How to Make Autonomous Vehicles a Reality with Arm

arm

Guide

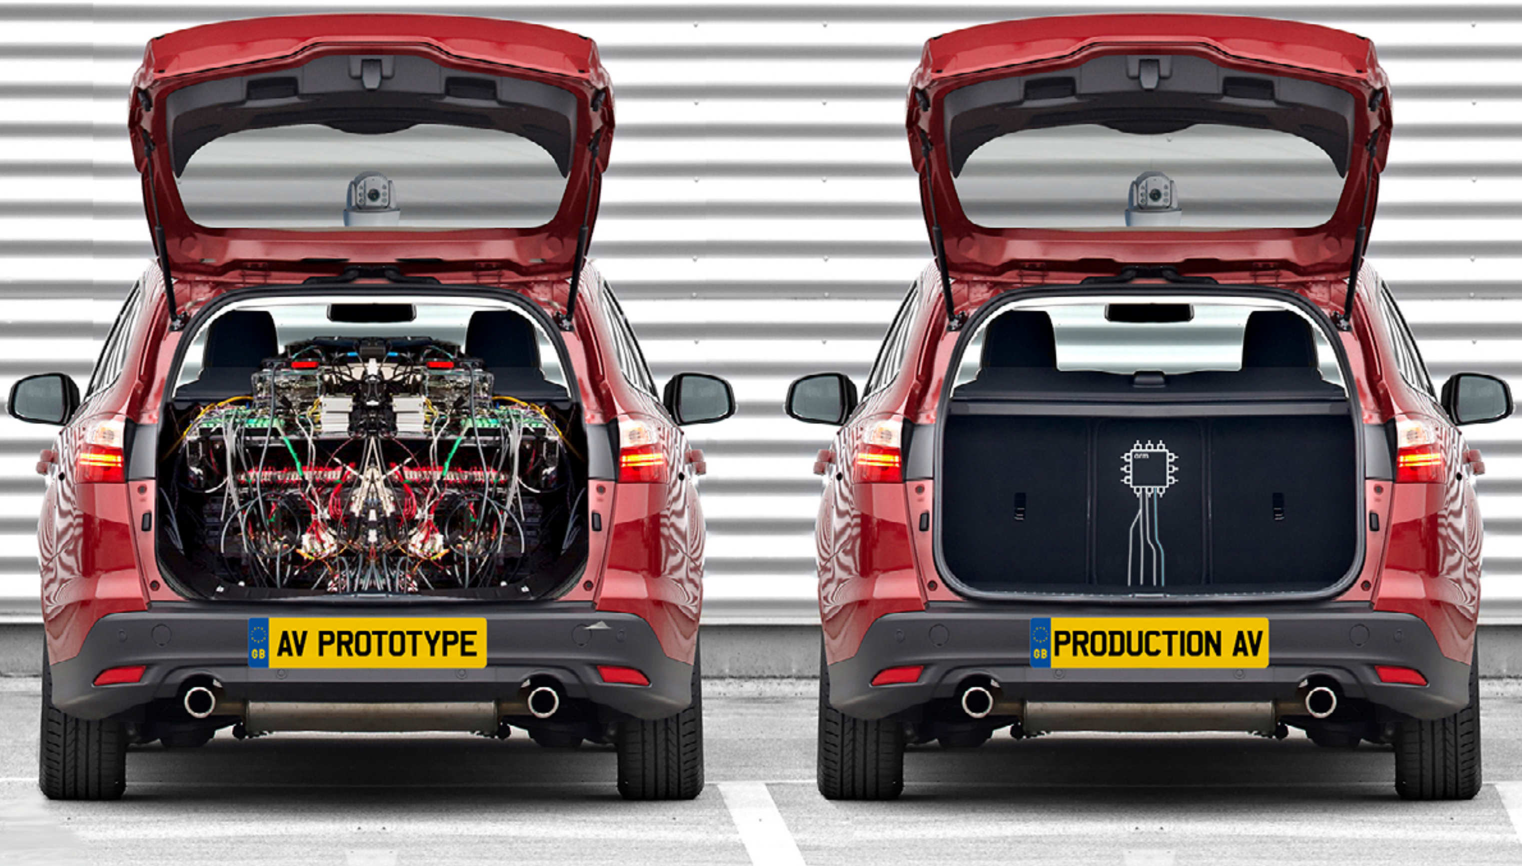

Challenges Impacting the Mass Deployment of Autonomous Vehicles

The automotive industry is currently looking at the technology innovation needed to move from today's prototype autonomous vehicles to deployable safe, self-driving solutions. This technology must be able to tackle the key challenges that are currently preventing us from reaching and producing safe Level 4 and Level 5 autonomous vehicles.

We're seeing the complexity of autonomous automotive systems growing at an unprecedented rate, and computational processing must keep pace with this growth without compromising the current challenges of power consumption, thermal properties, size and cost, safety and security. To add to these technology challenges, there are still many debates about consumer and regulatory acceptance to full autonomy. For example, a survey from the US-based AAA [1] showed seventy-three percent of Americans are too afraid to ride in fully self-driving vehicles. Another social and technological challenge is that it is very hard to have autonomous cars sharing the road with human drivers like us.

