



# JETSON AGX XAVIER AND THE NEW ERA OF AUTONOMOUS MACHINES



# WEBINAR AGENDA

## Intro to Jetson AGX Xavier

- AI for Autonomous Machines
- Jetson AGX Xavier Compute Module
- Jetson AGX Xavier Developer Kit

## Xavier Architecture

- Volta GPU
- Deep Learning Accelerator (DLA)
- Carmel ARM CPU
- Vision Accelerator (VA)

## Jetson SDKs

- JetPack 4.1
- DeepStream SDK
- ISAAC SDK

## Resources & Support

- Developer Site & Documentation
- Community Forums & Wiki
- Tutorials
- Quick Start Platforms

# BILLIONS OF AUTONOMOUS MACHINES



Industrial



Aerospace/Defense




Healthcare



Construction



Agriculture



Smart City



Retail



Logistics



Inventory Mgmt



Delivery

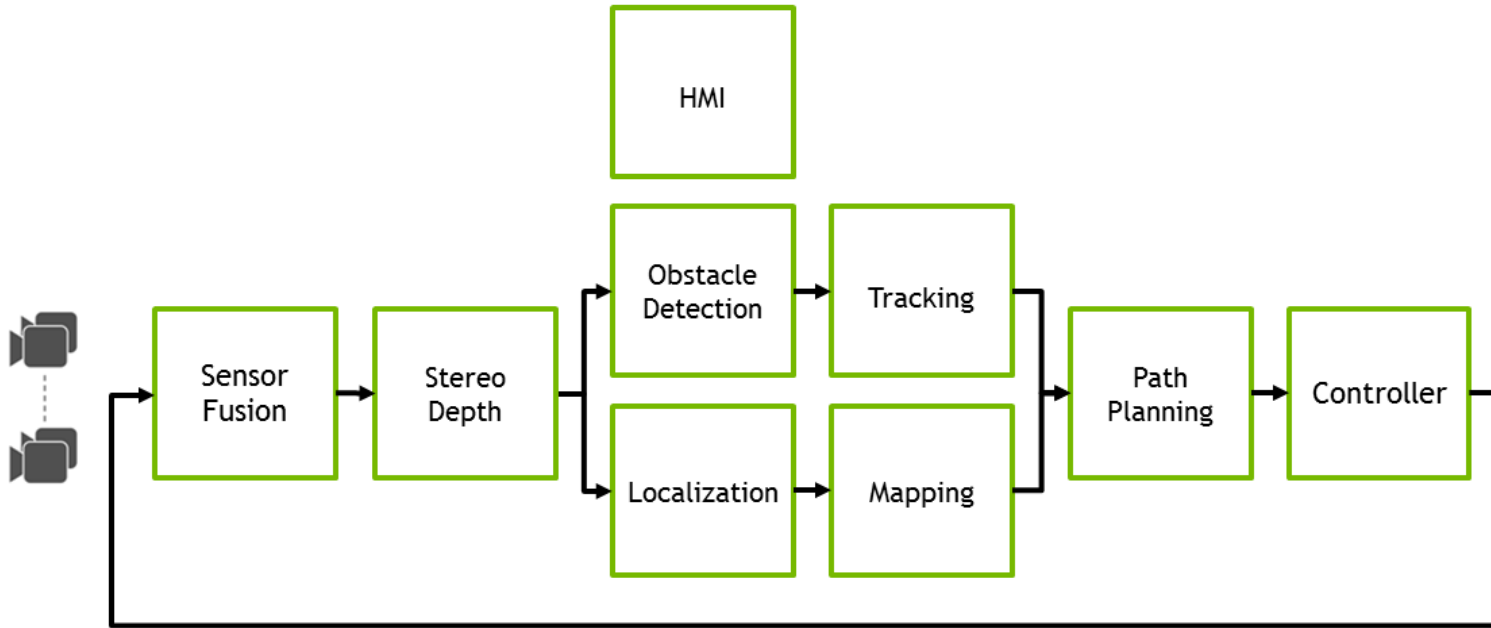


Inspection




Service

# EXAMPLE - AI DELIVERY



Total: 20-30 TOPS

# EXAMPLE - VIDEO ANALYTICS



Typical application: 30+ TOPS



# VISION NETWORKS

## Compute Demand

	Input Size	GOPs/Frame	GOPs @ 30Hz		Input Size	GOPs/Frame	GOPs @ 30Hz
<b>Image Recognition</b>				<b>Segmentation</b>			
MobileNet	224x224	.6	17	FCN-8S	384x384	125	3,750
AlexNet	227x227	.7	22	DeepLab-VGG	513x513	202	6,060
GoogleNet	224x224	2	60	SegNet	640x360	286	8,580
ResNet-50	224x224	4	120				
VGG19	224x224	20	600	<b>Pose Estimation</b>			
				PRM	256x256	46	1,380
<b>Object Detection</b>				Multipose	368x368	136	4,080
YOLO-v3	416x416	65	1,950				
SSD-VGG	512x512	91	2,730	<b>Stereo Depth DNN</b>	1280x640	260	7,800
Faster-RCNN	600x850	172	5,160				


# NVIDIA AGX SYSTEMS

EMBEDDED AI HPC

NVIDIA DRIVE AGX | Self-Driving Cars

NVIDIA Jetson AGX | Autonomous Machines

NVIDIA Clara AGX | Medical Imaging



# JETSON AGX XAVIER


World's first AI computer for  
Autonomous Machines

AI Server Performance in 30W • 15W • 10W

512 Volta CUDA Cores • 2x NVDLA

8 core CPU

32 DL TOPS





# COMPREHENSIVE HIGH PERFORMANCE I/O SUBSYSTEM

## PCIe

5 16GT/s gen4 controllers  
1x8, 1x4, 1x2, 2x1  
3x Root port + Endpoint  
2x Root port




## USB

3x USB3.1 (10 GT/s) ports  
4x USB2.0 ports



## ETHERNET

1x Gigabit Ethernet-AVB  
RGMII PHY  
PTP, WoL



## DISPLAY

3x DP/HDMI/eDP  
4K @ 60 Hz  
DP HBR3  
HDMI 2.0




## CAMERA

16 CSI2 lanes, 8 SLVS-EC lanes  
40 Gbps in DPHY 1.2 Mode  
109 Gbps in CPHY 1.1 Mode  
Up to 36 Virtual Channels



## OTHER I/Os

I2C I2S UFS  
CAN SPI SD  
UART GPIO



**Total I/O >650 Gbps**

# JETSON AGX XAVIER


## Compute Module

	JETSON TX2	JETSON AGX XAVIER
GPU	256 Core Pascal @ 1.3GHz	512 Core Volta @ 1.37GHz 64 Tensor Cores
DL Accelerator	-	(2x) NVDLA
Vision Accelerator	-	(2x) 7-way VLIW Processor
CPU	6 core Denver and A57 @ 2GHz (2x) 2MB L2	8 core Carmel ARM CPU @ 2.26GHz (4x) 2MB L2 + 4MB L3
Memory	8GB 128 bit LPDDR4 58.4 GB/s	16GB 256-bit LPDDR4x @ 2133MHz 137 GB/s
Storage	32GB eMMC	32GB eMMC
Video Encode	(2x) 4K @30 HEVC	(4x) 4Kp60 / (8x) 4Kp30 HEVC
Video Decode	(2x) 4K @30 12 bit support	(2x) 8Kp30 / (6x) 4Kp60 12 bit support
Camera	12 lanes MIPI CSI-2 D-PHY 1.2 30Gbps	16 lanes MIPI CSI-2   8 lanes SLVS-EC D-PHY 40Gbps / C-PHY 109Gbps
PCI Express	5 lanes PCIe Gen 2 1x4 + 1x1   2x1 + 1x4	16 lanes PCIe Gen 4 1x8 + 1x4 + 1x2 + 2x1
Mechanical	50mm x 87mm 400 pin connector	100mm x 87mm 699 pin connector
Power	7.5W / 15W	10W / 15W / 30W



