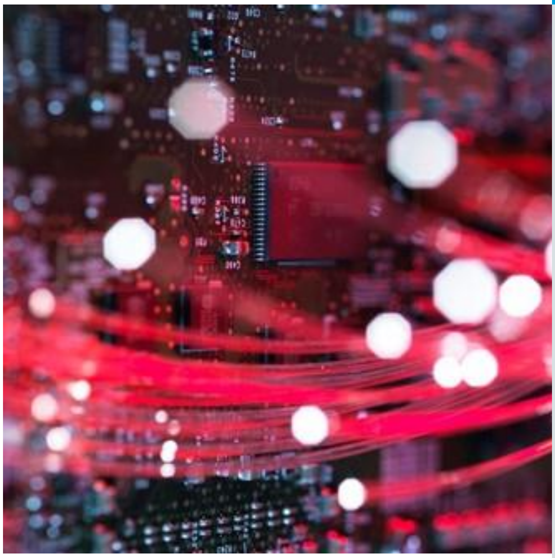

Study for the review of Commission Regulation 2019/424 (Ecodesign of servers and data storage products)

Task 4 Technologies – DRAFT v3

June 2024

Submitted to:

Davide POLVERINI

Policy Officer, DG GROW

Avenue d'Auderghem 45, 1040 Brussels, Belgium

Task 4 Technologies – Draft v3

A report submitted by ICF S.A.

Date: June 2024

Tom Lock
ICF S.A.
Avenue Marnix 17
Brussels
B-1000
Belgium
T +32 (0) 2 275 01 00
www.icf.com

Document Control

Document Title	Task 4 Technologies – Draft v3
Prepared by	Todd Leddy, John Clinger, Abhishek Jathar
Checked by	Eirini Passia, Tom Lock
Date	June 2024

Contents

Table of Tables

Table of Figures

4 Introduction to Task 4 Technologies

This report is produced in line with MEErP and serves to identify, retrieve, analyse data and report on a number of topics relevant to the Technologies of this study. These include the technical product description which also covers data on performance and impact of resources/emissions.

Furthermore, the report considers the production, distribution, and end of life aspects of the technologies.

Finally, the recommendations will be presented at the next stage of the report from a technical perspective with regards to product scope, barriers and opportunities for Ecodesign, along with the typical design cycle for the products and approximately appropriate timing of measures.

4.1 Technical product description

This section serves two main purposes which include the capacity building for the policy makers/stakeholders as well as a first assessment in a way of pilot/preview of the modelling work required as part of Task 6.

4.1.1 Existing products

The objective of this section is to outline the existing technology which will enable the work towards the definition of a Base Case(s).

4.1.1.1 Product overview

Servers and data storage products are equipment used for Information and Communication Technologies (ICT). Servers are different to data storage products as they offer additional functionalities, however they operate in an interconnected environment. Both are sold in a business-to-business market (B2B) environment with more complex procurement routes.

The ICT services offered are provided via connectivity and interoperability between the server and the data storage system. There are also additional hardware and software products that facilitate this connectivity and service. These are part of the network equipment products.

4.1.1.2 Servers

<u>Description of key components</u>

A server is a modular product (i.e., made up of more than housing element) which consists of the following key components. These remain unchanged from the previous preparatory study[1]:

- **Chassis**: A metal/plastic casing that encloses all components. It also incorporates mounting features for them.
- **Mainboard (server board)**: This is the main printed circuit board (PCB) of the server. It provides the firmware, processing units and points for mounting the

[1] Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC, DG ENTR Lot 9 - Enterprise servers and data equipment, Task 4: Technologies, July 2015 – Final report, Produced by bio by Deloitte and Fraunhofer IZM

memory modules and further active components. It also includes passive electronic components such as capacitors, resistors etc

- **Processor**: This is the central processor unit (CPU), the foundation of the server operation, which processes numerous software instructions to derive a required result. The processor works closely with memory, which both hold the software instructions and data to be processed as well as the results or output of those processor operations.
- **Memory (Random Access memory)**: RAM keeps the relevant instructions and data required by the processor. It also holds any output from the processor.
- **Storage devices/drives**: Hard disc drives (HDD) were the first ones to be produced as well as semiconductor-based solid state devices (SSD) that are integrated into a housing of a disk drive.
- **I/O control and network connectors**: Servers are remotely accessible via Ethernet connection and therefore feature multiple network interfaces and links (connectors) on the backside.
- **Cooling system**: The cooling of the active components (such as the CPU) is very important and is typically achieved by a combination of passive and active options. The active cooling includes in most cases a fan unit, while the passive cooling includes either a heat spreader attached to the processor, a large heat sink or heat pipes that distribute the thermal energy away from the processor. Both systems work in combination.
- **Power supply unit (PSU)**: This is typically a single unit or multiple unit coming in its own casing, usually metal.

The main components are discussed further below.

Chassis

The chassis (also typically called enclosure, case or housing) is in most cases, a simple metal box with a type of frame and mounting parts. The purpose of the chassis is to enclose and mount the main subassemblies including the server's printed circuit boards, storage devices, integrated fans and passive cooling devices, the power supply unit, and interfaces. The chassis may feature rails and cages (bays) for mounting different exchangeable devices such as disk drives or power supply units. The form factors, dimensions, and intended way of mounting dictates the design of the server's chassis. Servers can be stand-alone devices, where the chassis includes a pedestal. However, the majority of servers are mounted in 19-inch server racks. Each server is connected via cabling to the power distribution and network. The market shows a high diversity in terms of dimensions, form factors, and system configurations, however, with respect to the chassis, there are two main types:

- Integrated single servers (e.g. rack server). Within these, all components are in the man chassis of the server. The integrated server is fixed directly to the rack, while one rack fits multiple servers usually.
  The actual reck-mounted servers usually feature a display as user interface and openings at the front and back. Many of the key components such as power supply units and the disk drives are designed so that they can be exchanged while the server is in operation. The covers (top) are removable to provide easy access to the key components.
  The internal of the chassis incorporates frames/rails for mounting of the key components. The chassis is mounted on the rack with fasteners which are quick to open.

- Modular server systems (e.g. blade system). These systems consist of two housing elements which are the individual chassis of the server modules (motherboards with connectors for power and I/O) as well as the system enclosure. The individual servers are also call blade, cartridge or book and they are inserted into the larger enclosure. This also provides the shared resources for the blades eg. Network, storage, cooling and power supply. The modular system is also then mounted to a server rack. With regards to fastening mechanisms, these are in most cases base on rails and clips with the aim to ensure fast exchange.

Materials:

The materials used for the chassis vary depending on the dimensions, form and layout, however the main materials utilised are low-alloyed steel or chromium steel, brass and some plastic parts.

There are also several fasteners such as screws and clips which are not standard and vary from product to product.

Although there is a large variety of types and dimensions, there is opportunity for further recyclability of the materials used, due to the possibility of separating the different and main materials from each other as an end-of-life treatment.

The chassis contains ferrous and non-ferrous metals which are recyclable materials. It has been recommended in Section 2.8 of the Phase 1 report that rates of recyclability and recuperation of servers in the EU are investigated further. As within the scope of the Ecodesign regulation to improve the recycling rates of servers, we recommend investigating the inclusion of a more extensive product datasheet requirement which would track material content. This would mean that more information would be provided to recycler about the contents of the chassis, which will facilitate recovery and recycling activities.

Energy usage:

With servers being energy using products, the requirement to lower their energy consumption is a main consideration. In addition, optimising the PUE with managing the internal temperature and coordinate the air flow is a key consideration. Some chassis have baffle plates for better air flow. The chassis function sometimes as a heat spreader. The placement and utilisation of multiple servers in a rack also influences the surrounding thermal conditions.

Within the Phase 1 report the study team have discussed how the energy consumption of a data centre can be managed by maintaining the set temperatures within the server rooms. In Section 2.10 of the Phase 1 report, we recommend not to increase servers operating temperatures above 27°C, as this can result in an increase in total data centre consumption caused by an increase in individual server fan energy consumption. This follows the advice from ASHRAE which shows how for every degree increase in the air inlet temperature from 17.7°C, 4% can be saved on cooling costs. However, beyond 27°C, although cooling costs continue to decrease, and PUE value decreases, the energy consumption of the IT equipment increases, resulting in a total data centre energy increase.

Mainboard (server board)

This is the main printed circuit board (PCB) of the server. It provides the mounting for both the active components such as semiconductors devices (processor, memory, etc.), and the passive components (resistors, capacitors, inductions, etc.)

as well as various sockets (CPU socket, DIMM sockets, etc.) and connectors (Ethernet ports, USB ports, etc.). These together enable the computing and communication services.

The electronic devices are directly connected (soldered) on the mainboard mostly by Surface-Mount Technology (SMT) or Through-Hole Technology (THT). A serverboard is multi- layer printed with approximately 12 layers (plus/minus 4). The number of layers is not seen to be increasing due to the fact that many functional circuitries are nowadays realised as integrated circuits (ICs).

Materials used:

The mainboard is still made of the same materials which are FR-4 laminates and copper foils (cores).

FR-4 is a glass-reinforced epoxy laminate and its thickness is typically a few hundred micrometers (µm).

The copper cores are foils with a thickness of about 18 µm or 35 µm. The single copper layers are connected through small, copper-plated holes.

The mainboard's outer layer is coated with a protective surface finish which varies according to materials utilised. Common technologies remain unchanged since the last study and these are Electroless Nickel Immersion Gold (ENIG) and Immersion Silver (Imm Ag) in higher end products and Immersion Tin (Imm Sn), Organic Solder Preservative (OSP) or lead-free Hot Air Solder Levelling (HASL) in more economic products. Currently Cu with solderability preservatives (CuOSP) is regarded as the established technology.

The dimensions of the printed circuit board and the number of its layers are important parameters for the material efficiency considerations.

Energy usage:

There are energy intensive processes in the manufacturing of printed circuit boards (Galvanic processes) and the lead-free soldering process (high temperatures) which have an increased energy demand.

Processor and memory

Processor: The central processor unit (CPU) in conjunction with the employed memory (RAM) and chipset is generally providing the computing functionality of the enterprise server. There are only a few processor architectures utilised in the server market:

- The x86 instruction set architecture is the most widely distributed and supports Microsoft Windows, Unix, Linux, AIX and Solaris operating systems. It is also compatible with virtualization software and most other software stacks employed in the industry. This processor architecture goes back to 1978.
- ARM processors are extensively utilized in mobile devices like smartphones and tablets due to their RISC architecture, which simplifies instruction cycles, resulting in lower power consumption and heat generation by the CPUs. In the server realm, ARM CPUs can match the computing prowess of top-notch x86 processors. For instance, the Ampere® Altra® CPU, which is a single-socket solution, boasts an impressive configuration of up to 80 cores in a single CPU. This feature makes it suited for handling enormous amounts of data. On the other hand, x86 processors, being based on CISC architecture, excel in executing complex instructions that may span through multiple clock cycles.

Consequently, despite having fewer cores compared to ARM processors, x86 CPUs can still deliver exceptional performance by employing techniques like multithreading.

There are other architectures which include Intel Atom, Intel Xeon Phi and FPGA-type Intel Xeon processors, ARM 64-bit SoC, GPGPUs from AMD and NVIDIA remain current. Upcoming trends of cloud computing being used for gaming purposes and similar applications result in adapted server configurations and technology implementations that relate to different requirements compared to common server configurations.

Figure 4.1 Example of current node development[2]

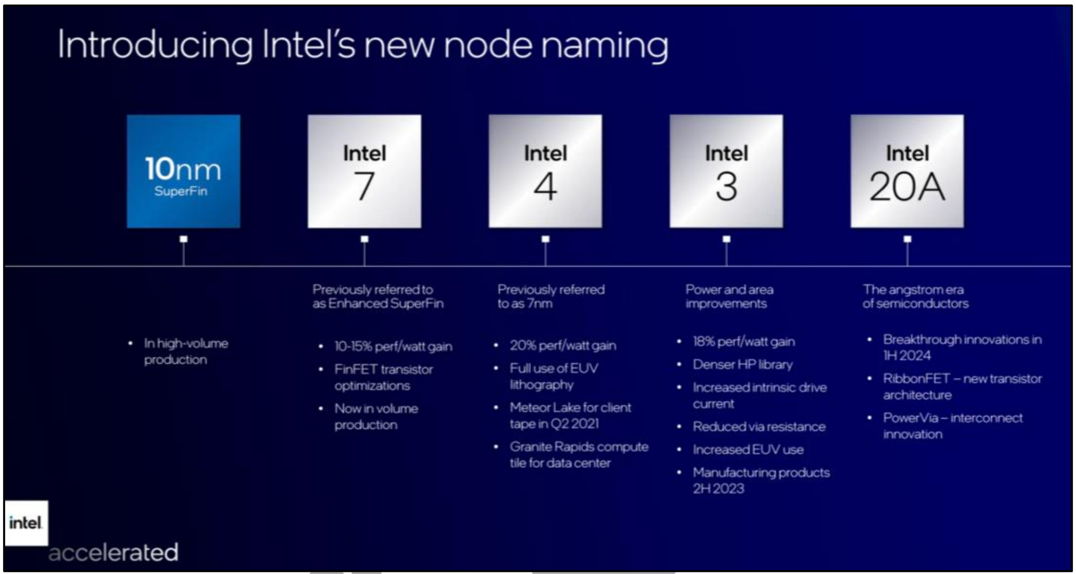

The CPU core serves as the central component where instruction cycles are executed. Within a single CPU, multiple cores can exist, each capable of performing its own set of instruction cycles. To handle larger workloads, parallel computing comes into play, enabling multiple cores to execute tasks simultaneously. While the number of cores significantly impacts the server processor's performance, factors like thread count, instructions per cycle/clock (IPC) and clock speed also come into play. Threads represent the smallest independent sequence of instructions that can be executed within each core. Leveraging multithreading, also known as hyperthreading, a core can effectively increase performance by roughly 30%. Clock speed, often referred to as clock rate, is a key metric for assessing a processor's speed. It denotes the frequency at which the processor generates clock signals, or "pulses," to synchronize operations internally. Clock speed is measured in hertz (Hz), indicating the number of clock cycles per second. Selecting the optimal processor for your server requires considering various factors, including throughput between processors and RAM, cache architecture, workload characteristics, and more.

