


# PCIe Gen5 GPU Solution

## Empowering Generative AI with Flexibility and Efficiency

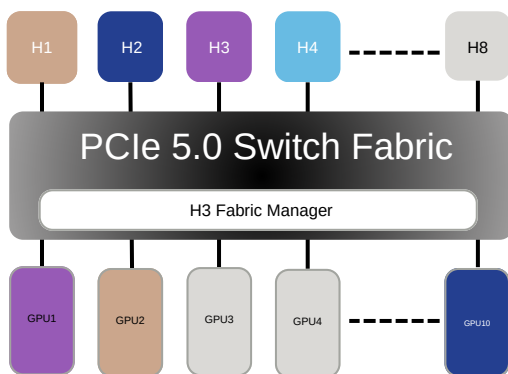


### Revolutionizing Generative AI: PCIe Gen5 GPU Solution Meets the Growing Demands

The PCIe Gen5 GPU solution revolutionizes the field of generative AI, satisfying the escalating demands by enabling the independent deployment of GPUs at a larger scale. Achieved through a composable architecture and the innovative approach of composable disaggregation, this solution significantly impacts data-intensive tasks. It enhances transfer speeds and overall performance while providing flexibility and scalability by isolating GPUs from the server. The simplified hardware management in independent GPU expansion further contributes to cost savings. Overall, the PCIe Gen5 GPU solution meets the substantial GPU demands in generative AI, delivering enhanced performance, flexibility, and cost-effectiveness.



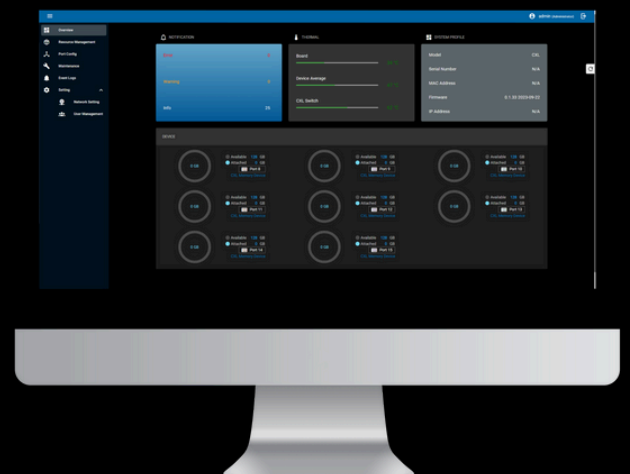
### Next-Level System Dynamics: Breakthrough GPU Functionalities



Revolutionizing system dynamics, Falcon 5012 introduces advanced GPU management, enabling up to 8 hosts to dynamically share 10 devices on-demand. This streamlines resource allocation, saving significant setup time. The solution allows seamless removal and addition of GPU resources between hosts, ensuring intelligent rearrangement. IT administrators benefit from the cohesive hardware and software solution, focusing more on gaining insights and less on component management. Falcon 5012 makes your workflow more efficient and effective.

### Comprehensive Resource Management with RESTful API and GUI

Falcon 5012 ensures optimal system performance through strategic resource management. Actively monitoring GPU health, real-time performance, and errors in both host and device ports, it facilitates quick issue identification and resolution for stable operation. Introducing flexible switch cascade topologies and configurable host virtual trees enhances resource management flexibility. The solution also provides intuitive resource allocation insights through the Host and GPU views, enabling intelligent resource deployment.



<b>Software Features</b>	<b>Features</b>	<ul style="list-style-type: none"> <li>• GPU composability</li> <li>• GPU hotplug</li> <li>• GPU peer to peer (GPU P2P)</li> <li>• PCIe port configuration to host/device ports</li> <li>• Power control of a single PCIe slot</li> <li>• GPU health information</li> <li>• Performance and error monitoring of host and device ports</li> <li>• Configurable host virtual tree</li> </ul>	<ul style="list-style-type: none"> <li>• Host view of assigned GPU</li> <li>• GPU view of assigned host</li> <li>• Temperature threshold setup</li> <li>• Device overheated protection</li> <li>• Firmware secure boot</li> <li>• Link capability and status of devices</li> <li>• User authentication management</li> <li>• 256bit encryption of sensitive data</li> <li>• Sha128 encryption</li> <li>• ELK central log server integration</li> </ul>
	<b>Management Interface</b>	Redfish®, RESTful API, GUI	
	<b>System Management</b>	<ul style="list-style-type: none"> <li>• H3 management center</li> <li>• Real-time GPU cluster topology</li> <li>• Dashboard for GPU utilization, performance and other information</li> <li>• Predictive health monitoring</li> <li>• Role-based authentication and access control</li> </ul>	
<b>Hardware Features</b>	<b>Model Name</b>	Falcon 5012	
	<b>BMC</b>	<ul style="list-style-type: none"> <li>• AST 2600, fully IPMI V2.0 compliant</li> <li>• AST 2500</li> </ul>	
	<b>mCPU</b>	Intel ATOM X86 CPU	
	<b>PCIe Switch</b>	Broadcom PEX 89144 PCIe 5.0 switch	
	<b>Device</b>	<ul style="list-style-type: none"> <li>• GPU, FPGA, NIC, and other standard PCIe device (add-in card)</li> <li>• Dual-slot width and 10.5" length</li> <li>• Support up to 700W GPU</li> </ul>	
	<b>Host Interface - Standard mode</b>	<ul style="list-style-type: none"> <li>• One (1) PCIe 5.0 x16 port for host connection</li> <li>• Need one HBA (host bus adapter) card on server host</li> <li>• 1m CDFP cable</li> </ul>	
	<b>Host Interface - Advanced mode</b>	<ul style="list-style-type: none"> <li>• Up to Four (4) PCIe 5.0 x16 ports for host connections</li> <li>• Need HBA (host bus adapter) cards on server hosts</li> <li>• 1m CDFP cables</li> </ul>	
	<b>Operating temperature</b>	10°C ~ 35°C (50°F ~ 95°F)	
	<b>Power</b>	Four (4) pieces 3200W (2+2 or 3+1 redundant), back removable, hot-swap	
	<b>Fan</b>	Hot-swappable 8 80x80mm dual rotor fans	
	<b>Ethernet Ports</b>	RJ45 of mCPU for GPU management/BMC for chassis management	
	<b>Dimension</b>	4U; 175(H) * 448(W) * 650 (D)mm	