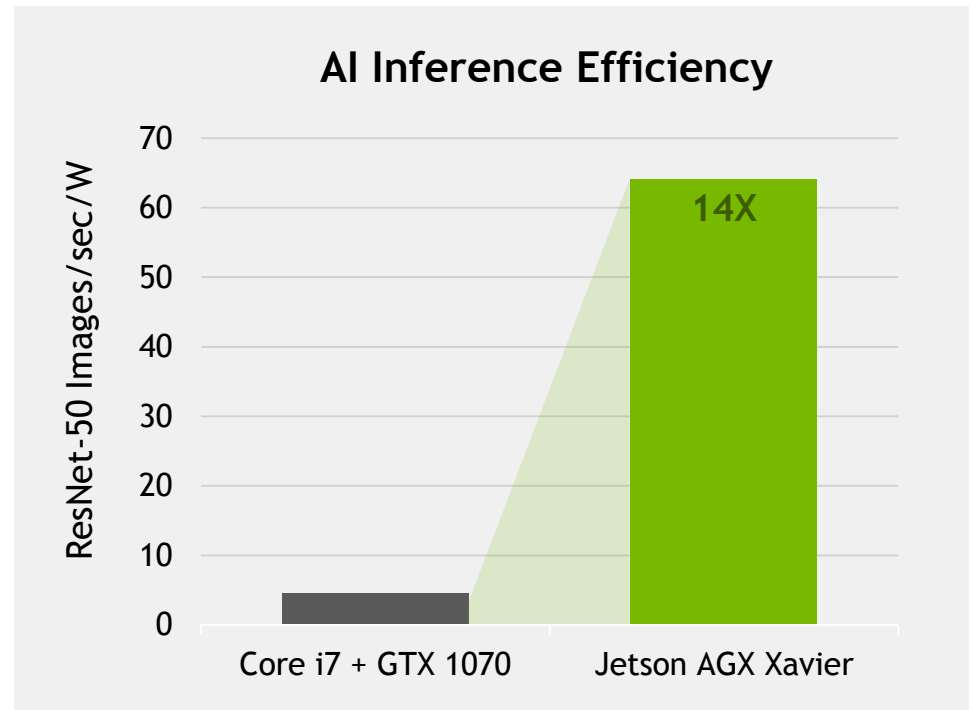

# JETSON AGX XAVIER

## 20x Performance in 18 Months





# JETSON AGX XAVIER

## GPU Workstation Perf • 1/10<sup>th</sup> Power



# JETSON AGX XAVIER

## Compute Module



Jetson Xavier  
Module  
Connector

Thermal  
Transfer  
Plate (TTP)



\$1599 (qty. 10+)  
\$1299 (qty. 100+)  
Coming Soon

# JETSON AGX XAVIER

## Developer Kit

I/O	
PCIe x16	PCIe Gen4 x8 / SLVS-EC x8
RJ45	Gigabit Ethernet
USB-C	(2x) USB 3.1   DisplayPort, Power Delivery
eSATAp + USB 3.0	SATA (Power + Data for 2.5" SATA) + USB 3.0
Micro USB	(1x) USB 2.0
Camera Header	(16x) CSI-2 lanes
M.2 Key M	NVMe storage
M.2 Key E	PCIe x1 (for Wi-Fi / LTE / 5G)
40-pin Header	UART, SPI, CAN, I2C, I2C, DMIC, GPIOs
HD Audio Header	High-Definition Audio
HDMI Type A	HDMI 2.0, eDP 1.2a, DP 1.4
uSD / UFS Card	SD / UFS
DC Barrel Jack	9V - 20VDC
Size	105mm x 105mm



\$2499 (Retail), \$1799 (qty. 10+)  
\$1299 (Developer Special, limit 1)  
Available Now, see NVIDIA.com

# JETSON AGX XAVIER

## Developer Kit


I/O	
PCIe x16	PCIe Gen4 x8 / SLVS-EC x8
RJ45	Gigabit Ethernet
USB-C	(2x) USB 3.1   DisplayPort, Power Delivery
eSATAp + USB 3.0	SATA (Power + Data for 2.5" SATA) + USB 3.0
Micro USB	(1x) USB 2.0
Camera Header	(16x) CSI-2 lanes
M.2 Key M	NVMe storage
M.2 Key E	PCIe x1 (for Wi-Fi / LTE / 5G)
40-pin Header	UART, SPI, CAN, I2C, I2C, DMIC, GPIOs
HD Audio Header	High-Definition Audio
HDMI Type A	HDMI 2.0, eDP 1.2a, DP 1.4
uSD / UFS Card	SD / UFS
DC Barrel Jack	9V - 20VDC
Size	105mm x 105mm



\$2499 (Retail), \$1799 (qty. 10+)  
\$1299 (Developer Special, limit 1)  
Available Now, see [NVIDIA.com](https://www.nvidia.com)

# JETSON AGX XAVIER

## Developer Kit





# JETSON AGX XAVIER ECOSYSTEM

## AI SOFTWARE



## QUICK START PLATFORMS



## RESEARCH



## SENSORS



## SYSTEM DESIGN



## SYSTEM SOFTWARE/TOOLS



## DISTRIBUTORS WORLDWIDE





# XAVIER ARCHITECTURE

Volta GPU

Deep Learning Accelerator (DLA)