The technical performance and features of a CPU is characterised by a number of factors including:

- Technology node, die dimensions and resulting number of transistors
- Number of cores per chip, threads per core, functional specialisation of the core
- Cache configuration and capacity, supported memory

[2] https://www.theverge.com/2021/7/26/22594074/intel-acclerated-new-architecture-roadmap-naming-7nm-2025

- Number, type and control of I/Os System-on-chip or further integrated functionality such as power and memory control
- Maximum operating frequency, IPC, and frequency scaling per core,
- Operating voltage, power scaling options, intelligent throttling
- Packaging and mounting on the mainboard including type of socket

The Intel Xeon Scalable processors were based on various technology nodes, including 14nm and 10nm. The 14nm technology node was utilized for earlier generations of Intel Xeon Scalable processors, such as the Skylake and Cascade Lake architectures. However, Intel has also been transitioning to its 10nm technology node with subsequent generations, such as the Ice Lake architecture. The Xeon Scalable processors represented a shift to the 10nm process, offering improved performance and power efficiency compared to the 14nm-based predecessors. The AMD EPYC processors were based on a 7nm technology node. AMD made significant advancements with its Zen architecture, which included the Zen, Zen+, Zen 2, and Zen 3 microarchitectures used in different generations of EPYC processors. These microarchitectures were manufactured using a 7nm process technology. The 7nm process technology offered improvements in power efficiency, performance, and transistor density compared to larger nodes, enabling AMD to deliver high-performance server processors with a greater number of cores and enhanced capabilities. As technology nodes shrink to smaller dimensions, they approach the physical limits of atomic scale and quantum effects. This introduces challenges in maintaining the desired performance, power, and reliability of the transistors and circuits. As feature sizes decrease, the risk of defects increases. Even minor defects can impact the functionality and yield of chips. Maintaining high yield rates becomes more challenging, leading to increased costs and potential delays in production. To address these challenges, the semiconductor industry invests in research and development, exploring new materials, transistor designs, and manufacturing techniques. Techniques like EUV (Extreme Ultraviolet Lithography) have been adopted to overcome some of the challenges in photolithography. Moreover, new architectures and technologies, such as 3D stacking and specialized accelerators, are being explored to enhance performance while mitigating the limitations of miniaturisation.

Figure 4.2 provides a simplified overview of existing variety of CPUs on the example of the current versions of Intel Xeon Scalable processors. The list differentiates the Xeon processor families.

Figure 4.2 Intel Xeon process overview[3]

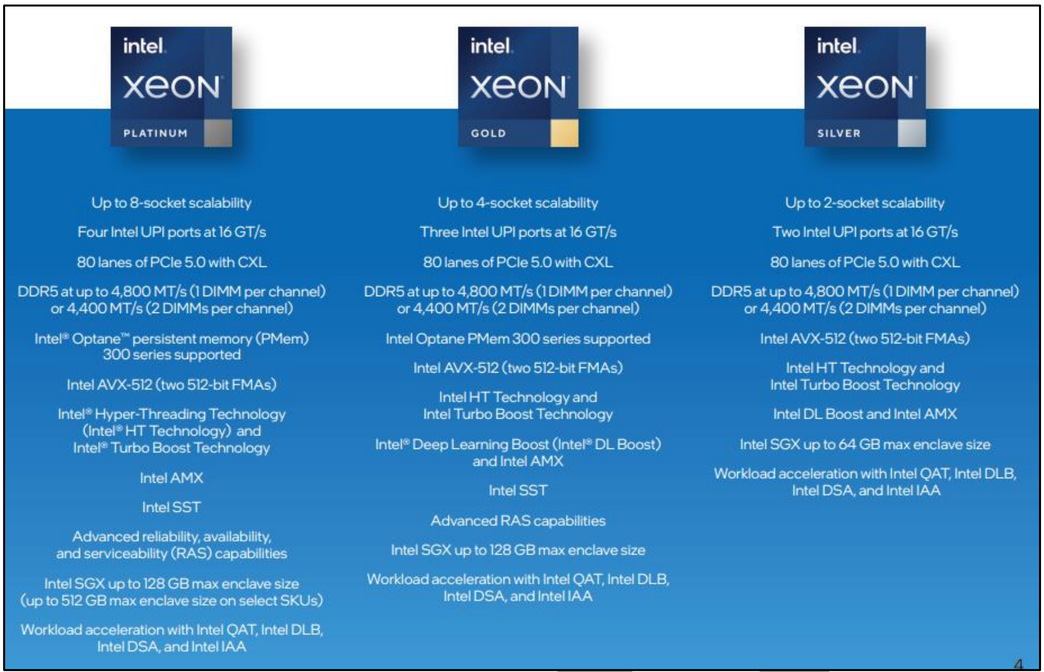

With respect to the performance, the elements remain unchanged and the clock frequency, number of individually addressable cores (and threads), cache capacity, input/output (I/O) speed and capacity are still very important design features of a CPU. In the past, CPU development addressed integrated GPUs. In order to further improve the performance of servers, two or more processors units are connected via a high speed bus or routing interface. Development since the last preparatory study are that there are few if any mainstream server CPUs that contain iGPUs, this space is nowadays better utilised for additional performance of its core function and/or ancillary on board accelerator chips.

Each processor consists of two or more cores allowing typically two threads per core. Multithreading aims to increase utilisation of a single core by using thread-level as well as instruction-level parallelism. The processor also features caches – memory buffers – on various levels (L1, L2, and L3) and other functional segments. Caches store information to avoid multiple computations and improve access times. The cache can be implemented as a hardware or software element. In a CPU, the cache answers many of the requests drastically lowering the effective load on the processor. Due to big and fast caches being rather uneconomic, there are usually several cache elements in a cache hierarchy. The smallest and fastest cache element is the first that attempts to answer a request; if the requested information is not stored, the request is handed to the next bigger and slower cache element. Caches utilise the faster SRAM to fulfil access time demands.

Static RAM (SRAM) retains data bits in its memory while power supply is connected.

Dynamic RAM (DRAM) stores bits by using capacitors and a transistor, that needs to be refreshed periodically. As SRAM does not have to be refreshed, it is much faster than DRAM but uses more parts resulting in a drastically lower memory per chip value. SRAM is also more expensive than DRAM.

[3] https://www.intel.com/content/www/us/en/products/docs/processors/xeon-accelerated/4th-gen-xeon-scalable-processors-product-brief.html

DDR SDRAM (Dual Data Rate Synchronous Dynamic Random Access Memory) is a DRAM package on a Dual In-line Memory Module (DIMM). DDR means that the data is transferred at both the rising and falling edge of the clock signal. SDRAM is different from Asynchronous RAM because it is synchronized to the clock of the processor and hence the bus. Today, virtually all SDRAM is manufactured in compliance with standards established by JEDEC, an electronics industry association that adopts open standards to facilitate the interoperability of electronic components. This makes DDR5 an important spec for any standard server.

DDR5 is the fifth major iteration of this standard. Compared to its predecessors, DDR5 provides higher bandwidth and increased bandwidth efficiency.

The core counts are growing with every new generation of CPU. DDR4 has reached its limit in terms of memory bandwidth and density. It can only support up to 16GB Density and 3200MT/s speed. This is where DDR5 technology offers solutions to meet customer needs for greater memory capacity per core, and bandwidth per core.

DDR5 offers a 50% increase in the bandwidth with 4800MT/s as compared to DDR4 with 3200MT/s. It also supports a maximum of up to 32Gb density (a density that is not available in the latest PowerEdhe generation launch), as compared to 16Gb in the previous generation. DDR5 also offers 2x the burst length, 2x bank groups, 2x banks, Decision Feedback Equalization, two independent 40-bit channels per DIMM, and optimized power management on DIMM.

Figure 4.3 Memory Bandwidth by DDR Generation[4]

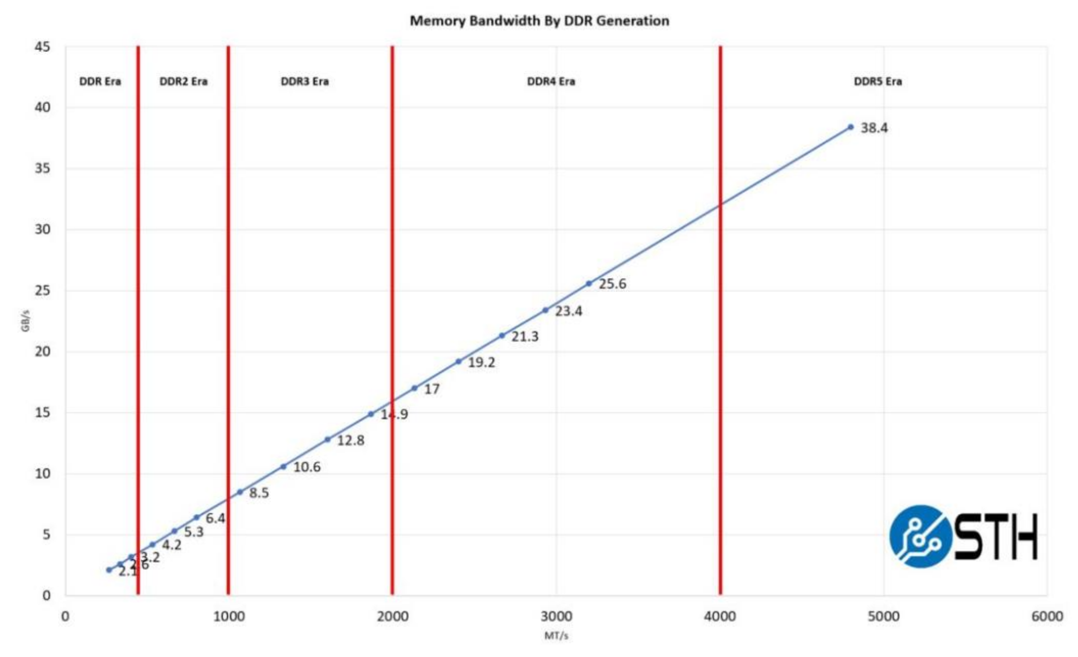

DDR DDR2 DDR3 DDR4 And DDR5 Memory Bandwidth By Generation 1

The following Table 4.1 provides information about the latest Dell PowerEdge portfolio for DDR5, including capacity, bandwidth, DIMM type, and Dell part numbers. Note that Dell does not support DIMM capacity mixing on the latest generation. These represent maximum bandwidth at ideal configurations. CPU vendors may reduce bandwidth capability based on their respective DIMM

4 https://www.servethehome.com/why-ddr5-is-absolutely-necessary-in-modern-servers-micron

population rules. The total system bandwidth is expected to vary between platforms based on population capability, such as on 8 x 1 DPC Intel® CPU- based platforms.

Table 4.1 Details about the attest Dell PowerEdge portfolio for DDR5

Table 1. Details about the latest Dell PowerEdge portfolio for DDR5

DIMM Capacity (GB)	DIMM Speed (MT/s)	DIMM Type	Dell PN*	Ranks per DIMM	Data Width	Density	Technology
16	4800	RDIMM	1V1N1	1	x8	16Gb	SDP
32	4800	RDIMM	W08W9	2	x8	16Gb	SDP
64	4800	RDIMM	J52K5	2	x4	16Gb	SDP
128	4800	RDIMM	MMWR9	4	x4	16Gb	3DS
256	4800	RDIMM	PCFCR	8	x4	16Gb	3DS

* Part numbers are subject to change. Additional part numbers may be required.

There are now two memory channels on the chip. Also, Error Correction Code (ECC) DDR5 unbuffered DIMM, is no longer compatible with platforms that support RDIMMs.

One of the biggest challenges with system scaling is the memory bandwidth. Having more cores and more PCIe devices is advantageous, but if portions of the system are sitting idle waiting for data, then they are being wasted. In the DDR3 and DDR4 generations, generally, one could use either unbuffered ECC (and often non-ECC) memory or RDIMMs in servers. That is no longer the case. The key reason for this is that the power delivery components are now on the DIMMs themselves. In servers, 12V power is supplied. In client systems, only 5V. This is converted to 1.1V for the DIMM and managed by the onboard power management IC, or PMIC. The PMIC moves a motherboard function to the DIMMs, but that means we have an added component on the modules. With DDR5, there are now two channels with a total of 80 bits. Each channel is half of that or 40 bits with 32 bits for data and 8 bits for ECC. The channels are split on the DDR5 DIMM with one channel on the left and one on the right. On one side, you will find the RCD or register clock driver. This is responsible for providing the clock distribution to the different chips on the memory module. DDR5 has on-chip ECC. As DRAM density increases and clock speeds increase as well, the folks designing the DDR5 spec realized that the combination has the potential to lead to more on-die errors. For those more familiar with flash technology, NAND has been going through this for years, and now we have higher bit density but also more error correction throughout the SSD. Still, the inherent on-chip ECC benefit is that issues that may not have been caught previously on consumer UDIMMs can potentially be caught.

