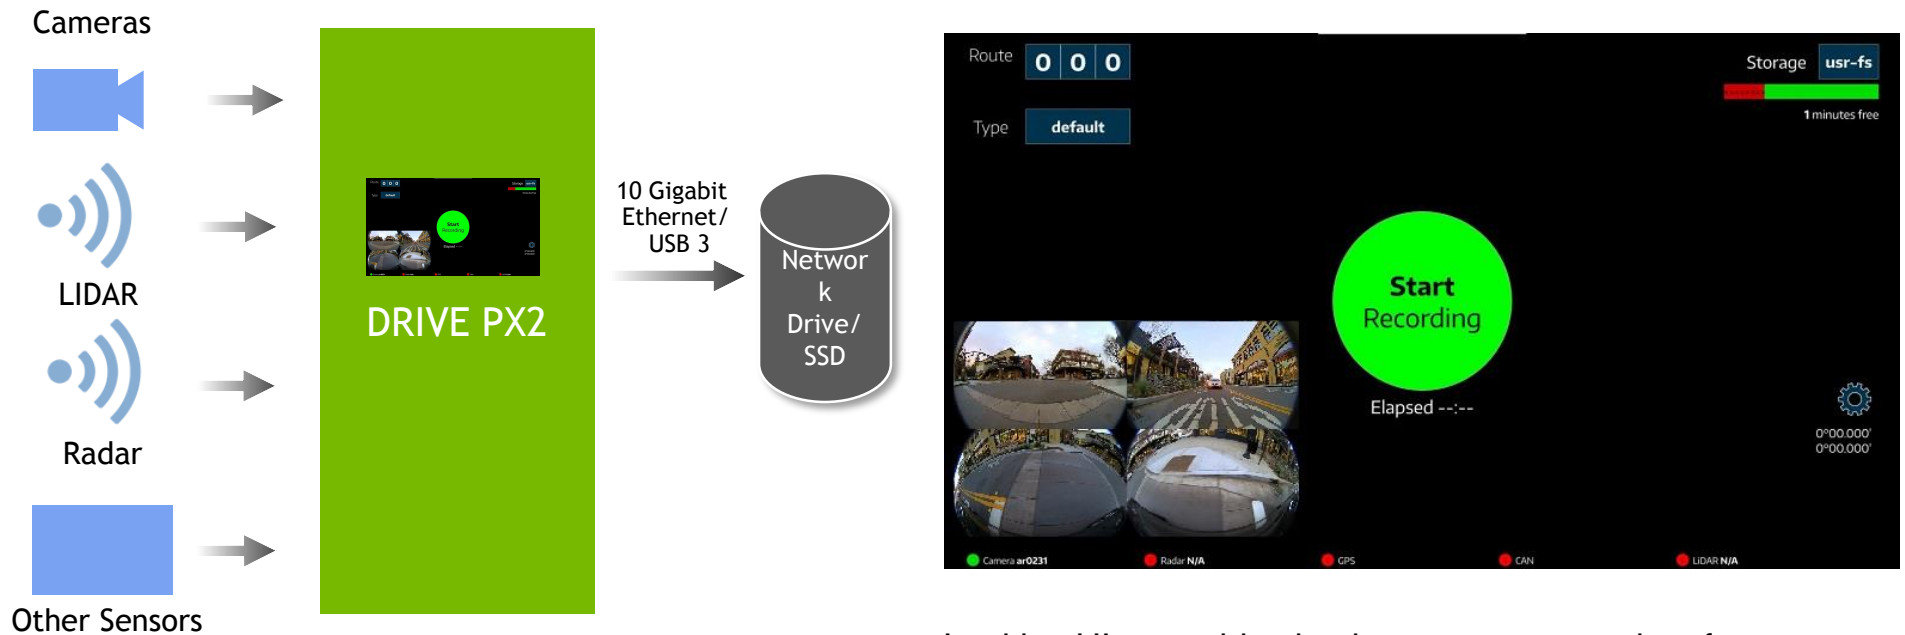
Free space detection

PilotNet



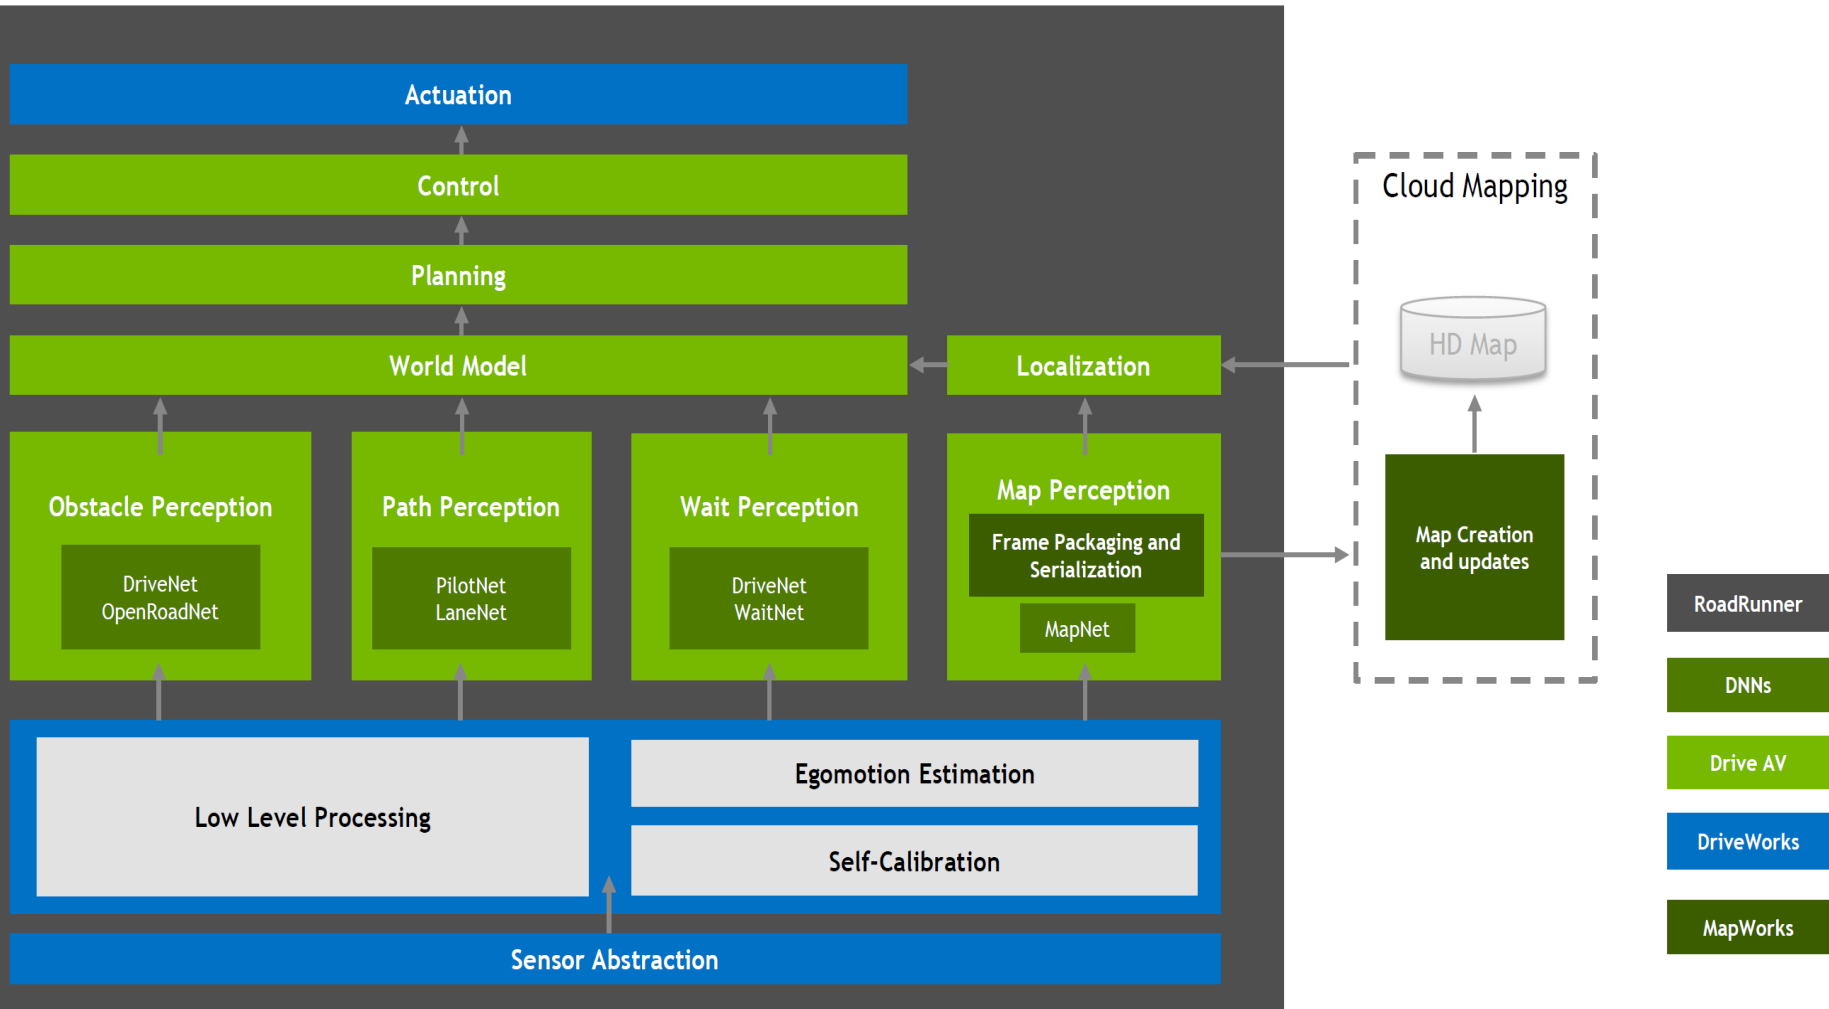
End to end in-lane driving

TOOLS - DATA ACQUISITION



Intuitive UI to enable simultaneous capture data from sensors

HIGH LEVEL ARCHITECTURE




NEXT GENERATION AD SOLUTION



INTRODUCING XAVIER

AI SUPERCOMPUTER SOC




- 7 Billion Transistors 16nm FF
- 8 Core Custom ARM64 CPU
- 512 Core Volta GPU
- New Computer Vision Accelerator
- Dual 8K HDR Video Processors
- Designed for ASIL C Functional Safety

THE NEXT GENERATION

DRIVE PX 2 IS THE TIME MACHINE TO XAVIER

Today

PASCAL GENERATION



2x PARKER + 2x PASCAL

Next

VOLTA GENERATION




XAVIER



INTRODUCING XAVIER

AI SUPERCOMPUTER SOC



DRIVE PX 2

2 PARKER + 2 PASCAL GPU | 20 TOPS DL | 120 SPECINT | 80W

XAVIER

20 TOPS DL | 160 SPECINT | 20W



Thank you!

Visit GENIVI at <http://www.genivi.org> or <http://projects.genivi.org>

Contact us: help@genivi.org

This work is licensed under a Creative Commons Attribution-Share Alike 4.0 (CC BY-SA 4.0)
GENIVI is a registered trademark of the GENIVI Alliance in the USA and other countries.
Copyright © GENIVI Alliance 2017.

