DIPL.ING.(FH)KLAUS ROCK

HTTP-QuSS

HTTP - QUANTUM SPEED AND SECURITY Ψ

February 14, 2022

AUTONOMOUS DRIVING SEGMENT

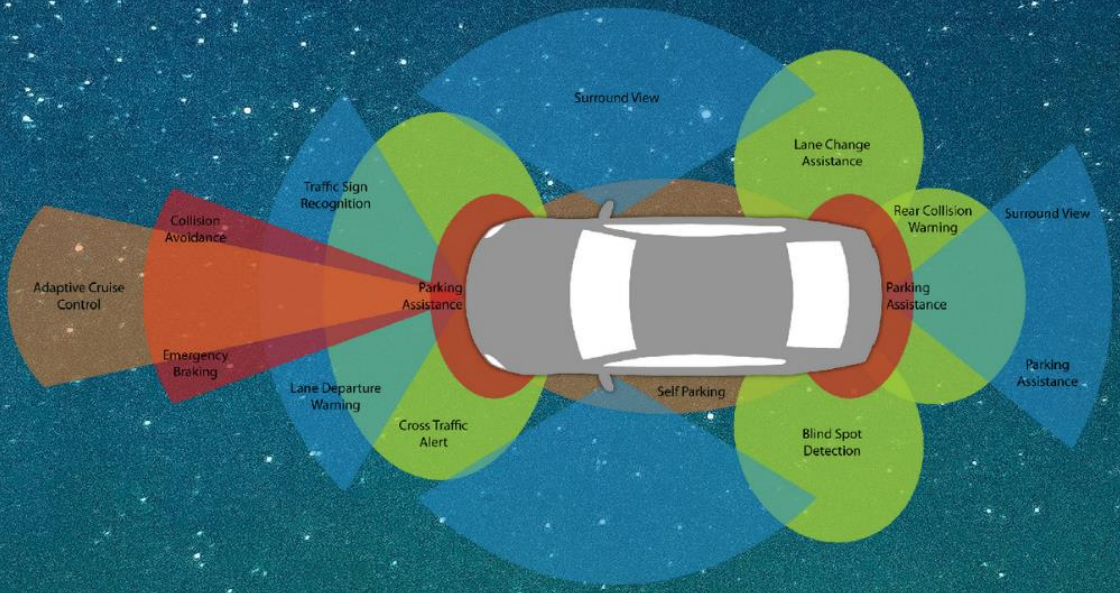

ROCK TECHNOLOGIES
Bonhoefferstr. 37 | 73432 Aalen | Germany | +49-7367-9222-958

Requirements for autonomous Driving

While the future of the automotive industry is still unwritten, here is what we know: Industry leaders will need to master connectivity to deliver the **V2X** (Vehicle-**T**o-Everything) capabilities fully autonomous driving promises.

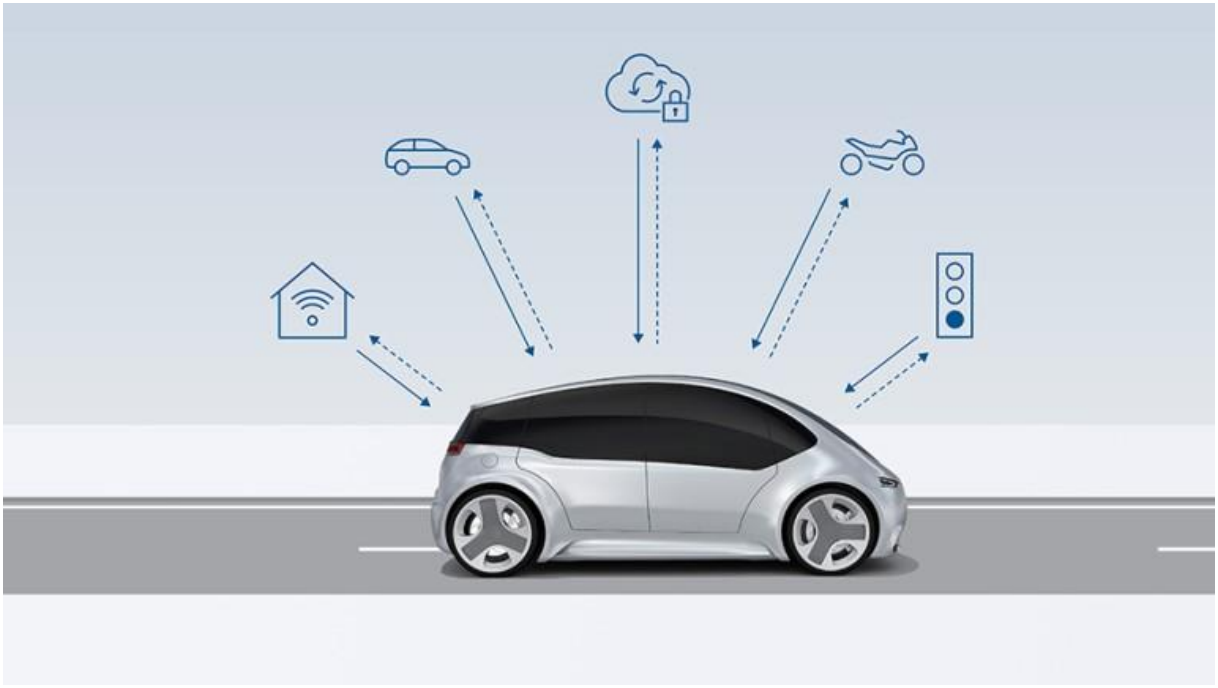

Autonomous driving is both an audacious visionary goal and a highly achievable feat of technical engineering. The race to launching the industry's first fully autonomous car is accelerating, as technology companies like Apple and Waymo battle car companies Audi, Ford, Tesla, Renault, Waymo and ride sharing companies Lyft and Uber to overcome technical challenges and enable an entirely new way of driving that will surprise and delight users.