Figure 4.4 Micron DDR5 RDIMM Two 40 Bit Channels[5]

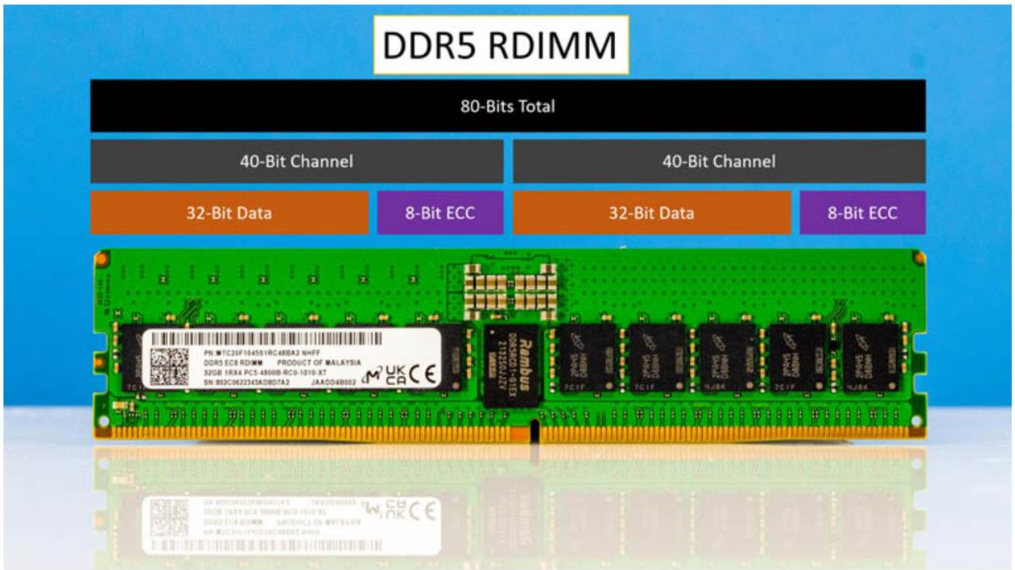

Micron DDR5 RDIMM Two 40 Bit Channels

Energy and Material use during the manufacturing phase: The CPU and memory are made of three major materials, which are silicon, copper and plastics. The steps of transforming these raw materials into the micro components and integrated circuits (IC) include a number of intensive processes. These require certain temperature and sanitation conditions, water, chemicals operated by complex machinery. Considerations for energy and resource efficiency during the manufacturing process is an important factor.

Figure 4.5 Silicon-to-chip process[6]

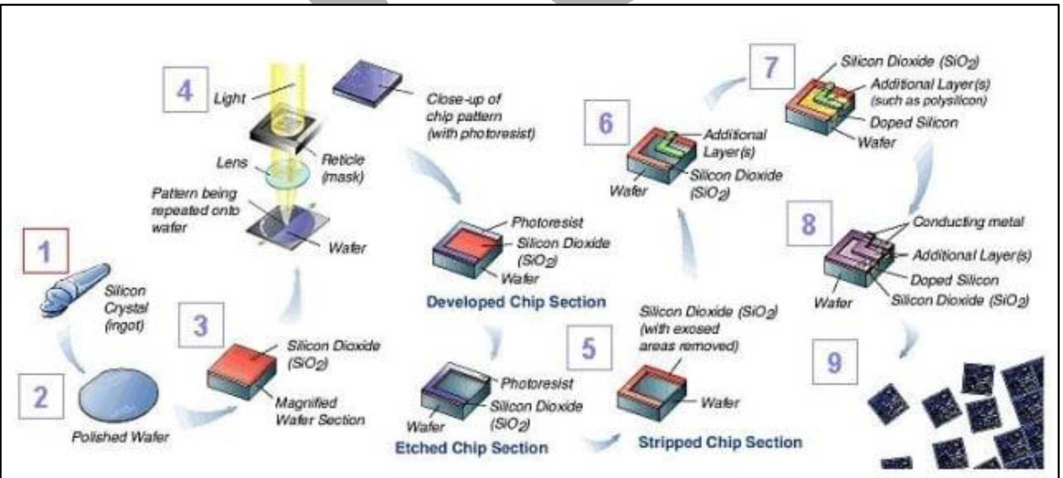

Energy demand during product use:

As discussed within Phase 1 section 2.4, Dynamic Voltage and Frequency Scaling (DVFS) is a technique used to reduce energy consumption in electronic devices by adjusting the operating voltage and frequency. DVFS allows for the dynamic

5 https://americas.kioxia.com/content/dam/kioxia/en-us/business/ssd/data-center-ssd/asset/KIOXIA_EDSFF_E3_Intro_White_Paper.pdf

6 What Raw Materials Are Used to Make Hardware in Computing Devices? | Engineering.com

adjustment of the supply voltage and clock frequency of a device based on its workload or performance requirements. By lowering the voltage and frequency during periods of low activity or idle states, energy consumption can be significantly reduced. This approach helps to optimise the power-performance trade-off and minimise unnecessary energy usage. By reducing voltage and frequency, the energy consumption and associated carbon footprint of electronic devices can be lowered.

The feedback from stakeholders suggested that servers should have processor power management enabled by default and at the time of shipment, with all processors capable of reducing energy consumption during low utilisation. Choosing one option (reducing voltage or frequency) was recommended instead of both, in order to avoid diminishing server resilience and increasing implementation complexity and cost. While P-states provide energy savings, it should be acknowledged that they introduce latency that may not be acceptable for certain customers.

Cooling system (active and passive)

The cooling system of enterprise servers is a combination of passive and active cooling elements. The aim of the system is to transport the heat generated by the CPU and other active components away from the devices in order to enable appropriate and reliable function. The cooling capacity is the main metric (W/cm²) and is the rate at which heat is removed from the space producing it for example the CPU. The technology, design, and material characteristic of the cooling system is determined by various factors such as:

- The number, type, and thermal design power of the CPU, memory and other electronic components.
- The type, form factor, actual dimensions, and modularity of the server.
- The external (ambient) temperature conditions, air intake design, dust filter mats, and other aspects of the rack design.

There are two basic types of cooling system technology (air cooling system, liquid and their combination).

Passive cooling elements are considered to be heat spreaders, sinks and pipes with the aim of radiating away the heat from the active components. Other design features are also part of the passive system and these can include round shaped air intake which is considered more effective, air baffles for channelling the airflow and air flaps/valves at the rear of the chassis that prevent the air leakage when modules are pulled out of the system.

Active cooling elements are considered an array of fans which take the cool air from the cold aisle in the data centre through air intake holes. The two basic types of fans are the centrifugal fan or blower, and the axial flow fan. Adjusting the fan rotation speed according to the current temperature conditions is a standard technology.

Another method of cooling is the liquid cooling technology which utilises water, enclosed in cooling loop channels coupled within heat spreaders. It is often used alongside air cooling systems. In contrast to passive cooling, the liquid cooling system includes active components too. The heat is transported through the pipes to a heat exchanger in the server chassis (or outside). Still, this is not a commonly used technology due to its cost.

Material use:

Aluminium and/or copper are typically the main materials of the passive cooling elements (heat sinks and heat pipes). The material selection and the design of the cooling elements are influenced by the specific thermal conductivity (W/m*K) and the related form factor and costs.

With respect to material efficiency, the selection of the passive and active cooling elements including its material composition should be carefully considered. There is a preference to have to have mono-material design, which helps the end-of-life and separation of materials for the recycling process.

Power supply unit

The power supply unit (PSU) provides the electric current with which the server operates. It typically receives alternating current (AC) and converts this into direct current (DC) and specific voltages for the server. There are also PSUs on the market that can receive DC directly and downrate the input voltage for use within the server, though this application remains niche at this time.

There is a variety of PSUs on power capacity, conversion efficiency rating, input and output power as well as redundancy, hot-swap capability and failure monitoring options.

Still since the previous iteration of the study, the typical type of power supplies used are switched-mode power supplies (SMPS) with power factor correction (PFC).

Energy use:
All types of servers have multiple configuration options with regards to power supply in order to provide a wide customer selection. The capacity of power units ranges from 250W to 1500W and can be achieved with either one large or two smaller units in the same system. In any case an essential power supply unit is installed for cases of hardware failure of the main one(s).

The PSU's rated output power and conversion efficiency in partial loads is influencing the baseline energy consumption of the server. The peak efficiency is typically somewhere between 50 and 100% loading.

Figure 4.6 shows how PSU performs under different loads and supply voltages.

Figure 4.6 An efficiency chart from Cooler Master for their V1300 Platinum PSU

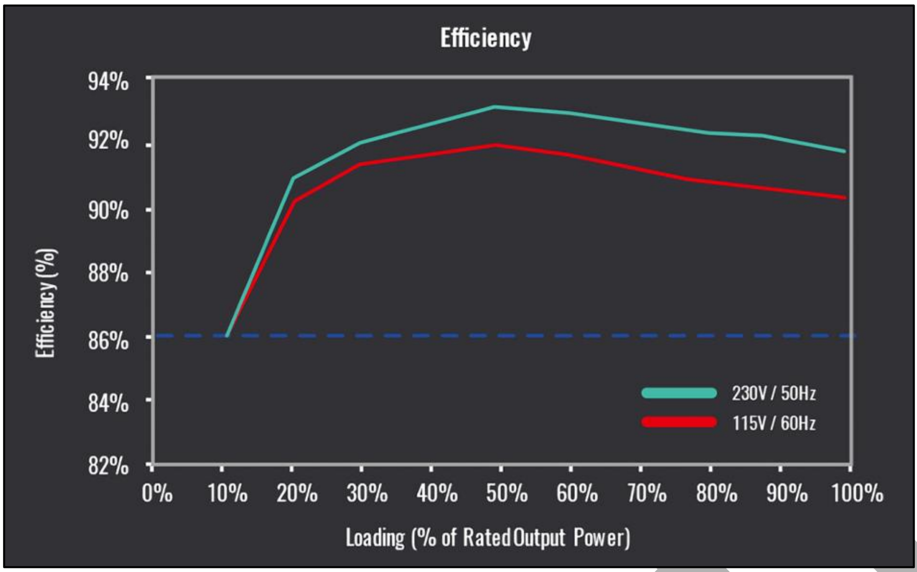

The current Regulation under review introduces minimum PSU efficiency and power factor requirements from 1 January 2023 as follows:

Figure 4.7 2019/424 introduced minimum PSU efficiency and power factor requirements from 1 January 2023

	Minimum PSU efficiency				Minimum power factor
% of rated load	10 %	20 %	50 %	100 %	50 %
Multi output	—	90 %	94 %	91 %	0,95
Single output	90 %	94 %	96 %	91 %	0,95

There are currently discussions on how to best enable that efficiency based on 80 PLUS certification scheme[7] as well as possible solutions that bring the best outcome[8] given the availability in the market of PSUs meeting the current Regulation's requirements.

Figure 4.8 demonstrates a typical layout of a PSU. There are many layouts, with different make and model of the parts inside, still serving the same functionality.

[7] 80 Plus Overview | CLEAResult

[8] https://www.clearesult.com/80plus/sites/80plus/files/news/230V%20ENTERPRISE%20PSU%20MARKET-READY%20EFFICIENCY%20STUDY_1.pdf

Figure 4.8 Parts comprising a PSU[9]

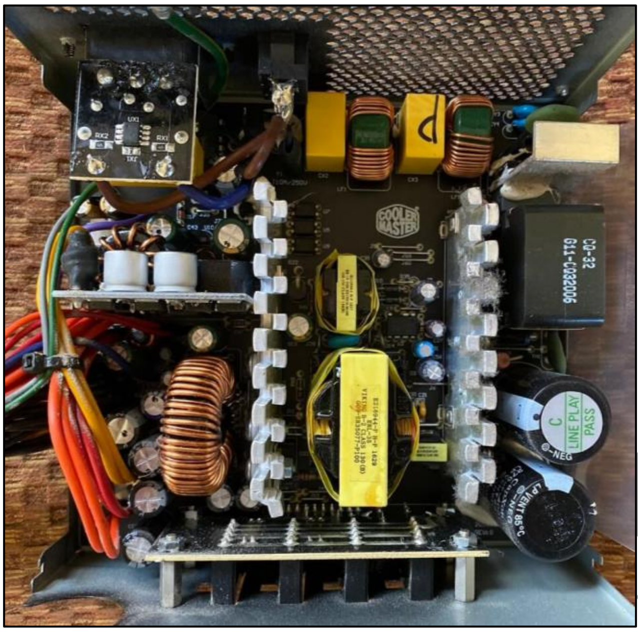

The main materials used are in the robust metal housing (steel alloy), the transformer and inductors (e.g. copper and ferrites), the passive cooling elements (aluminium), the fan (plastic) and the printed circuit boards (glass, epoxy, plastics, solder).

Prices: The sales price of a PSU is determined by certain factors such as:

- The number of output currents,
- connector configuration (pins),
- conversion efficiency (80plus certification),
- redundancy,
- hot plugging capability, and form factor.

The quality and performance are influenced by the component selection and thermal design of the PSU as well. Table 4.2 shows current price ranges for single PSUs according to various internet sources.

Table 4.2 Current price ranges for PSUs[10]

Single PSU	<400W	400-800W	850-1500W
80 Plus Gold	€ 99-115	€ 94- 200	€ 141-305
80 Plus Platinum	N/A	€ 153-282	€ 211-470
80 Plus Titanium	N/A	€ 294	€ 387-434
80 Plus Gold	€ 99-115	€ 94-200	€ 141-305

Please note prices ranges for this table were originally displayed in Pound Sterling. This was converted to Euros using Conversion factor extracted from https://www.exchangerates.org.uk/ on 23/05/2024. (1 GBP = 1.1742 EUR)

DC operated servers:

There is a perception that as the power supply unit has loses associated with the conversion from AC to DC, having DC operated servers which are provided with DC

[9] Anatomy of a Power Supply Unit (PSU) | TechSpot

[10] 80 PLUS TITANIUM Certified,1500W and above Power Supplies | Newegg.com

power, can save on the PSU material cost and energy loses. The DC servers still have PSUs, they are just DC-DC and step down the voltage, associated with some cost to that conversion into DC that can be used by the server board as well.

