CUTTING POWER BILLS AND EMISSIONS WITH A MODERNIZED AND EFFICIENT IT INFRASTRUCTURE

Authors:

Andrew Buss
Vladimir Kroa
Zuzana Kovacova

April 2023

An IDC Market Spotlight sponsored by AMD

IDC #EUR150512723

Cutting Power Bills and Emissions With a Modernized and Efficient IT Infrastructure

Introduction

The global economy is accelerating a transition to digital business engines. Goods and services increasingly depend on IT infrastructure, and with this the challenges and opportunities related to controlling electricity costs and greenhouse gas (GHG) emissions are compelling.

In 2022 alone, IDC forecasts that the global spend on digital will top $1.8 trillion, with $469 billion of this in EMEA. This will continue to grow, reaching $2.8 trillion by the end of 2025.

This relentless march of digital means that digital power consumption and emissions are also becoming a significant contributor to not only the overall emissions, but also the power bill. The growth in spend on digital continues unabated, requiring us to rethink how we buy, build, and operate our digital infrastructure to eliminate overhead and waste and reduce GHG emissions.

Europe has recognized this need for more sustainable digital transformation. The EU aims to make Europe the world's first climate-neutral continent by 2050. This is the objective behind the European Green Deal[1] (COM(2019) 640 final), an ambitious strategy that should drive European businesses to benefit from advancing efficient and more sustainable business strategies.

AT A GLANCE

WHAT'S IMPORTANT

- Rising power costs are leading companies to reevaluate the power efficiency of their IT infrastructure and the wisdom of running servers beyond five years.
- The growing extent of digital business models means that growing IT emissions are becoming a much greater concern in digitally transformed businesses.
- Upcoming efficiency and emissions regulations and reporting requirements will significantly impact datacenter investments in the EU.

KEY TAKEAWAYS

- A proactive infrastructure renewal that provides significantly more performance for the same power can help to reduce emissions and dramatically cut operational costs by lowering power bills.
- IT organizations should look to avoid overhead and waste by using automation and AIOps to enable dynamic workload management to run workloads on the fewest, most modern, and most efficient servers as possible, powering up less efficient servers only to handle burst requirements.
- Switching to renewable energy sources can dramatically reduce emissions, but the switch will not come cheap in the medium term, so companies need to increase infrastructure efficiency by using modern, energy-efficient server processors. This includes the latest generations of CPUs and workload accelerators such as FPGAs, GPUs, and DPUs to reduce power consumption.

