AMD EMBEDDED SOLUTIONS

PARTNER:
ADVANTECH

INDUSTRY:
DATACENTER

AMD TECHNOLOGY AT A GLANCE:

AMD EPYC™ EMBEDDED PROCESSORS

CASE STUDY

Hyperconverged Infrastructure (HCI): System Design Imperatives for Achieving Extreme Versatility and Scalability

Here's what to look for in the next generation of cost efficient, high performance HCI platforms

The continued proliferation of hyperconverged infrastructure (HCI) owes in equal parts to HCI's clear technology merits and the dynamic evolution of our global network from the datacenter to the edge. The efficiencies that HCI strategies and solutions provide – in cost, deployment and scalability – make them ideally suited to accommodate shifting network and computing requirements as compute infrastructure is distributed ever farther away from our central datacenter hubs, and ever deeper into regional and/or remote geographies.

HCI – particularly when deployed far from the datacenter core – brings data, content and AI-guided decision-making agility closer to the end users and end infrastructure where it's needed. This is achieved with an easy-to-scale model that prioritizes reliability and resiliency – and flexibility, most of all – at ultra-efficient cost structures.

HCI CORE BENEFITS

HCI is broadly defined as an approach to ICT architecture that combines computing, storage and networking in a single system, simplifying IT environments while simultaneously making them more adaptable in their ability to accommodate a wide range of applications and workloads. HCI prioritizes elasticity and scalability, and can be flexibly deployed and maintained in a simplified manner.

As ICT infrastructure and data traffic spreads progressively outward from our core datacenters, there will naturally be decreased bandwidth capacity, rack space, power and IT staff resources to service it. HCI is therefore optimized to maximize computing, storage and networking efficiencies in a dense, low power system profile that's highly modular, reliable and secure.

HCI is software-defined in nature and compute resources can therefore be readily repurposed and reallocated as requirements and/or use cases evolve. It's designed from the ground up for simplified, centralized management of consolidated workloads, and reduces the need for cost-intensive configuration, maintenance and support due to its 'building block' modularity, hot swapability and ease of serviceability.

HCI must also meet aggressive performance/cost requirements, providing the onboard processing horsepower to accommodate compute-intensive applications like AI and IIoT, with the connectivity/IO flexibility to maximize throughput efficiency and minimize latency for a wide range of cost-sensitive applications. In this model, multifunction HCI servers and hybrid server/storage systems can be leveraged in myriad functional combinations, flexibly configurable and seamlessly expandable to support NICs, storage capacity and GPU or FPGA accelerator cards – with a unified software layer to release the full value of highly-integrated HCI stacks.

HCI clusters can be scaled up or down quickly, with nodes added and/or redeployed on the fly – without the typical delays required for manual hardware configurations and integration processes. More broadly, HCI allows organizations to gradually and smoothly evolve and expand their infrastructure over time, without having to worry about interoperability issues and/or protracted and complex testing procedures common to heterogeneous technology environments.

For virtualized workloads, HCI can be leveraged to support more virtual machines (VMs) per server to improve CAPEX/OPEX, reducing the overall cost of VMs deployed. HCI built with high levels of compute and connectivity can be used to abstract compute resources for easier allocation and management, and consolidate workloads and storage capacity as needed.

MEASURING HCI PERFORMANCE & CAPABILITIES

When assessing and/or comparing HCI technology platforms, it's important to understand their respective abilities to deliver enterprise-class performance across four key capabilities.

NETWORKING – The ideal HCI server and storage platform is built to accommodate high I/O throughput supporting numerous lanes of high-speed PCIe® connectivity, and it should also include memory capacity enough for very large datasets. The HCI platform must be at home with advanced virtualization technologies, with access to highly parallelized CPU cores for optimal handling of network functions virtualization (NFV) and software-defined networking (SDN) capabilities, among others.

COMPUTE – HCI hardware platforms should satisfy the performance rigors of demanding AI and edge inferencing applications out of the box, and this typically manifests in high core counts and threads at the processing layer. The platform should be augmentable with GPU or FPGA cards where performance needs dictate, while providing high performance-per-watt compute efficiency, as well as high VM density and capacity.

STORAGE – Storage workloads from core to edge will continue to be as voluminous as they are varied, and HCI platforms must be able to scale capacity as readily as IOPs performance. This includes native support for SATA/SAS and NVMe interfaces to accommodate high-capacity HDDs as well as high-performance SSDs, respectively. High I/O bandwidth is essential for faster data/content loading and data collection at the edge, just as memory encryption is critical for overall data security.

INDUSTRIAL – So many attributes must converge in order for HCI to effectively service industrial applications, with cost and power/cooling efficiency chief among them. HCI server, storage and hybrid systems targeted for industrial apps must also support enterprise-class reliability, accessibility and serviceability (RAS), with robust security features sufficient to protect business critical data. HCI platform longevity is another hugely important consideration for industrial apps, and HCI system lifecycles often hinge on the supported lifetimes of the processors within.

SKY HIGH INNOVATION WITH AMD

HCI features and functionality are designed for maximum versatility, accommodating almost any workload from datacenter to edge – sophisticated in capability, yet simple to deploy and manage. This is an extremely careful balance to achieve – and it requires significant HCI experience and expertise.

Advantech has taken these and other key considerations into account in the design of its SKY-7260S3P server and storage platform for HCI, the latest entry in Advantech's SKY series of industrial server solutions spanning HCI, storage servers, edge servers and more. Based on AMD EPYC™ 7543 CPUs, Advantech's SKY-7260S3P is a fully featured, single socket server for virtualized workloads and HCI deployments.

The SKY-7260S3P supports high memory capacity up to 2TB for extraordinary memory throughput coupled with large I/O throughput leveraging up to 128 lanes of PCIe Gen4. This is complemented with powerful compute spanning up to 64 AMD "Zen 3" cores, high storage capacity up to 200TB (all HDD configuration), and networking flexibility via 10G SFP+ and 10GBase-T OCP mezzanine modules. Advantech's SKY-7260S3P delivers outstanding price/performance benefits and almost

ADVANTECH'S SKY-7260S3P

2X less power consumption than what was achievable with a competing x86 processing solution offered in a two-chip configuration, based on Advantech's own benchmark tests.[1]

Advantech's proprietary BMC management platform provides a reliable server and secure environment for next-generation AI-driven ICT infrastructure. This is complemented by AMD Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV) to provide state-of-the-art security features and a, reliable HCI server platform designed with enterprise-class robustness for the AI and IoT-driven compute requirements of tomorrow.

To learn more about Advantech's SKY-7260S3P HCI server and storage platform based on high-performance AMD EPYC™ processors, visit https://page.advantech.com/AMD

ABOUT ADVANTECH

Advantech, a global leader in industrial IoT, is fully committed to fulfilling its corporate vision of Enabling an Intelligent Planet, and its specific goals are aimed at partnering for smart city and IoT solutions. In recent years, in line with technology trends in IoT, big data and AI, Advantech's WISE-PaaS IoT cloud platform was launched to expand industrial application solutions across many different vertical markets. In addition, in order to promote vigorous IoT development, Advantech actively invites partners in the ecosystem to co-create solutions combining all their domain know-how. This will accelerate the completion of more smart applications and realize the vision of an intelligent planet. For more information, visit www.advantech.com*

ABOUT AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics and visualization technologies – the building blocks for gaming, immersive platforms and the datacenter. Hundreds of millions of consumers, leading Fortune 500 businesses and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit www.AMD.com

* Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

1. Compared to Intel Xeon Gold 6346 (two chip configuration)

AMD EMBEDDED SOLUTIONS