Who will win the race? And how will driverless cars transform society? While the future of the automotive industry is still unwritten, here is what we know: Industry leaders will need to master connectivity to deliver the V2X (vehicle-to-everything) capabilities fully autonomous driving promises. Here's an overview of six key connectivity requirements to make that happen.

Architectures Must Be Redundant and Real-time

Autonomous driving will increasingly demand more and more reliable network-based structures, requiring redundant, real-time architectures. These architectures will organize high performance clusters in functional domains and be connected via a central gateway in a high-speed data backbone structure. Group sensors and actuators will be organized hierarchically.

Functional domains for connected, autonomous vehicles with high speed in-vehicle networks.

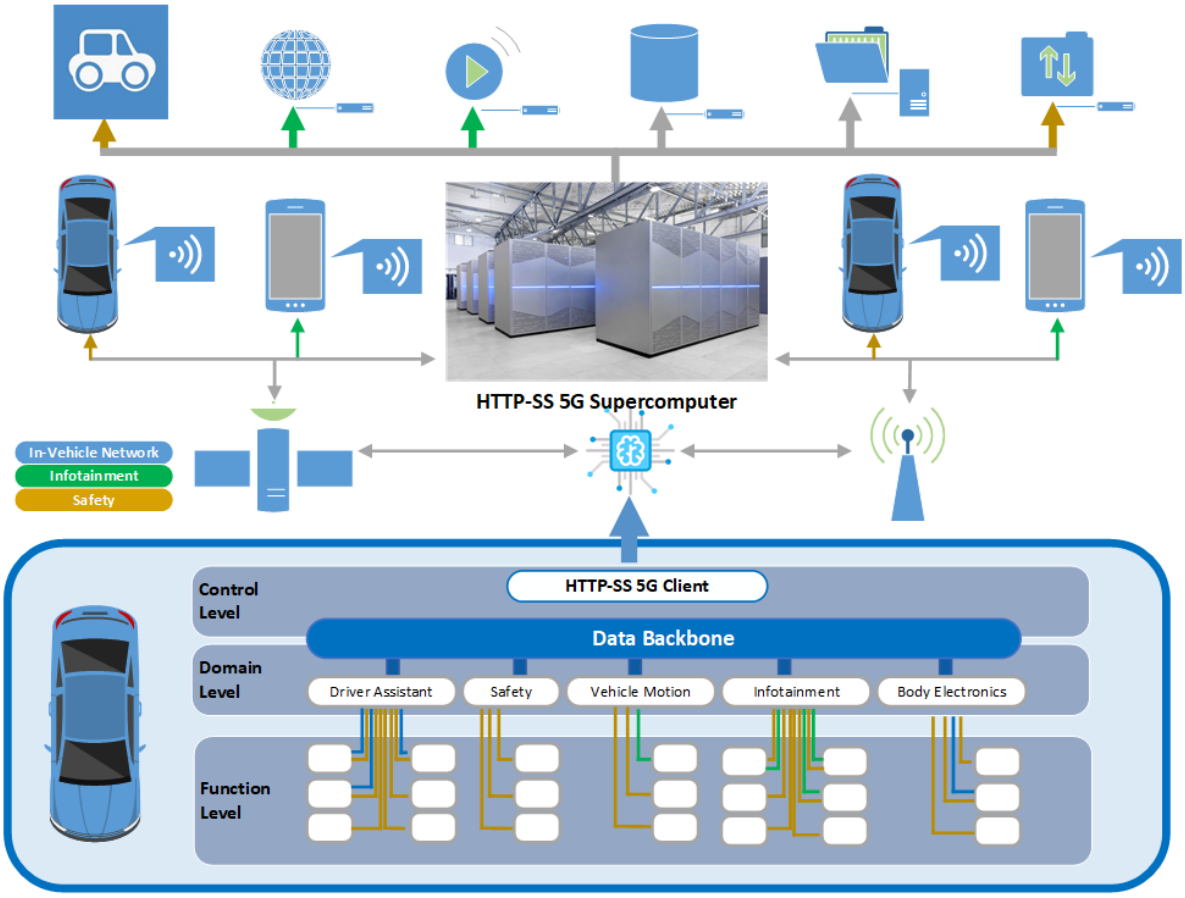

The Demand for High-Speed Data Will Only Increase

Driverless cars will have incredibly sophisticated systems, including high performance computers and an increasing number of advanced driver assistance system (ADAS) sensors, such as high-resolution stereo and/or mono cameras, RADAR, and LIDAR, as well as future human-machine interfaces (HMIs), such as large 4K/8K screens or head-up displays (HUDs).

These systems will generate vast amounts of data, which will require sophisticated electronic support, including high-speed data nodes, links, cables, and assemblies. The inside of the driverless car will literally be an information highway, and data streams will run in parallel or different directions. They'll be managed in switched networks with sufficient margin to ensure the car's safe operations.

Car manufacturers will be required to continuously increase bandwidths for both point-to-point data pipes and distributed network structures to meet new data demands. Level 5 cars will send 25 gigabytes of data to the cloud every hour.

Rock Technologies has developed a HTTP-QuSS gigabit-speed network technology that enable global automotive OEMs to meet these connectivity requirements. Rock Technologies helps automotive leaders achieve their goals, by providing connectivity solutions that overcome myriad mechanical, electrical, and software protocol and interface compatibility challenges.

Safety and Other Applications Require External Connectivity

V2X communications must continuously capture and interpret complete data concerning the surrounding environment. Fully autonomous vehicles require all of this information to correctly execute actions such as braking or accelerating without driver assistance.