Carmel ARM CPU

Vision Accelerator

# VOLTA GPU

Optimized for Inference

8x Volta SM @ 1377MHz

512 CUDA cores, 64 Tensor Cores

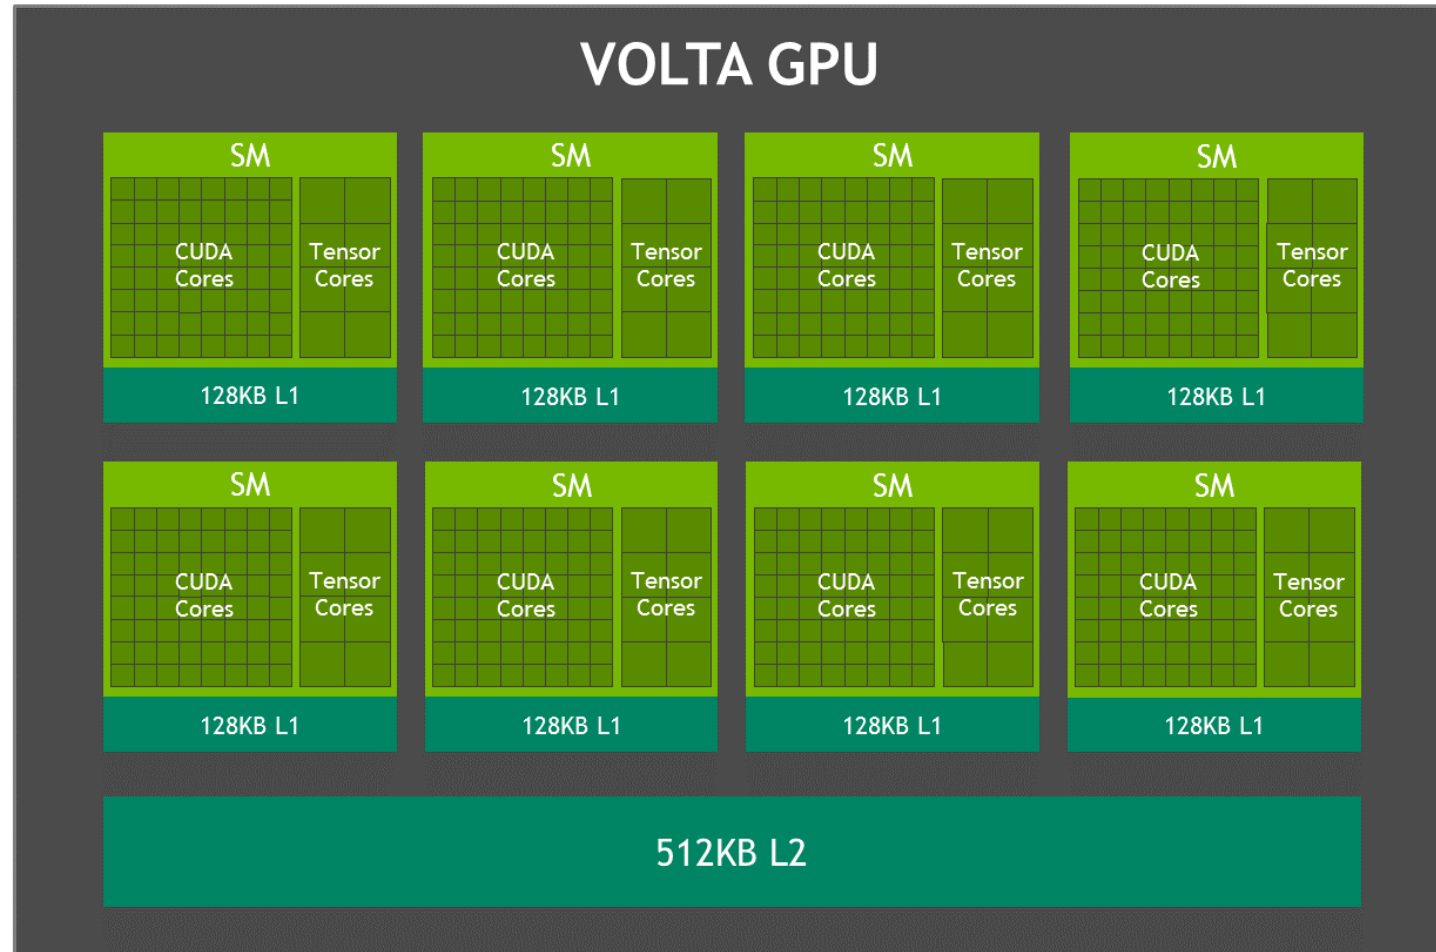
22 TOPS INT8, 11 TFLOPS FP16

8x larger L1 cache size

4x faster L2 cache access

4 scheduler partitions per SM

CUDA compute capability 7.2



# TENSOR CORES

## HMMA / IMMA

4x4 matrix processing array,  $D=A*B+C$


HMMA/IMMA FP16/INT8 Matrix Multiple Accumulate

Accumulation occurs in full precision with overflow protection

Each Tensor Core performs 64 floating-point or 128 integer ops per clock

Results can be composed to construct larger matrix multiplies & convolutions

Integrated with cuBLAS, cuDNN, TensorRT, and programmable through CUDA



$$D = \begin{pmatrix} A_{0,0} & A_{0,1} & A_{0,2} & A_{0,3} \\ A_{1,0} & A_{1,1} & A_{1,2} & A_{1,3} \\ A_{2,0} & A_{2,1} & A_{2,2} & A_{2,3} \\ A_{3,0} & A_{3,1} & A_{3,2} & A_{3,3} \end{pmatrix} \begin{pmatrix} B_{0,0} & B_{0,1} & B_{0,2} & B_{0,3} \\ B_{1,0} & B_{1,1} & B_{1,2} & B_{1,3} \\ B_{2,0} & B_{2,1} & B_{2,2} & B_{2,3} \\ B_{3,0} & B_{3,1} & B_{3,2} & B_{3,3} \end{pmatrix} + \begin{pmatrix} C_{0,0} & C_{0,1} & C_{0,2} & C_{0,3} \\ C_{1,0} & C_{1,1} & C_{1,2} & C_{1,3} \\ C_{2,0} & C_{2,1} & C_{2,2} & C_{2,3} \\ C_{3,0} & C_{3,1} & C_{3,2} & C_{3,3} \end{pmatrix}$$

HMMA FP16 or FP32  
IMMA INT32

FP16  
INT8 or UINT8

FP16  
INT8 or UINT8

FP16 or FP32  
INT32

# DEEP LEARNING ACCELERATOR (DLA)

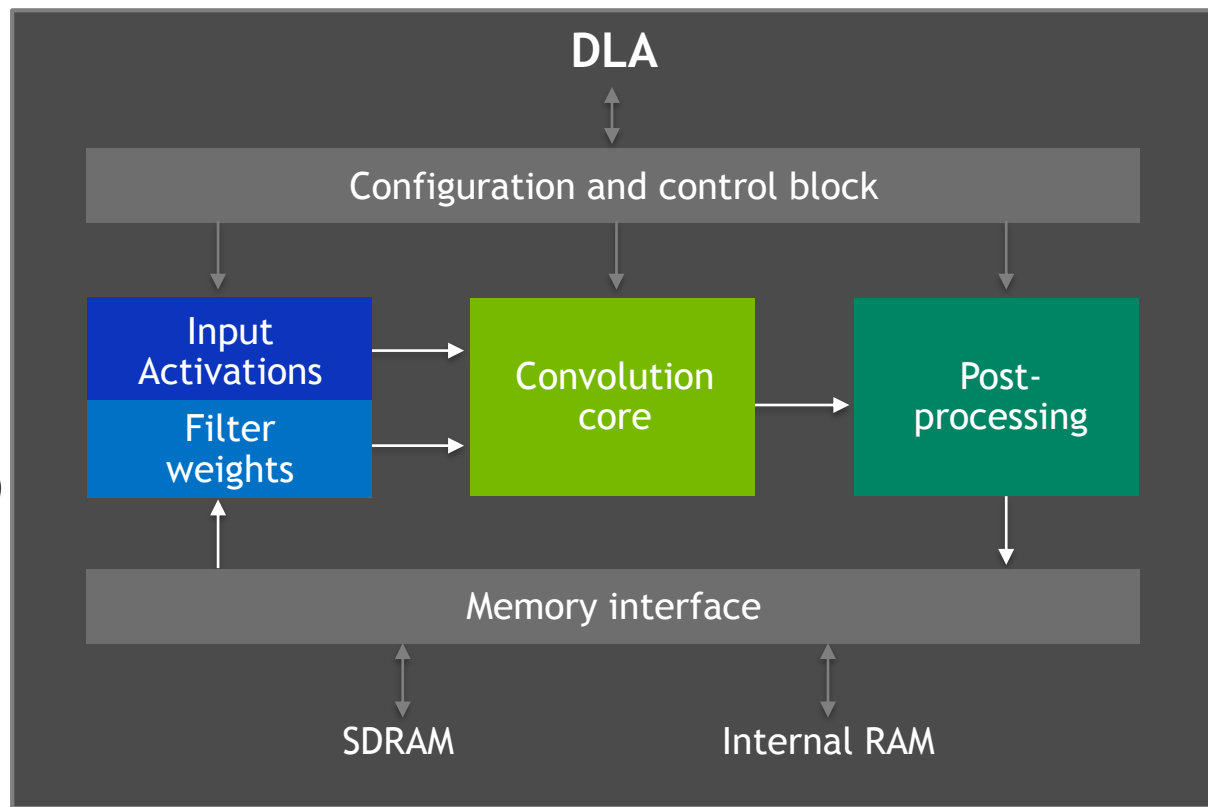
2x DLA engines per Xavier

5 TOPS INT8, 2.5 TFLOPS FP16 per DLA

Optimized for energy efficiency (500-1500mW)

Programmed with TensorRT 5.0


Supported layers include: Convolution, Deconvolution, Activations, Pooling, Normalization, Fully Connected



Open-source architecture: [NVDLA.org](https://nvidia.com/open-source/nvdl)


### AlexNet Inference Throughput

Jetson AGX Xavier / Jetson TX2




### GoogleNet Inference Throughput

Jetson AGX Xavier / Jetson TX2




### ResNet-50 Inference Throughput

Jetson AGX Xavier / Jetson TX2




### VGG19 Inference Throughput

Jetson AGX Xavier / Jetson TX2




### AlexNet Inference Energy Efficiency

Jetson AGX Xavier / Jetson TX2 (15W Power Modes)




### GoogleNet Inference Energy Efficiency

Jetson AGX Xavier / Jetson TX2 (15W Power Modes)




### ResNet-50 Inference Energy Efficiency

Jetson AGX Xavier / Jetson TX2 (15W Power Modes)



