WHITEPAPER

High Performance Computing Can Still Be Sustainable

The Register®

AMD | TYAN

EXECUTIVE SUMMARY

This white paper discusses the challenges faced by datacenter operators in providing processing power and storage for AI workloads while also meeting sustainability goals and energy regulations.

Readers will learn how they can use the latest datacenter infrastructure equipment to find the right balance between compute performance and energy efficiency for their organization.

We will examine the combination of 4th Gen AMD EPYC processors, along with rack and barebone servers from manufacturer Tyan, and how the two designs work together to boost performance while reining in energy usage and computing costs.

INTRODUCTION

Modern datacenters face a major sustainability challenge as high-performance computing (HPC), once strictly the domain of weather forecasters and medical researchers, has become more mainstream. HPC uses many silicon cores to run large computing jobs in parallel. From distributed data processing using tools like Hadoop to the rise of AI training and inference jobs, compute-intensive workloads have become increasingly common.

This trend will only intensify as the rise of the Internet of Things (IoT) and connected devices contributes to an explosion of data. HPC workloads are likely to proliferate across all datacenters, from hyperscale facilities to colocation and enterprise-owned buildings.

The rise in compute-intensive workloads will mean sustainability headaches for many datacenter operators and colocation customers. The International Energy Agency currently estimates global annual datacenter electricity usage at between 220-320TWh, or around 0.9-1.3% of total global electricity demand[1]. That produces around 300Mt of CO2, equating to around 0.6% of total global greenhouse gas emissions.

Increased power consumption in datacenters is self-propagating; the more power that server equipment consumes, the more heat it generates, calling for more cooling. The cooling system also consumes power, accounting for roughly a third of the total facility's energy usage. Together, server and cooling energy requirements drive up datacenter electricity consumption.

Most datacenters use evaporative cooling systems, in which water heated by hot air from the server room passes through a cooling tower that sprays it into a fill material. Some of the water evaporates into cooler, dryer outside air pulled through the fill material by fans. The rest, cooled by the process, falls back to the bottom of the tower and returns to capture more heat from the server room.

This cooling process uses lots of water - up to 8.55 gallons per minute for a 1MW facility, according to a 2016 analysis by the Uptime Institute[2].

Footnotes

1. International Energy Agency, "Datacenters and Data Transmission Networks", September 2022. https://www.iea.org/reports/data-centres-and-data-transmission-networks
2. Uptime Institute, "Ignore Data Center Water Consumption at Your Own Peril", June 2016. https://journal.uptimeinstitute.com/dont-ignore-water-consumption/

Regulators too are beginning to eye datacenter sustainability issues. Ireland for example sees a greater-than-average burden from datacenter energy use due to its popularity among tech firms, leading to the country's state-owned grid operator, EirGrid, limiting the development of datacenters in Dublin until 2028[3]. Singapore lifted a similar moratorium in 2022, but revealed new guidelines on datacenter sustainability[4].

In the EU, a forthcoming Corporate Sustainability Reporting Directive will take the region's Code of Conduct for Energy Efficiency in Datacenters - formerly a recommended set of guidelines - and make them mandatory.

Even in regions where energy efficiency is not regulated or where water consumption remains unmetered, pressures for sustainable operation are still rising. Water is becoming increasingly scarce in some areas as drought conditions increase, and energy prices are rising. So sustainable operations increasingly make economic sense too.

REDUCING THE SERVER FOOTPRINT

One of the best ways to manage power demands in the datacenter is to economize on server energy consumption. Many high-performance workloads bolster their performance using accelerator cards, with their enhanced floating point arithmetic and large core counts. However, CPUs have and will continue to play an important part in high-performance workloads. They are crucial in adding precision to large-scale calculations, and they coordinate other components such as the accelerators. A highly functional CPU is crucial, but it should not force unreasonable trade-offs in terms of energy consumption to achieve its goals.

AMD'S EPYC PROCESSOR

AMD has been driving datacenter sustainability for years with innovative chip designs, kicking the mission into high gear in 2017 with the launch of its AMD EPYC 7000 Series Processors, a family of chips built specifically for datacenter servers.

In 2021, the company announced a concrete goal to boost energy efficiency for its EPYC CPUs and Instinct acceleration cards by a factor of 30 between 2020 and 2025. It expects the power required for a single AI and HPC calculation to subsequently plummet by 97%[5].