FIGURE 1

From Digitally Transforming to Running a More Sustainable Digital Business

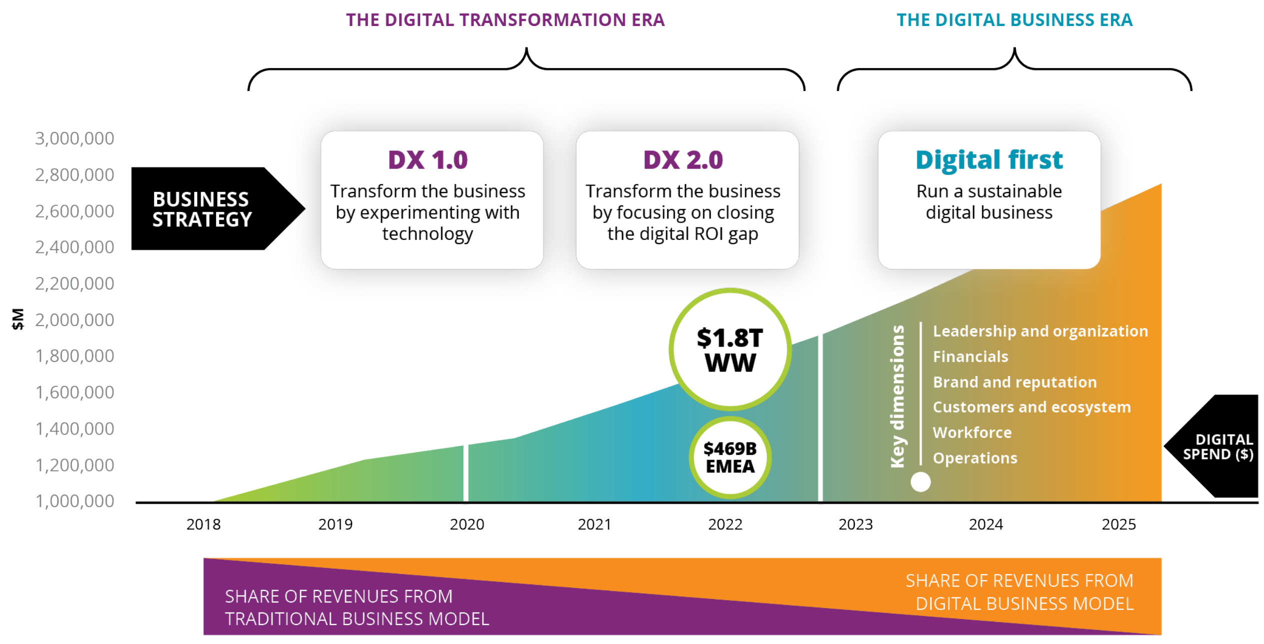

Source: IDC Worldwide Digital Transformation Spending Guide, 2021 V2

As part of this drive, the EU is making available almost €700 billion in Next-Generation EU funds, with at least 20% of RRF funding to accelerate the transition to digital business, and 37% which should support climate-related initiatives.

The Importance of More Sustainable Infrastructure

Environmental, social, and governance (ESG) is now a top concern for most European IT organizations[2], with 51% saying it is very important overall, rising to 69% in the top 25% most digitally transformed Digital Leaders.

What this means for many IT organizations is that increasingly they are getting dedicated sustainability funding to improve the efficiency and sustainability of their IT services. In the IDC *2022 European Sustainability and Technology Survey*, only 22% of companies dedicated more than 10% of their IT budget to sustainability- and efficiency-related initiatives in 2022. By 2024, 51% of companies will be doing this.

The importance of ESG in Europe is keenly felt in emissions-heavy industries such as manufacturing and resource industries such as mining, as well as construction, retail, and telecommunications. Despite driving much of the sustainability agenda in Europe, government tends to lag behind, as does banking, which, being performance heavy and physical assets light, tends to focus on ultimate performance over efficiency. But even here, the signs are that changes to improve efficiency are being adopted.

While ESG has been growing in overall importance over the past decade, the environmental side — and particularly energy efficiency — was given additional urgency in 2022 by the sharp rise in the cost of power due to geopolitical uncertainty and the resulting uncertainty of gas and energy supplies.

This has led to structural change in energy supplies in Europe with an accelerated shift toward investing significantly more in renewable energy from solar and wind. The impact of this is that the higher prices for energy seen in 2022 are likely to be sustained for some time as the region progresses on this energy supply rearchitecting.

The Four Key Pillars of IT Infrastructure Sustainability and Efficiency

IT infrastructure is a complex and interdependent set of suppliers, facilities, hardware, and software that all needs to work together to deliver a resource-efficient platform for digital innovation. This can be thought of as four key pillars that need to be considered both individually and collectively:

- Modernized and efficient infrastructure
- Intelligent workload management and infrastructure utilization
- Low-waste datacenter facilities
- Low- to no-carbon power supply

Making Best Use of Available Power With Modernized Infrastructure

IDC research[2] shows that European companies have tended to keep their IT infrastructure hardware in production for an extended period beyond the warranty and support term. Overall, 26% of active servers were three to five years old, and 22% were over five years. Prior to 2022, many companies routinely extended the operational life of servers to save additional capital spend on compute capacity with little thought to the cost of powering these servers as the relatively low and predictable cost of energy was not a significant pain point.

However, regulations are starting to increase regarding GHG emissions monitoring and reporting, and energy costs rose sharply in Europe in 2022. In some European markets, the non-household consumer cost of power increased threefold in a matter of months, with a cost per kWh in 2H21 of €0.15[3] rising to over €0.40[4] per kWh in 2H22 in markets such as Germany and the Netherlands. The result is a significant increase in the operational cost of a typical datacenter two-socket server as a result of power consumption. For companies operating hundreds or thousands of servers, this can be a very substantial increase.

Taking a reasonably configured two-socket server with a 1,200W power supply unit (PSU) operating at a sustained high load of around 80%, the power consumed by the server from the wall socket is around 660W, and over the course of a year would consume 5,780kWh of electricity. Table 1 shows 2H21 power costs compared to examples of increased 2H22 power costs in datacenters with a variety of PUE ratings.