V2X communications uses both sensor technology and radio-based communication. Car sensors help vehicle systems interact with their immediate surroundings. Radio systems enable vehicles to exchange information with other vehicles and with traffic infrastructure such as lights, signs and tolls. These surroundings actively communicate their state and changing conditions around them to the vehicle. Thus, a light communicates that it is about to change from red to green, a sign indicates that the next turn-off is five miles ahead, and a vehicle a lane over communicates that it is about to signal and turn right. All of these signals make driving safer. Information on driving conditions, such as icy roads ahead, traffic slowdowns, or tolls, can also enable vehicles to interpret data and make automated decisions, such as choosing another route.

Drivers make decisions in split-seconds, and fully autonomous cars will, too. To be safe, fully autonomous driving requires real-time data transmission. Current cellular radio standards such as LTE has a latency of 30 to 40 milliseconds, making it unsafe for driverless vehicles. 5G mobile communication, which will be available for deployment in 2019, will provide higher data rates of up to 10 Gbps with considerably lower latency than LTE, making it suitable for real-time safety applications within a 5G Basic

But it should be noted:

The 5G claims very short latency times of less than 1 ms but this is only possible under laboratory conditions. The total latency consists of several parts:

- The air interface, i.e. the connection from the mobile device to the base. The latency of the air interface can be realized at 5G under laboratory conditions under 1 ms.
- The latency of the data processing behind the base station towards the telecommunications network.
- Latency to the Internet: If the user accesses applications and data on the Internet, the latency of the participating servers (which are completely independent of the 5G standard) is added.
- The real end-to-end latency reached by US phone provider Verizon in Chicago in March 2019 is in the 30 ms range.

HTTP-QuSS will solve this hidden 5G disadvantage by eliminating this well-known **TCP Latency** problem which also destroys 1 Gbit/s Bandwidth by over 80 %.

Cybersecurity Threats?

Just like any other computer-enabled device, however, driverless cars are prone to cybercrimes. Criminals might be motivated to hack into the vehicles' operating systems and steal important passenger data, or else disrupt its operation and jeopardize the passenger's safety. Some of the possible cyber-precipitated scenarios that you can expect with the full adoption of driverless cars include:

- Criminals hacking cars for ransom before allowing the user either in or out of the car. This can happen when the car is parked or driving.

- Terrorists hijacking the network and taking control over a transport system in each area. Hacking a network can cause major crashes by disabling the light-detecting and ranging sensors, leading to endless confusion.

- Hacking the car's operating system remotely to destroy it could harm the user financially.

- As with any other hacking scenario, hacking into an autonomous car could expose a great deal of your personal data—including your destination. With this information, someone could potentially track the user with an aim toward robbery or assault. If hackers can gain access to the controls of the vehicle, it could also be possible to redirect the vehicle to a more convenient location for either of those scenarios.

- As the technology evolves, driverless cars will be able to turn on any smart device in your home, be it the TV, heater, garage door, or front gate, and everything programmable in the home. Hackers could use these features to gain access to your home.

The threats and benefits of autonomous cars will not be entirely clear until the technology is fully rolled out. The good thing is that even with the many probable cybersecurity threats the **HTTP-QuSS Supercomputer** has an integrated Active Cybersecurity which eliminates all these threads before reaching the car data backbone vie the HTTP-QuSS Client

What is HTTP-QuSS

HTTP-QuSS consists of an **Embedded Automotive Supercomputer** of the newest Generation stationed at an IP Backbone Site and a **S**ystem **o**n a **C**hip (**SoC)** integrated in **self-driving Cars**.

The **HTTP-QuSS Supercomputer** eliminates the Bandwidth destroying Latency Issue caused by Protocol Handshakes, Long Distances, many Hops, Data Losses, congested Networks by many Users etc. and therefore existing Wi-Fi and 4G Networks are able to provide 5G Requirements like low Latencies and Gbit/s Bandwidths. 5G Networks on the other Hand can even fulfil their own claimed Features when accessing the normal Internet and not only connected Cars within a Basic Cell.