### VGG19 Inference Energy Efficiency

Jetson AGX Xavier / Jetson TX2 (15W Power Modes)



# POWER MODES

Different power mode presets: 10W, 15W and 30W

Default mode is 15W

Users can create their own presets, specifying clocks and online cores in `/etc/nvpmode1.conf`

```
< POWER_MODEL ID=2 NAME=MODE_15W >
CPU_ONLINE CORE_0
CPU_ONLINE CORE_4 0
CPU_DENVER_0 MAX_FREQ 1200000
GPU_MIN_FREQ 0
GPU_MAX_FREQ 670000000
EMC_MAX_FREQ 1331200000
```

## NVIDIA Power Model Tool

```
sudo nvpmode1 -q    (for current mode)
sudo nvpmode1 -m 0  (for changing mode, persists after reboot)
sudo ~/tegrastats   (for monitoring clocks & core utilization)
```



# NVPMODEL CLOCK CONFIGURATION

Mode Name	EDP	10W	15W	30W	30W	30W	30W
Power Budget	n/a	10W	15W	30W	30W	30W	30W
Mode ID	0	1	2	3	4	5	6
Online CPU	8	2	4	8	6	4	2
CPU Maximal Frequency (MHz)	2265.6	1200	1200	1200	1450	1780	2100
GPU TPC	4	2	4	4	4	4	4
GPU Maximal Frequency (MHz)	1377	520	670	900	900	900	900
DLA cores	2	2	2	2	2	2	2
DLA Maximal Frequency (MHz)	1395.2	550	750	1050	1050	1050	1050
Vision Accelerator (VA) cores	2	0	1	1	1	1	1
VA Maximal Frequency (MHz)	1088	0	550	760	760	760	760
Memory Maximal Freq. (MHz)	2133	1066	1333	1600	1600	1600	1600

The default mode is 15W (ID:2)

# CARMEL CPU COMPLEX

Full ARMv8.2 including RAS support

8 NVIDIA Carmel cores @ 2.26GHz

500-1500mW power per core

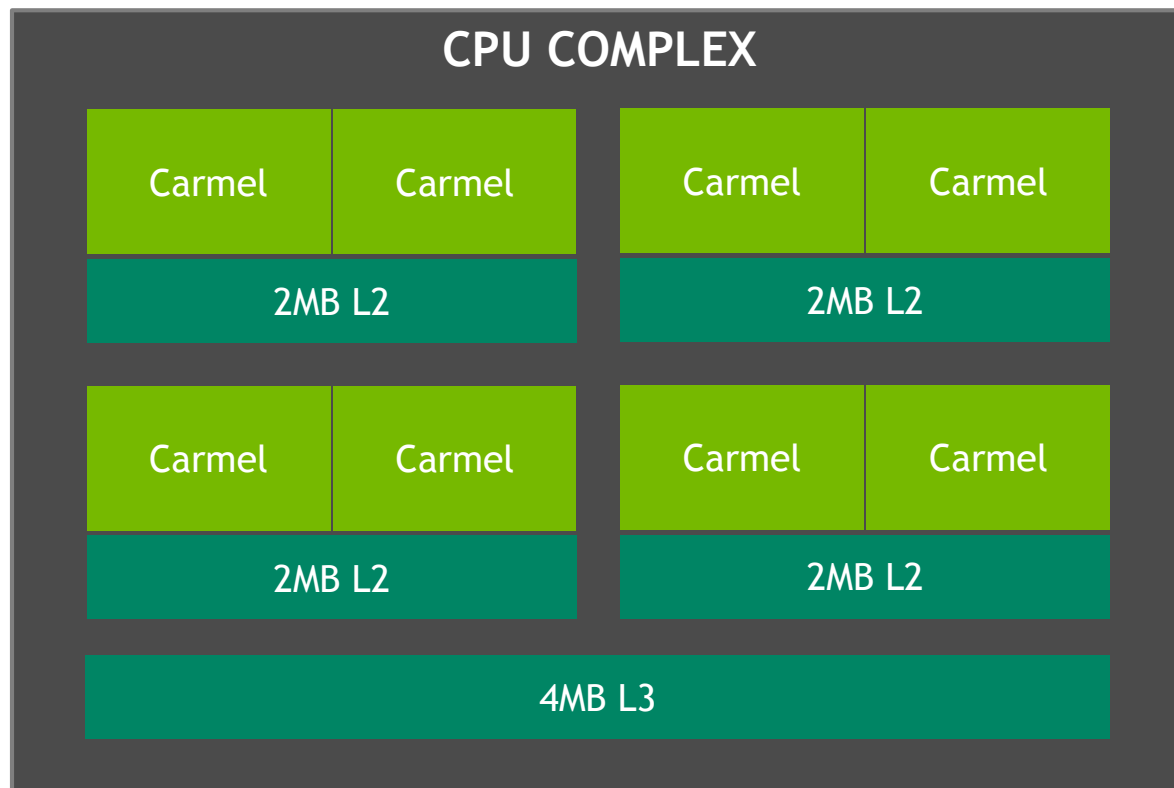
2 cores + 2MB L2 per cluster

Cache Coherent Across CPU Complex

NVIDIA Dynamic Code Optimization

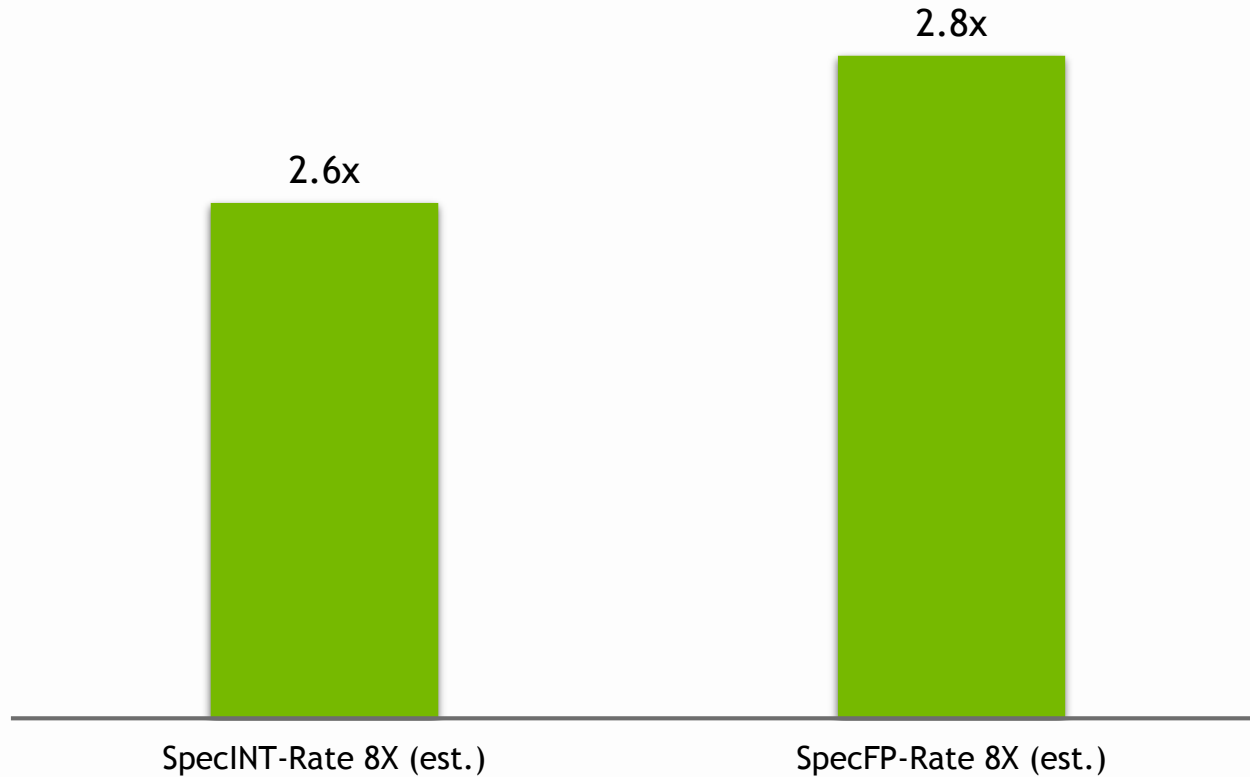
I/O Coherent Memory

4MB Exclusive L3 cache



# CPU BENCHMARKS

Speed-up of Xavier over TX2



# VISION ACCELERATOR

2x Vision Accelerator engines

Optimized offloading of imaging & vision algorithms - feature detection & matching, stereo, optical flow