[1] – https://newsroom.aaa.com/2018/05/aaa-american-trust-autonomous-vehicles-slips/

Challenges to Deployment

1. Price of Autonomy

It has been suggested that if produced in 2020, a Level 4 and Level 5 car could cost between an extra $75,000 to $100,000 compared to a regular car. Truthfully, this figure may even be too low as the total cost is likely to exceed $100,000 when considering the number of sensors needed to achieve Level 4 and Level 5 autonomy. To make the purchasing of these vehicles feasible, the price must come down dramatically to make it affordable for consumers. It is likely that this high price means that the first real deployments of autonomous vehicles will be part of Mobility-as-a-Service (MaaS), ride-sharing or robotaxi fleets. By replacing the cost of a human driver, and by driving a much higher utilization of the vehicles than consumers, these entities could build a business model that can support these more expensive vehicles.

2. Is Level 3 a Deployable Reality?

Fig 1: SAE levels of autonomy showing shift from ADAS to Autonomous.

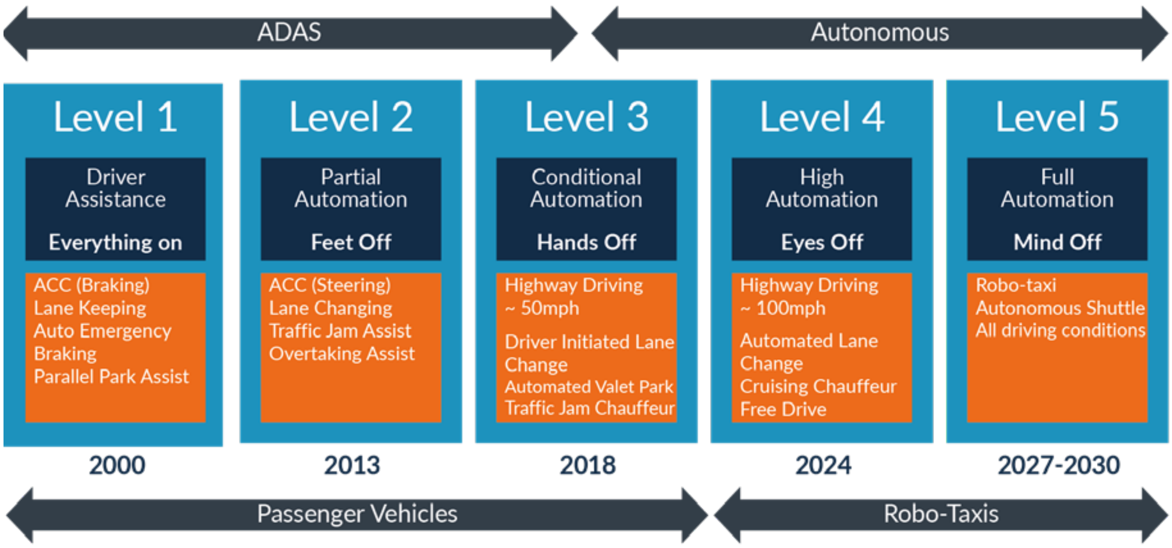

As shown in the diagram above, Level 3 is the first step in the move from advanced driver assistance systems (ADAS) to autonomy. However, there is currently some debate over Level 3 autonomy and the requirements put on both the vehicle and the driver. Successfully deployed Level 3 autonomy requires the driver to be alert when the vehicle's self-driving functions are active. However, we as drivers will instinctively assume that as soon as we take our hands off the wheel, we no longer need to pay attention, and can quite happily do our email, send texts etc, which takes both our eyes and our minds off the road. However, with Level 3 the car can ask you to take back control of the vehicle at any time. This raises the issue of how quickly a distracted driver can return to the wheel and take back control of the vehicle to manage the situation that the autonomous car couldn't handle.

Some car manufacturers are currently discussing skipping Level 3 as a way to overcome this challenge. Also, from a liability perspective, skipping Level 3 would make it easier to determine if the driver is controlling the vehicle or if the vehicle is self-driving. There are also discussions about advanced driving monitoring systems using in-cabin cameras and advanced software algorithms to determine whether the driver is alert, and fit to take back control, and if not, activate the appropriate warning to bring the driver back to full readiness. Even if car manufacturers decide to skip this level, the technology complexity required to get from Level 3 to Level 4 is much greater.

3. Processing Needs of Dramatically Increased Sensor Complexity

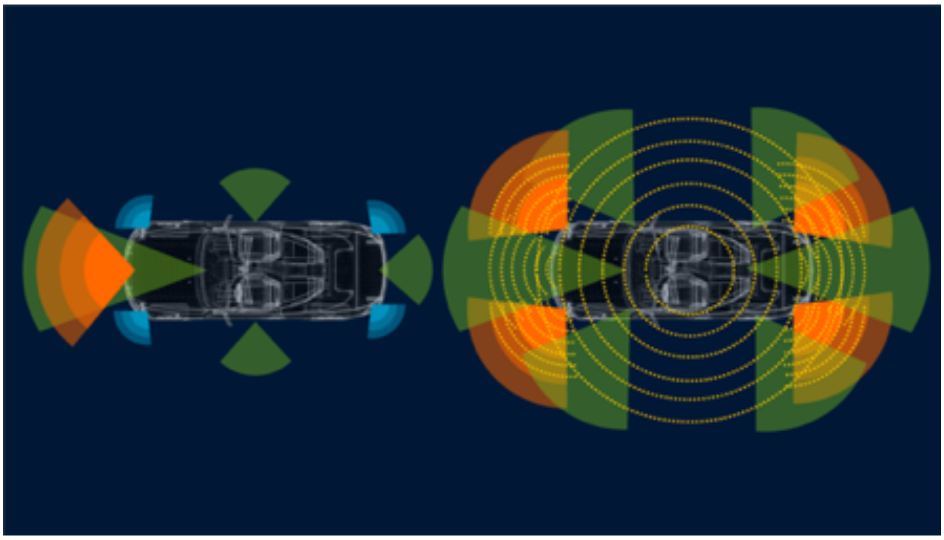

Fig 2: Shift from ADAS to Autonomous requires greater sensor awareness.