It is noted however that although savings can be made at the PSU level, efficiency loses are typically passed on to another location in the data centre. Indeed, as utility power is primarily in AC, the AC-DC conversion will need to happen for the ICT equipment, which may mean a shift of the problem at the PSU server level to a centralised data centre conversion unit.

SPEC SERT currently does not support testing of DC servers.

Configuration options

The market offers a humongous variety of server products and configuration options. These translate into the technical and environmental performance of the server as well as in respective purchasing costs (CAPEX) and cost of utilisation (OPEX).

The various options for configurations are still relevant since the past study and outlined here:

- Type, dimensions, weight, and mounting options of the server chassis
- Mainboard type and dimensions
- CPU options including type, cores, threads, I/O, frequency, type and number of sockets, etc.
- Memory options including type, capacity, interfaces, number and type of memory sockets, etc.
- Internal storage (drives) including type, capacity, interfaces, number and storage bay space, etc.
- Connectivity (input/output) including network controller, connectors, expansion slots, etc.
- KVM (keyboard, video, mouse) interfaces, monitoring and service options, etc.
- Power requirements including rated power, power supply unit configuration, redundancy, etc.
- Power consumption values including maximum, sometimes idle or thermal design power of CPU (power-to-performance benchmarks e.g. SPEC are typically only offered on request)
- Active cooling elements such as fans and fan control
- Operational conditions and allowances including min and max temperatures, humidity in room
- Operation system and other software specifications including licenses

With respect to the current server market there are two rack-server main systems and blade systems as outlined in the section covering Chassis above.

4.1.1.3 Data storage products

Regulation 2019/424 provides the definitions for data storge products as a fully functional storage system that supplies data storage services to clients and devices either directly or through a network. The components and subsystems that are an integral part of the data storage product architecture are considered to be part of the data storage product. Whereas components that are not associated with a storage environment at the data centre level are not considered to be part of the data storage product. A data storage product may be composed of integrated storage controllers, data storage devices, embedded network elements, software and other device.

Data storage products provide non-volatile data services to connected servers and/or to remote computing devices via a network connection. A data storage product provides the server with more capacity, redundancy and flexible data management. With a data storage product there are large number of storage media and a controller than handles the input and output requests. Data storage products also need the capability to connect to networks for data transmission. They are connected to servers or clients by means of direct connection or network connections. These connections are achieved via the use of the following four form factors: direct attached storage (DAS), storage area networks (SAN), network attached storage (NAS) and content addressed storage (CAS).

Data storage systems have the following sub-systems:

- **Storage media and devices:** hard drive disc (HDDs), solid state devices (SSDs)
- **Storage controller:** Processors and other electronics that automatically process the input and output (IO) requests directed to the storage device.
- **Storage elements:** configuration of the data storage product. Such as a redundant array of independent disks (RAID) or robotic tape library with a number of storage devices and integrate storage controller for handling I/O requests.
- **Connectivity and network elements:** connectivity is based on various technologies such as: Serial Advanced Technology Advancement (SATA), Serial Attached SCSI (SAS), Fibre Channel (FC), Infiniband (IB) or Ethernet (TCP/IP)
- **Connectors and cables:** transmission between servers and data storage products requires interface controllers, connectors and cables.

Storage system form factor and configuration

Similar to servers, there is a very large number of possible configurations for the customer to choose, which are dependent entirely on the specific use case and workload profiles required. Therefore, the efficiency and product performance is distinctive to each product configuration that is chosen. There are several factors which impact performance and efficiency these include the application of software-based functionality including compression, de-duplication, software defined storage management systems, and workload tiering. Utilising these software functions enables data storage products to significantly reduce the number of storage systems required to manage specific workloads. At product level it may increase the energy consumption, resulting in in a lower efficiency as measured by SNIA Emerald. However, it will reduce the overall energy consumption required to perform a workload as a less storage devices and systems are required to service a given workload.

Storage controllers are either installed internally or externally as an extra controller enclosure (CE) with attached disk enclosures (DE) which can provide extra scalability. These systems are capable of organising hundreds to thousands of attached HDDs and SDDs.

The configuration that determines the type and number of attached drives can differ depending on the requirements of the customers. Typically, systems are sold with various basic configurations which differentiating the number of controllers and drive size (2.5 and 3.5 inch) while modern systems usually support all of the common interfaces like SAS, iSCSI, FCoE or FC. Common supported RAID levels are 0, 1, 3, 5 and 6. An entry-level example system is the FUJITSU Eternus AF150, the technical specification for this product are displayed below in Table 4.3.

Table 4.3 Fujitsu AF150 technical specifications[11]

Technical specifications	
Type	All-Flash Array
Latency	140µsec (Read), 60µsec (Write)
Sequential access performance	10,000 MB/s (128KB Read) 5,000 MB/s (128KB Write)
Random access performance	390,000 IOPS (8KB Read) 270,000 IOPS (8KB Write)
Host interfaces	Fibre Channel (16 Gbit/s, 32Gbit/s) iSCSI (10Gbit/s [10GBASE-SR, 10GBASE-CR, 10GBASE-T])
Maximum disk drives	24
Support Pack Options	Available in major metropolitan areas: 9x5, Next Business Day Onsite Response Time 9x5, 4h Onsite Response Time (depending on country) 24x7, 4h Onsite Response Time (depending on country)
Maximum Storage Capacity	92 TB
No. of host interfaces	4/8 ports
Max. no. of hosts	1,024
Supported RAID levels	0, 1, 1+0, 5, 5+0, 6
Drive interface	Serial Attached SCSI (12 Gbit/s)
Automated Storage Tiering	Yes
Remote Copy functionality	Synchronous and asynchronous
Recommended Service	24x7, Onsite Response Time: 4h

Modern data storage system designs are becoming increasingly case-specific and require products that can adapt to optimisation in performance or capacity, or even a hybrid of both. Advancing technologies such as IoT and IIoT (e.g., edge computing, machine to machine communication), along with faster communication protocols and increasingly robust interfaces (5G, WiFi, and GPON) has meant that storage media based on 2.5 inch drive formats are becoming increasingly limited[12]. Data storage products with these hard drive formats are not able keep pace with the increased demands of modern servers such as servers based on PCIe® 5.0 and 6.0 technologies. This is because the 2.5 inch drive is not optimal for flash memory packaging or optimised for flash memory channels[13]. Since performance is expected to continue to increase, so too will the flash memory and all of the dies, therefore, the energy consumption of the flash memory and interface will increase alongside this technology advancement. Therefore, in the near future a form factor will be

[11] https://www.fujitsu.com/global/products/computing/storage/all-flash-arrays/eternus-af150-s3/index.html#specs

[12] Introducing the EDSFF E3 Family of Form Factors (kioxia.com)

[13] Introducing the EDSFF E3 Family of Form Factors (kioxia.com)

required which can scale power, PCIe speed increases, and wider PCIe link widths to enable full-throttle input/output operations per second (IOPS) performance per terabyte (TB) capacity[14].

Hard disk drives (HDD)

Hard disk drives are a non-volatile data storage device. HDDs are a mature technology that is the most commonly uses volume storage media for high capacity block storage devices. HHDs are more cost efficient than SSDs. Within the HDD the data is stored on fast rotating disks by magnetic remanence. The most common form factors for HDDs in servers are the following:

- Large form factor - 3.5-Inch HDD for very high capacity;
- Small form factor - 2.5-Inch HDDs for faster access but smaller capacity

The HDD form factor refers to the physical size or geometry of the data storage device[15]. These follow industry standards that govern their length, width and height as well as their position and orientation of the host interface connector.

The performance requirements for HDDs are storage capacity, average latency, type of connectivity or interface, the maximum sustained transfer rate, spin speed, aerial density, cache, encryption, format, maximum operating shock, form factors, average failure rate, operating temperature, and average operating power[16]. The operating temperature for HDD is typically 5°C to 40°C. however, current product data sheets indicate even higher operating temperatures of up to 60°C.

HDDs are often differentiated for various applications using the mean time between failures ratings. These failures often occur due to vibration and thermal cycling which causes the spindle's swivel to gall. Workload is determined by the amount of time the HDD is in active mode. Input output per second (IOPS) is one of the main performance indicators for HDDs. It represents the average rotational latency and read/write seek latency. It can be defined by the following equation:

- IOPS = 1000/ (seek latency + rotational latency).

As demonstrated in Figure 4.9 basic HDDs contain several disk platters, which are made from either aluminium, glass, or ceramic. These platters spin with a motor which is connected to the spindle. The read and write heads magnetically record the information to and from tracks on the platters using the magnetic head.

[14] Introducing the EDSFF E3 Family of Form Factors (kioxia.com)

[15] https://www.techtarget.com/searchstorage/definition/hard-disk-drive

[16] https://www.seagate.com/gb/en/products/enterprise-drives/

Figure 4.9 Internal components of a HDD[17]

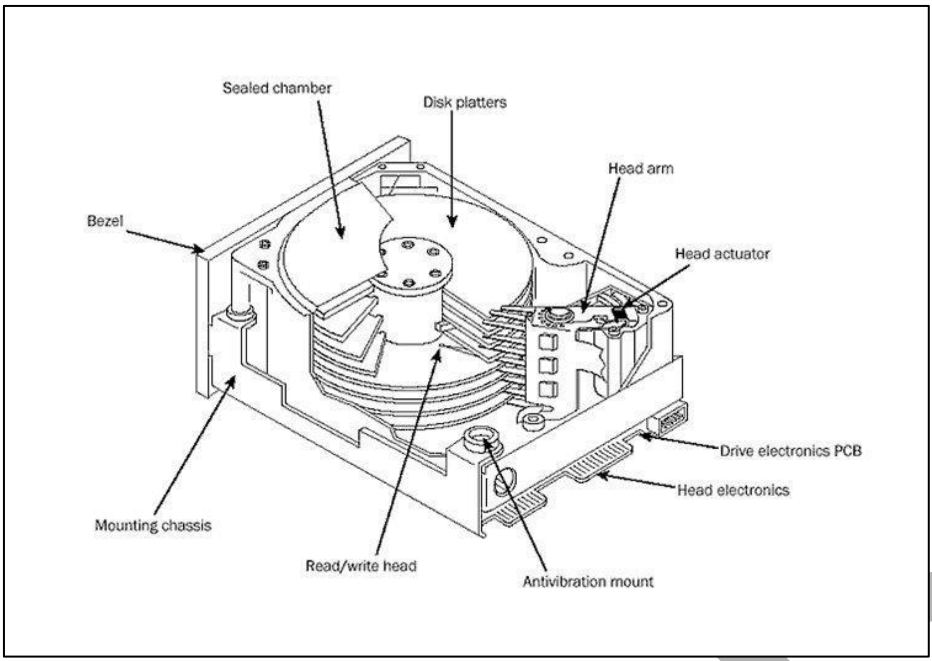

The general hardware elements that are shown in 0 are shown below:

- Platters – these are the storage media
- Electromechanical read and write systems
- Network and power connectors
- Aluminium cast housing (90% of the total weight)

Solid state drives (SSD)

Solid state drives (SSD) are another type of storage device used by servers. They are non-volatile storage medias storing persistent data on solid-state NAND flash memory chips. They are packaged as single-chip or multi-chip modules on a printed circuit board. Since SSDs are not defined as drives, they do not have to meet specific industry requirements which HDDs do. SSD are controlled by a firmware, an I/O interface controller and respective connector (bus). SSDs are offered in different form factors including PCIe cards and HDD standard housing.

An SSD reads and writes the data it collects on underlying NAND flash memory chips made from silicon. These chips can be stacked in a grid in order to achieve different densities. The flash memory chips use floating gate transistors, allowing it to hold an electric charge. Even after the device is disconnected from the power supply. Each FGT has a single bit of data, these are defined by a 1 for a charged cell or a 0 if the cell has no electrical charge. Every block of data is accessible at a constant speed. However, it should be noted SSDs are fastest when they write onto empty blocks, writing onto blocks with data already on will reduce the performance.

There are several types of SSDs, these include the following[18]:

[17] https://www.techtarget.com/searchstorage/definition/hard-disk-drive

[18] https://www.techtarget.com/searchstorage/definition/SSD-solid-state-drive#:~:text=An%20SSD%2C%20or%20solid%2Dstate,functions%20as%20a%20hard%20drive

- Single-level cells (SLC) – hold one bit of data per cell. They are the most expensive but are the fastest and most durable.
- Multi-level cells (MLC) – hold two bits of data per cell and have a larger amount of storage space for the same amount of physical space as a SLC.
- Triple-level cells (TLC) – hold three bits of data per cell. They are the cheapest and deliver more flash capacity. However, they are write at slower speeds and are less durable than SLC and MLC.
- Enterprise multi-level cell (eMLC) – this is a MLC flash with a controller and software that helps remedy the faults for MLC.

SSDs have an I/O interface which contains the I/O controller and I/O connector. The I/O controller is a microprocessor that creates the link between the storage devices (SSD) and the server (host). There are many interfaces available on the market, the most common are SSDs that support Serial ATA (SATA), Serial Attached SCSI (SAS) and non-volatile memory express (NVMe). Figure 4.10 below illustrates where in the SSD these interfaces are found and demonstrates simply the internal structure of an SDD.

Figure 4.10 Internal components of an SSD[19].