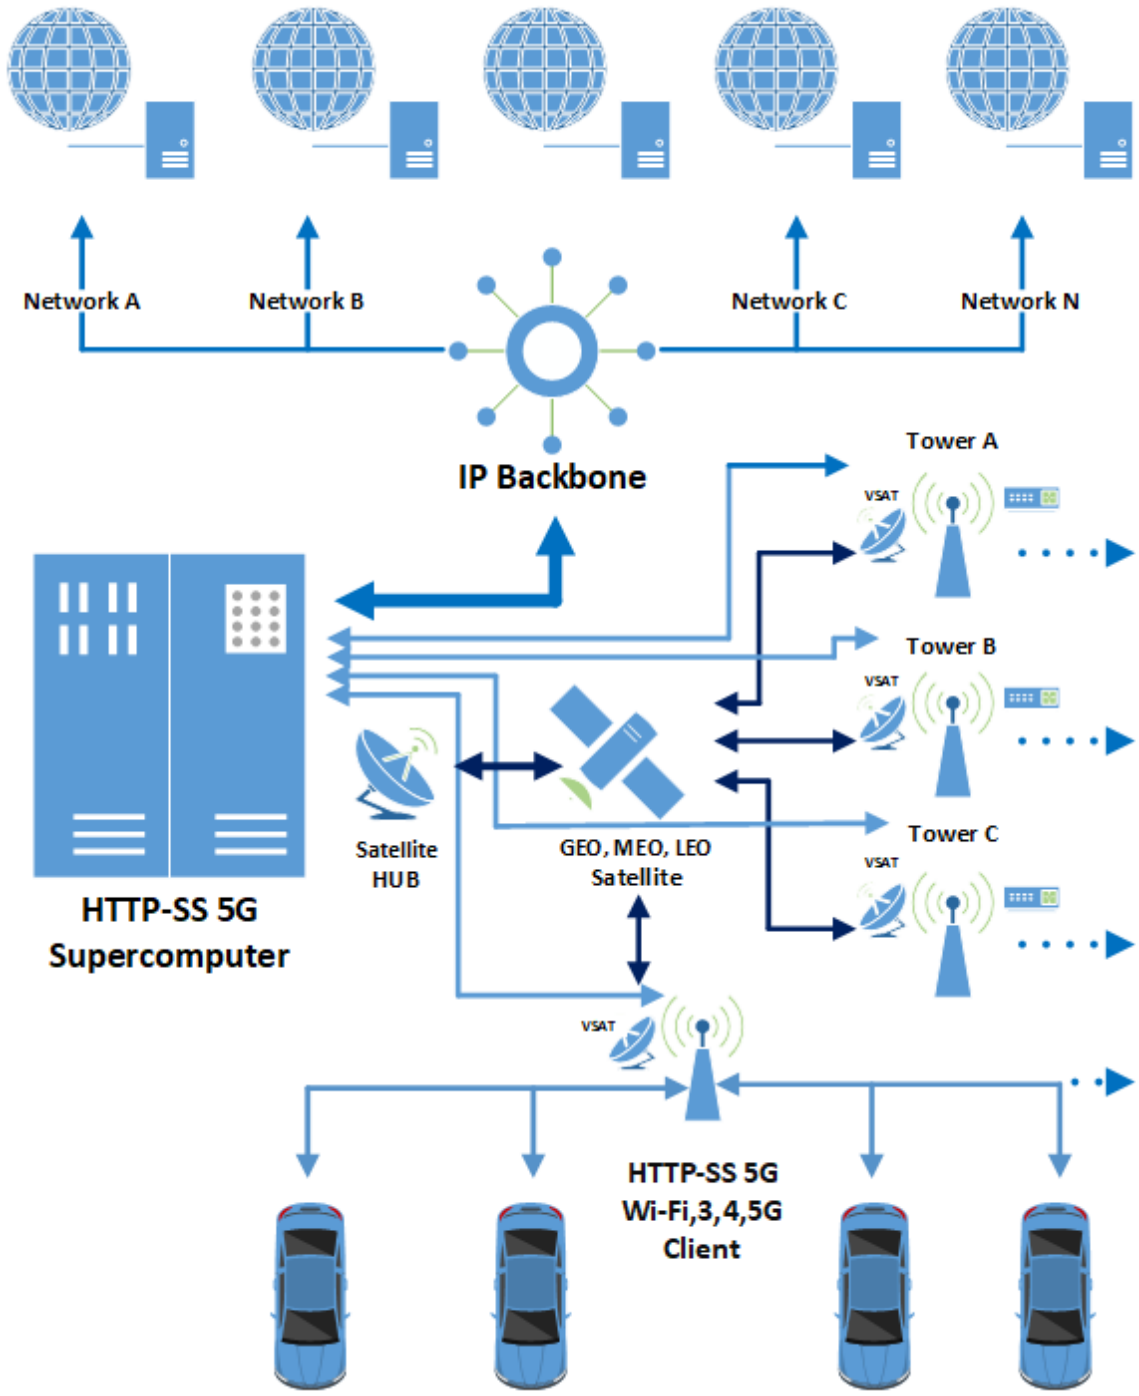

Unique Selling Proposition

- **Total Elimination of the TCP Latency Problem and related Bandwidth Losses Within Wi-Fi, 3,4,5G Communication Networks**
  - ✓ Through a new **HTTP 5G S**ingle **S**tream Architecture and Technology
  - ✓ By using the next Generation of Embedded Server Supercomputer Hardware
  - ✓ By Artificial Intelligence supported Data Transmission and highly Parallelized Process Chains
  - ✓ Integrated Dynamic Bandwidth Shaping and Slicing for max Bandwidth Usage
- **Highly reduced Data Transmission for all WEB Objects and Files**
  - ✓ By AI supported WEB Object Push and File Descriptor Delta Data Algorithm
  - ✓ And 1 Round Trip Protocol Handshake
  - ✓ and therefore 90 % reduced secure Data Transmission
  - ✓ All WEB- and TCP Applications are supported
- **5G Performance and Latency for existing 4G and Wi-Fi Towers**
  - ✓ With its own highly efficient Process and Processor Management
  - ✓ Integrated parallel Processing Architecture without Site Effects
  - ✓ Smart parallel Process Chains with highly efficient Inter Communication
  - ✓ And much more …
- **Quantum Secure Cyber Security**
  - ✓ Through keyless 2 Level Data Encryption
- **Pure Client Software Solution for legacy Devices**
  - ✓ No need of special Hardware on Client Site
  - ✓ No need of special Browser or Proxy Settings on Client Site
- **Fast and transparent Wi-Fi, 3,4,5G Tower Installation and Integration**
  - ✓ No need of Reconfigurations in existing Infrastructure
  - ✓ Fully transparent for User
  - ✓ Supports all Devices PC''s, Laptops, Smartphones (Android, iOS), DSL- and VSAT Routers etc.etc.

Major Advantages for autonomous Driving Apps

- Best use of limited Resources (Bandwidths)
- No Bandwidth Losses through TCP Latency Issue
- Fulfils Realtime Automotive Requirements even beyond 4/5G Basic Stations
- More efficient and faster Data Transmission

  - ✓ Faster WEB Page Loading
  - ✓ Significantly better Performance
  - ✓ All TCP/UDP Sensor/Security and Infotainment Apps are supported

- Transparent and fast Integration into existing Car Infrastructure
  - ✓ Part of Car Automotive Operating System
  - ✓ No need of extra Hardware
- No need of 5G Contract
- Can reduce his monthly Costs dramatically because Data Volume will be reduced by **90 %**
- No upgrade of existing Provider Contracts
- In Satellite, WLAN / Wi-Fi, and Mobile Networks for no Latency Issues
- Much better Performance for Cloud Applications
- Fast Cloud Data Backup
- Broadband Availability, even in rural Areas