The move from ADAS to autonomous demands a much greater awareness of everything around the car. To accomplish this, the number of sensors on the car are dramatically increasing, with multiple LiDAR, camera and radar sensors required to essentially replace and enhance human sight and situational awareness. Not only are these sensors expensive, but the processing required to understand what they are "seeing" and the situation evolving outside the car is dramatically different to the compute required by simpler ADAS functions such as adaptive cruise control or emergency braking.

4. Increased Software Complexity

The majority of autonomous vehicles being prototyped right now are essentially testing the increased sensor complexity and the software algorithms needed to process the large amount of information coming into the car, make the right decision about what to do and then action it. This processing requires a substantial amount of software, with current estimates around 1 billion [2] lines of code to power a fully autonomous car. The compute requirements to execute this large amount of software are more akin to server performance than traditional automotive embedded processing. This is driving a trend towards a consolidation of much more powerful clusters of application processors and accelerators in more performant multicore SoCs rather than discrete CPUs. This consolidation requires a dramatic change in software architectures, and can also cause a dramatic increase in the software footprint.

[2] – https://www.talismanexecutive.com/blog/needs-happen-make-autonomous-vehicles-reality

The software application complexity is far greater than even the most advanced passenger jets already brimming over with autonomous functionality, because autonomous cars will have to deal with highly chaotic roads full of unpredictable human drivers and pedestrians vs. the relatively empty skies full of professional pilots and a mature global network of air traffic control. This leads to a large amount of algorithmic processing that must happen in real time to understand everything that is happening around the car. It also requires a huge software stack for all of the autonomous compute components to make the right decisions and execute them safely. This greater complexity lends itself to a common and unified platform architecture on which to build an easily upgradable and portable software stack.

Fig 3: Safety is a key part of the mass deployment and acceptance of self-driving cars.

5. Safe Deployment of Autonomous Vehicles

As stated earlier, recent statistics have shown that 73% of American drivers are too afraid to ride in fully self-driving vehicles and, astonishingly, 63% of US adults would feel less safe sharing the road with self-driving vehicles while walking or cycling. This raises a new and interesting challenge of how we gain consumer trust, both as a passenger in an autonomous car, and as someone sharing the same environment with the car.

Safety is a key part of many automotive systems, and rigorous safety standards and certifications are applied to any functions that need to work reliably when requested by the driver, such as brakes and steering. When we increase autonomy in a car, we are essentially replacing the safe decision making of a human driver with a complex computer system comprising many heterogenous compute elements and, as discussed earlier, up to a billion lines of code. How can we guarantee that this hugely complicated compute system executes to the highest levels of passenger and environmental safety?

With the consolidation of functions onto powerful multi-core SoCs, there will also be the need to support mixed-criticality applications on a single SoC. This is where some applications will require the highest levels of functional safety as they are executing life critical functions, mixed with applications operating at a lower criticality level. It would be impossible to try to take all the software to the highest level of functional safety, and so a compute and software architecture must support these different safety levels without having to dedicate a separate SoC for each application.

6. Prototype to Production

The compute systems going into today's autonomous prototypes are typically based on off-the-shelf server technology. The challenge with server technology is that the size, power consumption, and thermal properties are not suitable for cars. There must be a significant reduction in all of these current attributes. The common belief is that power consumption must be reduced by 10x, size by 5x, and if both of these can be achieved then there will be a significant reduction in cost and dissipated heat, which also leads to simpler and more reliable cooling methods. These improvements will lead to the true deployment of self-driving cars, both in the consumer space and robotaxis.

Fig 4: The shift from driver to passenger will change the cabin environment for self-driving cars.

7. Enhanced In-Car Passenger Experiences

There is a growing trend within the cabin for consumers to demand a more enhanced and enriched in-car experience. As we get to higher levels of autonomy, vehicle occupants will turn from drivers to passengers and their requirements for information, entertainment and connectivity will be more akin to their home and office.

Before we arrive at full autonomy, an interesting hybrid of driver and environmental information will fuse with entertainment and productivity features. This will pose the interesting challenge of mixing safety into the existing feeds, and ensuring that the driver safety information is not compromised by the other forms of information being displayed.

As we move into the next 5 to 7 years, to a more autonomous world, there will be different kinds of information delivered to the screens including driver information from autonomous systems, media experiences, driver monitoring systems, sensors facing inside the car, all of which will be helping to deliver a more personalized in-car experience. This requires high throughput capability for delivering data to screens, high bandwidth connectivity, and enhanced safety, especially for critical information such as driver warning information.

How Arm Can Help with Your Autonomous Systems

Automotive OEMs and Tier 1s increasingly recognize the need for a strong technology partner to help them address the challenge of getting autonomous vehicles from prototype to mass deployment and view Arm and its broad automotive ecosystem as the right partnership to make this happen.