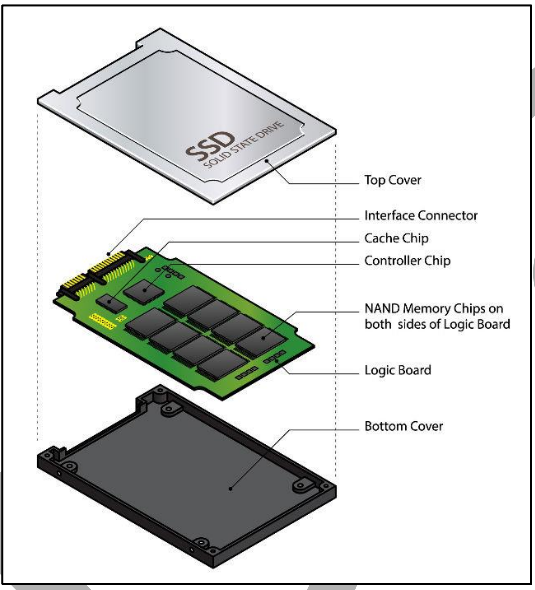

SSD are often used to replace small form factor 2.5 inch HDDs because they are significantly faster. The main difference is that HDDs contain a spinning disk with read/write head on a magnetic arm, HDDs read the data magnetically. However, this can lead to mechanical breakdowns. SSD on the other hand have no moving parts, therefore, are less susceptible to breakdowns. Although SSDs have no moving parts they do have a wear-out mechanism which can be caused by too many writes on the SSD, this not an issue for HDDs. Nevertheless, having no moving parts is the reason why the access time is much faster for SSDs. Additionally, this means that the power consumption is also lower, and they can operate under slightly higher temperatures. The operating temperature of SSD is between 0°C to 70°C and therefore considerably wider in comparison to HDDs. Although standard SSDs may have a lower power consumption than comparable 10k RPM HDDs, a NVMe (non-

[19] https://uwaterloo.ca/arts-computing-newsletter/winter-2018/feature/everything-you-need-know-about-solid-state-drives-ssd

volatile memory host controller) SSDs will consume higher amounts of power than the equivalent HDD.

However, SSDs are significantly more expensive than HDDs, HDDs also have much larger capacities than SSDs. With SSDs technology continuing to advance it is expected that they will soon be similarly priced. This has even started to happen in some instances. For example, SSDs are replacing the 10k RPM disk drives due to their better performance and cost. SSDs are cost competitive with HDDs for high performance systems[20]. Nonetheless, for some products the price of a HDD remains the more competitive. For example, a high capacity 7.2K RPM HDDs cost and power efficiency are competitive with the comparable SSDs and continue to be used in applications that do not require the higher performance of SSDs[21].

Direct attached storage (DAS)

DAS is a low cost, short distance storage solution which includes a fixed wired, direct connection between the server and the data storage product. DAS can be connected to one, two or four servers depending on the users needs. There are also options to create a DAS-system that is based upon Fibre Channel which can allow for the servers and data storage product to be separated by a larger distance. However, these are not common and are often more expensive than other wireless options.

Storage area network (SAN)

Storage area network (SAN) is a network (fabric) attached storage product, which provides fast, secure and lossless communication between the server and data storage product. SAN have a significant advantage over DAS storage units because they can connect to much larger number of servers via the fabric (network). This storage architecture will support the mutual administration of all storage elements, which improves the utilisation storage capacity. Which can then be accessed by the servers like a local storage device through the servers operating system. **Error! Reference source not found.** illustrates the SAN's architecture, which consist of the Host Bus Adapter (HBA) situated in the case of Fibre Channel (FC), a network element such as switch or router, and arrays of storage devices including the RAID systems. In addition, **Error! Reference source not found.** demonstrates the basic topology and elements of a SAN product.

Typically, SANs are used for block I/O services rather than file access services. However, they are very sensitive to packet drops. SAN utilises high latency, lossless protocols and network interfaces such as fibre channels and InfiniBands. Development of SANs have included Fibre Channel over Ethernet (FCoE) and internet Small Computer System Interface (iSCSI) which both allow communication through TCP/IP.

Network attached storage (NAS)

Network attached storage is a more economical solution to increase the storage capacity and performance of a server than SAN. NAS are specialised systems that are attached via a local area network (LAN), which makes it far more cost efficient than SANs. As a result, NAS is very popular among smaller business networks,

[20] Stakeholder feedback

[21] Stakeholder feedback

whereas SAN is usually preferred by larger businesses. This is because the user can install more storage capacity to a NAS system simply by installing more devices, even though each NAS system will operate independently. Contrastingly, SANs are utilised to handle high speed file transfers or many terabytes of centralised file storage. As SANs have a high-performance disk array but they do require their own hardware and interfaces.

The NAS architecture is displayed in 0 and consists of an engine that implements the file services to one or more devices, on which the data is stored. Typically, a NAS product will contain its own operating system. NAS use distinctive remote file access protocols such as network file service (NFS) in Linux/Unix environments or Common Internet File Services (CIFS) in Windows environments. The NAS data storage product can be tested by the SNIA Emerald test standard, as it fits within the Near Online category.

Content addressed storage (CAS)

Content addressed storage (CAS) stores data by assigning a content address to each object, this means that the data can be retrieved based on its contents, rather than its name or location. This high-speed storage device ensures that the retrieval of data is extremely fast. This is because CAS minimises the storage space consumed by data backups and archives by assigning a retention period to each object and avoiding duplicates. Since, there is only one copy of each object, data retrieval is much faster than SAN, DAS or NAS[22].

CAS systems function by passing data within a file through a cryptographic hash function, which generates a unique key, known as the content address. The CAS will not store duplicate files because the duplicate file will create the same key, which will not be stored as it is already in the device. Today CAS systems continue to be utilised across numerous emerging technologies such as cryptocurrencies[23]. CAS have suited these technologies because every time an item is accessed it must be verified with a hash value to prove that it is the same item that was stored[23]. Feedback has suggested that their current market share is expected to be maintained moving forward[24].

4.1.1.4 Firmware

Firmware is defined within the current Regulation 2019/424 as system, hardware, component, or peripheral programming provided with the product to provide basic instructions for hardware to function inclusive of all applicable programming and hardware updates.

Firmware is installed permanently into a device's read only memory and dictates how the device (hardware) should operate. It enables the servers/storage systems to function.

Throughout the life of a component, there is a need for firmware updates which would enable better functionality, cybersecurity and improve compatibility. It is however a potentially risky maintenance operation which needs to be carried out carefully by the Administrators of those products. This is because often the new

[22] What is content-addressed storage (CAS) and how does it work? (techtarget.com)

[23] Design-And-Implementation-Of-Msha256-On-Blockchain-Using-Content-Addressable-Storage-Patterns.pdf (ijstr.org)

[24] Stakeholder feedback.

versions have increased performance requirements, which make the hardware unable to keep up and run. This can lead to limitations in product life expectancy.

The current Regulation already ensures that products have access to latest firmware updates for eight years after the last product was placed on the market. However, the availability of previous versions could be helpful for refurbishers to extend life expectancy. The impacts of making these previous versions available shall be investigated, without lessening cybersecurity norms.

There is also a barrier on releasing firmware as open-source data, which is likely to be a breach of manufacturer Intellectual Property and is therefore difficult to implement.

4.1.1.5 Network Protocols

Small computer system interfaces (SCSI)

SCSI is the industry standard that specifies the physical connection and data transport between the host (servers) and storage devices. It is a buffered peer-to-peer interface, and a parallel bus-system that comes with a 68-pin (single) with large connectors which outlines the limitations of this SCSI. Of particular importance is the fact that SCSI works independently from the hosts operating systems.

Table 4.44 provides the technical specifications of some of the latest generations of SCSI.

Table 4.4 SCSI specifications[25]

Technology Name	Maximum cable length (metres)	Maximum speed (MBps)	Maximum number of devices
SCSI-1	6 Single Ended	5	8
Fast SCSI	3 Single Ended	10	8
Fast Wide SCSI	3 Single Ended	20	16
Ultra SCSI	1.5 Single Ended	20	16
Ultra SCSI	3 Single Ended	20	4
Wide Ultra SCSI	-	40	16
Wide Ultra SCSI	1.5 Single Ended	40	8
Wide Ultra SCSI	3 Single Ended	40	4
Ultra2 SCSI	12 LVD/ 25 HVD	40	8
Wide Ultra2 SCSI	12 LVD/ 25 HVD	80	16
Ultra3 SCSI (Ultra160 SCSI)	12 LVD	160	16
Ultra320 SCSI	12 LVD	320	16
Ultra640 SCSI	10 LVD	640	16

However, it's worth noting that SCSI technology, including Ultra 640 SCSI, has largely been replaced by newer interface standards such as Serial Attached SCSI (SAS) and Serial ATA (SATA) for most mainstream applications. These newer interfaces offer comparable or higher speeds while being more cost-effective and easier to implement.

[25] What is SCSI (Small Computer System Interface)? | Definition from TechTarget

Serial attached SCSI (SAS)

SAS are serial protocol that create a point-to-point connection instead of the bus-topology that is used by SCSI. It specifies all aspects from the physical layer up to the link, transport and application layer. With each generation of SAS the data throughput improves. SAS-1, the first version of the SAS standard released in 2003, supported data bandwidth of 3 gigabits per second (Gbps). SAS-2 supports 6 Gbps, SAS-3 supports 12 Gbps and the current version, SAS-4, supports 24 Gbps. SAS-5 is under development and will support bandwidth greater than 24 Gbps. Within an SAS device there is one transceiver with multiple transmitter and receiver for duplexing. For example, a HDD consists of two transceivers and a host bus adapter 4 or 8. Please see Figure 4.11 for an illustration of this multi-node cluster application.

Figure 4.11 SAS multi-node cluster applications[26]

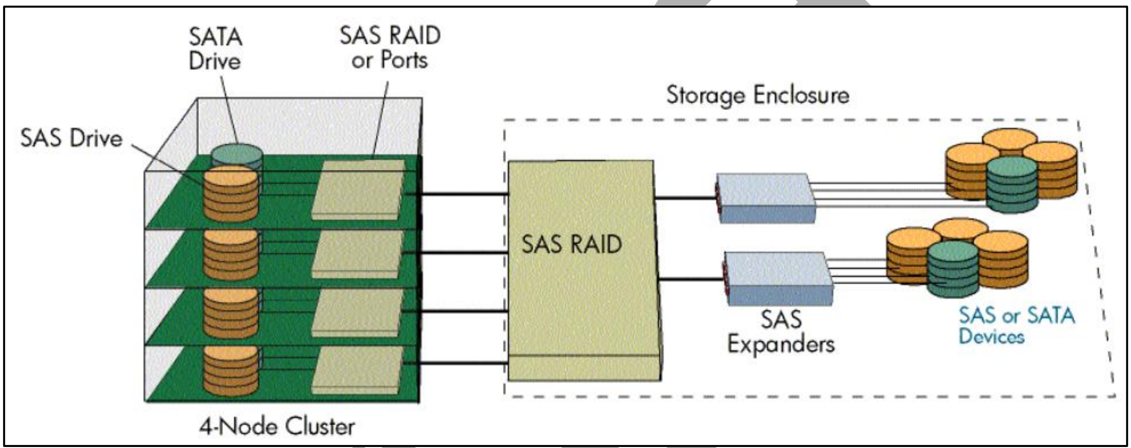

The standard also specifies the various transport protocols for command-level communication with SATA devices (STP), SCSI devices (SSP), and for managing the SAS network (SMP). SAS expanders are an additionally element that links many SAS devices to a switch. Expanders contain eight or more external expander ports. SAS supports both point-to-point and daisy-chain configurations, providing flexibility in connecting multiple devices to a single SAS controller. It allows for the connection of up to 65,536 devices in a SAS domain, providing extensive scalability.

There are various small form factor (SFF) connector types available for SAS devices. One should distinguish between the internal and external connectors of a SFF. External connectors facilitate the interpolation with multiple SAS devices (SSF 8484) and with other standards including InfiniBand (SFF 8470) and SATA (SFF 8482). A small form factor allows the use of smaller storage devices such as the 2.5 inch HDD but also increase the possible number of connectors on the server side.

Serial Advanced Technology Attachment (SATA)

SATA is most recent version of the parallel ATA interface (bus) that features a bit-serial, point-to-point connection for storage devices with the server. SATA is available only via short distance connection. It provides a quick, simple, economical, small form factor, and hot-pluggable technology. The external SATA interface competes with USB and Firmware. The physical SATA is made of a 7 line flat band,

[26] http://h10032.www1.hp.com/ctg/Manual/c00256909.pdf

8mm cable, which is up to 1 metre long. In addition, SATA utilises low voltage differential signalling (LVSD) similar to SCSI to avoid signal interference.

The following bullet points cover all the revisions of SATA that have occurred. These have been extracted from the webpage TechTarget[27]:

- SATA Revision 1. These devices were widely used in personal desktop and office computers, configured from PATA drives daisy chained together in a primary/secondary configuration. SATA Revision 1 devices reached a top transfer rate of 1.5 Gbps.
- SATA Revision 2. These devices doubled the transfer speed to 3.2 Gbps with the inclusion of port multipliers, port selectors and improved queue depth.
- SATA Revision 3. These interfaces supported drive transfer rates up to 6 Gbps. Revision 3 drives are backward-compatible with SATA Revision 1 and Revision 2 devices, though with lower transfer speeds.
- SATA Revision 3.1. This intermediate revision added final design requirements for SATA Universal Storage Module for consumer-based portable storage applications.
- SATA Revision 3.2. This update added the SATA Express specification. It supports the simultaneous use of SATA ports and PCI Express (PCIe) lanes.
- SATA Revision 3.3. This revision addressed the use of shingled magnetic recording
- SATA Revision 3.5. This change promoted greater integration and interoperability with PCIe flash and other I/O protocols.