- Supports upcoming Low Earth Orbit Satellite Networks by avoiding heavy 1 Gbit/s Bandwidth Losses caused by Round Trip Times > 30 ms

Transparent Wi-Fi, 3,4,5G, Satellite Network Integration

- No need to change Infrastructure by fully **transparent Network Integration**
- Supports **all common Browsers** and **TCP Application**

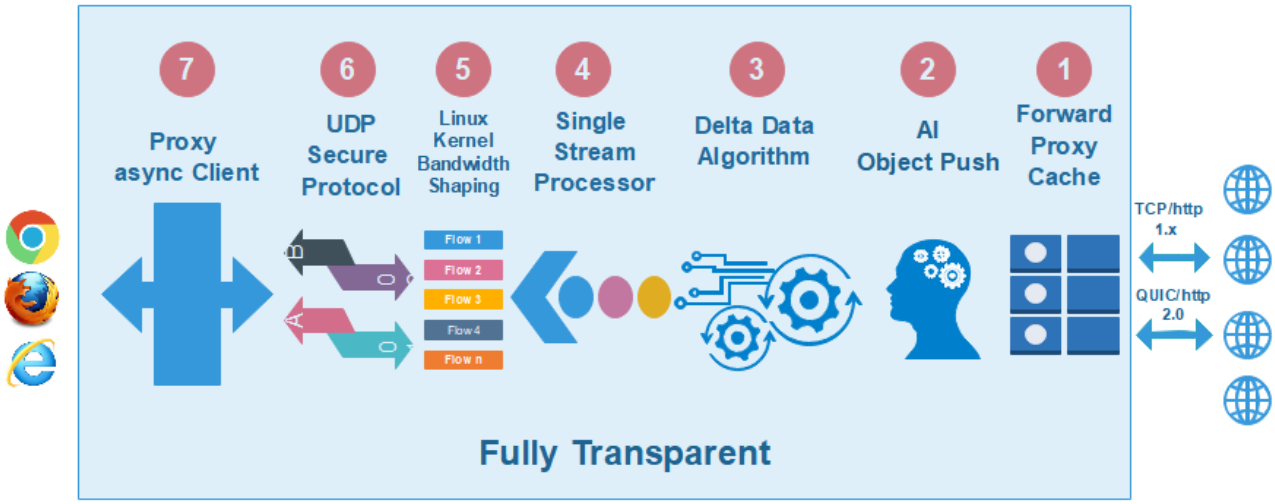

No Latency to Bandwidth Dependency

- **No Dependency** between **R**ound **T**rip **T**ime and **TCP Bandwidth**
- **Breakthrough** in higher bandwidths Regions even at very long Round Trips