TABLE 1
Energy Cost Example

TIME	ENERGY COST (€/KWH)[3, 4]	DIRECT SERVER POWER COST	PUE 1.1	PUE 1.4	PUE 1.7
2H21	€0.15	€870	€957	€1,218	€1,479
2H22	€0.40	€2,310	€2,541	€3,234	€3,927
2H22 power and cooling cost increase for 100 servers			€158,400	€201,600	€244,800
2H22 power and cooling cost increase for 1,000 servers			€1,584,000	€2,016,000	€2,448,000
2H22 power and cooling cost increase for 10,000 servers			€15,840,000	€20,160,000	€24,480,000

Source: IDC, 2023

CPU technologies continue to rapidly advance, and this is continuing with new approaches to designing and packaging CPUs, with technologies such as chiplets enabling new designs with unprecedented core counts and flexibility. CPUs are also continuing to advance rapidly on a performance/watt basis even in a short space of time. The initial generation of AMD EPYC server CPUs launched in 2017 featured up to 4 compute die and 32 cores per socket on a 14nm process node. Per socket performance increased in 2019 with the launch of the second-generation EPYC CPUs with 8-core CCD chiplets manufactured on a 7nm node, resulting in a doubling of core count, with up to 64 cores per socket. Architectural changes in 2021 boosted performance with the third generation of EPYC CPUs, and in 4Q22 AMD launched the fourth generation with an all new design with 6nm IO die and up to 12 5nm 8-core CCDs, for a maximum of 96 cores per socket. The combination of increased core count and enhanced CPU cores, coupled with the smaller process nodes, has doubled the performance per watt per socket over the intervening five years.

This means that running servers for many years without tracking and monitoring their workload output and power consumption may now be a significant but hidden operational cost. IDC research[2] shows that many companies realize the most effective way to cut IT operations costs and reduce emissions is to refresh the server estate to remove the servers with low performance per watt. In addition, increased per socket core density and improved performance, coupled with advanced reliability, availability, and servicing (RAS) features, enable server and socket consolidation to reduce overall server footprint without compromising IT infrastructure resilience.

Another area to consider is performance-intensive computing (PIC) workloads, such as high-performance computing (HPC), artificial intelligence (AI), and Big Data and analytics (BDA). These workloads put significant load on servers and infrastructure, and can often benefit from using workload accelerators (such as datacenter GPUs, FPGAs, or ASICs) to increase performance and cut power consumption. This is achieved by having hardware support for accelerating specific software routines, enabling them to scale to high performance without dramatically increasing the power budget too.

Using the Minimum With Intelligent Workload Management

All the above can help with the structural aspects of making IT infrastructure more efficient and sustainable. To be really efficient, workloads need to be intelligently managed to run on as few servers as possible. This means running servers with as high a utilization as feasible while still being able to deliver according to the service-level agreement or performance requirements of the workloads.

With intelligent workload management, as performance starts to approach the agreed limits, additional capacity can be started automatically and workloads started on or moved to the additional servers to increase scale.

The challenge to achieving this is that only the 24% most digitally transformed European companies — the Digital Leaders — have embraced the extensive automation and orchestration and AIOps that is needed to drive dynamic workload management[2] (see Figure 2).