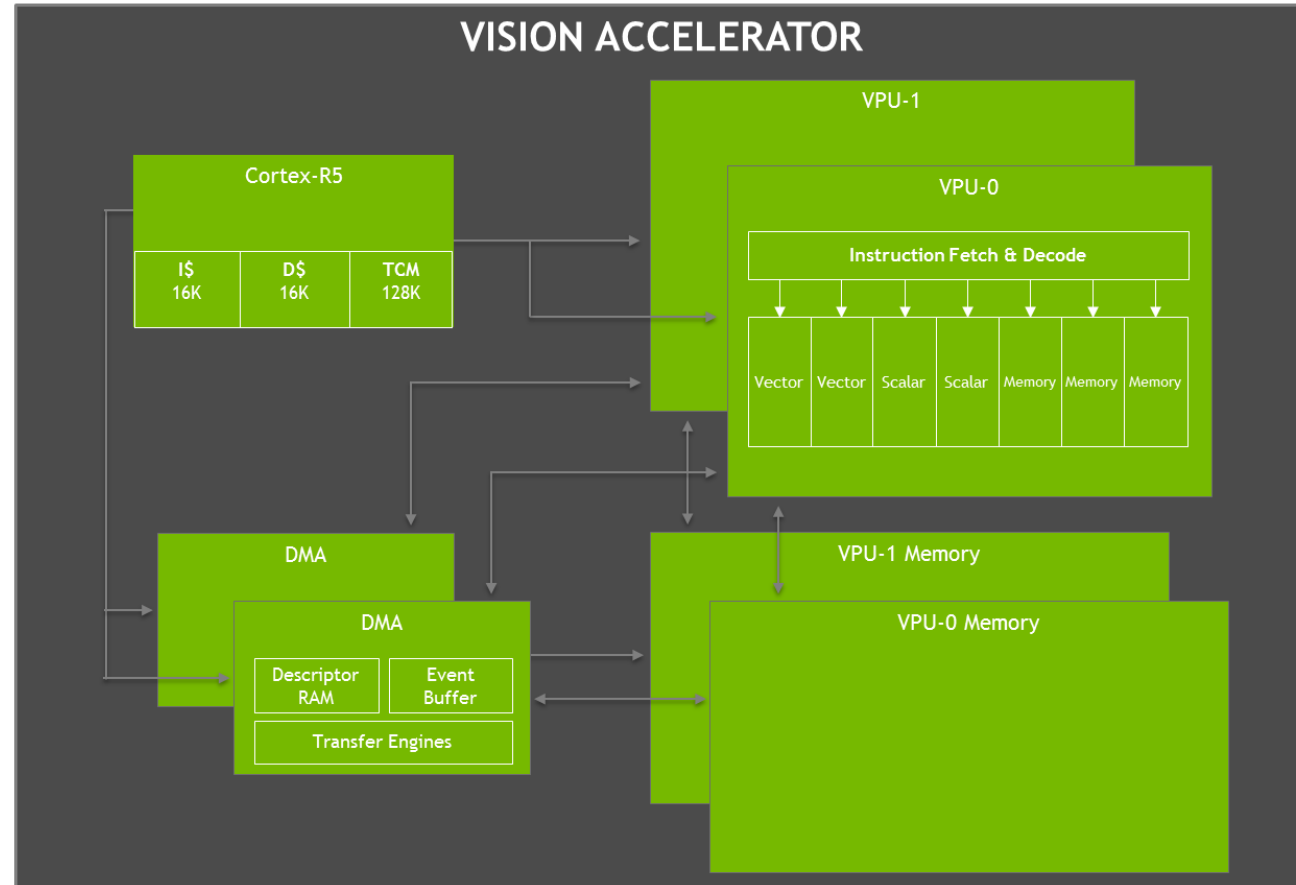
SW support enabled in future JetPack

Each Vision Accelerator includes:

- Cortex-R5 for config and control

- 2x 7-way VLIW Vector Processing Units

- 2x DMA for data movement to/from internal/external memories





# JETSON SDKs

JetPack - AI at the Edge

DeepStream - Intelligent Video Analytics (IVA)

ISAAC - Robotics & Autonomous Machines

# JETSON SDKs




DEEPSTREAM SDK  
FOR VIDEO ANALYTICS



ISAAC SDK  
FOR AUTONOMOUS MACHINES

JETPACK SDK  
FOR AI AT THE EDGE



JETSON AGX XAVIER


# JETPACK SDK

## for AI at the Edge

Sample Code

Nsight Developer Tools

Multimedia API




TensorRT  
cuDNN  
TF, PyTorch, ...

Deep Learning



VisionWorks  
OpenCV  
NPP

Computer Vision



Vulkan  
OpenGL  
EGL/GLES

Graphics



libargus  
GStreamer  
V4L2

Media

CUDA, Linux For Tegra, ROS



Jetson AGX Xavier: Advanced GPU, 64-bit CPU, Video CODEC, DLAs

# JETPACK 4.1

## DEVELOPER PREVIEW EARLY ACCESS



Available Now For Jetson AGX Xavier  
[developer.nvidia.com/jetpack](https://developer.nvidia.com/jetpack)

## Package Versions

L4T BSP	31.0.2
Linux Kernel	4.9
CBoot	1.0
Vulkan	1.1.1
OpenGL	4.6.0
OpenGL-ES	3.2.5
EGL	1.5
GLX	1.4
X11 ABI	24
Xrandr	1.4
Multimedia API	31.1
Argus Camera API	0.97
GStreamer	1.14
Nsight Systems	2018.1
Nsight Graphics	1.0
Jetson OS	Ubuntu 18.04
Host OS	Ubuntu 16.04 / 18.04

CUDA	10.0
cuDNN	7.3
TensorRT	5.0 RC
VisionWorks	1.6
OpenCV	3.3.1
NPP	10.0


Install TensorFlow, PyTorch, Caffe,  
ROS, and other GPU libraries





# NVIDIA TensorRT

## Production Inferencing



Compile and Optimize Neural Networks  
Support for Every Framework  
Optimize for Each Target Platform

# NVIDIA TensorRT 5

## Deep Learning Inference Optimizer and Runtime

New support for Jetson AGX Xavier in TensorRT 5:

- Volta GPU INT8 & Tensor Cores (HMMA/IMMA)
- Early-Access DLA FP16 support
- Updated samples to enabled DLA
- Fine-grained control of DLA layers and GPU Fallback
- New APIs added to IBuilder interface:

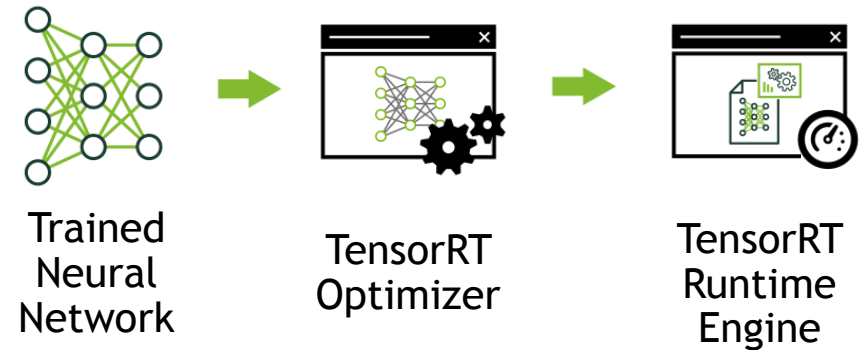
```
setDeviceType()
```

```
canRunOnDLA()
```

```
getMaxDLABatchSize()
```