AMD is delivering on its promises. It had already reached a 6.79x improvement over its 2020 energy consumption as of mid-2022 with its 3rd Gen AMD EPYC Processor range[6], again focused on datacenter servers. In November that year the company advanced its sustainability mission a step further with the launch of the 4th Gen AMD EPYC Processors.

The 4th Gen AMD EPYC processor was one of several significant steps in a broader sustainability journey for AMD.

Footnotes

3. Datacenter Dynamics, "EirGrid pulls plug on 30 Irish data center projects", May 2022. https://www.datacenterdynamics.com/en/news/eirgrid-pulls-plug-on-30-irish-data-center-projects/
4. Lexology, "Building Green Datacenters - Singapore Lifts Moratorium on New Datacenters, Introduces Environmental Sustainability Standards", May 2022. https://www.lexology.com/library/detail.aspx?g=254ea4fc-98b1-465f-9904-2f4c4c3496fc
5. AMD, "AMD Announces Ambitious Goal to Increase Energy Efficiency of Processors Running AI Training and High Performance Computing Applications 30x by 2025", September 2021. https://www.amd.com/en/press-releases/2021-09-29-amd-announces-ambitious-goal-to-increase-energy-efficiency-processors
6. AMD, "Data Center Sustainability", Accessed July 2023. https://www.amd.com/en/corporate/corporate-responsibility/data-center-sustainability.html

This was one of several significant steps in a broader sustainability journey for AMD. In June 2023, the company launched its 4th Gen AMD EPYC 97x4 Processors, which target cloud servers and uses smaller cores, along with its AMD EPYC 9004 Series Processors with AMD 3D V-Cache technology, a flagship processor for compute-intensive applications. AMD CEO Lisa Su said that this chip would deliver 80x extra performance per watt compared to its closest rival.

The EPYC Zen 4 core at the heart of AMD's chiplets is inherently sustainable by design.

EPYC's sustainability benefits by the numbers

4th Gen AMD EPYC processor power-performance improvements produce some impressive numbers. AMD tested 2000 virtual machines running on 11 EPYC 9654 chips, each sporting 192 threads on 96 cores. It ran the same setup on 17 Intel Xeon 8490H servers.

The results of the test showed a 29% drop in energy consumption, along with an estimated saving of 35 tons of CO2 annually.

CHIPLET DESIGN

AMD switched away from a monolithic core design to smaller chiplets in 2015. This enabled it to expand the per-processor core capacity which can be embedded in datacenter servers. The smaller chiplets are easier to produce with fewer on-silicon errors, also improving its yields.

The chiplet design positions a collection of Zen 4 architecture cores around a single I/O die which handles communication between them.

The EPYC Zen 4 core at the heart of these chiplets is inherently sustainable by design. It is highly integrated containing enough on-board functions to qualify as a system on chip (SoC). Fewer external chips means more CPU and server power efficiency.

Separating the I/O and core chiplet designs allows the two components to evolve at different rates. The I/O die, sporting a separate memory controller for each core, uses a 6nm process size. It communicates with each core using AMD's ultra-fast Infinity Fabric.

AMD has pushed its EPYC 9004 Series Processor's CPU cores to an even smaller 5nm die size. This increases the processor's compute density, enabling it to push out more calculations across its innovative chiplet design. Competitors are still struggling to move beyond a 10nm die size.

The chiplet design also allows AMD to keep pushing up the CPU core count in datacenter servers. The 4th Gen AMD EPYC processors contain multiple CPU dies, each having several cores. This version of the chip has up to 12 dies, each containing eight cores, for an impressive upper limit of 96 cores, supporting 192 threads . This shatters the core count on the previous AMD EPYC 7003 Series Processors, which feature a maximum 64 cores. The 4th Gen AMD EPYC 97x4 processor design released in June, has eight dies with 16 cores each for a maximum 128 cores, with 256 threads per socket. This far exceeds the number of CPU cores offered by competing vendors and server manufacturers using different chips, increasing compute density and keeping AMD ahead in terms of the power/performance ratio it can offer for datacenter servers.

