Orientation Guide
to Power Density and Load Determination of Servers, Data Cabinets and Data Centres

eco
ASSOCIATION OF THE INTERNET INDUSTRY

Table of contents

About the authors

Ulrich Terrahe
Engineer: Dipl.-Ing. (FH),
Dipl.-Wirtsch.-Ing. (FH)
Managing Director dc-ce
DC-Beratung GmbH Co. KG

Ulrich Terrahe is a respected figure in the data centre industry with extensive expertise in concept development as well as air conditioning and cooling technology for data centres. After studying at the technical universities in Giessen and Berlin, he gained valuable professional experience at renowned companies such as Lennox Industries in England, Kraftwerksunion, Rudolph Otto Meier, and Raab Karcher Wärmetechnik. In 1997, he started in the data centre industry at Schnabel AG.

As founder and Managing Director of dc-ce DC-Beratung GmbH & Co KG, he now plays a key role in shaping the planning and operation of data centres. Up until 2019, he organised the annual industry meeting 'Future Thinking' and launched the German Data Centre Award, which recognises innovations in the data centre industry. With more than 20 years of experience, Terrahe has been involved in projects of various sizes and was honoured with the Datacentre Award in 2007. From 2008 to 2011, he was a member of the jury for this award.

His focus is on developing innovative solutions for data centres and integrating digital processes into data centre planning. Terrahe is committed to efficiency and sustainability and is in demand as a speaker at national and international conferences. He has also published numerous specialist articles on current topics in the data centre industry.

www.linkedin.com/in/ulrich-terrahe

www.dc-ce.de

Peter auf dem Graben
Engineer: Dipl.-Ing. (FH) Mechanical Engineering
Senior Data Centre Consultant

Peter auf dem Graben began his professional career in the development, implementation, and integration of microprocessor-controlled systems after completing his degree in mechanical engineering in 1992. During this time, he focussed on developing IT networks and systems for office communication and industry, profiting from his sound knowledge of mechanical engineering, electronics and IT.

Since 2004, he has worked at DXC-Technology (formerly CSC), where he was involved in various international projects as an IT engineer, project manager, consultant and technical architect. As part of these activities, he supported data centre operations in conversion, expansion and migration projects and has dedicated the last 14 years exclusively to the design, planning and virtualisation of high-performance data centre systems.

In 2021, auf dem Graben moved to dc-ce DC consulting and took on the role of data centre consultant. In this position, he uses his many years of experience to develop individual data centre solutions for customers and bridge the gap between IT and data centre organisations.

His current focus is on strategic planning, developing concepts and optimising data centre infrastructure. Thanks to his extensive knowledge in this field, he is able to quickly familiarise himself with complex IT infrastructures and implement this knowledge in his projects in a targeted manner.

www.linkedin.com/in/peter-auf-dem-graben-8342275a

www.dc-ce.de

Marcus Pump
Head of Strategic IT Consulting Division
SVA System Vertrieb Alexander GmbH, Wiesbaden

Marcus Pump has been working in the IT industry since 1994 and has extensive experience in consulting, project implementation and the operation of IT systems and data centres with on-premise and cloud architectures. His professional career began with various IT service providers and system integrators in the field of network and IT infrastructure, where he gained his first valuable experience as a systems engineer and consultant.

From 2004, Pump took on the role of Project Manager, Consultant and Managing Director in the Outsourcing division at Hamburg-based IT service provider akquinet. He played a key role in the planning, construction and operational launching of data centres and supported the development of IT architectures and operational organisations for customers who were transitioning to managed outsourcing environments. Since 2020, Pump has headed the Strategic IT Consulting division at SVA System Vertrieb Alexander GmbH in Wiesbaden. Here, he advises companies on the development of forward-looking IT strategies, service-orientated multi-sourcing architectures and IT target operating models.

In addition to strategy development, he supports companies in the design of sourcing and governance models as well as modern service and architecture models for on-premise and cloud services. With the support of the system integrator SVA, these concepts can be both strategically developed and fully implemented

www.linkedin.com/in/marcus-pump-300594143

www.sva.de/strategisches-it-consulting

www.sva.de

All three experts share many years of experience, extensive industry knowledge and a strong focus on innovation and efficiency in data centre operations.