Fibre Channel (FC)

Fibre channel (FC) protocol is a standard interface for fast and lossless storage area networks. The FC standard defines the connection architecture, physical components including the connectors and cables (not only glass fibre but also copper), data throughput, distances, coding, management, and services. One of the main characteristics of FC is its high availability. It is able to avoid frame drops through link flow control mechanism based on credits that are called buffer to buffer credits. FC contain their own hardware element, the Host Bus Adapter (HBA), and the workload for the central processor unit of the server.

There are three topologies with respective port types these include point-to-point (P2P), arbitrary loop, and switched fabric. FC are comparable to SCSI hot-pluggable devices. The specifications for FC's are shown below in Figure 4.12.

[27] https://www.techtarget.com/searchstorage/definition/Serial-ATA

Figure 4.12 Fibre Channel Development[28]

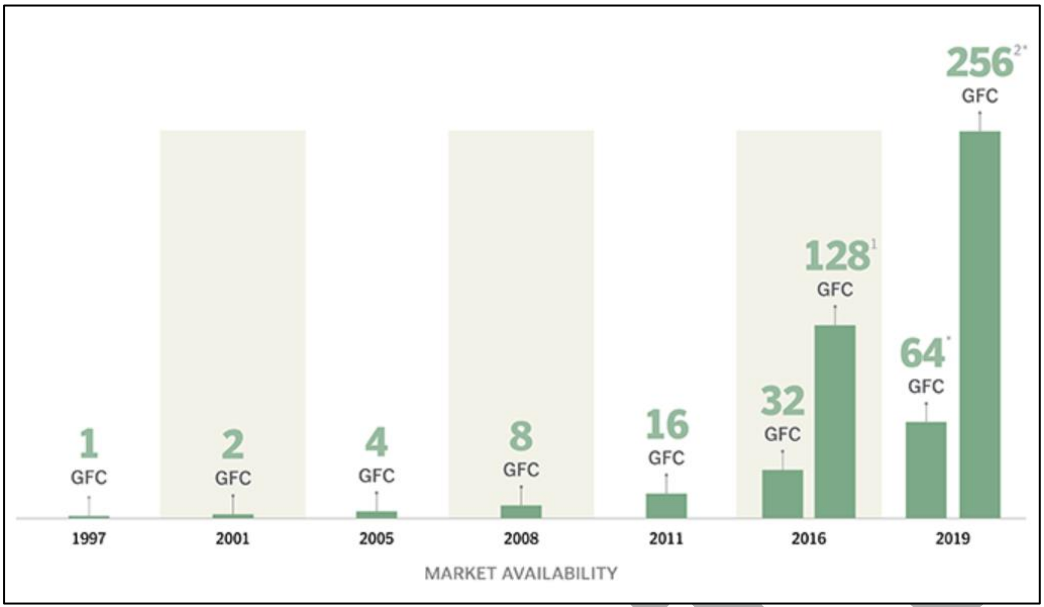

Fibre Channel over Ethernet (FCoE)

FCoE is an extension of FC it covers fibre channel frames that use an Ethernet networks that are 10 Gigabit or faster. It is usually used in high performance applications with a lot tolerance for latency, packet loss and high throughput demands. FCoE is usually chosen by larger enterprises because it runs on FC hardware that is typically already implemented in high performance data centres. In comparison to iSCSI, FCoE has a significantly lower package overhead which allows for greater data transfer rates.

FC usually uses a stack with five layers (FC-0 to FC-4) and the Ethernet uses Open Systems Interconnection Model (OSI-model) with seven layers. Whereas, FCoE transfers the FC layer to FC-2 to Ethernet which enables higher FC-layers (FC-3 and FC-4) to be implemented in Ethernet infrastructures as well. This does however, limit the FCoE to layer 2 domains within the data centre infrastructure. Since, it is often very expensive to change system to contain an FCoE, because it requires a lot of hardware. Therefore, FCoE are usually installed in existing FC infrastructure or considered in the design of a new data centre.

Internet small computer interface (iSCSI)

The iSCSI standard is an IP-based SCSI network which operates over TCP. As servers utilise the same interface for both LAN and SAN, this allows iSCSI to merge the parallel network infrastructures to a convergent network, thus reducing cabling and the need for adapters. Since no new hardware is required for iSCSI it is relatively inexpensive, however, the performance is far lower than that of the more expensive FCoE. Nevertheless, the lower price of iSCSI makes it a popular choice amongst small and medium size enterprises. Table 4.5 describes the similarities and difference between iSCSI and FCoE, demonstrating which standard has the upper hand for reliability, performance, cost and flexibility.

[28] https://www.techtarget.com/searchstorage/definition/Fibre-Channel?Offer=abMeterCharCount_var3

Table 4.5 Fibre channel (FC) SAN vs. iSCSI SAN[29]

	FC SAN	iSCSI SAN	Supporting points
Reliability	Yes	No	▪ FC has been a ratified standard for 9 years longer than iSCSI (1994 vs. 2003). ▪ FC SANs are isolated from non-storage traffic, thus, creating a safe haven for critical data. ▪ FC SANs are much simpler than iSCSI SANs. Simple = reliable
Performance	Yes	No	▪ FC requires host bus adapter initiators. ▪ FC can offer line rate, non-blocking speeds for every device. ▪ FC ensures storage traffic have 100% of available network resources.
Cost	No	Yes	▪ iSCSI doesn't require host bus adapter or storage specific network devices ▪ iSCSI leverages IP network knowledge base reducing manpower. ▪ iSCSI storage arrays are cheaper than FC and aimed at SMEs.
Flexibility	No	Yes	▪ iSCSI doesn't limit initiator to target distance like FC does, as it runs on IP. ▪ iSCSI initiators, networks and targets have less interoperability challenges than FC. ▪ iSCSI are well suited for production, development, disaster recover and virtual environments.

Infiniband (IB)

Infiniband is a high speed, low latency technology use to interconnect servers, storage and networks in a datacentre. It has a low latency networking with delays around 20 microseconds. The InfiniBand roadmap details 1x, 2x, 4x, and 12x port widths with bandwidths reaching 600Gb/s data rate HDR in the middle of 2018 and 1.2Tb/s data rate NDR in 2020. The roadmap is intended to keep the rate of InfiniBand performance increase in line with systems-level performance gains.

[29] https://www.techtarget.com/searchstorage/definition/iSCSI

Figure 4.13 InfiniBand specifications[30]

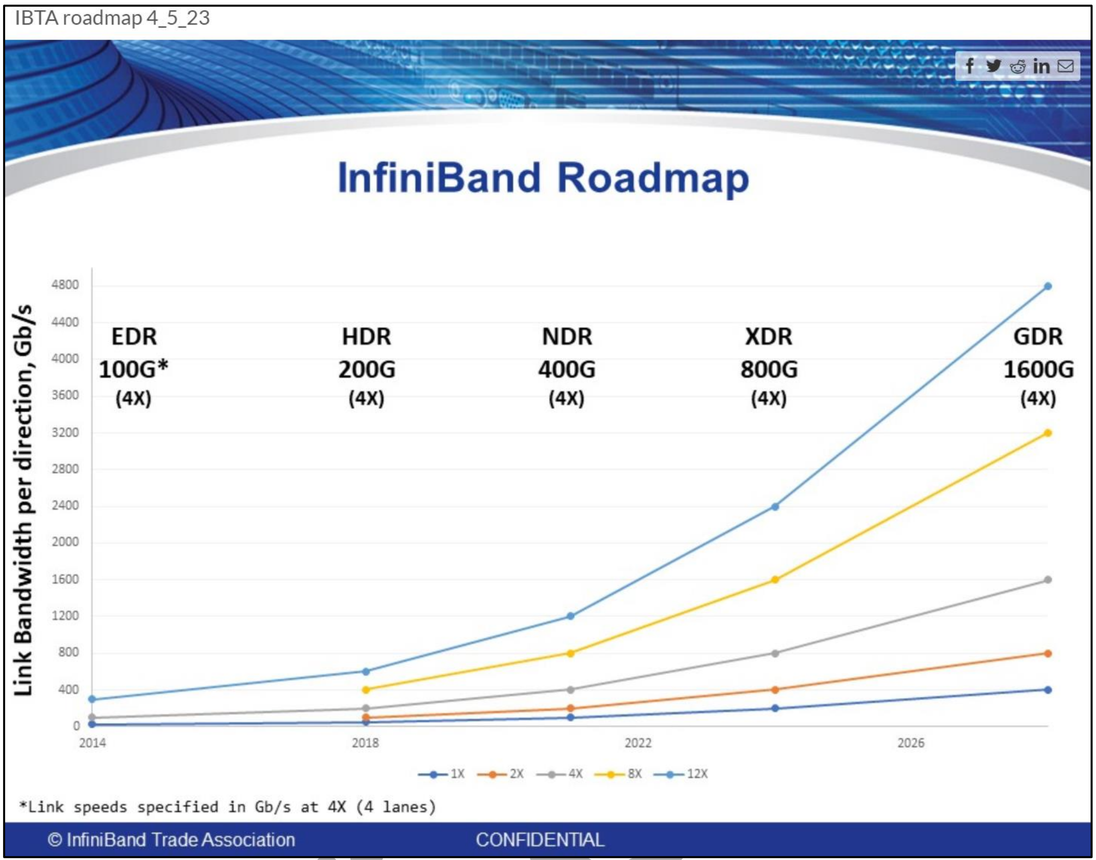

IB is usually more expensive and complex than comparable solutions such as FCoE or iSCSI. This is because it often requires new hardware. Therefore, it is primarily used in applications where other solutions cannot match performance requirements.

4.1.1.6 Material efficiency- disassemblability and repair

Material efficiency requirements for disassembly are included in the current Ecodesign regulation for servers and data storage products in Annex II, 1.2.1. This requirement sets out that: *joining, fastening or sealing techniques do not prevent the disassembly for repair or reuse purposes of the following components, when present: data storage devices, memory, CPU, motherboard, expansion card/graphic card, PSU, Chassis, Batteries.*

The following section discusses the requirements in the scope of the following components (if included): data storage devices, memory, CPU, motherboard, graphic card, PSUs, chassis, batteries, fans, integrated switch, RAID controllers and Network Interface Cards.

According to EN 45554, there are four classifications related to disassembly requirements: the fastener types, the necessary tools, the working environment, and the skill level of the disassembler. These are described below.

Skill level:

[30] https://www.infinibandta.org/infiniband-roadmap/#iLightbox[01ab527150f12ce7d56]/0

- ***Layman (Class A):*** *person without any specific repair, reuse or upgrade experience or related qualifications.*
- ***Generalist (Class B)****: repair, reuse or upgrade process cannot be carried out by layman (class A) but can be carried out by a person with a general knowledge of basic repair, reuse or upgrade techniques and safety precautions.*
- ***Expert (Class C):*** *person with specific training and/or experience related to the product category concerned.*

Most disassembly can be done with generalist skills of **class B**. However, there is feedback that specific tasks currently require **class C** skills. For example when completing a refurbishment, there are server reconfigurations that need to be done. These require technical expertise to deliver but these are software based.

<u>*Working environment:*</u>

- ***Use environment (Class A):*** *If a repair, reuse or upgrade process can be carried out in the environment where the product is in use without any working environment requirements.*
- ***Workshop environment (Class B):*** *If a repair, reuse or upgrade process cannot be carried out in the environment where the product is in use (class A) but does not require a production-equivalent environment.*
- ***Production-equivalent environment (Class C):*** *If a repair, reuse or upgrade process can only be carried out in an environment that is comparable with the environment in which the product was manufactured.*

Class B workshop environment is likely to be the most apt description of working environment requirements. Many repairs are performed on data centre premises, **Class A** is also relevant.

<u>*Necessary tools:*</u>

- ***Class A:*** *feasible with no tool; a tool supplied with the product or spare part; or with basic tools (screwdriver, hex key, pliers, spanner)*
- ***Class B:*** *product group specific tools*
- ***Class C:*** *other commercially available tools*
- ***Class D:*** *proprietary tools*
- ***Class E:*** *not feasible with any existing tool*

The Intel study[31] revealed that the PSU, motherboard and drives are the components which are most likely to fail. Class A requirement of “tools supplied with the product or spare part” should be avoided, as many maintenance and repair activities use harvested parts that will not have the tools supplied.

<u>*Fastener types:*</u>

- ***Reusable (class A):*** *An original fastening system that can be completely reused, or any elements of the fastening system that cannot be reused are supplied with the new part for the repair, reuse or upgrade process.*
- ***Removable (class B):*** *An original fastening system that is not reusable, but can be removed without causing damage or leaving residue which precludes*

[31] IT@Intel: Green Computing at Scale, August 2021

reassembly (in case of repair or upgrade) or reuse of the removed part (in case of reuse) for the repair, reuse or upgrade process.

There is currently a mix of reusable and non reusable fasteners used. It is understood that

- fasteners should be reusable or removable. Feedback is provided that
- the number of different types of screws should be reduced and that
- no proprietary fastening systems be used in order to facilitate disassembly.

Easy separation of the product facilitates increased reuse and recycling. Not only does the disassemblability allow for access to parts for reuse, but also ensures material is not lost during recycling, as making it easier to access circuit boards will ensure precious metal connections are not lost through shredding.

4.1.1.7 Power over Ethernet (PoE)

Power over Ethernet (PoE) is a technology that enables both data and electrical power to be transmitted over a single Ethernet cable. This technology simplifies the installation and deployment of devices like IP phones, wireless access points, security cameras, and other networked devices by eliminating the need for separate power cables. PoE offers convenience, cost savings, and flexibility in network design. PoE involves injecting direct current (DC) electrical power onto the unused pairs of wires in an Ethernet cable while still transmitting data over the other pairs. Power Sourcing Equipment (PSE), such as PoE-enabled switches or injectors, provides the power, while Powered Devices (PDs) receive power from the Ethernet cable.