A CONFIGURABLE BALANCE OF POWER VS PERFORMANCE

AMD offers several SKUs for its 4th Gen AMD EPYC processors with varying levels of memory and core counts. This enables server manufacturers like Tyan to right-size processors to match different customer workloads by varying the cores and power requirements. Across different categories:

- *Balanced:* The AMD EPYC 9004 Series core chiplets each contain eight Zen 4 cores that are reliable workhorses on balanced enterprise workloads such as software development and business applications.
- *License cost challenged:* Some enterprise software vendors bolster their revenues by charging on a per-core basis for virtualized workloads. AMD can help to drive down those revenues using lower core counts operating at higher clock speeds.
- *Memory intensive:* AMD EPYC 9004 Series processors with AMD 3D V-Cache technology provide access to more L3 cache for on-chip calculations serving applications with high memory requirements.
- *Compute intensive:* Finally, high-performance computing workloads hum along nicely with EPYC's high core counts. The AMD EPYC 97x4 Processor is a winning proposition here thanks to its maximum 128 cores per processor, which offers 256 cores and a massive 512 threads in a dual-socket datacenter server configuration.

While server manufacturers and customers can buy different SKUs, the AMD EPYC 9004 Series processor is designed for flexibility due to its configurable thermal design power (TDP). The TDP, measured in watts, represents the power consumption of the chip when under maximum load.

It's lowest TDP is 45W higher than that found in the previous 3rd Gen AMD EPYC CPUs, while the highest is up by 120W. The important metric here is the performance it delivers per watt. AMD says that it offers 1.7 times more performance per watt than its predecessor. That means the 4th Gen AMD EPYC processor is using a little more power but doing a lot more work, which has big implications for datacenter server compute density.

Its TDP is designed with workload flexibility in mind. The company has allowed server manufacturers and customers to configure the TDP for different workloads for a while now. The configurable TDP enables those buyers to save power on different workloads by pulling the power consumption limit down to a lower target in the BIOS settings. That could be cumbersome in a large fleet of servers, but Tyan offers TSM, a configuration management tool that enables administrators to adjust and automatically push out updated BIOS settings across a large group of EPYC-based servers.

MORE CACHE FOR BETTER PERFORMANCE

AMD has also boosted the amount of cache available to those cores. An increase in cache memory is critical to server performance because it puts short-term 'scratchpad' memory close to a core's on-board compute circuitry, reducing the latency and bottleneck that some other chips suffer when accessing off-chip system memory. In short, the more cache on the processor, the faster it is.

AMD EPYC 9004 Series processors feature 1MB of private L2 cache per core to drive better server compute performance compared to the 512KB found on AMD EPYC 7003 Series processors. They also have up to 1152MB of L3 cache, compared to a maximum of 768MB on the previous generation's silicon.

We can thank AMD's 3D V-Cache feature - also available on the AMD EPYC 7003 Series chips - for this increase in maximum L3 cache. The company developed a copper-to-copper hybrid bonding that fits more memory onto a core by stacking it directly onto the silicon. It claims to offer 200X the interconnect density of regular 2D cache connections, which also helps to boost chip and server performance.

It isn't just the cache memory support that got a boost with the 4th Gen AMD EPYC processors. This is the first of the EPYC chips to support DDR5, a new, faster memory architecture that offers powerful performance benefits by doubling the throughput from system memory, levelling up performance for high-bandwidth applications. It also supports up to 6Tb of DDR5 memory per socket across 12 memory channels.

NEW CONNECTIVITY IMPROVEMENTS

The AMD EPYC 9004 Series processor comes with some other firsts, not least of which is its support for PCIe 5, or Gen 5 – the interconnect standard which links different server components to ensure maximum data transfer speeds. At up to 32Tb/sec, this is twice as fast as its predecessor, providing faster data throughput to external devices connected to the server system board.

The latest chip generation also increases the number of PCIe lanes available. Whereas AMD EPYC 7003 Series processor featured up to 128 PCIe 4 lanes, the 4th Gen chip offers up to 160 on a dual-socket configuration. It can also dedicate 32 lanes to SATA connections, which is especially useful for server configurations that need to draw in lots of data from faster off-chip storage, such as NVMe SSDs.

AMD also enables server manufacturers to dedicate 64 of those PCIe 5 sockets to something else new in its 4th Gen AMD EPYC processors: Compute Express Link (CXL). This provide cache-coherent connectivity between the processor and off-chip devices, including memory devices. CXL offers hardware support for both the chip and the devices to share the same cache memory, effectively allowing users to address remote memory as local memory.