Foreword

Dr. Béla Waldhauser
Head of the Data Center Infrastructure Competence Group

'Quo vadis, power density?' – Here we go again! After 13 years, eco is now publishing the second white paper on load determination and power density in data centres. An important guide that is just as relevant today as it was then. Ulrich Terrahe has once again taken the lead in producing the paper, this time with the active support of Peter auf dem Graben and Marcus Pump.

In 2012, eco invited experts to a workshop for the first time, which the two authors, Ulrich Terrahe and Marc Wilkens, moderated together with me. I still remember the extremely good atmosphere at this workshop quite clearly. At that time, the topic was still relatively new for data centre deployers and the development was dynamic. Above all, energy-efficient IT hardware and software were still in their infancy – and with them, power densities. Is the industry still so much characterised by change after thirteen years? The answer is an emphatic YES, and more than ever.

Back then, we were still talking about 3 to 6 kilowatts (kW) per rack, whereas today it is more like 5 to 25 CW per rack. For AI or high-performance applications, it is often even 50 kW, 80 kW or even well over 100 kW per server cabinet. In my opinion, it is not yet possible to say how these values will develop. Exactly 60 years ago, Gordon Moore (co-founder of Intel) postulated that processor performance would double every 18 months – better known as Moore's Law. Coincidence? There has been speculation and debate about the end of this law for years now. However, researchers and developers have repeatedly managed to increase CPU performance. I am curious to see whether Moore's Law will live to see its 70th, 80th or even 100th anniversary. In any case, it is clear that the worldwide increase in digitalisation – both professionally and privately – has led to a sharp increase in the demand for data centre capacity over the last 13 years and continues to do so. Donald Trump has just announced the Stargate programme with a budget of 500 billion US dollars to build huge data centres for AI systems.

Despite all the euphoria, business economics cannot be ignored. This is because the construction costs for data centres have also increased significantly. In this respect, it is more important than ever to consider the correct dimensioning of a new data centre – or the efficient modernisation of an existing data centre.

Now back to the motivation for this white paper. In 2012, as today, despite all the information available, both on-premise data centres and rented performance in co-location data centres are usually significantly oversized. A capacity underutilisation of 50 per cent is not uncommon. With production costs of around 12,000 Euro per kW plus the running costs for maintenance and repairs, this is more than surprising. Incidentally, for customers in colocation data centres, the rent is also around 12,000 Euro per kW. However, here it is spread over a period of only five years.

Since the end of 2023, another important factor has come into play in Germany: the German Energy Efficiency Act, which includes some challenging key figures for the efficiency of the power and cooling infrastructure in a data centre. For example, from mid-2027, data centres with a non-redundant net connected load of at least 300 kW must achieve a Power Usage Effectiveness (PUE) value of 1.5. With an average of 1.52 for all German data centres in 2022[1] , this seems easily achievable. However, this does not take into account the fact that the vast majority of new data centres today are already well below 1.5. Older ones currently tend to have PUE values of 1.8 to 2.0 or above. So significant efforts are needed here to achieve the 1.5 target in two years. It will then get really exciting from mid-2030, when the law sets a PUE limit of 1.3. The rising temperatures caused by climate change will make it challenging to achieve the target for many existing data centres. A number of things need to be done here. Starting with significantly higher temperatures in the cold aisle. 27 degrees Celsius should become the standard (based on ASHRAE recommendations). However, this requires a significantly higher utilisation of the data centres in order to operate the air conditioning efficiently. Today's oversizing must then become a thing of the past.

So, will the problem solve itself in the next few years? Neither the authors of this white paper nor I can really see that happening. Therefore, this white paper, like its predecessor 13 years ago, is absolutely necessary. We would like to see the findings of this guide given significantly more attention in the future – for business reasons, on the one hand, and for the sake of sustainable data centres, on the other.

Happy reading!