FIGURE 2

Digital Platform Rankings and Element Adoption

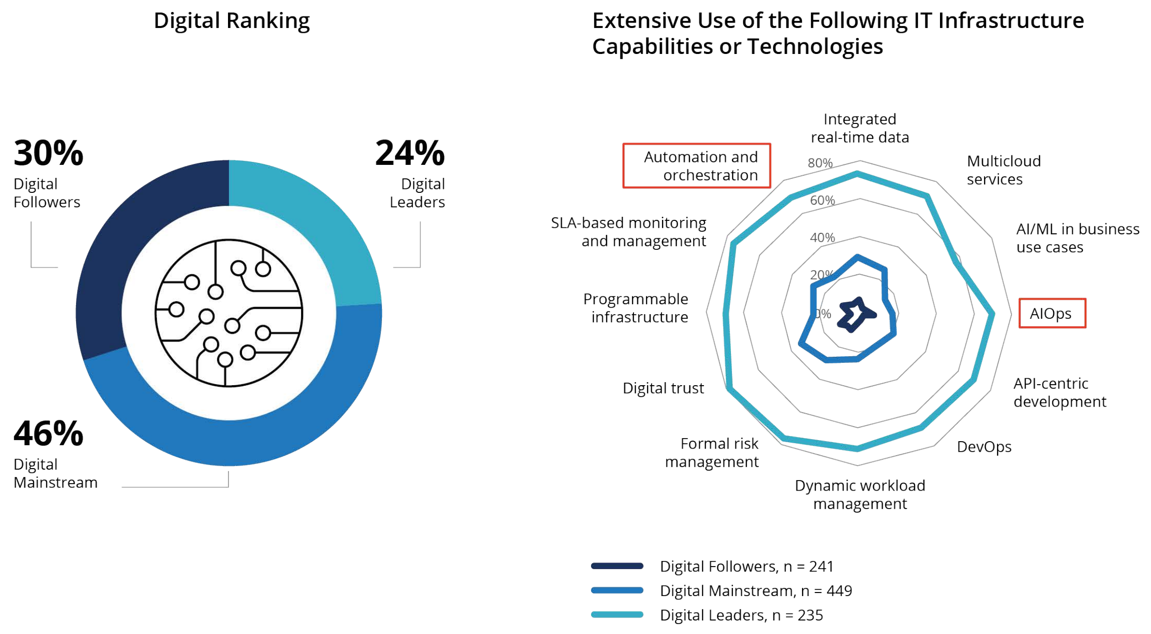

Source: *IDC European Infrastructure Survey 2022* (n = 928)

Investing in integrated monitoring and automation tools is critical for efficient and cost-effective infrastructure operations. This can be hard for many companies to achieve on their own, so working with advisors and partners that have experience of designing, implementing, and operating automation stacks can help accelerate an automation initiative while also cutting the costs and reducing the associated risk. Digital Leaders are also leading the shift to programmable infrastructure such as hyperconverged infrastructure (HCI), which is built for cloud-native automated operations. Increasingly, full cloud stacks including automation and billing engines such as Microsoft Azure Stack HCI, Google Anthos, and VMware Cloud are being used to drive this shift to intelligently automated IT infrastructure.

Reducing Overhead and Waste With More Efficient Datacenter Cooling

When it comes to being more sustainable, the typical enterprise datacenter has many questions to answer. It's interesting to note that in the era of hyperscale datacenters, such as Microsoft Azure[5], a class-leading PUE of 1.1 is now the norm, and colocation providers such as Equinix[6] can achieve a very good PUE of 1.4, while enterprise datacenters may range from efficient at PUE 1.7, to significantly worse at PUE 2.0 or higher[7].

Improving the efficiency of datacenters remains a secondary priority for many companies[2]. Instead, the top priority for improved sustainability is modernizing the IT infrastructure. However, the scale of datacenter facility operations power overhead and waste — where the energy required to deliver power and cool the infrastructure can be more than double that required to actually run the infrastructure — is leading to higher costs as power costs spike and GHG emissions increase.

While updating enterprise datacenters midcycle is not typically feasible, for datacenter refreshes or new builds, and for workload refresh and migration, there are a number of elements that can help to drive down emissions and power costs:

- Where workloads are being modernized or refreshed, Digital Leaders are increasingly moving workloads to either public cloud or colocation-based infrastructure to lower overall power consumption and emissions.
- Digital Leaders are also focusing on improving the efficiency of their datacenter cooling with new approaches such as free-air cooling, higher-temperature operations, and closed-loop liquid cooling.
- Reutilization of waste heat for in-building or district-based hot water and space heating.
- Adoption of datacenter infrastructure management (DCIM) automation and tooling to dynamically monitor and manage the environment of the datacenter facility for optimum efficiency. Achieving this in practice can be a challenge, as IDC's research shows this is more a facilities operations role (rather than IT), and the facilities role is often disconnected from intelligent workload management.

Investing in more modern cooling approaches can also help to improve the resilience and reliability of digital operations, as companies have frequently suffered outages due to HVAC outages or local hotspots.

Decarbonizing IT With a Rich Mix of Renewable and Clean Energy Supplies

No matter how efficient infrastructure can become over time, it will always require energy to power it, and selecting the energy source can be one of the biggest levers to reduce overall GHG emissions. Energy that is from 100% renewable energy resources emits no net new emissions, meaning that overall IT infrastructure operations emissions would be close to zero.

While many hyperscalers, colocation providers, and service providers have moved to negotiate energy supply contracts that are significantly or even fully composed of renewable energy, many European enterprises still rely on fossil fuels for much of their overall datacenter energy needs.