Arm has been working closely with the automotive industry to understand the challenges of delivering autonomous vehicles and is now providing new solutions to help power the production of fully autonomous vehicles at scale.

1. All Processing Needs are Covered

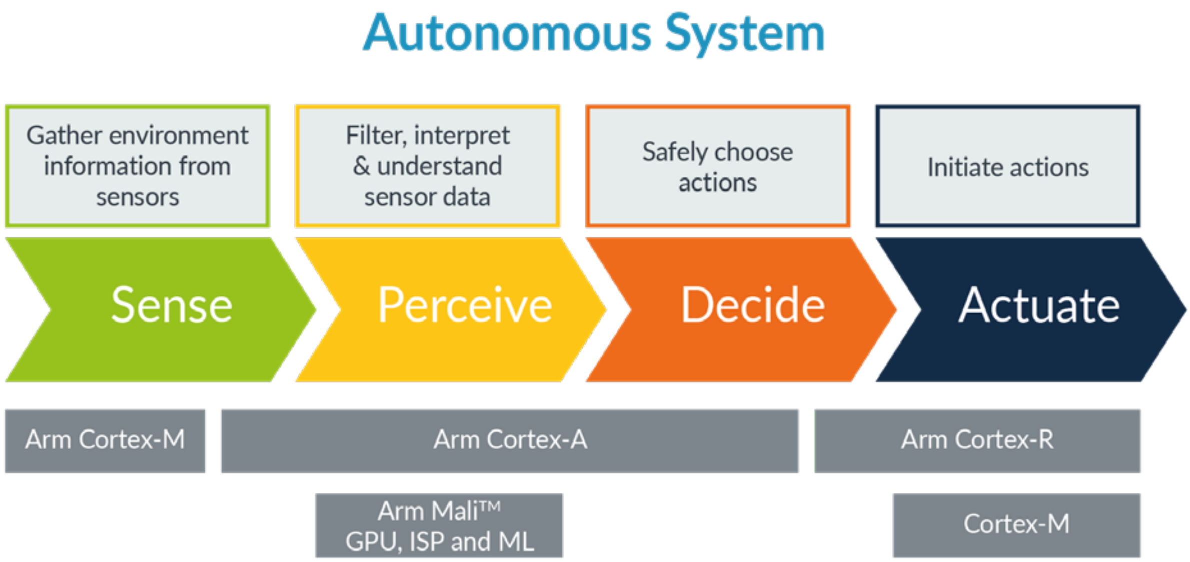

Fig 5: Arm-the foundation for autonomous systems.

The unparalleled range of Arm CPUs and other IP elements such as GPUs, ISPs and NPUs allows Arm-based solutions to be used throughout the whole vehicle, with the broadest set of automotive-grade SoCs being offered by Arm's semiconductor partners. This range of application processors (Cortex-A), real-time processors (Cortex-R) and small, low-power, microcontrollers (Cortex-M) fits across all the phases of an autonomous system, as shown above. As Arm's semiconductor partners bring more of these compute elements onto single heterogeneous SoC platforms, this will help to meet the processing requirements, and at the same time help reduce the power, price, size and thermal properties.

2. Safety Ready

Safety Ready

The Arm Safety Ready portfolio encompasses Arm's existing and future products that have been through a rigorous functional safety process, including systematic flows and development in support of ISO 26262 and IEC 61508 standards. Safety Ready is a one-stop shop for software, tools, components, certifications, and standards to simplify and reduce the cost of integrating functional safety for Arm partners. By taking advantage of these portfolio offerings, partners and car makers can be confident their SoCs and systems incorporate the highest levels of functional safety required for autonomous applications.

Arm has recently introduced it's Functional Safety Partnership Program, which introduces a broad group of safety partners that specialize in software and tools, design services and training, to help the automotive industry bring safe deployable autonomy to market.

3. Delivering Performance and Safety

Fig 6: Split-Lock Technology: Flexible Operation Modes.

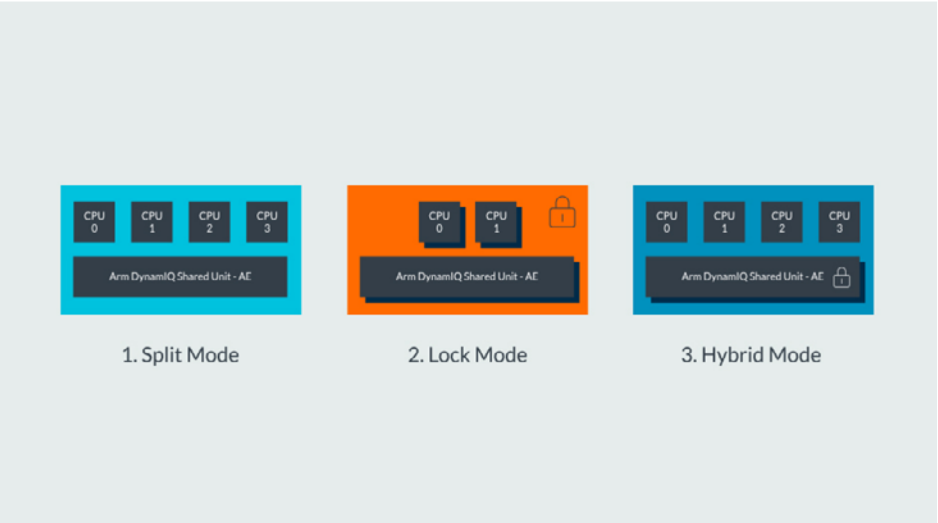

Arm has also taken another huge step forward in meeting the innovation needs by adding AE features to some of its key technologies. One of these features is Split-Lock, first introduced in the Cortex-R52 real-time processor, this feature is now a mainstay of Arm's mission-critical application processor line-up, currently comprised of Cortex-A78AE, Cortex-A76AE and Cortex-A65AE, targeting the high-performance safety requirements for autonomous sensing and perception processing.