Dr. Béla Waldhauser

Leader of the eco Data Centre Infrastructure Competence Group

Source: https://www.borderstep.de/facts-and-figures/rechenzentren-2022/

Management Summary

The rapid consolidation of performance and the transformation of the IT landscape require a rethink in data centre planning. This white paper provides guidance for planners, installers and deployers of data centres, especially for small and medium-sized performance classes, in order to master the challenges of IT transformation. Poor planning, whether over-sizing or under-sizing, leads to significant cost risks – up to 8,000 - 10,000 Euro per kW – and endangers economic and legally compliant operation.

A continuous dialogue between management, IT and data centre planning is essential. This is the only way to integrate corporate strategy, IT strategy, technological progress, service and sourcing models as well as regulatory requirements (KRITIS, GDPR, NIS-2, IT Security Act) and sustainability goals into planning. Taking these factors into account enables a 360-degree view and leads to valid, needs-orientated planning.

The planning process itself is based on a 5-step approach: analysis of existing IT, determination of future space requirements including migration and growth areas, consideration of availability and security requirements, optimisation of the operating strategy through automation and cloud integration, and monitoring of new technologies. This iterative process, which weighs up calculable and incalculable factors (80/20 ratio), enables flexible adaptation to changes and reduces the risk of making costly mistakes. Recognising errors at an early stage is crucial, as subsequent corrections are significantly more expensive.

The white paper provides a detailed guide to determining power density and load requirements, from servers and data cabinets to the entire data centre. It supports planners and deployers in designing and operating future-proof, efficient and legally compliant data centres. The paper recommends involving qualified IT consultants to adequately map the complexity of the IT landscape.

Target group

This white paper is aimed at planners, installers, builders and deployers of data centres and server rooms.

The aim of this white paper is to provide a sound basis for data centre planning.

It is aimed at both private and public companies and focuses on server rooms and data centres with a small to medium performance class.

In the data centres referred to here, IT services are provided on different system architectures. These include classic server and storage systems as well as hyper-converged infrastructure systems and legacy platforms that are at the end of their life cycle.

This includes typical basic services, such as directory and network services, but also data backup, safety and support systems for IT operations.

It also includes application systems, such as ERP and CRM applications, and often special process-supporting systems, such as manufacturing and production control or laboratory systems.

The explanations described also apply in principle to IT service providers such as colocation providers or (cloud) service providers. However, it should be noted that the same laws of physics and the same laws and regulations apply to the operation of co-location data centres, but the data centre operator only has limited influence on the IT systems operated there. In the case of (cloud) service providers, on the other hand, the sphere of influence can vary depending on the business model. These different spheres of influence must be taken into account in the methods presented in the document.

A fundamental challenge in data centre planning

The main problem when planning a server room and data centre (DC) is that rooms, buildings and system technology are designed for long-term use and tend to be static in their type and design.