Over time, heightened investment in renewables means that they will grow to make up a larger proportion of the overall energy supply.

For many companies that are not yet digitally transformed, this will be sufficient for their IT infrastructure GHG emissions reduction plans, but for companies making heavy use of digital technologies for their overall business and operations, it's time to proactively manage energy supply contracts to accelerate the proportion of renewable energy in the energy supply mix.

FIGURE 3

The Most Important Elements to Improve Energy Efficiency and Sustainability

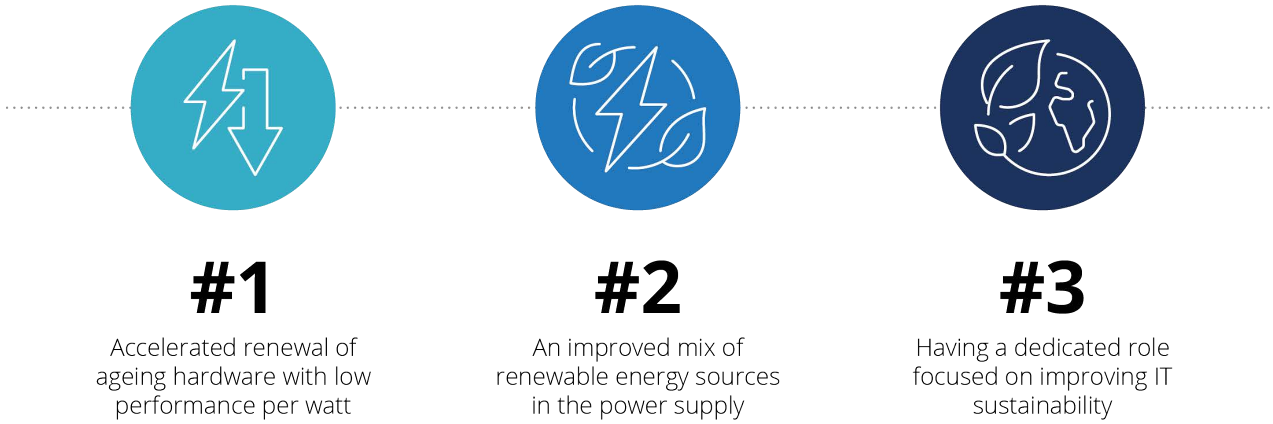

Source: *IDC European Infrastructure Survey 2022*, n = 928

Another area to consider is to migrate toward using energy storage technologies such as large-scale battery installations that can help supply power during peak loads and recharge when loads are lower to balance out power use and make use of off-peak discounting energy pricing where available.

Looking more to the future, a key consideration for decarbonizing infrastructure operations will be understanding the optimum sites for new datacenter builds. With modern global internetworking capabilities, companies can now operate all or significant parts of their IT infrastructure in countries that enable more sustainable datacenter builds. Countries with lower temperatures can vastly reduce the need for active cooling and therefore wasted energy use and overhead, while proximity to predictable renewables such as hydroelectric and geothermal power sources can deliver carbon-neutral energy.

Considering AMD for Your Sustainable IT Infrastructure

Over the past five years AMD has reemerged as a potent force in datacenter compute. AMD launched the first generation of the innovative multidie EPYC server CPU platforms in 2017, and subsequently introduced the 2nd, 3rd, and 4th generations in 2019, 2021, and 2022 respectively. The generational improvements have been substantial. The 2nd generation EPYC line grew the core count from 32 to 64, spread across 8 CCD chiplets. The introduction of a dedicated IO die solved some of the key issues of the 1st generation EPYC line, enabling more scalable and uniform performance. The 3rd generation introduced redesigned CCD chiplets, giving all 8 cores unified access to the L3 cache, and new 3D V-Cache variants helped accelerate memory-intensive workloads. The 4th generation EPYC platform — launched in November 2022 — has moved the needle yet again with up to 96 physical cores and an improved IO die with PCIe 5.0 and 12 channels of DDR5 memory.

These multiple generations of EPYC server CPUs have disrupted the market and accelerated the pace of innovation. As a result, between 2Q17 and 2Q22, AMD tripled its share of the x86-64 server CPU market in EMEA[8]. It has seen particular success driving new capabilities into new markets, such as growing the 1-socket server space and achieving a 27% share in EMEA in this space in 2Q22 as high core counts, large memory support, and expansive PCIe connectivity enabled multisocket servers to be replaced by scalable single-socket machines.