Split-Lock enables clusters of processors to be split for performance and locked together for higher levels of safety. In addition, Arm has introduced a Hybrid mode for the Cortex-A78AE, where the processor cores remain split for high performance, but the control common logic around the CPU cluster, the DynamIQ Shared Unit, runs in lock mode. The choice of Split, Hybrid or Lock mode can be made at boot time, bringing maximum flexibility for different use cases of the same device.

The current second generation of split lock offers the following advantages:

- Support up to 4 lock-stepped application threads ("lock") or up to 8 regular threads ("Split")
- Support intermediate levels of Safety with the newly introduced Hybrid mode
- Achieve all safety modes with no impact to silicon integrators or application software developers
- Comply with development and error-reporting requirements set forth by standardsISO26262/IEC61508
- Software Test Library support for targeted safety deployments
- Post-silicon choice of mode of operation, enabling multi-modal deployment of SoC.
- CPU-cluster-level independence of mode of operation for mixed safety criticality

This innovation will help enable Arm's semiconductor partners build safety-capable SoCs that can be configured at deployment to have mixed-criticality elements on a single SoC, and also use the same SoC for multiple applications in a vehicle. This helps greatly with the deployment of mixed-criticality systems, such as an ECU running both ADAS and IVI functionality and also helps keep the costs down as single parts can be used across a wide set of automotive applications.

4. Performance, Safety and Reduced Power

The first generation Cortex-AE CPUs offered an alternative to server-class chipsets for autonomous processing while freeing system developers from the burden of safety around the compute complex. The second generation Cortex-A78AE follows from where the Cortex-A76AE left off by offering a 30% uplift in performance at the same frequency and process node. The CPU achieves this feat at comparable or lower power efficiency thereby enabling automotive customers to dramatically lower their investments thru application consolidation and optimized system design. Coupled with the ability to be paired in SMP/AMP clusters with Cortex-A65AE, Cortex-A78AE opens the door wide open to deployable autonomy.

Fig 7: Cortex-A78AE

Efficient single thread performance for integrated safety.

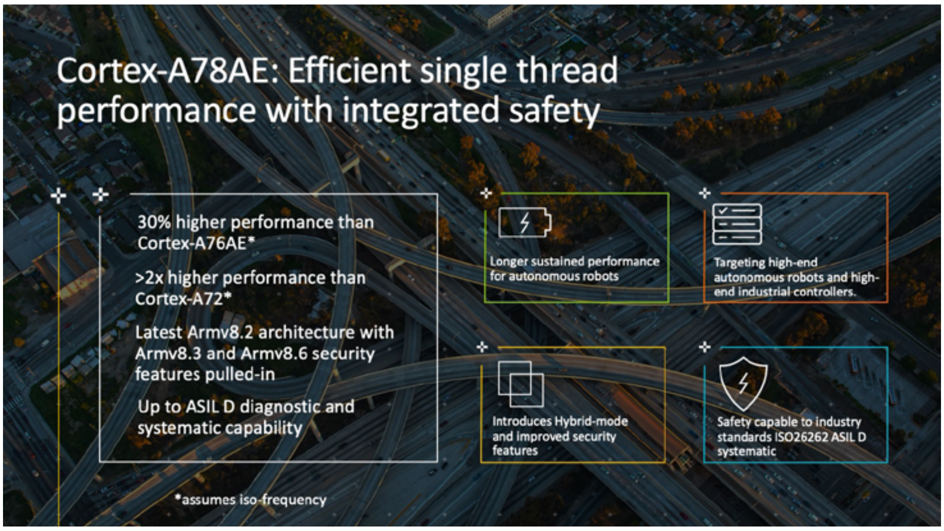

The first AE application processor was the Cortex-A76AE launched in September 2018. This processor offers the performance required for autonomous compute, but at a much lower power level than traditional server chips. Coupled with the Split-Lock capability, this processor allows for truly deployable safe autonomous compute.

Fig 8: Cortex-A76AE

World's first autonomous-class processor with integrated safety.