CXL also provides more support for CPUs to share memory with accelerator cards, providing better performance and simpler memory management for customers exploring HPC and AI applications.

AN UPGRADED INSTRUCTION SET

Finally, the 4th Gen AMD EPYC processor also included more support for Intel's 512-bit wide advanced vector extensions (AVX-512), which is an extended instruction set better enabling software developers to conduct simultaneous operations across a range of data items. This is especially important for some HPC and AI applications, such as the Tensorflow AI toolset, which support it.

AMD has taken a sustainable approach with its support for this instruction set, too, by choosing a 256-bit bus to pass instruction data, retrieving it in two cycles. This allows for a smaller die size while also offering better power management. On chips using a 512-bit data path for AVX instructions, the extra power requirement sometimes forces power throttling to avoid overheating on the chip, but there are no such worries here.

The sum total of all these new capabilities is that datacenter servers based on AMD EPYC processors can be optimized to deliver more performance per system and per rack – critical to delivering capacity growth in hosting facilities with constraints on power costs and availability. And providing the best possible workload performance per watt also means datacenter servers can do more while using less power per unit of work, making them better placed to offer the type of density needed to support efficient, scale out cloud-native workloads.

TYAN: A SERVER LINEUP BUILT FOR AMD EPYC

AMD's chips are an important part of the equation, but nothing works without an appropriate high-quality server for them to go in. Tyan's experience here is critical. The company is an original design manufacturer (ODM) that custom-produces servers directly for large clients. It regularly submits RFPs for custom fleet builds in the 5-10,000 server range.

Tyan, which is a server brand owned by MiTAC Group, also has a channel operation to help regulate inventory. Served by local sales and support offices in North America, Asia Pacific and Europe, these products bring its efficiency of scale and manufacturing expertise to a wider range of customers.

Transport CX TD76-B8058

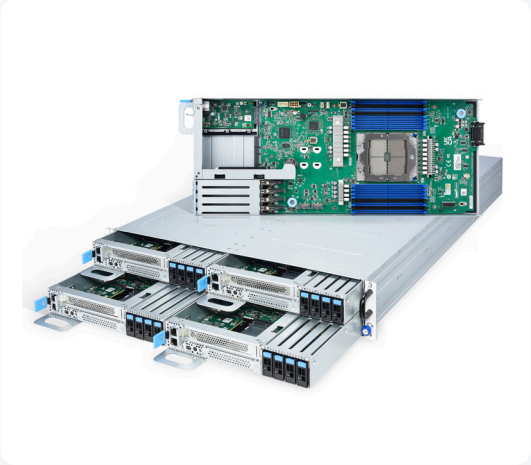

These systems included the Transport CX TD76-B8058, a cloud-focused 2U 4-node rack server with single-socket configuration, heavily focused on computing density and power efficiency.

The CX TD76-B8058 is a single-socket device. Tyan has invested heavily in single-socket servers because this design offers multiple benefits. Many customers want more easily configurable workloads that don't have to grapple with non-unified memory architecture (NUMA) resource handling (although the 4th Gen AMD EPYC processor is certainly well-equipped to handle dual-socket configurations).

Single-socket devices are appealing from a sustainability perspective because the server doesn't have to devote power to an inter-processor bus. It also doesn't have to contend with heat from one chip passing over the next, making it easier to cool.

Thanks to AMD's powerful multi-core options, single-socket designs don't mean a trade-off in computing power. This server is compatible with EPYC devices all the way up to the 4th Gen AMD EPYC 97x4 Processors, providing a maximum of 512 cores per box.

Tyan has carried EPYC's sustainability ethos through to the design of its TD76-B8058 server chassis. It chose front I/O ports for this server, placing them all at the front of the rack. In traditional servers all the I/O is at the back, but this means that the front-to-back airflow blows hot air onto the network card. This design protects the network card from the CPU heat, and works with the wide chassis design to make the server easier to vent into the hot aisle.

This rack-mounted server also features extended volume air cooler (EVAC) heat sinks, which have pipes connecting the heatsink fins on top of the CPU to remote outrigger sinks along the chassis. Finally, a built-in temperature monitoring mechanism for the CPU and other critical system components dynamically adjusts the fan controller to minimize the power expenditure spent on shifting air through the unit.