Along with the increases in performance that AMD introduced between 2017 and 2022 with the EPYC server CPU generations, it also moved production of the CPU and IO die silicon to more advanced process nodes, starting with Global Foundries 14nm in 2017 for the 1st-generation EPYC, to a combination of TSMC's leading-edge 5nm and 6nm nodes for the CPU CCDs and IO die respectively for the 4th generation in 2023. This progression naturally enables significantly more transistors to be included, increasing the core count and the scale of IO capabilities of the platform. Crucially it also reduces the power consumed by each transistor. The end result is a significant increase in the overall performance per watt.

We can see clear examples of the significant generational improvements in overall performance that have come in two years by looking at the performance per watt of some key benchmarks of the 64-core 3rd-generation EPYC 7763 compared to the 96-core 4th-generation EPYC 9654 (see Table 2).

TABLE 2

Performance per Watt Comparison

EPYC CPU MODEL/ BENCH-MARK	EPYC 7763 SCORE	EPYC 7763 DEFAULT SOCKET POWER	EPYC 7763 SOCKET-LEVEL PERF/ WATT	EPYC 9654 SCORE	EPYC 9654 CONFIG-URED SOCKET POWER	EPYC 9654 SOCKET-LEVEL PERF/ WATT	SOCKET-LEVEL PERF. IMPROVE-MENT	SOCKET-LEVEL PERF/WATT IMPROVE-MENT
SPECrate 2017_int-base	861 (9)	280	3.1	1,790 (10)	400	4.5	108%	46%
SPECrate 2017_fp_base	634 (11)	280	2.3	1,480 (12)	400	3.7	133%	63%
SPECjbb 2015 MultiJVM max-jOPS (critical-jOPS)	420,774 (13)	280	1,502.8	815,459 (14)	400	2,038.6	94%	36%

Source: SPEC POWER 2008

Customers can project and compare potential server performance, power consumption, and resulting greenhouse gas emissions and TCO by using AMD EPYC online tools, available at https://www.amd.com/en/processors/epyc-tools.

To optimize application performance and energy efficiency, AMD continuously works with industry-leading operating system and application software vendors, cloud service providers, and server vendors. Server power efficiency can be tracked by the SPEC POWER 2008 benchmark, updated regularly with the latest server test results, submitted by server OEM and ODM manufacturers.

SPEC POWER 2008 is an independent industry-standard benchmark for measuring performance per watt. It uses a server-side Java workload and performs a series of 10 runs, incrementing the power at 10% intervals and then providing an aggregate view of the performance per watt across all the runs. The overall work over power consumed is a good indicator of energy efficiency with many application workloads. This is critical in today's modern datacenters where power and cooling, along with real estate, are at a premium.

By January 2023, AMD EPYC processors led[15] in a number of SPEC Power categories, including:

- The overall result
- 1-socket and 2-socket standard rack configurations
- 1U/2U/blade form factors
- Linux and Windows operating systems

AMD EPYC server processors can deliver leading performance while reducing energy consumption[16]. As a result, this can help meet the most challenging application performance demands with fewer physical servers, leading to a smaller datacenter footprint, lower cooling requirements, lower hardware spend, lower TCO, and crucially lower power bills and significantly reduced greenhouse gas emissions.

The path forward with EPYC looks even brighter: AMD is committed to increasing energy efficiency 30-fold for EPYC processors and AMD Instinct accelerators for AI training and HPC applications by 2025. By the middle of 2022, AMD was on track to achieve 30 x 25, having reached 6.8x improvement in energy efficiency from the 2020 baseline using an accelerated compute node powered by one 3rd-gen AMD EPYC CPU and four AMD Instinct MI250x GPUs[17].

To accelerate performance-intensive applications that are not optimal for running on CPUs and compute GPUs, AMD also offers a diverse range of AMD Xilinx FPGAs and adaptive computing SoCs. AMD's latest addition in its datacenter portfolio is Pensando DPU, a CPU offload engine to accelerate and filter datacenter traffic at line rate with high performance and efficiency.

Challenges

The technology industry has suffered from component shortages and supply chain disruption since the pandemic. AMD relies on third-party manufacturing, primarily with TSMC and Global Foundries, to manufacture its portfolio of CPUs and GPUs, and its ability to supply the market is tightly coupled to its ability to secure manufacturing capacity on the appropriate process nodes in these foundries. AMD has been executing well and growing strongly over the past five years, particularly in the datacenter space with AMD EPYC CPUs, which are built on a flexible chiplet architecture to enable more diversity in product binning and targeting during assembly rather than at the early stages of manufacturing.