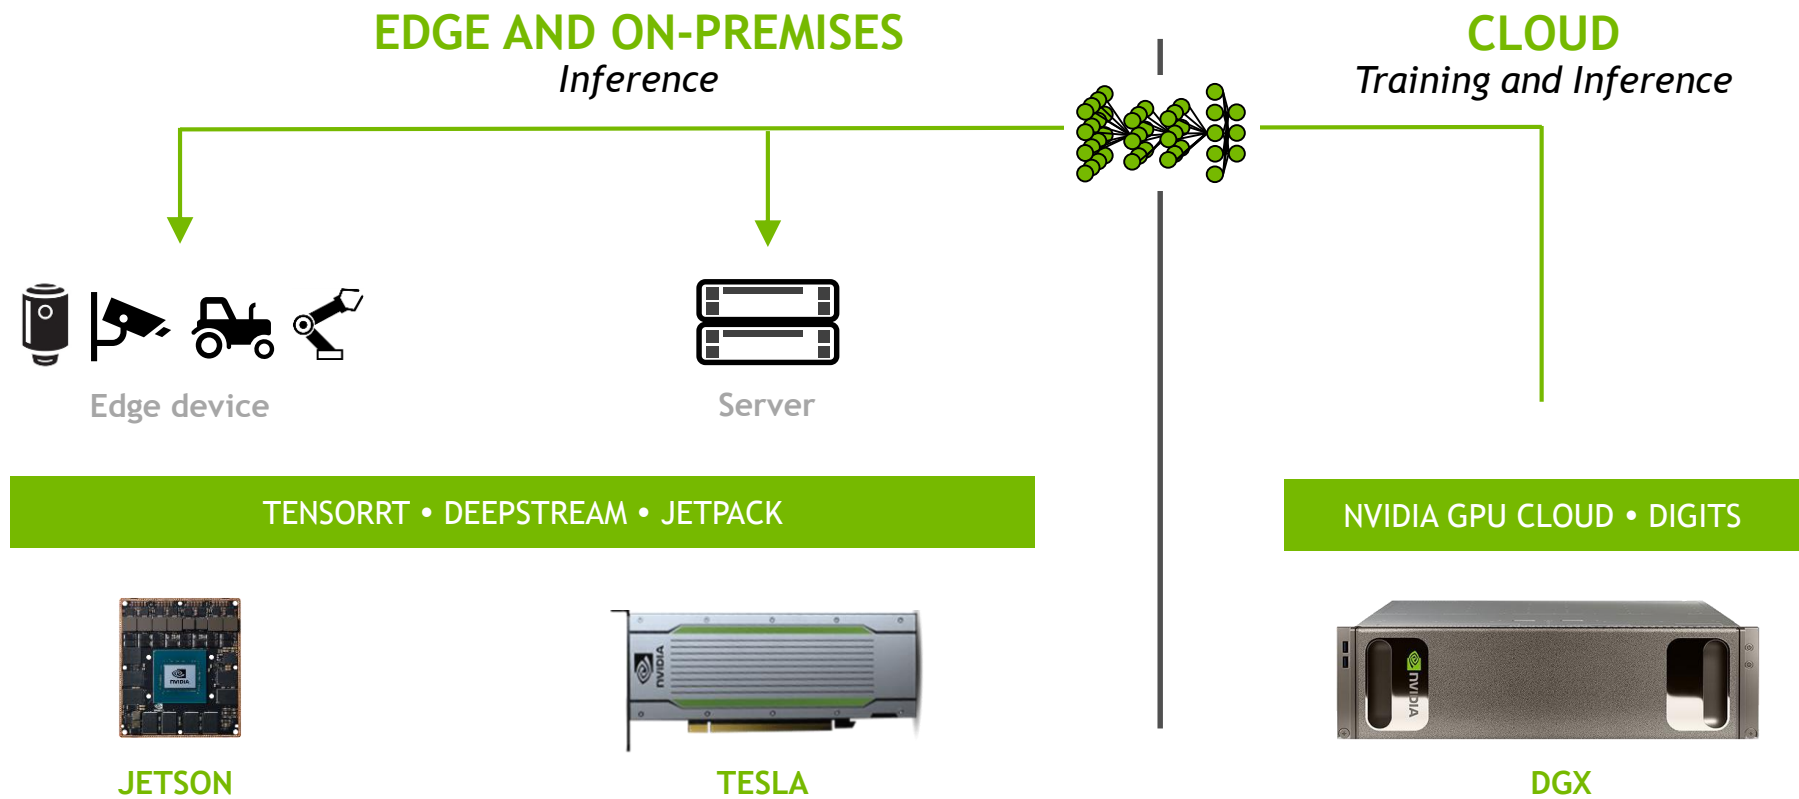
```
allowGPUFallback()
```

[developer.nvidia.com/tensorRT](https://developer.nvidia.com/tensorRT)



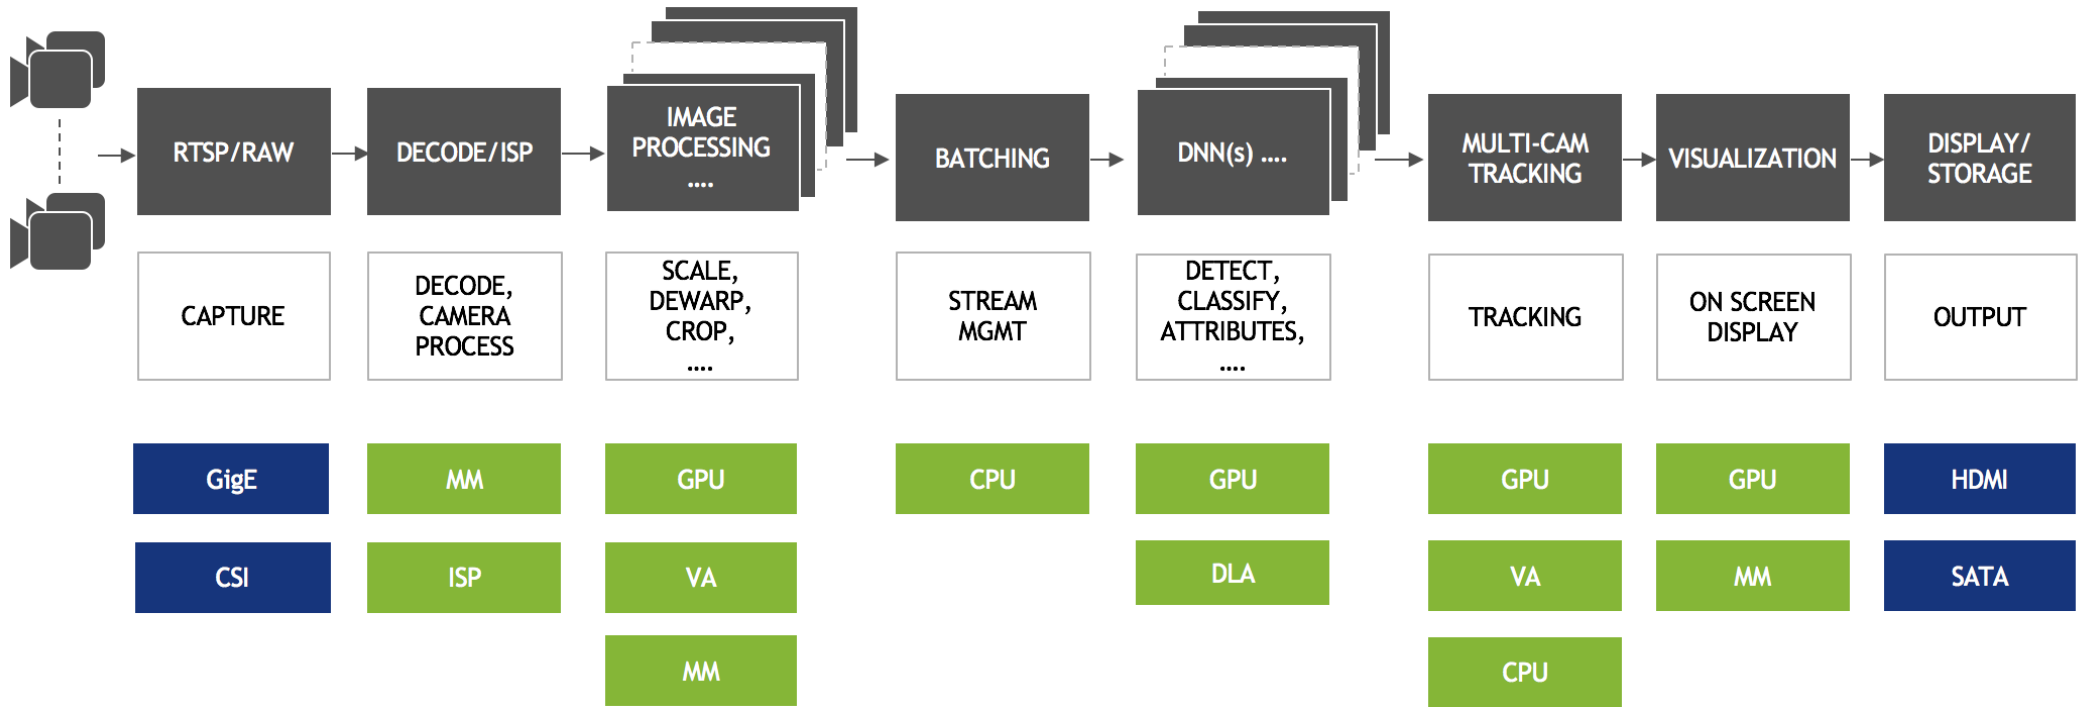
- Fuse network layers
- Eliminate concatenation layers
- Kernel specialization
- Auto-tuning for target platform
- Select optimal tensor layout
- Batch size tuning
- Mixed-precision INT8/FP16 support

