

# Execute Multiple Workloads on a Single Hardware Platform

Real-time hypervisors — the central technology for workload consolidation — enable the safe execution of multiple workloads on a single hardware platform.

Embedded systems become increasingly hardware-intensive — and that’s a problem. Every time new hardware is bolted on, it brings extra costs, more points of failure, greater maintenance complexity, and a more difficult path for future upgrades.

[Virtualization](#) and [hypervisors](#) offer a way out. By replacing physically separate platforms with virtual machines (VMs), multiple subsystems can be combined into a single hardware platform — a technique known as workload consolidation.

This approach solves many of the challenges posed by fragmented hardware. And now, with platforms that pre-

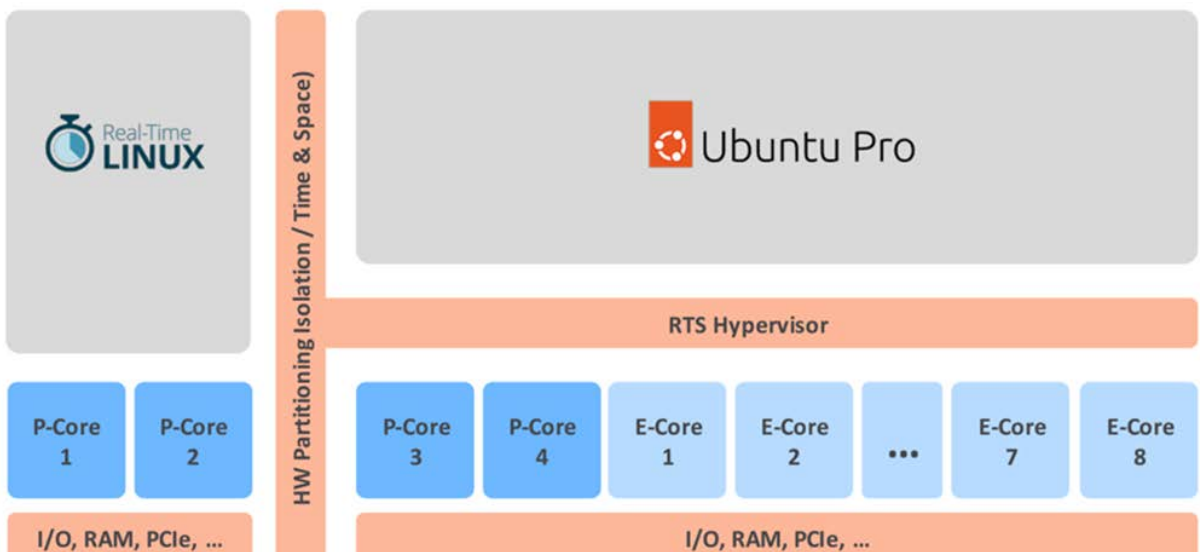
integrate industrial-grade virtualization with embedded hardware, adoption is easier than ever.

## Workload Consolidation: One Device Fits All

The most immediate advantage of workload consolidation is space and cost savings. With fewer hardware components, systems consume less power, require less maintenance, and occupy less physical space. For example, motion control and HMI functions, traditionally deployed on separate platforms, can run side by side on a single device.

Reliability also improves. Consolidating workloads reduces the total system count, which in turn increases mean

## conga-aCOM/cRLP | System Consolidation



1. A real-time hypervisor partitions hardware into virtual machines with deterministic performance. (Credit: congatec)

time between failures (MTBF). Fewer components — especially connectors and other common failure points — translate into less downtime and simpler maintenance.

Workload consolidation also enhances system flexibility. Once the capability to run multiple tasks on a single platform is established, additional functions, such as AI analytics or predictive maintenance, can be integrated alongside existing workloads. This allows OEMs to upgrade equipment functionality without redesigning the underlying computing architecture.

### Understanding Virtualization

Virtualization is the foundation that makes this all possible. Virtualization uses an abstraction layer known as a hypervisor to transform a single piece of physical hardware into multiple virtual machines (VMs). Each VM is an isolated software environment that can host its own operating system (OS).

In embedded context, it's typically used to create a heterogeneous system, i.e., a platform that hosts a mix of real-time operating systems (RTOSs) and general-purpose OSs (*Fig. 1*).