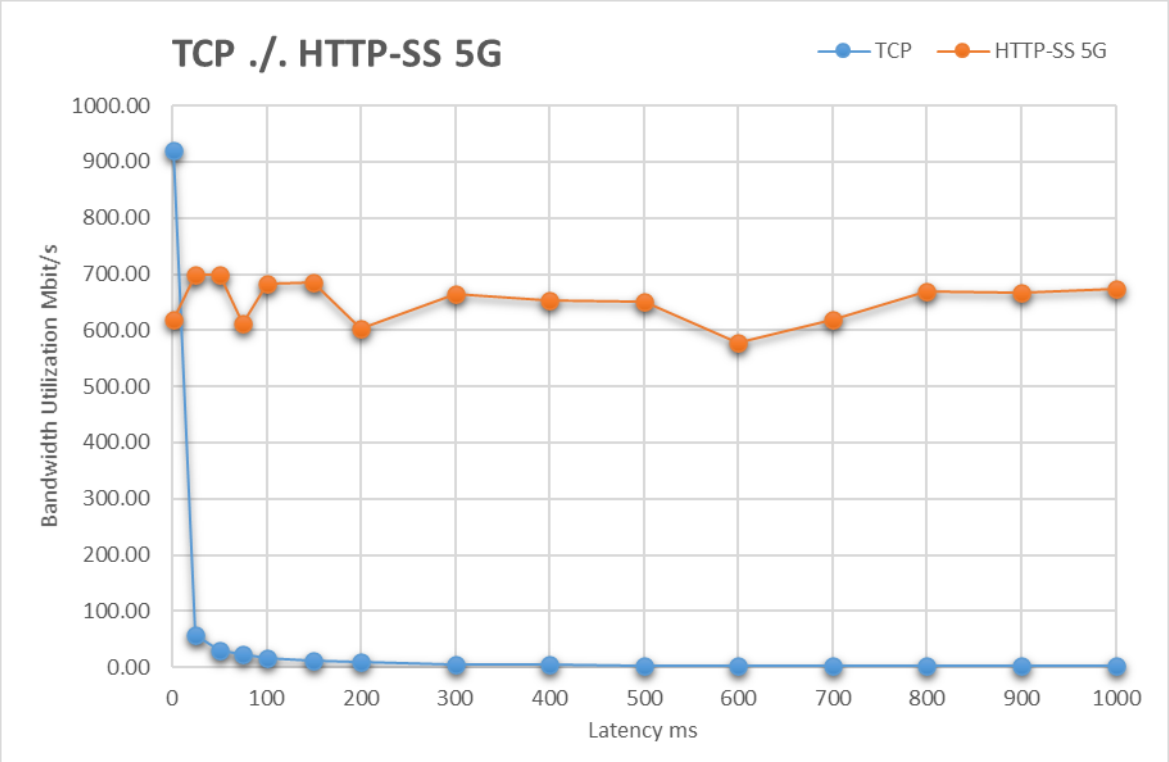

HTTP-QuSS - Automotive Backbone Supercomputer

NVIDIA - DGX A100

Yesterday

Today with NVIDIA A100 Supercomputer

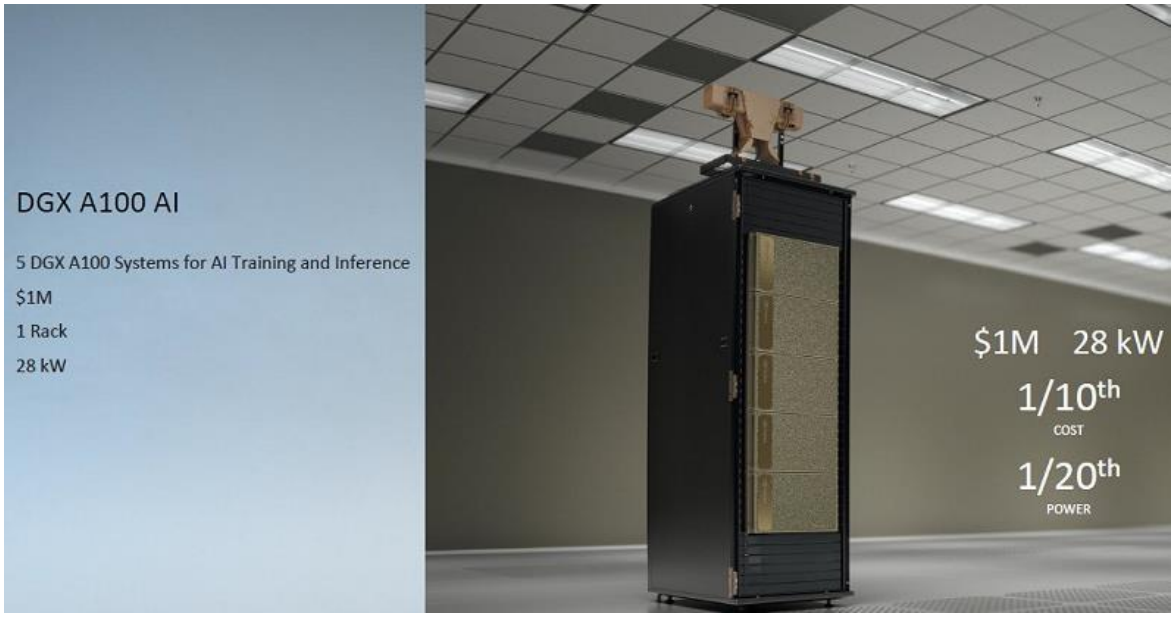

NVIDIA - A-100

The NVIDIA A100 Tensor Core GPU delivers unprecedented acceleration at every scale for AI, data analytics, and HPC to tackle the world's toughest computing challenges. As the engine of the NVIDIA data center platform, A100 can efficiently scale up to thousands of GPUs or, using new Multi-Instance GPU (MIG) technology, can be partitioned into seven isolated GPU instances to accelerate workloads of all sizes. A100's third-generation Tensor Core technology now accelerates more levels of precision for diverse workloads, speeding time to insight as well as time to market.