Although a long-established and significant player in the semiconductor market, AMD is still a relative newcomer to the modern datacenter infrastructure with EPYC CPUs, Instinct GPU accelerators, and Xilinx FPGAs. It takes time for software providers to certify their software and solutions on new platform systems, and there may be a more limited choice of AMD systems that are certified to run various applications or software stacks. However, the testing and validation that occurs on modern enterprise server and infrastructure platforms means the risk of incompatibilities is much reduced compared to a decade ago, and the potential advantages in performance and overall power efficiency will likely justify extra efforts in testing to ensure the hardware can run the PIC workloads reliably and consistently.

IDC Conclusion and Recommendations

Upcoming emissions reporting regulations coupled with actively managing power use and eliminating waste and overhead are driving a rethink of how to build and operate our IT platforms. As digital processes and services become more engrained, the power and emissions associated with this transformation are also increasing, and doing things how we've always done them is unsustainable.

How we buy, build, and operate our IT infrastructure needs to change, so that we can intelligently reduce and remove overheads and waste wherever possible.

Looking to the short term, over the next 6 to 12 months, you should aim to prioritize a number of objectives:

- Scrutinize your energy supply and contracts, and look to what can be done in three to six months to increase the use of renewables in the mix
- Audit your infrastructure energy efficiency and carbon footprint

- Identify and decommission the equipment with the worst performance/watt, and proactively replace this with modern, highly efficient servers and storage
- Look to your automation capabilities and how to implement them effectively
- Investigate the possibility of migrating power-hungry workloads to colocation facilities or to the public cloud

In the longer term, you should look to:

- Improve your datacenter facilities with new approaches to cooling
- Deploy AIOps across your datacenter facilities, IT infrastructure, and workload estate to continuously monitor and optimize workload management for maximum efficiency
- Focus on power supply diversification
- Consider investing in onsite energy generation such as solar or wind
- Replace diesel generators with hydrogen fuel cells
- Develop life-cycle assessments leading to carbon intensity optimization for all scopes of your IT-related emissions

References

1. European Green Deal, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Renewable_energy_statistics#:~:text=In%202020%2C%20renewable%20energy%20accounted,increasing%20from%2011.7%20%25%20in%202004
2. IDC's 2022 European Infrastructure Survey (n = 928)
3. Eurostat, Electricity Prices for Non-Household Consumers, Second Half 2021, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=File:Electricity_prices_for_non-household_consumers,_second_half_2021_v2.png
4. Electricity Prices Increased in Most of Europe, GMK Center, https://gmk.center/en/news/electricity-prices-increased-in-most-of-europe-in-december-2022/
5. Microsoft Azure PUE, https://azure.microsoft.com/en-us/blog/how-microsoft-measures-datacenter-water-and-energy-use-to-improve-azure-cloud-sustainability/#:~:text=Understanding%20Power%20Usage%20Effectiveness%20(PUE)&text=The%20closer%20the%20PUE%20number,efficient%20the%20use%20of%20energy
6. Equinix Datacenter PUE, https://cloud.kapostcontent.net/pub/fae282ef-8d29-4266-96cd-b1bc91164299/sustainability-report-fy2021?kui=nKl648GVlEjEXFAjNIzzuw
7. IDC's 2021 Datacenter Operational Survey (n = 400)
8. IDC EMEA Server Tracker, November 2022
9. 2P AMD EPYC 7763 (861 SPECrate 2017_int_base, 128 total cores, http://spec.org/cpu2017/results/res2021q4/cpu2017-20211121-30148.html)
10. 2P AMD EPYC 9654 (1790 SPECrate 2017_int_base, 192 total cores, www.spec.org/cpu2017/results/res2022q4/cpu2017-20221024-32607.html)
11. 2P AMD EPYC 7763 (634 SPECrate 2017_fp_base 128 cores, http://www.specbench.org/cpu2017/results/res2021q2/cpu2017-20210524-26504.html)
12. 2P AMD EPYC 9654 (1480 SPECrate 2017_fp_base, 192 cores, https://www.spec.org/cpu2017/results/res2022q4/cpu2017-20221024-32605.html)
13. 2P AMD EPYC 7763 (420,774 SPECjbb 2015 MultiJVM max-jOPS, 165211 SPECjbb 2015 MultiJVM critical-jOPS, 128 total cores, http://www.spec.org/jbb2015/results/res2021q3/jbb2015-20210701-00692.html)
14. 2P AMD EPYC 9654 (815,459 SPECjbb 2015 MultiJVM max-jOPS, 356204 SPECjbb 2015 MultiJVM critical-jOPS, 192 total cores, http://www.spec.org/jbb2015/results/res2022q4/jbb2015-20221019-00861.html)
15. MLN-094A: SPECpower_ssj 2008 overall ssj_ops/watt comparison based on highest system results published as of 07/06/2021; configurations: 2x AMD EPYC 7763 (23505 overall ssj_ops/watt,

