






The Architect's Checklist for **GPU Virtualisation**

November 2025





Ensure performance, security, and efficiency in your SoC design

Evaluating **GPU virtualisation** solutions is critical for modern SoC architectures in automotive, data center, and embedded systems. This checklist guides you through the key technical and strategic considerations to select the right IP for your project.



Performance & Architecture

Near-Native Performance:

Does the solution allow GPU workloads to be sent directly from the guest OS to the hardware, bypassing the hypervisor to minimise latency and overhead?

Hardware-Level Implementation:

Is the virtualisation logic implemented in hardware (e.g., a dedicated firmware CPU and tags like OSID), rather than a software-only model?

Independent Operation:

Can the GPU continue functioning independently, even if the host OS reboots or crashes? (A key feature for safety-critical applications).

Multi-GPU Load Balancing:

Does the architecture support assigning multiple virtualised GPUs to a single VM, enabling dynamic load-balancing for optimal resource utilisation?

Number of Supported Virtual Machines:

Does the GPU virtualisation solution support a high number of concurrent virtual machines (VMs) per GPU, scaling to meet your workload density and market requirements (e.g., multiple OS instances in automotive or multi-tenant workloads in data centers)?



Security & Isolation

Hardware-Enforced Domains:

Does the solution provide hardware-level isolation (e.g., memory protection via IOMMU, separate register banks) between virtual machines, ensuring one misbehaving VM cannot impact another?

Secure Resource Partitioning:

Is memory and processing time for each VM securely carved out and protected at the hardware level?

Use Case Alignment:

Is the isolation level sufficient for your target market (e.g., ASIL-B in Automotive, multi-tenancy in Data Centers, sandboxing in Mobile)?



Use Case & Application Fit

Primary Market:

- **Automotive:** Consolidating Instrument Cluster (safety-critical OS) and Infotainment (rich OS) on a single SoC.
- **Data Center:** Cloud Gaming, AI/ML workloads, Virtual Desktop Infrastructure (VDI).
- **Mobile & IoT:** Running secure applications in a trusted execution environment (TEE) isolated from the main OS.

Workload Type:

Is the solution optimised for your primary workload (graphics, compute, or a hybrid mix)?



Platform & Software Compatibility

Hypervisor Support:

Is the solution compatible with your required hypervisor type?

- **Type 1 (Bare Metal):**

e.g., Xen

- **Type 2 (Hosted):**

e.g., Linux KVM, QNX Hypervisor, Jailhouse

- **None (AMP):** Asymmetric Multiprocessing for “hypervisor-less” designs.

Bus Topology:

Does the IP work seamlessly in your chosen bus architecture (e.g., SoC-integrated, PCI Express with SR-IOV)?

OS & Driver Support:

Are production-ready drivers available for your target guest operating systems (Linux, QNX, Integrity, Android, etc.)?



Advanced Features & Scalability

Live Migration:


Does the solution support moving a VM between physical GPUs or servers for maintenance, load balancing, and energy savings?

Multi-Tenancy:

Can the system reliably and securely serve multiple untrusting clients or applications (e.g., in a cloud environment)?

Container Integration:

Can virtualised GPUs be assigned directly to containers (e.g., Docker, Kubernetes) for finer-grained application isolation?



Production & Ecosystem

Proven in Silicon:

Is the IP already shipping in production-scale devices? (e.g., Imagination’s technology has a decade of proven use in automotive).

Development Tools:

Is a robust software stack, debugging tools, and documentation provided to accelerate your development cycle?

Vendor Expertise:

Does the IP provider offer expert technical support and architectural guidance throughout the design-in process?


How to use this checklist

Use this list during your initial research and vendor discussions. A "Yes" to most of these questions, particularly in the Performance & Architecture and Security & Isolation sections, indicates a mature, capable, and future-proof GPU virtualisation solution.


This checklist is based on insights from Imagination Technologies' webinar **"A Deep Dive into GPU Virtualisation."**

Ready to See How Imagination Stacks Up?


Imagination's HyperLane technology is the foundation for our hardware-based GPU virtualisation, designed to meet every critical item on this checklist.



Download
White paper



Contact
an expert



Download our technical deep-dive whitepaper or contact our experts to see how HyperLane can power your next design.



www.imaginationtech.com

[Contact us now](#)