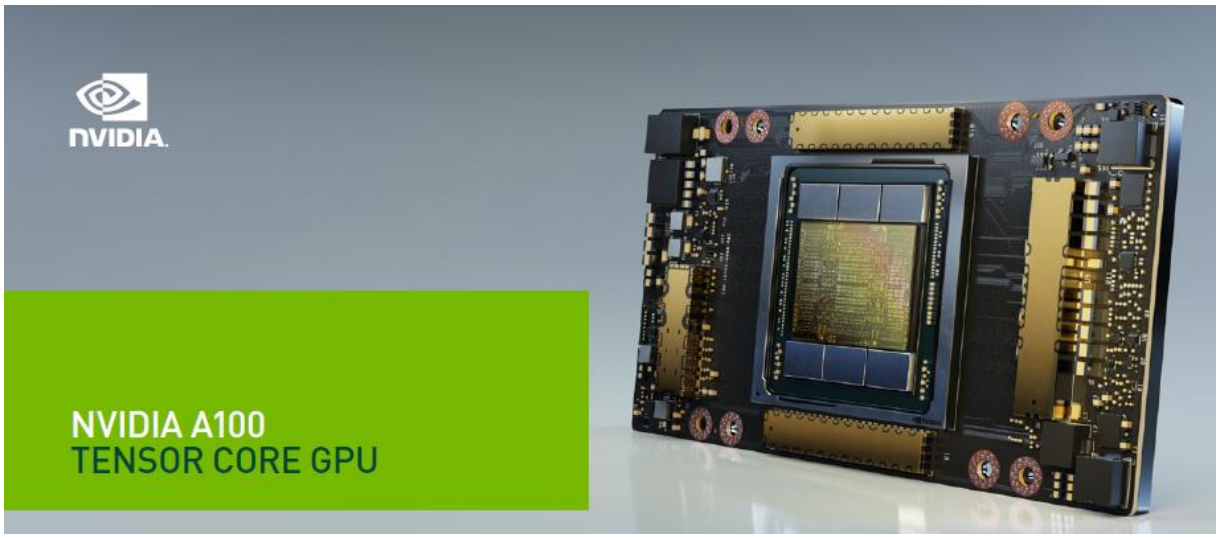

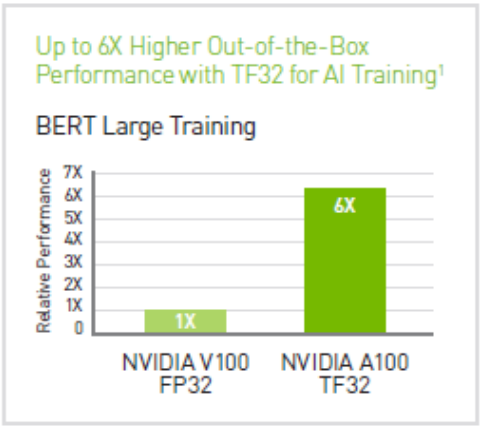

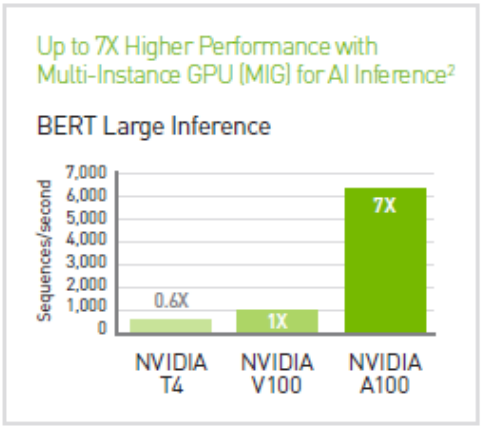

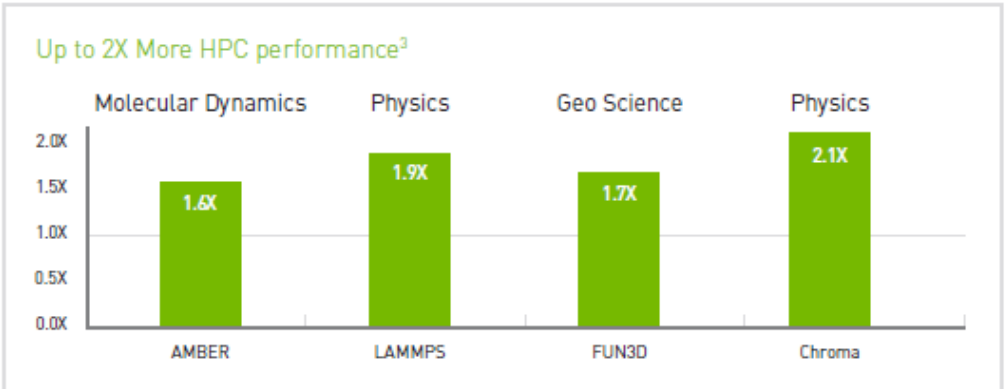

SYSTEM SPECIFICATIONS (PEAK PERFORMANCE)

	NVIDIA A100 for NVIDIA HGX™	NVIDIA A100 for PCIe
GPU Architecture	NVIDIA Ampere	
Double-Precision Performance	FP64: 9.7 TFLOPS FP64 Tensor Core: 19.5 TFLOPS	
Single-Precision Performance	FP32: 19.5 TFLOPS Tensor Float 32 (TF32): 156 TFLOPS | 312 TFLOPS*	
Half-Precision Performance	312 TFLOPS | 624 TFLOPS*	
Bfloat16	312 TFLOPS | 624 TFLOPS*	
Integer Performance	INT8: 624 TOPS | 1,248 TOPS* INT4: 1,248 TOPS | 2,496 TOPS*	
GPU Memory	40 GB HBM2	
Memory Bandwidth	1.6 TB/sec	
Error-Correcting Code	Yes	
Interconnect Interface	PCIe Gen4: 64 GB/sec Third generation NVIDIA® NVLink®: 600 GB/sec**	PCIe Gen4: 64 GB/sec Third generation NVIDIA® NVLink®: 600 GB/sec**
Form Factor	4/8 SXM GPUs in NVIDIA HGX™ A100	PCIe
Multi-Instance GPU (MIG)	Up to 7 GPU instances	
Max Power Consumption	400 W	250 W
Delivered Performance for Top Apps	100%	90%
Thermal Solution	Passive	
Compute APIs	CUDA®, DirectCompute, OpenCL™, OpenACC®	

* Structural sparsity enabled

** SXM GPUs via HGX A100 server boards; PCIe GPUs via NVLink Bridge for up to 2 GPUs

GROUNDBREAKING INNOVATIONS

NVIDIA AMPERE ARCHITECTURE

A100 accelerates workloads big and small. Whether using MIG to partition an A100 GPU into smaller instances, or NVLink to connect multiple GPUs to accelerate large-scale workloads, A100 can readily handle different-sized acceleration needs, from the smallest job to the biggest multi-node workload. A100's versatility means IT managers can maximize the utility of every GPU in their data center around the clock.

THIRD-GENERATION TENSOR CORES

A100 delivers 312 teraFLOPS (TFLOPS) of deep learning performance. That's 20X Tensor FLOPS for deep learning training and 20X Tensor TOPS for deep learning inference compared to NVIDIA Volta™ GPUs.

NEXT-GENERATION NVLINK

NVIDIA NVLink in A100 delivers 2X higher throughput compared to the previous generation. When combined with NVIDIA NVSwitch™, up to 16 A100 GPUs can be interconnected at up to 600 gigabytes per second (GB/sec) to unleash the highest application performance possible on a single server. NVLink is available in A100 SXM GPUs via HGX A100 server boards and in PCIe GPUs via an NVLink Bridge for up to 2 GPUs.

MULTI-INSTANCE GPU (MIG)

An A100 GPU can be partitioned into as many as seven GPU instances, fully isolated at the hardware level with their own high-bandwidth memory, cache, and compute cores. MIG gives developers access to breakthrough acceleration for all their applications, and IT administrators can offer right-sized GPU acceleration for every job, optimizing utilization and expanding access to every user and application.

HBM2

With 40 gigabytes (GB) of high-bandwidth memory (HBM2), A100 delivers improved raw bandwidth of 1.6TB/sec, as well as higher dynamic random-access memory (DRAM) utilization efficiency at 95 percent. A100 delivers 1.7X higher memory bandwidth over the previous generation.

STRUCTURAL SPARSITY

AI networks are big, having millions to billions of parameters. Not all these parameters are needed for accurate predictions, and some can be converted to zeros to make the models "sparse" without compromising accuracy. Tensor Cores in A100 can provide up to 2X higher performance for sparse models. While the sparsity feature more readily benefits AI inference, it can also improve the performance of model training.