http://www.spec.org/power_ssj2008/results/res2021q2/power_ssj2008-20210324-01091.html), SPEC and SPECpower_ssj are registered trademarks of the Standard Performance Evaluation Corporation, see www.spec.org for more information

16. EPYC-028: as of 2/2/22, of SPECpower_ssj 2008 results published on SPEC's website, the 55 publications with the highest overall efficiency results were all powered by AMD EPYC processors; more information about SPEC is available at http://www.spec.org; SPEC and SPECpower are registered trademarks of the Standard Performance Evaluation Corporation
17. EPYC-030: Calculation includes 1) base case kWhr use projections in 2025 conducted with Koomey Analytics based on available research and data that includes segment-specific projected 2025 deployment volumes and datacenter power utilization effectiveness (PUE) including GPU HPC and machine learning (ML) installations, and 2) AMD CPU socket and GPU node power consumptions incorporating segment-specific utilization (active vs idle) percentages and multiplied by PUE to determine actual total energy use for calculation of the performance per watt. 6.79x = (base case HPC node kWhr use projection in 2025 x AMD 2022 perf/watt improvement using DGEMM and typical energy consumption + base case ML node kWhr use projection in 2025 *AMD 2022 perf/watt improvement using ML math and typical energy consumption)/(2020 perf/watt* Base case projected kWhr usage in 2025); for more information on the goal and methodology, visit https://www.amd. com/en/corporate-responsibility/data-center-sustainability

MESSAGE FROM THE SPONSOR

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses, and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work, and play.

4th-gen AMD EPYC is ideally placed to help enterprises gain a competitive edge. Datacenter computing powered by AMD EPYC processors delivers leadership, application performance, TCO, and security features to help your business scale.

Find out more about AMD datacenter solutions at amd.com/en/processors/epyc-9004-series.

About the Analysts

Andrew Buss, Senior Research Director, IDC

Andrew Buss is senior research director for IDC's European Infrastructure Strategies research program. He is responsible for driving IDC's research covering present and future trends impacting servers, storage and networking, and IT service delivery. Central to this is understanding how on-premises IT is evolving under the emergence of open source, software-defined enterprise, multicloud adoption, and cloud-native development practices, and how this will impact everything from the low-level silicon underpinnings and system design, through to the design and integration of the different infrastructure components, up to the platform management and service delivery.

Vladimir Kroa, Associate Vice President, IDC

Vladimir Kroa leads the IT and business services delivery practice and related research for Central and Eastern Europe, the Middle East, and Africa (CEMA). He is responsible for the direction and development of the program, which includes support services, systems integration, outsourcing services, and business services and solutions. He also manages IDC's Cloud and Datacenter Services Practice for Central and Eastern Europe (CEE). He has been involved in projects ranging from market forecasts for the CEE and Middle East and Africa IT and cloud services markets, to strategic consulting involving market entry, competitive analysis, and end-user studies. He participates in local and international projects and delivers presentations to conferences and clients.

Zuzana Kovacova, Program Manager, IDC

Zuzana Kovacova is a program manager for IDC's European Services program and IDC Europe's latest Technology for Sustainability and Social Impact Practice, helping technology providers measure and maximize the impact of their own sustainability actions.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

IDC UK

5th Floor, Ealing Cross,
85 Uxbridge Road
London
W5 5TH, United Kingdom
44.208.987.7100
Twitter: @IDC
idc-community.com
www.uk.idc.com

Global Headquarters

140 Kendrick Street,
Building B
Needham,
MA 02494
+1.508.872.8200
www.idc.com

Copyright and Restrictions