1. POE Standards:

There are different PoE standards that define the power levels and requirements:

- IEEE 802.3af: Provides up to 15.4 watts of power per port.
- IEEE 802.3at (PoE+): Provides up to 30 watts of power per port.
- IEEE 802.3bt (Type 3 and Type 4): Provides up to 60 watts or 100 watts of power per port.

Each new standard provides for backward compatibility with all previous standards and notes the minimum power per port to be supplied to the powered devices. This minimum requirement accounts for power loss over the length of the cable, with a maximum length set at 100m.

With each power increase, cabling requirements also increased, with Cat 5 cabling being the minimum requirement for Type 3 (60W) and Type 4 (90W) PoE.

2. Benefits of POE

The use of PoE can provide several advantages during its installation that make it an attractive investment. These benefits are displayed below in Figure 4.14.

Figure 4.14 List of advantages of POE during installation[32]

Time and cost savings	PoE can reduce the time and expense of having electrical power cabling installed. Network cables do not require a qualified electrician to install. Reduction of power outlets required per installed device saves money. It also reduces energy costs by allowing for centralized control over lighting, window shades, and heating and cooling.
Flexibility	Without being tethered to electrical outlets, devices such as monitors, security cameras, and wireless access points can be positioned in ideal locations and be easily repositioned if moved.
Safety	Power delivery using PoE is designed to intelligently protect network equipment from overload, underpowering, and incorrect installation. It also eliminates the danger of working with or around dangerous high-voltage power sources.
Reliability	PoE power comes from a central and universally compatible source and not from a collection of distributed wall adapters. It can be backed up by an uninterruptible power supply (UPS), allowing for continuous operation even during power failures. PoE also allows for devices to be easily disabled or reset from a centralized controller.
Scalability	Having power available on the network means that the installation and distribution of network connections are simple and effective.

4.1.2 Standard improvement options

This section will discuss the current state of servers in the market today and where the technology currently is. At present the servers and data storage market is still a very fast paced industry demonstrating constant progress in a number of simultaneous areas such as new component development, product configuration, system architecture and all in short time intervals within months. Therefore, the baseline for products with standards improvement options is challenging to identify. Below is a list of areas of potential improvement:

- New component development
- Storage and utilisation
- Power management
- Data processing
- Virtualisation
- Modularity in the design
- Efficiency in power supply
- Thermal management

4.1.2.1 New component development

New and improved components provide much better performance compared to their predecessors. There is also significant energy consumption reduction.

On the chip-level this is achieved by techniques such as voltage scaling, multi-core designs, new materials, and single core power management as well as power gating.

On the board-level energy consumption this is achieved by integrating passive component functionality into semiconductor ICs. This reduces the number of actual components that need to be powered and generating thermal energy loss.

[32] https://www.cisco.com/c/en/us/solutions/enterprise-networks/what-is-power-over-ethernet.html

An example is the latest DDR5 memory generation. It debuted in 2021 and already featured a 50% increase in bandwidth, and ~20% less energy consumption compared to DDR4[33].

Configuration of the system is another aspect that helps on the energy and resource efficiency: The use of the appropriate configuration for the intended application can contribute significantly to savings.

4.1.2.2 Storage and utilisation

The available SSD and HDD options in the market, offer different advantages depending on the application and therefore storage device selection is important. Performance improvements are seen in storage worklets in SPEC SERT as well as in the comparison of maximum and idle power measurements on otherwise identical systems (all SSDs or all HDDs).

There are also software based options that enable optimisation of the storage capacity. This is achieved through managing redundant data.

4.1.2.3 Power management

There are a number of power management options available and some main examples of those are:

- Dynamic voltage and frequency scaling (also known as CPU Throttling). This function automatically reduces the processor's frequency (or clock speed) and voltage when idle.
- Core Parking. This feature dynamically disables CPU cores in an effort to conserve power when idle.
- Power profiles. These are a number of pre-defined groups of power management settings. Although varying from brand to brand, they offer three basic options of balanced, maximum performance and maximum energy efficiency.
- Power metering & budgeting. This function offers the setting of power limits, or caps, for individual servers based on monitoring of power use of the server[34].

ICF has created reference material for the U.S. EPA on how users can take full advantage of power management in servers titled "How to Save Idle Energy in Computer Servers"[35]

As discussed in the above section Processor and memory, Dynamic Voltage and Frequency Scaling (DVFS) is a technique used to reduce energy consumption in electronic devices by adjusting the operating voltage and frequency. DVFS allows for the dynamic adjustment of the supply voltage and clock frequency of a device based on its workload or performance requirements. By lowering the voltage and frequency during periods of low activity or idle states, energy consumption can be significantly reduced. This approach helps to optimise the power-performance trade-off and minimise unnecessary energy usage. There are currently no processor power management function requirements in Ecodesign. However, there is interest

[33] DDR5 Memory Standard: An introduction to the next generation of DRAM module technology - Kingston Technology

[34] https://www.energystar.gov/products/utilize_built_in_server_power_management_features

[35] ENERGY STAR Ask the Experts | Products | ENERGY STAR

to review this technique as by reducing voltage and frequency, the energy consumption and associated carbon footprint of electronic devices can be lowered.

Implementation of DVFS and activating reduced power states for idle cores or sockets, has already proven effective in the ENERGY STAR program.

Another example of power management is the latest DDR5 modules for example which feature on-board Power Management Integrate Circuits (PMIC), which help regulate the power required by the various components of the memory module (DRAM, register, SPD hub, etc)[36] .

4.1.2.4 Data processing

There are a number of methods utilised that have enabled effective data processing. These include having multiple cores to the CPU which allow multiple single threaded activities on the same machine. Others include muti-threading and parallel computing. These are both software activities that help accelerate processing for larger complex problems sets.

The latest processors are designed to accelerate performance across the fastest-growing workloads in artificial intelligence (AI), data analytics, networking, storage, and high-performance computing (HPC).

These processors incorporate the highest number of built-in accelerators. Built-in acceleration is an alternative, more efficient way to achieve higher performance than growing the CPU core count. With built-in accelerators and software optimizations, processors have been shown to deliver leading performance per watt on targeted real-world workloads. This can result in more efficient CPU utilization, lower electricity consumption[37]. However, it should be noted that built-in accelerators are not yet included in the SERT methodology. They are a side processor that performs more work, which is offloaded to it, which the CPU core cannot handle.

4.1.2.5 Virtualisation

Virtualisation allows multiple virtual machines to run on a single physical server, enabling better utilisation of hardware resources. Cloud computing services offer scalable infrastructure, allowing organisations to dynamically provision and deprovision resources as needed. Both virtualisation and cloud computing help reduce the likelihood of having idle or underutilised servers.

Over the last decade or so, a significant number of hypervisor vendors, solution developers, and users have been equipped with Virtualisation technology, which is now serving a broad range of users in the consumer, enterprise, cloud, communication, technical computing sectors.

Virtualisation is used in[38]:

- CPU virtualisation
- Memory virtualisation
- I/O virtualisation
- Graphics Virtualisation
- Virtualisation of Security and Network functions

---

[36] DDR5 Memory Standard: An introduction to the next generation of DRAM module technology - Kingston Technology

[37] 4th Gen Intel® Xeon® Scalable Processors

[38] Intel® VT: Intel® Virtualization Technology - What is Intel® VT? |...

It is worth noting however, that as discussed in Task 2, section 2.2, it was expected virtualisation to increase utilisation rates (and therefore lower sales). However, stakeholder feedback indicates that utilisation rates have not increased as much as expected, primarily due to end-user behaviour.

4.1.2.6 Modularity in the design

Modular server architecture comes as an answer to the search for more compute choices which can be provided faster. This is a result of rapid digital transformation which has ignited an eruption of intelligent devices and cloud-based services. Modular server architectures let suppliers use standard building blocks, flexibly configure systems and deliver innovative solutions.

In early 2022, six companies, Intel, Dell, HPE, Google, Meta, and Microsoft (and recent addition AMD) started a collaboration in the Open Compute Project (OCP) called Data Center Modular Hardware System (DC-MHS).

DC-MHS documents (through six base specifications) several types of Host Processor Modules (HPMs) and how they interface with building blocks needed to configure a full server. Host Processor Modules (with CPU and memory) and other modular, standardized elements allow a wider range of systems for Enterprise, Hyperscale, Cloud, Telecom and Edge sooner and with reduced engineering investment. The same module is deployed in multiple configurations at higher scale[39]. The benefits are deemed to be multiple and include reduced costs, reduced custom designs, reusable building blocks across multiple designs and markets, faster time to market, module replacement, easier reuse or recycling.

Figure 4.15 Demonstrating DC-MHS Modularity[40]

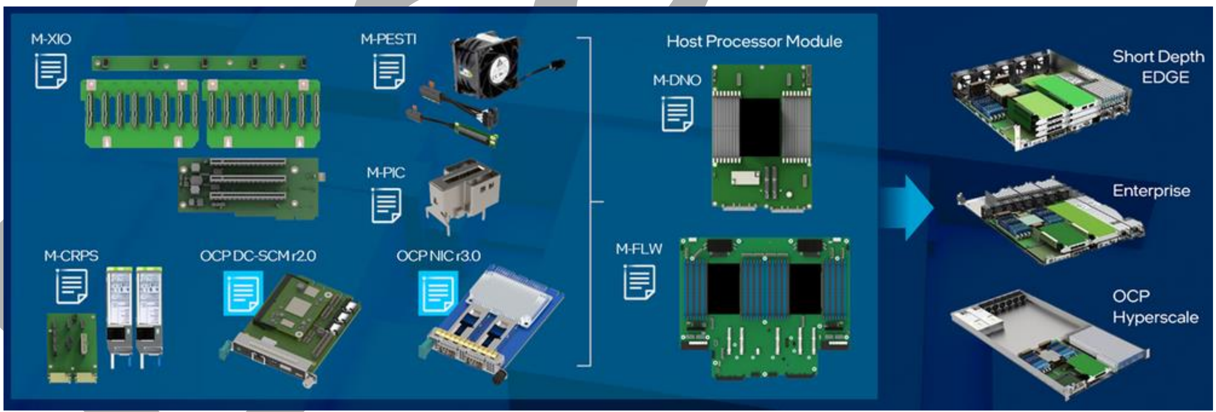

4.1.2.7 Efficiency in power supply

As discussed in section Power supply unit above, there are continuous discussions on how to best increase its efficiency.

80 PLUS® PSU certification scheme[41] continues to be sought after and a good practice to demonstrate meeting the current regulation requirements uplifted from 1 January 2023 on PSU as indicated in Figure 4.7.

[39] Modularity and Standardization Transforming Server Architecture (intel.com)

[40] As above

[41] 80 Plus Overview | CLEAResult

A white paper from ITI's The Green Grid published in 2021[42] studied the 230V Enterprise PSU market-ready efficiency and identifies barriers and solutions to the uplifted current Regulation's requirements from January 2023. It is worth noting that the required efficiency baseline requires 96% efficiency of the PSU, power supply, at 50% of rated power. This is equivalent to the 80 PLUS® 230V Data Center Titanium Certification, the highest Certification currently offered. It is claimed that currently there are very limited products that can offer this as shown in the graph below with a recommendation being on having PSU under 750W to use a baseline of 94% efficiency at 50% load (equivalent to the 80 PLUS Platinum Certification), while working with the power supply manufacturers to meet the current level.

Figure 4.16 Chart showing quantity of power supply units that meet Titanium Certification for wattages up to 2000 watts, in increments of 50 watts

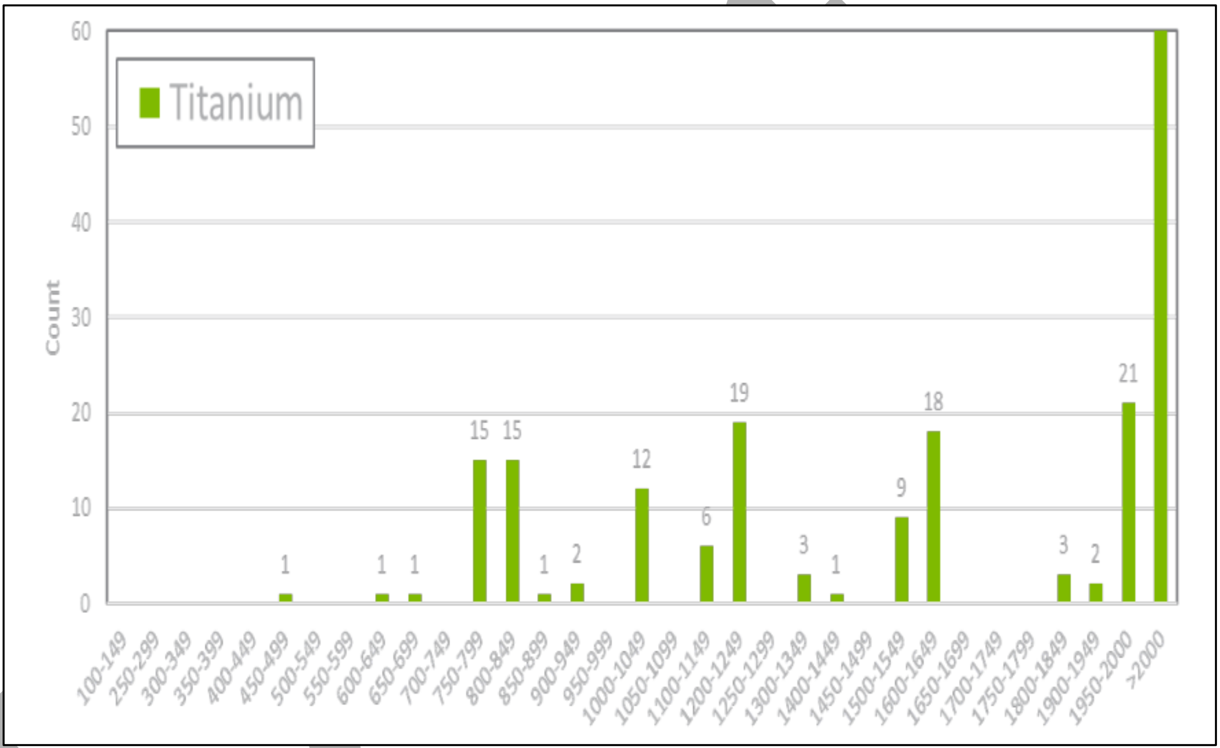

4.1.2.8 Thermal management

Thermal Management has two aspects and relates to the server level and the data centre level.

