



Separation Kernel – Basis for Certifiable Applications and Systems

© erdikocak | iStock

Functional security and cybersecurity are among the most important issues in the development of complex connected vehicle systems. The certification of individual systems will become increasingly important in the future. Here the automotive industry has some backlog to do and can learn from certification solutions from the aviation industry, including working according to the Safety-and-Security-by-Design principles. Real-time operating systems based on a separation kernel enable new approaches, as Sysgo describes.

AUTHOR



Chris Berg
is Solution Architect Automotive at
the Sysgo AG in Klein-Winternheim
(Germany).

CHALLENGES

Already today, many functions in automobiles are fully or partially automated; driver assistance systems such as distance alarms, rear-view cameras or parking aids entered series production even of small and medium-sized vehicles. Fully autonomous passenger cars, buses and even trucks are in test operation on public roads; the first production vehicles can move at least semi-autonomously. The car of the future will require even more electronics and computing power – on the one hand to increase safety (driver assistance systems) and on the other to meet passengers’ demands for comfort, entertainment and communication. The Internet of Things, which provides devices within the vehicle with the necessary connectivity, brings with it considerable new security challenges. Safety aspects are therefore increasingly determining software design, and certification standards will soon play a similarly important role in automotive engineering as they do today in the aircraft industry. In addition to functional safety, aspects of IT security and data protection also come to the fore.

SAFETY AND SECURITY

Safety is concerned with functional safety and thus the prevention of system failures that can lead to damage to health and the environment. IEC 61508 is an industry-independent basic standard for the functional safety of electrical, electronic and programmable systems with a safety reference. IEC 61508 distinguishes between four criticality levels: SIL-4 to SIL-1 (Safety Integrity

Level). Based on this standard, ISO 26262 specifically defines the functional safety requirements for vehicles in road traffic. Based on the safety integrity level of IEC 61508 and depending on the tolerable failure frequency, it defines four Automotive Safety Integrity Levels (ASILs) D to A, where D stands for the most critical systems, **FIGURE 1**.

IT security is primarily concerned with protecting critical systems against unauthorized access and thus against manipulation, but also against unauthorized access to personal data. There are also international standards for security-critical systems. The most important worldwide is ISO 15408, better known as Common Criteria (for Information Technology Security Evaluation). It introduces seven Evaluation Assurance Levels (EAL 1-7) for trustworthiness. In the automotive industry, SAE International (formerly Society of Automotive Engineers) is currently working on a series of standards for various aspects of IT systems in automobiles; in Draft J3061, published in 2016 and based on the SILs, it is introducing the term Automotive Cybersecurity Integrity Level (ACsIL). ISO is also working on ISO 21434 (Road Vehicles – Cybersecurity Engineering). Once these standards have been adopted, it is to be expected that many products will require corresponding certification. Some products already have Common Criteria (CC) certification. However, this does not automatically mean that they are safe – in any case, detailed attention must be paid to exactly what has been certified here.

Safety and security are easiest to ensure if all IT systems are strictly

separated from each other and therefore cannot influence each other. Today, however, more and more safety-critical and non-critical applications are being operated in vehicles, and in order to save costs and weight, a single hardware serves as a platform for multiple functions, including different levels of criticality. This trend harbours considerable risks. If, for example, an attacker gains access to a comparatively insecure entertainment system based on Android and can also access security-critical systems via this gateway, this can have very serious consequences. It is therefore essential to strictly separate applications with different criticality levels, even if they run on the same hardware. Here, the automotive industry can benefit from the experience gained in the aerospace industry with its extreme requirements, where hypervisor-based solutions have already established themselves.

MULTIPLE PARTITIONS INSTEAD OF MULTIPLE CPUs

A hypervisor can host different functions on a controller in multiple partitions that previously required separate CPUs. However, it must be absolutely ensured that the software that provides the hypervisor functionality actually guarantees a strict separation between the partitions. Otherwise, one has a unified hardware platform, but possibly interactions between critical and non-critical applications such as audio systems and brakes. ASIL-D and ISO 26262 certification ensures that functions in different partitions are actually separated as if they were running on different CPUs.