Transport CX GC68A-B8056

Fast data transfer is a big requirement for cloud applications. That's why Tyan announced this cloud-focused I/O server to coincide with the launch of the AMD EPYC 9004 Series processors in November. This 1U system supports the Open Compute Project's NIC mezzanine card for high-speed networking.

The CX GC68A-B8056 features 24 DDR5 DIMMs to support EPYC's interoperability with this new memory type. It also sports a pair of PCIe 5.0 x16 expansion slots and two 10GbE onboard Ethernet ports in a 1U configuration. The server's twelve 2.5" tool-free drive bays support NVMe U.2 SSDs, making it a good choice as an all-flash array platform to support high data throughput in a datacenter environment.

Transport SX TS70A-B8056

For enterprises needing even more all-flash storage, the SX TS70A-B8056 offers double the capacity. This is a 2U single-socket storage server with the same 24 DDR5 DIMM slots and an extra PCIe 5.0 slot. It also offers double 26 2.5" NVMe U.2 hot-swap drive bays.

NVMe is known for its fast throughput. All those extra IOPs can end up choking network data transfer for very demanding applications unless the server internals are boosted to support it. To avert this issue, Tyan has kicked the I/O up a notch with its Transport SX TS70A-B8056 unit, offering up to 400Gbit in Ethernet throughput.

Transport HX TN85-B8261

Tyan might offer plenty of support for single-socket customers but it hasn't skimped on dual-socket offerings. Dual-socket designs are especially good for accelerated computing, which is why the company used that design for this 2U server with 24 DDR5 slots and eight NVMe trays.

This unit uses registered DIMM (RDIMM) slots rather than regular DIMM. This design, which attaches a register between the CPU and the DRAM chip, increases transmission efficiency. It also supports buffered and high-performance register modes for better stability than other forms of dual in-line memory modules.

The TN85-B8261 is Tyan's mid-range server for handling GPU cards and accelerators, supporting up to four double-wide GPU cards and two half-height PCIe 5.0 x16 slots. Tyan built this design with HPC and deep learning performance in mind, but there will also be a non-GPU variant enabling customers to focus on storage rather than small AI training jobs.

Transport HX FT65T-B8050

Some of these single socket designs make their way outside of the server rack. The FT65T-B8050 is a rack-convertible pedestal unit that can operate as a desk-side server designed for applications like video editing and smaller AI training/inference models that focus on local data, for example.

The unit features eight DDR5 DIMM slots, eight 3.5" SATA and two 2.5" NVMe U.2 hot-swap, tool-free drive bays. The FT65T-B8050 supports up to two double-wide PCIe 5.0 x16 professional GPUs along with two additional high-speed networking adapters for parallel clustered workloads.

CONCLUSION

Both Tyan and AMD have something in common: They make their products for large customers whose focus is on driving down server power consumption per watt to get the maximum efficiency per square foot of datacenter space.

We see a similar focus on sustainability in their designs. AMD has taken a commodity chip architecture - the x86 - and modified its design to provide unique advantages over the competition. Tyan has done the same in the server space thanks to innovations in areas like I/O and cooling. Together, these companies will help make incremental improvements in an operating environment where every watt and instruction cycle counts.

Sponsored by Tyan and AMD

30x Projected power efficiency increase in AMD server chips 2020-2025

1.7x Performance-per-watt increase from 3rd Gen AMD EPYC processors to 4th Gen AMD EPYC processors

97% Projected reduction in power-per-instruction for AMD server processors

About TYAN

TYAN is a world class provider of server motherboards and server platforms. Founded in California's Silicon Valley in 1989, TYAN has been bringing innovation and reliability to the server industry for over 30 years. As a leading server manufacturer owned by MiTAC Computing Technology Corporation, TYAN is a core business unit of the MiTAC Group.

As a leading manufacturer, TYAN provides systems integrators and value added resellers with market leading products to build and customize their servers worldwide. By offering a full range of server products and accessories, TYAN enables customers to apply their expertise and bring value to their solutions. TYAN prides itself in supplying the channel with building block products.

Many customers have had their TYAN-based products recognized within their industries, covering a wide range of markets such as HPC, hyper-scale/data center, server storage and security appliances.

TYAN has locations spanning the globe, from North America to Asia Pacific and Europe, including R&D, distribution, local support, inventory and sales.

Contact us:
www.tyan.com
United States: sales_america@tyan.com
EMEA: johnny.cy.wang@tyan.com