The hypervisor makes this multi-OS setup possible by mediating access to resources such as CPU, memory, and I/O. This can be done by giving a VM dedicated resources — for example, a set of CPU cores — or by managing shared resources in a way that avoids interference between VMs.

There are two types of hypervisors:

- Type 1 hypervisors (bare-metal) run directly on hardware.
- Type 2 hypervisors (hosted) run on top of a host OS.

Only Type 1 hypervisors can deliver true determinism, as they eliminate the overhead imposed by the host OS. This

makes them the only viable option for safety-critical systems. More broadly, the reliable performance of Type 1 hypervisors is advantageous for all kinds of embedded systems.

A key benefit of virtualization is that it improves resource utilization. Modern processors often have an abundance of resources that can't be fully exercised by a single workload, such as multiple processor cores and a wide array of I/O. By running multiple workloads on a single platform, these otherwise wasted resources can be put to work.

Virtualization can also enhance security and reliability. Because VMs operate in isolation from one another, a security breach or software fault may be contained to a single VM without compromising the entire system.

### The Role of Real-Time Hypervisors in Workload Consolidation


A longstanding concern in workload consolidation is that running multiple applications on a shared platform may degrade real-time performance. Developers typically worry that general-purpose workloads might interfere with deterministic tasks, introducing delays or jitter that undermine system reliability.

Real-time hypervisors address these concerns through temporal and spatial isolation. Temporal isolation ensures that real-time processes meet their deadlines without interference from other workloads. Spatial isolation enforces physical separation by assigning dedicated CPU cores, memory, and I/O devices to specific VMs. Together, these techniques prevent resource contention and maintain deterministic performance.

This enables, for example, an RTOS handling motion control to coexist with a human-machine interface (HMI) running on a general-purpose OS. Here, the hypervisor can partition the system so that motion control is given absolute priority.

Real-time hypervisors can further minimize jitter and ensure low-latency response times by leveraging hardware-assisted virtualization technologies such as Intel VT-x and AMD-V. These technologies accelerate context switching, reduce virtualization overhead, and support direct device assignment (known as “passthrough”) for improved I/O performance.

Hardware assistance also enhances fine-grained control over system resources, e.g., CPU core allocation, device passthrough, and scheduling policies. As a result, high-performance, time-critical applications can run reliably on shared platforms.



2. Platforms like the conga-COM/mRPL pre-integrate hardware and hypervisor (Credit: congatec)

### Enabling Consolidation with Solution Platforms

While virtualization facilitates workload consolidation, implementing it from scratch can be complex and time-consuming. OEMs must integrate hardware, hypervisors, operating systems, and industrial I/O, each with its own configuration requirements and potential pitfalls.

To streamline adoption, solution platforms are available that pre-integrate these components into a single, validated system (Fig. 2). Such platforms combine hardware, hypervisors, and software stacks designed specifically for virtualized workloads, ensuring compatibility and reducing integration effort. For OEMs, this allows the focus to shift from low-level system integration to high-value application development.

Furthermore, working with pre-integrated solution platforms accelerates time-to-market and reduces the risk of performance bottlenecks or integration errors. By partnering with experienced providers, OEMs gain access to expert support and scalable designs that are ready for future expansion.

### Cutting Costs and Complexity with Workload Consolidation

Modern workload consolidation offers OEMs a powerful way to reduce costs, simplify system architectures, and enhance flexibility by running multiple applications on a single platform. Real-time hypervisors play a critical role in making this possible, ensuring that performance and determinism are maintained even in complex, mixed-workload environments.

To fully realize these benefits, it's essential to select hardware and software platforms that are designed to work seamlessly together. Pre-integrated solution platforms provide a proven path forward, minimizing integration challenges and accelerating time-to-market.

As industrial systems grow more complex and demand increases for new functionality, OEMs that embrace workload consolidation will be better positioned to deliver scalable, efficient, and competitive solutions.



*Andreas Bergbauer is Manager Solution Management at congatec. He has over 15 years of professional experience in product management, IT project management, and software development. After graduating with a degree in Business Informatics from Deggendorf Institute of Technology, he held leadership positions at msg and ConVista.*

*With his expertise in product strategy and technological innovation management, he's now involved in developing application-ready solution platforms in embedded computing.*