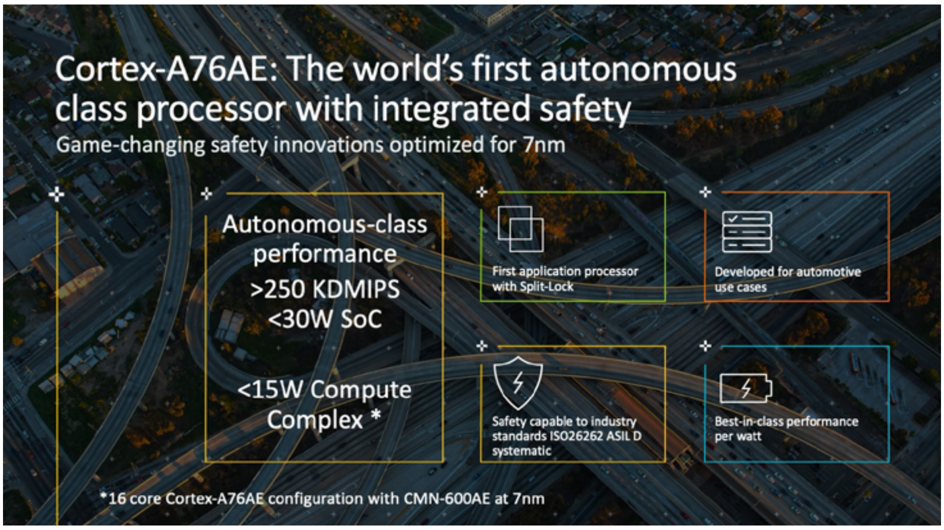

The AE application processor, the Cortex-A65AE, is aimed at safe sensor and information processing and was announced in December 2018. The Cortex-A65AE is Arm's first multithreaded processor and the industry's first with integrated safety designed for high-throughput applications such as sensor processing in autonomous vehicles. Its multithreaded capability is layered on top of Split-Lock, allowing developers to configure and prioritize safety or performance, and it delivers 3.5x higher throughput than prior generations with better power efficiency. Historically, a separate processor was needed for each sensor stream, but now with the Cortex-A65AE, two streams of sensor information can be processed per core, improving efficiency and latency. Both of these new application processors can be coupled with each other, and with other compute elements to complete a truly deployable autonomous computing complex, all based on a common architecture. This helps automakers build their ideal compute topography and optimize their software stack across these elements without having to change tools or architectures. As core safety locking is now performed in hardware, no software changes are required to accommodate it.

Fig 9: Cortex-A65AE

Arm's first multithreaded processor with integrated safety.

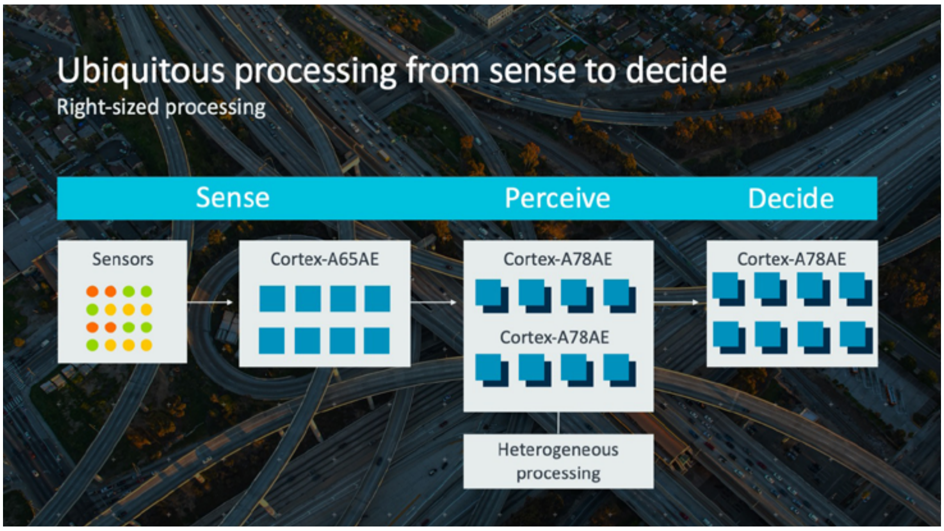

Fig 10: True heterogeneous compute with flexible safety features to provide deployable autonomous compute.

5. Broad Group of Ecosystem Partners

Arm has always attracted a broad and healthy ecosystem to bring key technologies to market, and this is particularly true in the automotive industry where multivendor ecosystem support for both hardware and software is a key factor for deployment for both OEMs and Tier 1s.

Of the top 20 semiconductor suppliers to the automotive industry, 15 are Arm licensees. The breadth of Arm's support allows ecosystem partners to build SoCs for all parts of the car, from powertrain and body through to cabin and connectivity and finally in the move from ADAS to autonomous systems.

However, as noted previously, one of the broadest challenges to the deployment of new complex automotive innovation is the software that enables it. No company can write up to a billion lines of code to win the race to full autonomy, and hence automakers are relying on a software ecosystem with the building blocks to get there. Some of those building blocks can be based on open source, but for much of the real-time and safety-critical software stack, commercial offerings are often preferred for their safety pedigree.

Arm has undertaken to support both our open-source communities and commercial software entities to make a broad range of software solutions available across all the vehicle systems optimized for the Arm architecture. This is a place where these companies can thrive and make their products easily available to the automotive industry.

The Road to AV Mass Deployment

Arm has spent a considerable amount of time working with the automotive industry to fully understand the challenges and pain points standing in the way of deploying the next wave of automotive innovation. Read this insightful report from Forrester to learn about the challenges that need to be overcome in order to deliver safe and affordable autonomous vehicles at scale.

Recent technology innovations from Arm help make those deployments a reality, and cut power, size and cost without compromising performance and safety. Arm is working with key OEMs, Tier1s and the broader ecosystem to help simplify and accelerate the path to the real deployment of autonomous vehicles which will redefine our concept of mobility and enable a new era of automotive innovation.

Arm Technology and Products

The Arm Safety Ready Portfolio

The Arm Safety Ready portfolio is a collection of products across the Arm range of IP that have been through various and rigorous levels of functional safety systematic flows and development. It includes our latest collection of AE IP specifically designed for automotive applications. Customers can be confident that they are receiving reliable functional safety through consistent documentation, processes, and features. The Arm Safety Ready portfolio helps reduce risk and accelerate time to market while fast-tracking the project's certification phase.