# AI - EDGE TO CLOUD



# NVIDIA DEEPSTREAM

## Zero Memory Copies



Typical multi-stream application: 30+ TOPS



# ISAAC

## Isaac SDK: Simulation to Reality


Simulate



Develop



Deploy



### World model

Warehouse · Office  
· Store · Home

### Robot model

Carter · URDF loader

### ML

TensorRT · CUDA  
· TensorFlow · ...

### Gems

Optimizers · Algebra  
· EKF's · Depth · ...

### Drivers

Lidar · Camera · IMU ·  
Robot Base · ...

### Jetson

Fully integrated with  
TX2 and Xavier

### Simulation Engine

Photo-realistic Graphics · Physics · Soft bodies ·  
· Procedural Generation · Massive parallelism  
· Unreal Engine 4 / Unity 3D

### Isaac Framework

Codelets · Behaviors · 3D Poses · Distributed  
· Messaging · Synchronization · Record & Replay  
· Configuration · Visualization

### Unified Message API

Use the same messages for simulation,  
actual hardware and across all apps

Virtual Sensors

Virtual Actuators

Sensor Processing

Actuator Control

HW Sensors

HW Actuators



Reality



Simulation


# NVIDIA ISAAC GEMS



Global Localization



LQR Path Planner



Depth Estimation




Human Pose Estimation



Object / People Detection



Map Editor



Visual Odometry



Physical Simulation



Gesture Recognition



ASR





# RESOURCES & SUPPORT

Developer Site

Documentation

Forums & Wiki

Tutorials

Quick-Start Platforms

# JETSON DEVELOPER SITE

End-to-end development  
from idea to final product

JetPack and Isaac SDKs

Developer tools

Design collateral

Developer forum

Training and tutorials

Ecosystem

[developer.nvidia.com/jetson](https://developer.nvidia.com/jetson)

The screenshot shows the NVIDIA Embedded Computing website. At the top, there's a navigation bar with links for Develop, Downloads, Community, Learn, SDKs, and Buy. A main banner features the text "Two Days to a Demo" with a "Get Started Now" button and an image of a futuristic human head. Below this, a section titled "Meet Jetson, the Platform for AI at the Edge" includes a paragraph about NVIDIA Jetson's capabilities and a "Learn More" button, accompanied by an image of a Jetson module. A horizontal menu contains icons for PARTNERS, ECOSYSTEM, FORUM, DOWNLOAD, FAQ, and TUTORIALS. Further down, there's a video player for "NVIDIA Jetson Xavier Develop..." and a section for "Jetson Xavier Developer Kit: Introduction" with a "Learn More" button. The footer includes a "Recent Tweets" section with a tweet from @MathWorks and a "Follow" button.

# GETTING HELP

## Jetson Community

### Developer Forums devtalk.nvidia.com

**NVIDIA ACCELERATED COMPUTING** Downloads Training Ecosystem Forums  [Join](#) [Login](#)

Home > CUDA ZONE > Forums > Accelerated Computing > Jetson & Embedded Systems > Jetson AGX Xavier

[+ Create Topic](#)

## Jetson AGX Xavier

	Activity	Started By	Last Comment
<a href="#">TensorFlow wheel for JetPack 4.0 !!</a>	6 Replies 198 Views	AastaLLL 5 days ago	rebotnix 4 hours ago
<a href="#">JetPack 4.0 Developer Preview Early Access for Jetson AGX Xavier</a>	21 Replies 775 Views	dusty_nv 2 weeks ago	prush 22 hours ago
<a href="#">Links to Jetson Xavier Resources &amp; Wiki</a>	8 Replies 723 Views	dusty_nv 3 weeks ago	dusty_nv 1 day ago
<a href="#">NVIDIA Webinar — Jetson AGX Xavier and the New Era of Autonomous Machines</a>	1 Replies 65 Views	dusty_nv 3 days ago	GeForceX 3 days ago
<a href="#">NVIDIA Announces Jetson AGX Xavier</a>	26 Replies 1,088 Views	dusty_nv 1 month ago	dusty_nv 1 week ago
<a href="#">Jetson AGX Xavier Developer Kit — Now Shipping!</a>	11 Replies 486 Views	dusty_nv 4 weeks ago	dusty_nv 3 weeks ago
<a href="#">nsight not working with me</a>	10 Replies 117 Views	sato5312 4 days ago	sato5312 2 hours ago
<a href="#">Setting up VNC server</a>	1 Replies 7 Views	hackbin 22 hours ago	dusty_nv 8 hours ago
<a href="#">Xavier Carrier Board Development started.</a>	4 Replies 80 Views	rebotnix 1 day ago	linudev 11 hours ago
<a href="#">Default Boot Resolution?</a>	3 Replies 39 Views	GeForceX 15 hours ago	linudev 13 hours ago
<a href="#">Intel 9260 adapter not working</a>	15 Replies 192 Views	fnunes 1 week ago	linudev 15 hours ago
<a href="#">xavier fan question</a>	7 Replies 83 Views	guo.tang 2 days ago	linudev 1 day ago
<a href="#">Building Driver for ASUS USB-AC-53 Nano Wi-Fi Adapter</a>	1 Replies 41 Views	nicolas.capens 1 day ago	linudev 1 day ago
<a href="#">Ubuntu 18.04 has experienced an internal Error</a>	8 Replies 69 Views	S4WRXTTCS 3 days ago	danpollack 1 day ago

### eLinux Wiki eLinux.org/Jetson

Page Discussion  [Read](#) [View source](#) [View history](#)  [Search eLinux.org](#)

## Jetson AGX Xavier

NVIDIA Jetson AGX Xavier<sup>®</sup> is an embedded system-on-module (SoM) from the NVIDIA AGX Systems<sup>®</sup> family, including an integrated Volta GPU with Tensor Cores, dual Deep Learning Accelerators (DLAs), octal-core NVIDIA Carmel ARMv8.2 CPU, 16GB 256-bit LPDDR4x with 137GB/s of memory bandwidth, and 650Gbps of high-speed I/O including PCIe Gen 4 and 16 camera lanes of MIPI CSI-2. Useful for deploying computer vision and deep learning to the edge, Jetson AGX Xavier runs Linux and provides 32 TeraOPS of compute performance in user-configurable 10/15/30W power profiles. Jetson AGX Xavier is currently available as the Jetson AGX Xavier Developer Kit<sup>®</sup>, with the OEM compute module becoming available in the future. See the wiki of previous Jetson's [here](#).