Function	Hazard	ASIL-A	ASIL-B	ASIL-C	ASIL-D
Driving	Sudden start		████████████████████		
	Abrupt acceleration		████████████████████		
	Loss of driving power		████████████████████		
Braking	Maximum 4-wheel braking			████████████████████	
	Loss of braking function			████████████████████	
Steering	Self-steering			████████████████████	
	Steering lock			████████████████████	
	Loss of assistance		████████████████████		

FIGURE 1 Different ASILs are required depending on the function; D is the most stringent level (© Sysgo)

Especially on multicore systems, the use of hypervisors is a suitable way to meet the challenges of system design. Although such CPUs are primarily used for performance reasons, they can also support the required separation of individual functions. However, the certification of multicore systems is very complex, and many certified systems actually use only one core. If, however, different functions are bundled in a single software that runs under a Real-time Operating System (RTOS) on only one CPU core, interference between the functions can easily occur – strict separation is not guaranteed. For example, the impact of one application on the runtime behavior of another application can lead to security issues, such as exceeding deadlines for real-time applications. Similarly, timing effects due to the sharing of system resources, such as caches and memory buses, can lead to hidden information channels that violate the confidentiality requirements of the application. These potential problems stem mainly from the fact that most hypervisors are added to a RTOS, which does not support such a separation due to its own design. Especially in safety-critical applications, it is important that the RTOS is already designed specifically for

the separate execution of different functions, that means it is more of a separation kernel than a simple RTOS.

SPATIAL AND TEMPORAL SEPARATION

Such a separation kernel enables the spatial and temporal separation between applications and provides the partitions for the execution of user applications. Time separation is achieved by time partitioning, where CPU time is divided into time partitions during configuration. Spatial separation is achieved by resource partitioning, in which system resources such as main memory, files, devices, secure communication channels and cores are partitioned and containers, known as resource partitions, are assigned statically. User applications are executed in the context of a resource partition.

An example of a software environment with strict separation of applications is Sysgo’s PikeOS, which has already proven itself in the aircraft industry. PikeOS offers both a full real-time operating system (Hard RTOS) and a virtualization and partitioning system to support the specific requirements of automotive applications, **FIGURE 2**. The

basis of the PikeOS platform is a small microkernel that provides a virtualization infrastructure. This makes it possible to place different applications and resources in secure, individual partitions. These can run in various environments – from Posix to Linux and Android to Autosar or Genivi, **FIGURE 3**. Thanks to the integrated time partitioning, it is irrelevant whether the applications are real-time applications or not. The separation kernel of PikeOS 4.2.2 is currently the only one in the world to be certified according to CC EAL 3+ (for the popular platforms X86_64, ARMv7 and ARMv8).

Especially in mixed environments with relatively poorly secured operating systems such as Android and critical environments such as Autosar, a separation kernel can also perform an important security function to aggravate attacks. This kernel, which acts as a hypervisor for the different guest operating systems, is able to intercept privileged system calls from the guest systems and first check for the required permissions before they are actually executed. While common desktop operating systems have all device drivers integrated into the kernel, it is possible to create an environment in which only a