Fig 11: Safety Ready Products.

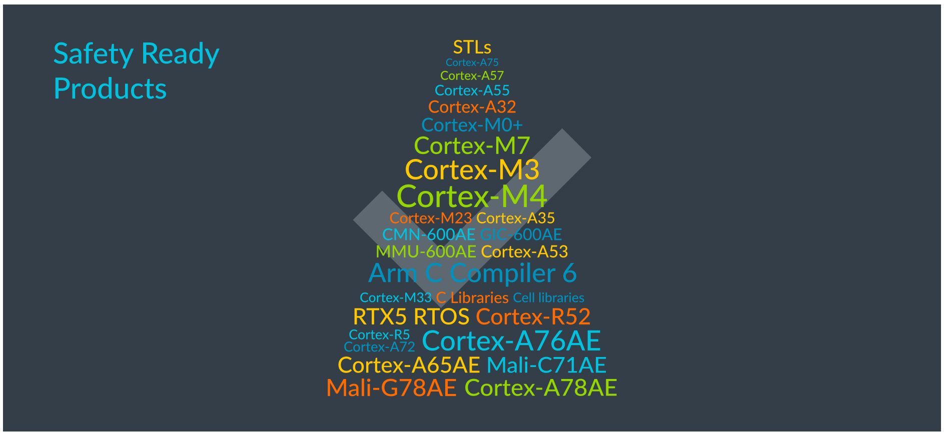

[Learn more](#)

Arm AE CPUs

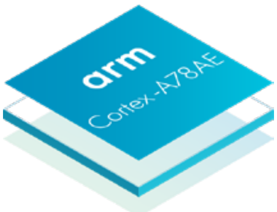

Arm Cortex-A78AE

Cortex-A78AE elevates Arm's AE processing lineup with enhanced safety features and breakthrough performance at higher levels of efficiency. For the emerging class of L2+ ADAS and beyond, A78AE is the processor of choice.

[Learn more](#)

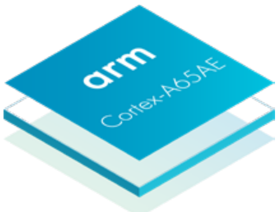

Arm Cortex-A65AE

Cortex-A65AE is Arm's first multithreaded AE Cortex-A CPU that delivers the highest levels of safety with DCLS. It can scale up to eight cores in a single cluster and supports Split-Lock for mixed-criticality automotive applications.

[Learn more](#)

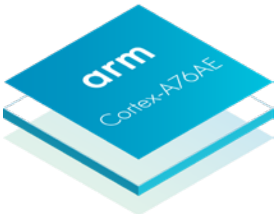

Arm Cortex-A76AE

Cortex-A76AE brings highest levels of safety with Dual Core Lock-Step (DCLS) capabilities, while delivering uncompromising performance and thermal efficiency. It is the processor of choice for next-generation ADAS and Autonomous Driving systems.

[Learn more]

Arm AE System IP

Arm CoreLink CMN-600AE

The Arm CoreLink CMN-600AE Coherent Mesh Network is designed for high performance automotive systems across a wide range of applications.

[Learn more]

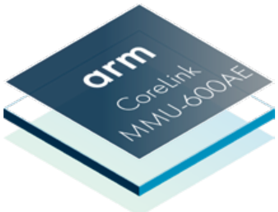

Arm CoreLink MMU-600AE

Software compatible with MMU-600. Adds additional safety features to meet safety requirements for building high-performance ASIL B to ASIL D systems.

[Learn more]

Arm CoreLink GIC-600AE

Arm’s newest Interrupt controller, the GIC-600AE is fully software compatible with the GIC-600 and is engineering with safety mechanisms to meet demanding safety requirements up to and including ASIL-D/SIL 3. The GIC-600AE supports the GIC v4 standard and is the interrupt controller for choice for the Cortex-A78AE.

Learn more

Arm AE GPU

Arm Mali-G78AE

Mali-G78AE is the first Mali GPU developed specifically to meet the high performance requirements of next generation automotive systems, including support for ASIL B.

Learn more

Automotive Related IP

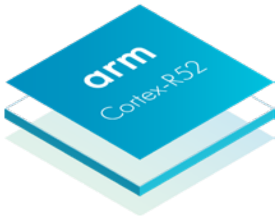

Arm Cortex-R52

Cortex-R52 is the most advanced processor in the Cortex-R family delivering real-time performance for functional safety.

Learn more

Arm Mali-C71AE

Mali-C71AE was developed for the smart automotive systems market, delivering key visual information to both computer vision systems and the driver display for clear and convenient viewing.

Learn more

Fig 12: Autonomous drive - performance with safety.

