



WHITE PAPER

AMD UNLOCKS OPEN-SOURCE EFFICIENCY

**PLATFORM CHOICE AND COMPETITIVE
DIFFERENTIATION FOR BAIDU'S
APOLLO AUTONOMOUS DRIVING INNOVATION**



System designers can implement high-performance AMD GPUs and the open AMD ROCm ecosystem in a flexible, portable framework for impressive cost efficiencies.

The advent of autonomous driving and artificial intelligence (AI) has positioned automotive OEMS – from legacy brands to upstarts – for a new era of competition that leverages CPU and GPU compute power within the vehicle to achieve breakthrough, differentiating capabilities. Innovations in machine learning (ML) and sensor fusion are playing a major role in this seismic technological shift, with untold implications for tomorrow's transportation, commerce and infrastructure.

The evolution from L2 (partial automation) to L2+/L3 (conditional automation) and finally L4/L5 (full automation) will profoundly transform modern transportation models and vehicle ownership trends. In so doing it will hasten the mainstream adoption of new Mobility as a Service (MaaS) offerings designed to maximize transport efficiency leveraging fully automated robo-taxis and delivery vehicles.

Deploying safe, autonomous vehicles at scale is a colossal undertaking that requires an expansive ecosystem for cultivating advanced vehicle technologies. Baidu is leading this effort with technical collaboration on unprecedented scale, and AMD is proud to play an enabling role by delivering open-source development complemented with the continued innovation and leadership by AMD in x86 CPU and GPU processing performance.

APOLLO ASCENDING

Baidu's Apollo initiative is an open-source autonomous driving platform developed by Baidu for commercial transport applications. It's targeted to provide a complete hardware and software service solution that includes cloud data services, software, vehicle and hardware platform.

The Apollo project provides an open, reliable and secure software foundation for its partners to develop their own autonomous driving systems, including software and hardware tools optimized for obstacle perception, trajectory planning, vehicle control, and more.

Since its launch in 2017, development work centered on Apollo has grown significantly. Baidu has established an autonomous driving business unit and has partnered with several automakers to develop autonomous vehicles.¹

With extensive field testing and implementation already underway, Baidu's Apollo is a leading open-source autonomous driving software platform. It offers complete Level 4+ autonomous driving support, comprised of over 750,000+ lines of code and 9,000+ repository forks.

1. [China's Baidu is Building an Entire 'Mobility as a Service' Ecosystem](#)



AN EXCITING OPEN ALTERNATIVE

The GPUs onboard the autonomous vehicles of today and tomorrow shoulder enormous computing workloads. Guided by ML, they're tasked with processing enormous volumes of data from sensors (camera, radar, lidar) throughout the vehicle to the sensor fusion hub, powering critical object detection and recognition capabilities that inform the real-time prediction and planning functions.

Until now, supported GPU models under Baidu's Apollo program were limited to NVIDIA offerings, with the attendant constraints of proprietary CUDA® development tools. Today, AMD is opening up the Baidu Apollo development ecosystem even wider by enabling the Apollo software workloads originally written for CUDA to be ported to innovative open-source models – and powered by high-performance AMD GPUs.

To achieve this, AMD has ported the full CUDA ecosystem software stack into the AMD ROCm™ software ecosystem. AMD ROCm™ is an open software platform comprised of drivers, development tools and APIs enabling GPU programming from the low-level kernel to end-user applications.

ROCm today spans several domains including general-purpose computing on graphics processing units (GPGPU) and high-performance computing (HPC). It's enabled in part by the AMD Heterogeneous-computing Interface for Portability (HIP), an OSS C++ GPU programming environment and its corresponding runtime.

The ROCm platform is built on the foundation of open portability, supporting environments across multiple accelerator vendors and architectures. With the AMD ROCm open software platform designed for flexibility and performance, automotive OEMs and Tier Ones can now gain access to open compute languages, compilers, libraries and tools designed to accelerate their code development – and get to market faster.

The AMD collaboration with and contributions to the open-source community are a driving force behind ROCm platform innovations. This industry-differentiating approach to accelerated compute and heterogeneous workload development gives users unprecedented flexibility, choice and platform autonomy. Tools, guidance and insights are shared freely across the ROCm GitHub community and forums.

SEAMLESS ONE-TO-ONE PORTING

The typical GPU workload for Baidu Apollo is comprised in large part of ML inference capabilities needed for obstacle perception, traffic light perception, etc. – each represents ML models inferred in the GPU. Image pre- and post-processing is another important element in the GPU workload, and in all cases Baidu found the porting from NVIDIA to AMD GPU solutions was graceful and readily achieved. CUDA kernels were likewise converted, leveraging HIP.



WORKLOAD	NVIDIA SOLUTION	AMD SOLUTION
MACHINE LEARNING INTERFACE	TensorRT + cuDNN + cuBlas	MIGraphX + MIOpen + hipBlas
IMAGE PROCESSING	NPP (Nvidia Performance Primitive)	RPP (Radio Performance Primitive)
GPU KERNELS	CUDA	HIP

Apollo uses Caffe, Tensorflow, PyTorch and Paddlenet Deep Learning frameworks to train their ML models. These models are inferred using TensorRT and PyTorch on NVIDIA GPUs. AMD has ported the TensorRT inference engine to MIGraphX for use with AMD GPUs. PyTorch has been supported in ROCm for some time, upstream to the PyTorch open-source community.

AMD has accomplished one-to-one mapping for all requisite libraries and APIs to unlock the full innovation in Baidu Apollo for the open-source development community, while offering automotive system designers more choice when it comes to selecting high-performance GPUs. AMD graphics and acceleration performance is renowned throughout the technology industry, and AMD GPUs are at the forefront of innovation in automotive, including in-vehicle infotainment (IVI) and autonomous driving applications.

Designers of autonomous driving systems based on Baidu's Apollo platform can now benefit from open-source development efficiency, greater platform choice and competitive differentiation in the fast-moving automotive market. The full Apollo source code has been ported to work on select AMD GPUs with ROCm support. The list of supported AMD GPUs is here: <https://github.com/RadeonOpenCompute/ROCm/tree/rocm-4.5.2#hardware-and-software-support>. The AMD source code has been uploaded to the Apollo github repository master branch available here: <https://github.com/ApolloAuto/apollo>

About AMD

For more than 45 years AMD has driven innovation in high-performance computing, graphics and visualization technologies – the building blocks for gaming, immersive platforms and the datacenter. Hundreds of millions of consumers, leading Fortune 500 businesses and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit www.amd.com