The NVIDIA A100 Tensor Core GPU is the flagship product of the NVIDIA data center platform for deep learning, HPC, and data analytics. The platform accelerates over 700 HPC applications and every major deep learning framework. It's available everywhere, from desktops to servers to cloud services, delivering both dramatic performance gains and cost-saving opportunities.

HTTP-QuSS Client - integrated in NVIDIA - DRIVE AGX Orin

Advanced, Software-Defined Platform for Autonomous Machines

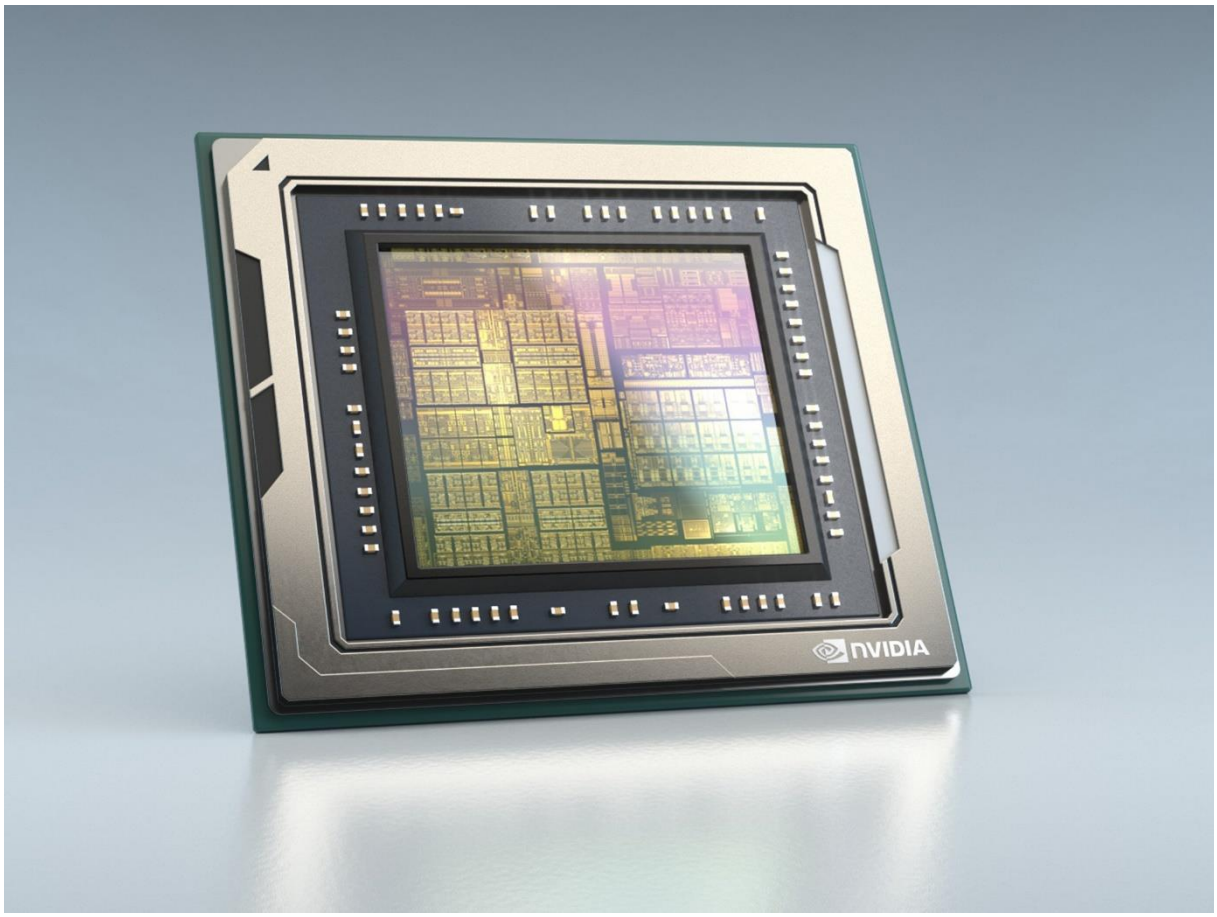

NVIDIA today introduced NVIDIA DRIVE AGX Orin™, a highly advanced software-defined platform for autonomous vehicles and robots.

The platform is powered by a new system-on-a-chip (SoC) called Orin, which consists of 17 billion transistors and is the result of four years of R&D investment. The Orin SoC integrates NVIDIA's next-generation GPU architecture and Arm Hercules CPU cores, as well as new deep learning and computer vision accelerators that, in aggregate, deliver 200 trillion operations per second—nearly 7x the performance of NVIDIA's previous generation Xavier SoC.

Orin is designed to handle the large number of applications and deep neural networks that run simultaneously in autonomous vehicles and robots, while achieving systematic safety standards such as ISO 26262 ASIL-D.

Built as a software-defined platform, DRIVE AGX Orin is developed to enable architecturally compatible platforms that scale from a Level 2 to full self-driving Level 5 vehicle, enabling OEMs to develop large-scale and complex families of software products. Since both Orin and Xavier are programmable through open CUDA and TensorRT APIs and libraries, developers can leverage their investments across multiple product generations.

“Creating a safe autonomous vehicle is perhaps society’s greatest computing challenge,” said Jensen Huang, founder and CEO of NVIDIA. “The amount of investment required to deliver autonomous vehicles has grown exponentially, and the complexity of the task requires a scalable, programmable, software-defined AI platform like Orin.”

“NVIDIA’s long-term commitment to the transportation industry, along with its innovative end-to-end platform and tools, has resulted in a vast ecosystem — virtually every company working on AVs is utilizing NVIDIA in its compute stack,” said Sam Abuelsamid, principal research analyst at Navigant Research. “Orin looks to be a significant step forward that should help enable the next great chapter in this ever-improving technology story.”

The NVIDIA DRIVE AGX Orin family will include a range of configurations based on a single architecture, targeting automakers’ 2022 production timelines.