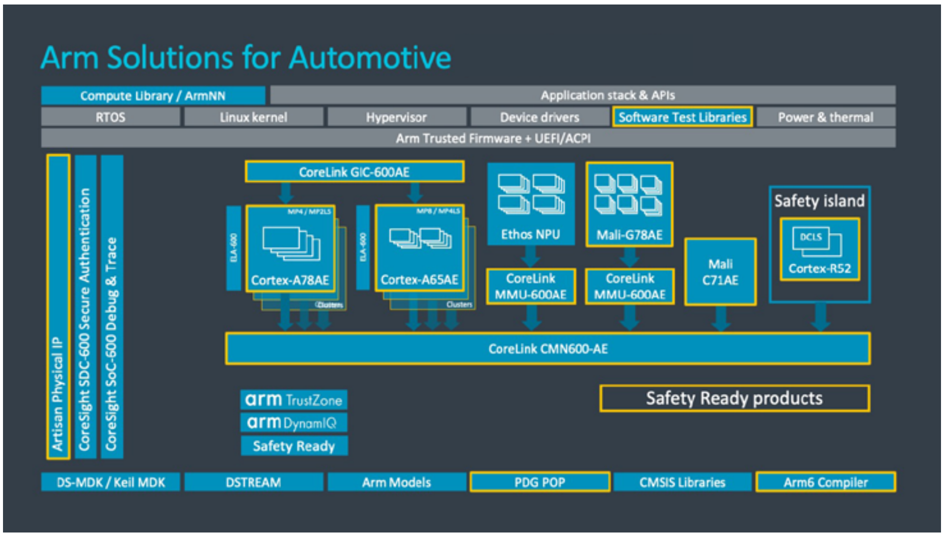

Software and Tools

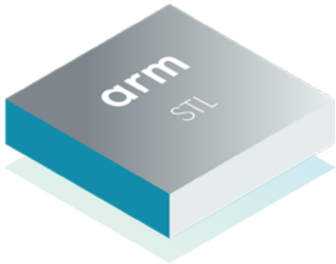

Arm Software Test Libraries

Arm Software Test libraries (STL) complement Arm's functional safety technology which supports systematic capability for ISO26262 ASIL D. STLs are efficient assembly language routines that test for the presence of faults when executing on Arm-based processors at startup and during run time. A single straightforward C language API lets developers quickly and easily scale test functions across multiple CPU systems.

Learn more

Swift Starling Positioning Engine

The Swift Navigation Starling Positioning Engine natively processes measurements from commercially available Arm-based GNSS (Global Navigation Satellite System) receivers to provide precision positioning and integrity capabilities.

Learn more

Development Solutions

Alongside its processor technology, Arm also creates professional software development tools, simulation models, and critical software components to accelerate innovation in the automotive segment. With 2020's cars likely to embed over 10x as many lines of code as those produced previously, it has become success defining for the entire supply chain to be able to go from design through to safety certification in less time. Arm's development solutions are highly optimized for the entire range of Arm processors and have been externally assessed by TÜV SÜD for use in applications up to ASIL D where applicable. Together with Arm's vast ecosystem of third-party software, tools and service providers, these solutions enable shorter product cycles on Arm-based systems.

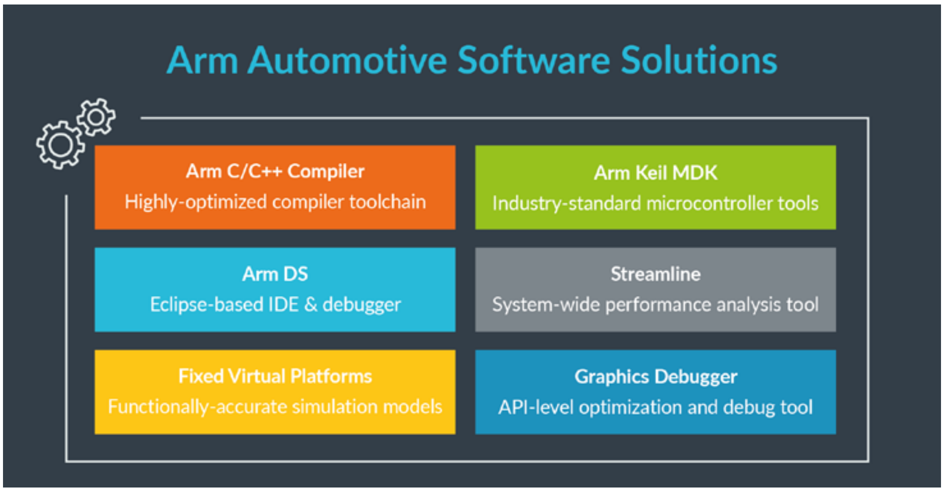

Fig 13: Arm automotive software solutions.

Arm Automotive Ecosystem

Automotive and Functional Safety

Arm Automotive Developer Community

Arm partners offer technologies to service the compute needs of the whole car, including ADAS and autonomous systems. These partners provide solutions for silicon and software, facilitating the efficient development of automotive solutions.

Learn more about the Arm Automotive Developer Community and how to become a member.

Learn More

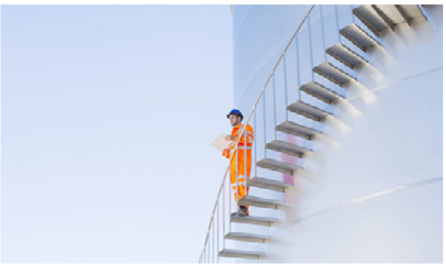

Arm Functional Safety Partnership Program

Arm's Functional Safety Partnership Program showcases a range of functional safety partners who specialize in the areas of software and tools, design services, and training services.

Arm Functional Safety Partner
Software & Tools

Arm Functional Safety Partner
Design Services

Arm Functional Safety Partner
Training Services

Learn more about the Arm Functional Safety Partnership Program and how to become a member.

Visit the Arm Functional Safety Partnership Program