On this section it is taken to be the transport of the heat away from the active components of the server to the warm aisle of the data centre. Design of the Server board must determine how the unit will be cooled and how that will affect the overall board.

Latest technology enables monitoring of equipment and components to improve performance, monitor temperature and predict failures. The sensors also help load the server more efficiently, such that workloads from one server rack can be transmitted to the other sever racks that may be idle (and therefore cool).

[42] 230V ENTERPRISE PSU MARKET-READY EFFICIENCY STUDY_1.pdf (clearesult.com)

Ensuring that thermal properties are carefully monitored will help to reduce downtime and integrating modern solutions like artificial intelligence into thermal management systems is the likely next step forward for this field[43].

4.1.3 Best Available Technology (BAT)

This section will discuss where the servers and data storage market is progressing towards and how it is expected to develop. A lot of the innovations are happening at the component level. Architecture of the servers and storage products are not evolving as rapidly.

4.1.3.1 Microprocessors and memory

The most advanced microprocessors being mass produced on the market currently employ TMSC's 4N feature size, which is a modified version of TMSC's 5nm node process[44]. These smaller features allow for additional transistor density, generating additional performance and improved efficiency over previous generations of CPUs and GPUs. This latest technology is available on the market in certain high end GPU applications.

While DDR5 is relatively new on the market and is currently the most advanced form of server memory available, additional iterations on the technology will cause its CAS latency ratings to fall over time. This will unlock additional performance and access the full potential of DDR5's upgraded speed and bandwidth over previous generations of DDR.

4.1.3.2 Hard disk drives and storage

The item will be finalised in the second iteration of this report once further insight has been obtained.

4.1.3.3 Modular Servers and Microservers

Modular servers have continued to evolve and incorporate newer CPU, GPU and advanced SSD components that can be configured to service a specific workload type more effectively. They now employ more advanced features including improved fabric and bus technologies such as CXL, higher speed connectors such as silicon photonics, and support higher maximum power limits[45].

Innovation in the microserver space has lagged since the last preparatory study was finalized as the market for these products did not materialize to the degree initially expected. There are a handful of microservers still available on the market, but they are largely a niche application.

4.1.3.4 Advanced cooling

As discussed in Phase 1, the best available advanced cooling technologies for ICT equipment are currently CPU cold plate and immersion.

43 Circuit Protection – Server Thermal Management | Arrow.com

44 https://www.tsmc.com/english/dedicatedFoundry/technology/logic/l_5nm

45 https://www.cisco.com/c/en/us/solutions/computing/what-is-modular-computing.html

- CPU cold plates are attached directly to the CPU, replacing traditional fan technology and runs circulating cooling water directly to the plate to cool the CPU.
- Immersion cooling involves submerging the entire ICT product in a dielectric fluid which allowing for efficient cooling of very high-density electronics. Two-phase immersion cooling in particular allows for high efficiency heat removal as it uses a dielectric liquid with low boiling point in a sealed system that generates vapor when boiling that then phase changes back into liquid and removing heat when interacting with the chilled water condenser coil within the system.[46]

These are highly efficient and allow for higher outlet temperatures, 60 – 75 degrees C, which allows for more efficient waste heat recovery applications and the use of additional free cooling. These have low maintenance and low total cost of ownership. There are some installation concerns, but the main drawbacks are in the design of the specialised server to accommodate for the liquid heat exchange. Internal analysis from Asetek shows that the electrical consumption for a per-node cooling system can be reduced by 40%. Examples of such technologies and companies are Denmark based Asetek, and Canada based CoolIT.

These technologies are currently employed in HPC applications, but their use may increase in high rack density enterprise applications in the future.

4.1.4 Best Not yet Available Technology (BNAT)

4.1.4.1 Standby-Readiness for Servers

As discussed in the Phase 1 report, section 2.5, there has been interest in servers to gaining the ability to enter deep sleep state and recover from them in a rapid fashion to give data centre operators the option to power down groups of servers in a behaviour similar to “core parking” at the CPU level. This functionality is not yet available on the market.

There is little demand for it currently and concerns were raised that this type of behaviour can negatively impact reliability and increase jitter and latency which can negatively impact certain workload types within the data centre more broadly.

4.2 Production, Distribution and End-of-Life

4.2.1 Bill of Materials

4.2.1.1 Rack- mounted server

The following Bill of Materials is maintained based on the previous preparatory study[47] and reiterated here for discussion and further input on materials and quantities involved.

[46] https://datacenters.lbl.gov/liquid-cooling

[47] Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC, DG ENTR Lot 9 - Enterprise servers and data equipment, Task 4: Technologies, July 2015 – Final report, Produced by bio by Deloitte and Fraunhofer IZM

Figure 4.17 Bill of materials of an average rack server

Component	Material	Weight (g)	Component	Material	Weight (g)
Chassis	Metal Body	12 265	PSUs (2*400W)	Low-alloyed steel	1 027
	Plastics	348		Chromium steel	66
	Plastics	282		Brass	42
	Aluminium	249		Copper	9
	Copper	179		Zinc	7
	Electronic components	131		Aluminium	491
Fans (4)	Steel	386		High Density Polyethylene	184
	Copper	78		Polyvinylchloride (PVC)	92
	Iron based	55		Paper	50
	Plastic (PBT-GF30)	206		Electronic components	1 101
	Plastic (PCABSFR40)	21		Solder	31
	Plastic (undefined)	200		PCB	326
HDDs (4)	Steel	12	CPU Heat Sink	Copper	442
	Low alloyed steel	222		Steel	140
	Aluminium	1 335	Mainboard	Controller board	1 667
	PCB	179	Memory	PCB	97
				IC	38
ODD	Low alloyed steel	115	Expansion Card	PCB	349
	Copper	7	Cables	Brass	7
	Aluminium	1		Copper	81
	High Density Polyethylene (HDPE)	28		Zinc 0.166 kg	96
	Acrylonitrile-Butadiene-Styrene	12		High Density Polyethylene	104
	Polycarbonate (PC)	7		Polyvinylchloride (PVC)	145
	Electronic components (capacitors,	8		Polyurethane (PUR)	2
	Solder	2		Synthetic rubber	35
	PCB	9	Packaging	Cartons	3629
CPUs (2)	Copper	31		HDPE/ unspecified plastics	78
	Gold	0,4		GPPS/ Styrofoam	1 026
	PCB	21			
	IC	2			
Total weight of BC-1: 27 748 g					

The mainboard does not include memory and CPUs, which are listed and evaluated as separate items.

4.2.1.2 Rack-blade server

The following Bill of Materials is again maintained based on the previous preparatory study[48] and reiterated here for discussion and further input on materials and quantities involved by the stakeholders.

[48] Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC, DG ENTR Lot 9 - Enterprise servers and data equipment, Task 4: Technologies, July 2015 – Final report, Produced by bio by Deloitte and Fraunhofer IZM

Figure 4.18 Bill of materials of a blade system with 8 servers

Enclosure		
Component	**Material**	**Weight (g)**
Chassis	Steel	87 000
Fans (6)	Steel	964
	Copper	194
	Iron based	137
	Plastic (PBT-GF30)	515
	Plastic (PCABSFR40)	52
	Plastic (undefined)	499
PSUs (4)	Low-alloyed steel	4 981
	Chromium steel	319
	Brass	202
	Copper	43
	Zinc	32
	Aluminium	2 384
	High Density Polyethylene (HDPE)	894
	Polyvinylchloride (PVC)	447
	Paper	245
	Electronic components	5 343
	Solder	149
	PCB	1 581
Total weight of enclosure		**105 981 g**

Figure 16: Illustration of a blade system

8 Blade Servers					
Component	**Material**	**Weight (g)**	**Component**	**Material**	**Weight (g)**
Chassis	Steel	33 600	**HDDs (16)**	Steel	47
CPUs (16)	Copper	244		Low alloyed Steel	888
	Gold	3		Aluminium	5 341
	PCB	170		PCB	717
	IC	15	**Packaging**	Cartons	14 969
CPU Heat Sinks	Copper	1 688		HDPE/ unspecified plastics	321
	Steel	560		GPPS/ Styrofoam	4 233
Memory	PCB	773	**Mainboards**	Controller Board	6 451
	IC	307			
Total weight of 8 Blade Servers: 70 327g					
Total weight of BC-2: 176 308 g					

4.2.1.3 Resource considerations for rack and blade mounted systems

As discussed in Task 3, section 3.3.6, Most hardware manufacturers have end-of-life mechanisms to recuperate products when customers no longer require these (whether due to failure or upgrade). These mechanisms will aim to reuse (or resell) the product, repair (or refurbish), recuperate functioning components and recycle remaining materials.

Table 4.6 below shows the breakdown of material usage, recycling, energy recovery and landfill for enterprise servers provided by a European waste disposal stakeholder in 2023. This data shows how 81% of collected server mass is either re-used or materially recycled. Including waste heat recovery, up to 99% of server materials value is recovered. Only 0.16% of material by mass is incinerated or sent to landfill.

Table 4.6 Inputs in the end-of-life phase of collected enterprise servers[49]

	Plastics	Metals	Electronics	Misc.
Mass ratio within server	0,97%	67,88%	30,88%	0,28%
Re-use	0%	0%	1%	0%
Material recycling	98%	98%	43%	50%
Heat recovery	0%	2%	56%	0%
Non-recovery incineration	0%	0%	0%	50%
Landfill	2%	0%	0%	0%
Total	100%	100%	100%	100%

4.2.1.4 Storage system

The following Bill of Materials is based on the previous preparatory study[50] and put forward for discussion and further input on materials and quantities involved by the stakeholders.

Figure 4.19 Bill of materials of a controller

Material	Weight (g)	Material	Weight (g)
Controller		**PSU Controller**	
steel	14900	Mainboard	1650
stainelss stell	3360	Cables	40
sluminum sheet	574	Chassis andbulk material	1778
copper	1040		
ABS	1020	**PSU Fans**	
PET	78	Steel	220
HDPE	174	Copper	130
PP	36	iron based	26
PC	62	Nylon 6	18
Nylon 6	10	PC	70
PVC	170	ABS	38
Other plastics	24		
Printed circuit board	1154	**Packaging**	
		Cartons	7258
		HDPE/unspecified plastics	156
		GPPS/ Styrofoam	2052

Total weight: 36038g

[49] Feedback from European electronics waste disposal stakeholder 2023

[50] Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC, DG ENTR Lot 9 - Enterprise servers and data equipment, Task 4: Technologies, July 2015 – Final report, Produced by bio by Deloitte and Fraunhofer IZM

Figure 4.20 Bill of materials for a Disc Array Enclosure (DAE)

Material	Weight (g)
Chassis	
PC	203
ABS	46
Steel sheet part	7687
Zinc Part	149
Steel Machined Part	1.5
PSUs in DAE (2)	
mainborad	2108.5
cables	52
chassis and bulk	2273
Fans in PSU (4)	
Steel	281.5
Copper	166
iro based	33
Nyylon 6	23.5
PC	88.5
ABS	47.5
Controller cards (2)	1154
Mid plane board	460

Total weight: 14774g

Figure 4.21 Bill of materials for an average storage media mix

Storage Media Mix (40.1 TB)					
3.5 HDD (19.35)	Steel	58	2.5 HDD (14.01)	Steel	278
	Low Alloyed Steel	1 103		Low Alloyed Steel	211
	Aluminium	6 637		Aluminium	2 562
	PCB	890		PCB	123
				ABS	7
SSDs (2.86)	Electronic components	172			
	IC	5			

Total weight: 12050g

4.2.1.5 Resource considerations for storage systems

Table 4.7 shows the inputs in the end-of-life phase of the storage systems assumptions from the 2015 Preparatory study for further discussion and considerations and input from stakeholders.

Table 4.7 Inputs in the end-of-life phase of storage systems[51]

	Plastics	Metals	Electronics	Misc.
Re-use	25%			
Material recycling	5%	70%	50%	68%
Heat recovery	69%	0%	24%	1%

[51] Source: Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC, DG ENTR Lot 9 - Enterprise servers and data equipment, Task 3: User, July 2015: Final report, bio by Deloitte & Fraunhofer IZM

	Plastics	Metals	Electronics	Misc.
Non-recovery incineration	0.5%	0%	0.5%	5%
Landfill	0.5%	5%	0.5%	1%
Total	100%	100%	100%	100%

4.2.2 Packaging

As discussed in Task 3, Product packaging has been categorised into three main elements:

- The cardboard packaging.
- Foam protection of the product (including bubble wrap), these are usually plastics based.
- Other plastics (films, wrappings)

The main packaging quantities for rack and blade servers are outlined in Figure 4.17 and Figure 4.18 above. Stakeholder feedback has suggested that these categories are appropriate. Additionally, it was mentioned that there is a trend to use greater amounts of cardboard for product production, replacing the use of foam.