Where else to find us  
Google+ Community  
Twitter (@elinux)  
#elinux on Freenode  
Facebook (@elinux.org)  
Mailing Lists

Tools  
What links here  
Related changes  
Special pages  
Printable version  
Permanent link  
Page information

### Contents [hide]

- Jetson AGX Xavier Module
  - Processing Components
  - I/O Interfaces & Ports
  - Form Factor
  - Software Support
- Jetson AGX Xavier Developer Kit
  - What's Included
  - Ports & Connectors
  - Getting Started
  - Availability
- Platform Documentation
  - Upcoming Documents
- Guides and Tutorials
  - System Tools
  - Deep Learning
- Ecosystem Products & Cameras
- Getting Help

## Jetson AGX Xavier Module

The Jetson AGX Xavier compute module contains all the active processing components. The I/O ports are broken out through a carrier board via a 699-pin board-to-board connector. Below is a partial list of the module's features. Please see the Jetson AGX Xavier Module Datasheet<sup>®</sup> for the complete specifications.

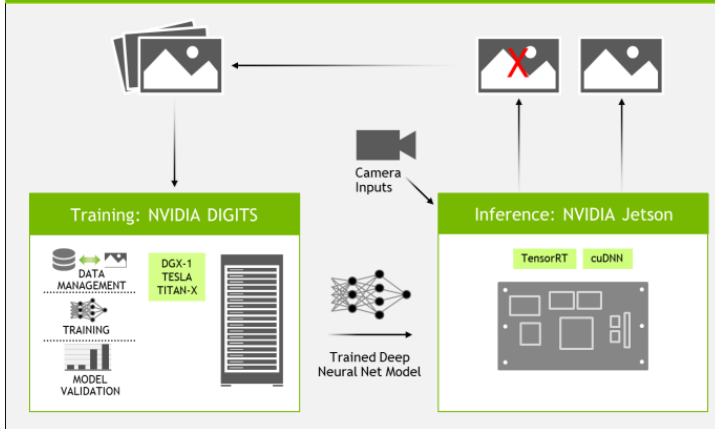
### Processing Components

- Octal-core NVIDIA Carmel ARMv8.2 CPU @ 2.26GHz
- 512-core Volta GPU @ with 64 Tensor Cores
- Dual Deep Learning Accelerator<sup>®</sup> (DLA) engines
- 16GB 256-bit LPDDR4x @ 2133MHz (137GB/s)
- 32GB eMMC 5.1
- Vision Accelerator engine
  - (4x) 4Kp60 H.264/H.265 video encoder
  - (2x) 8Kp30 / (6x) 4Kp60 H.264/H.265 video encoder

# TWO DAYS TO A DEMO

## Getting Started with Deep Learning

### AI WORKFLOW



Train using DIGITS and cloud/PC  
Deploy to the field with Jetson

### TRAINING GUIDES

This block shows a screenshot of the NVIDIA Jetson website's training guides. It features sections for 'Deploying Deep Learning' and 'Locating Object Coordinates using DetectNet'. The 'Deploying Deep Learning' section includes a 'DEEP VISION TUTORIAL' and 'NVIDIA\_JETSON\_TK2'. The 'Locating Object Coordinates using DetectNet' section includes a 'DETECTNET' tutorial and 'NVIDIA\_JETSON\_TK2'. The page also includes a 'Deploying Deep Learning' section with a 'DETECTNET' tutorial and 'NVIDIA\_JETSON\_TK2'.

All the steps required to follow to train  
your own models, including the datasets.

### DEEP VISION PRIMITIVES

This block displays various deep vision primitives. It includes 'Image Recognition Classification' (97.07% ice bear, polar bear), 'Object Detection Localization' (people in a street scene), 'Segmentation Free Space' (a street scene with a green overlay), '3D ShapeNet' (a grid of 3D object models), and 'Registration' (a 3D point cloud of a tree).

Image Recognition, Object Detection  
and Segmentation

# TWO DAYS TO A DEMO

## Reinforcement Learning Edition



### OpenAI Gym



Test environments and games for research and verification


### RL Algorithms

```
Do you want to continue [Y/n]? y
Get:1 http://archive.ubuntu.com/ubuntu/ lucid/universe python-keybinder
[12.2kB]
Get:2 http://archive.ubuntu.com/ubuntu/ lucid/universe terminator 0.93
[190kB]
Fetched 202kB in 5s (37.2kB/s)
Selecting previously deselected package python-keybinder.
(Reading database ... 129972 files and directories currently installed)
Unpacking python-keybinder (from .../python-keybinder_0.0.4-1_i386.deb)
Selecting previously deselected package terminator.
Unpacking terminator (from .../terminator_0.93-0ubuntu1_all.deb) ...
Processing triggers for desktop-file-utils ...
Processing triggers for python-gmenu ...
Rebuilding /usr/share/applications/desktop.en_US.utf8.cache...
Processing triggers for man-db ...
Processing triggers for hicolor-icon-theme ...
Processing triggers for python-support ...
Setting up python-keybinder (0.0.4-1) ...

Setting up terminator (0.93-0ubuntu1) ...
update-alternatives: using /usr/bin/terminator to provide /usr/bin/x-terminal-emulator (x-terminal-emulator) in auto mode.
```

DQN, A3C, Actor Critic using PyTorch

### Robotic Simulation



Observation from vision  
Pixels-to-actions

### Transfer Learning



Adapt network to real robot  
Online learning in the field

[github.com/dusty-nv/jetson-reinforcement](https://github.com/dusty-nv/jetson-reinforcement)

# TENSORFLOW

## Accelerated Performance with Jetson AGX Xavier

Download PIP Wheel installers from Jetson Download Center

Follow tutorials for popular vision tasks like object detection

Optimize for deployment with NVIDIA TensorRT (UFF/TFTRT)

[developer.nvidia.com/embedded/downloads](https://developer.nvidia.com/embedded/downloads)

[github.com/NVIDIA-Jetson/tf\\_to\\_trt\\_image\\_classification](https://github.com/NVIDIA-Jetson/tf_to_trt_image_classification)


[github.com/NVIDIA-Jetson/tf\\_trt\\_models](https://github.com/NVIDIA-Jetson/tf_trt_models)



Golden Retriever

Miniature Poodle

Toy Poodle



# JETSON QUICK-START PLATFORMS



Toyota HSR



Clearpath Robotics - Jackal UGV



JetsonHacks RACECAR/J




Aion Robotics - R1 UGV



NVIDIA - Redtail UAV

# Thank you!



Developer Site

[developer.nvidia.com/jetson](https://developer.nvidia.com/jetson)

Download JetPack

[developer.nvidia.com/jetpack](https://developer.nvidia.com/jetpack)

2 Days To a Demo

[github.com/dusty-nv](https://github.com/dusty-nv)

DevTalk Forums

[devtalk.nvidia.com](https://devtalk.nvidia.com)

Visit the Wiki

[eLinux.org/Jetson](https://eLinux.org/Jetson)

Q&A: What can I help you build?

A screenshot of the NVIDIA Developer Blog website. The page features a green header with the NVIDIA logo and navigation links. The main content area displays an article titled "NVIDIA Jetson Xavier Opens New Era of AI in Robotics" by Dustin Franklin, dated September 7, 2018. The article includes a sub-headline "New Jetson Delivers 32 TeraOPS in Energy-Efficient Embedded Module" and several paragraphs of text describing the capabilities of the Jetson Xavier. Below the text is a photograph of the Jetson Xavier embedded compute module with a Thermal Transfer Plate (TTP) attached. A caption below the photo reads: "Figure 1. Jetson Xavier embedded compute module with Thermal Transfer Plate (TTP), 100x87mm." The page also includes a search bar, login options, and social media links.

Dev Blog *NVIDIA Jetson AGX Xavier  
Opens New Era of AI in Robotics*