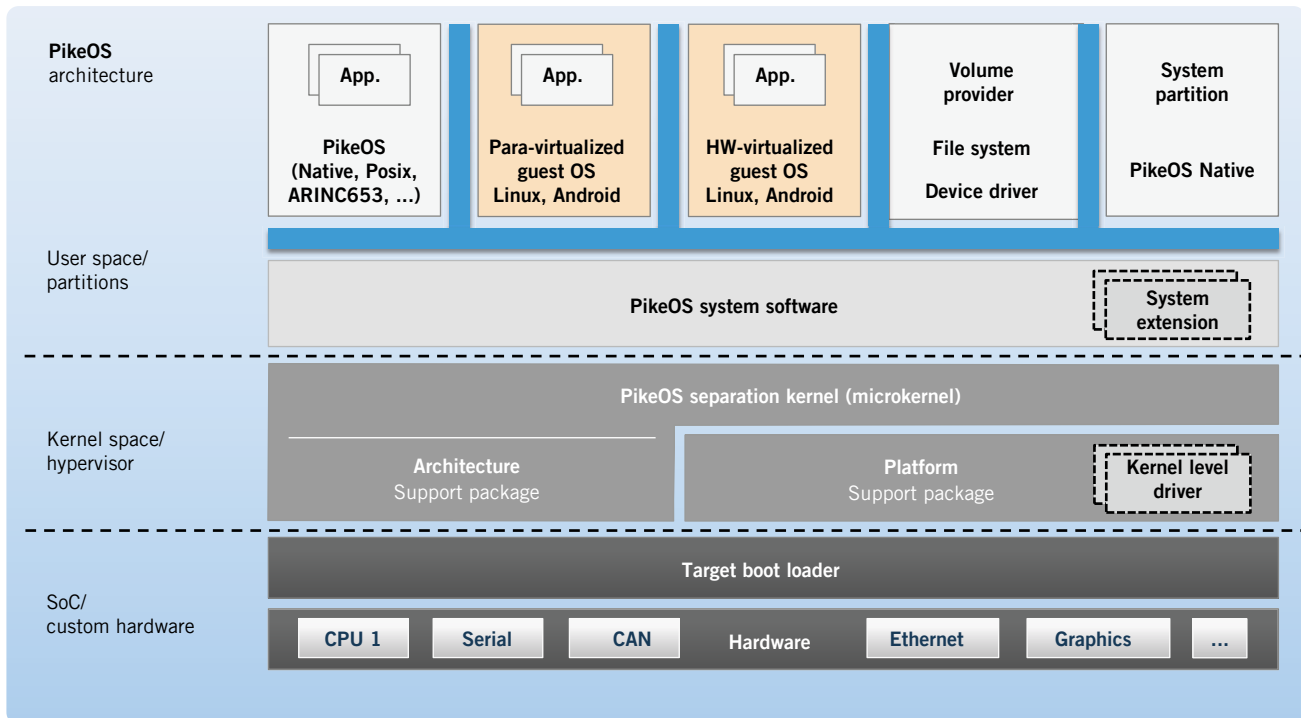


FIGURE 2 The separation kernel of PikeOS provides a secure separation of the different partitions (© Sysgo)

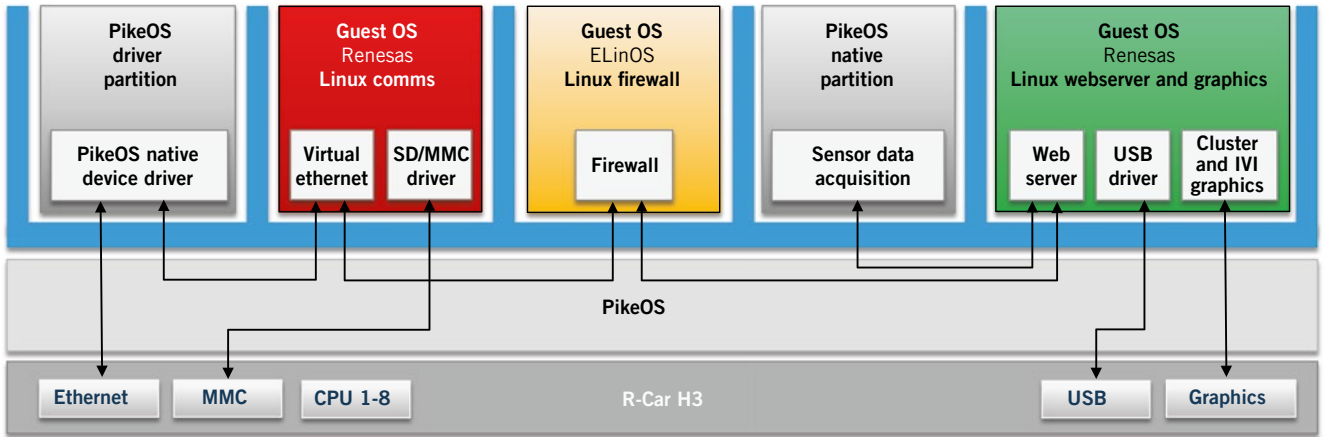


FIGURE 3 PikeOS in an automotive application based on the Renesas R-Car H3 (© Sysgo)

very small set of services in “privileged mode” runs in the kernel – such as scheduling, context switching, process communication and synchronization, and interrupt handling. Device drivers are then executed without privileges in normal user mode like any other application code. In this way, the attack surface of the entire system is significantly reduced.

SECURITY AND CERTIFICATION

The PikeOS Hypervisor itself is certified to the highest industry standards, providing a suitable foundation for critical systems where both functional and IT security must be ensured. The protection

mechanisms are essentially based on two principles: strict separation of applications through time and resource partitioning and control of communication channels. The individual applications within the overall system can have different criticality levels.

Due to these protection mechanisms of PikeOS, certification according to industry-specific safety and security standards can be carried out separately for each application – an essential feature to keep costs under control. PikeOS was also the first platform to receive SIL 4 certification in multicore environments.

To meet the requirements of ISO 26262, PikeOS is optionally offered with an auto-

motive certification kit, which incorporates Sysgo’s many years of comprehensive certification expertise. The certification kit contains an ISO 26262 Part 6 compliant PikeOS Hypervisor as well as comprehensive documentation aids for development and testing. Additional safety information can also be provided to achieve ISO 26262 compliant systems. Important components of these certification kits are a safety manual with guidelines for the use of PikeOS in safety-critical system designs and a case study with characteristic functional safety requirements according to the Automotive Safety Integrity Levels required.