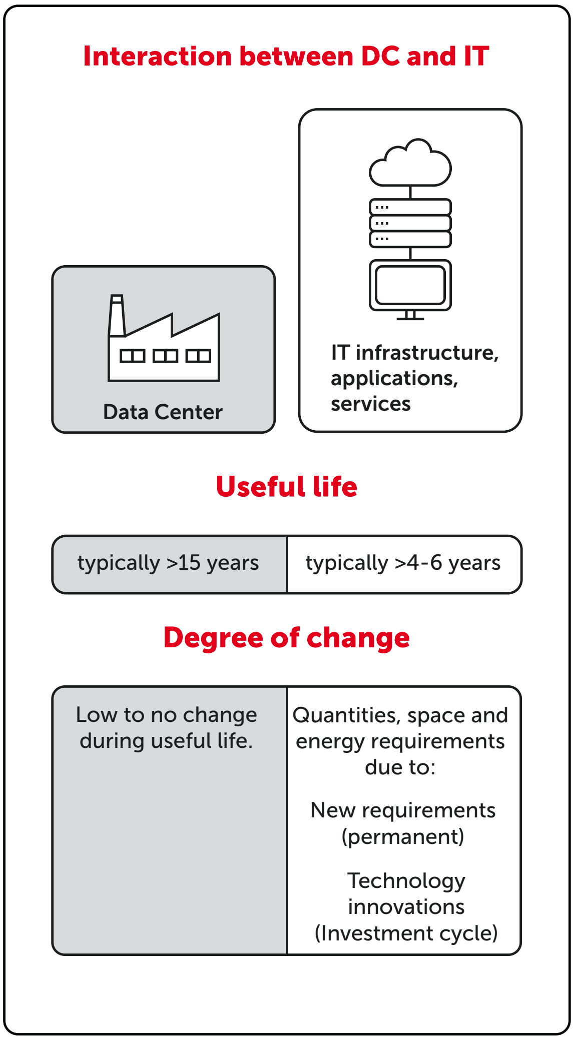

Figure 1: Illustration of the fundamental challenge in DC planning

In practice, DC systems are therefore only adaptable to a limited extent, i.e. they can only be replaced with great effort and expense.

Changes to IT systems due to various influencing factors, such as technological advances and changing usage scenarios, occur much more quickly.

Therefore, they increase the risk of bad investments in server rooms, DC buildings, and system technology and may significantly impact operating costs.

Typical planning elements for dimensioning

There are three main dimensions for DC planning: performance, space and availability (including safety).

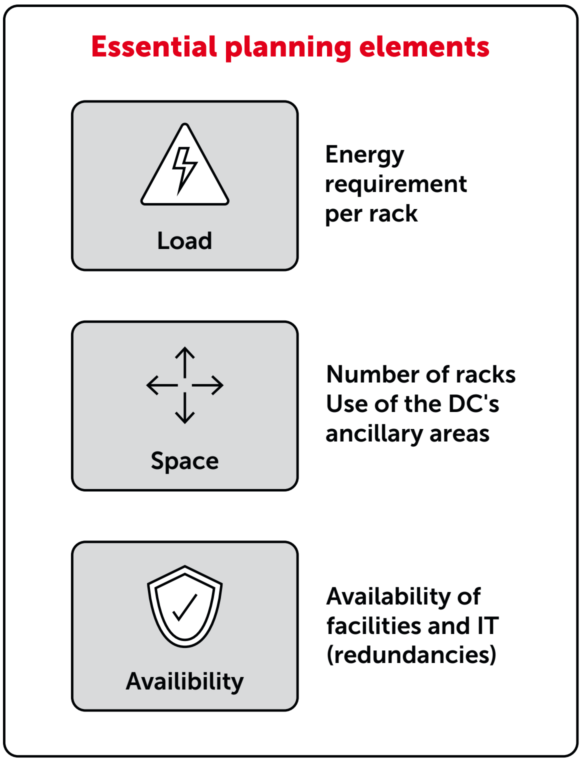

Figure 2: Illustration of the main dimensions in DC planning

The central planning element and key cost-determining factor for a data centre are the performance values set for the 19" racks and the quantity of racks required.

For this reason, the following explanations and calculation examples focus primarily on this element.

- Power specifications from 5 KW to 30 KW per racks are being discussed on the market.
- For special applications, such as high-performance computing, also significantly higher power specifications of up to 100 KW per rack.

Problem definition

Which specific consumption values can be assumed in reality can often not be answered by reviewing performance data from system manufacturers for servers, storage and networks or hyperconvergent systems. Here, empirical values from deployers can provide good support for planning.

Changes in IT systems and the assumptions of higher performance and new trends such as AI and edge computing are also causing uncertainty as to what performance per rack will be required in 3, 5 or 10 years' time.

The result of these uncertainty factors is usually a miscalculation of the performance actually required.

- In the event of oversizing, more space is required for ICT and building technology.
- The technical systems such as transformers, UPS, diesel and switching systems, batteries, air conditioning units and refrigeration systems are getting bigger.
- The energy efficiency requirements may not be met.

Alternatively, a misjudgement of the performance values can also lead to undersizing, so that the data centre reaches its performance limit too quickly and expansions are no longer possible. However, this rarely happens.

Cost impact in the event of a miscalculation

As a guideline, around € 8,000 to € 10,000 (currently 2024) can be estimated for each KW of oversizing of the required IT performance for a data centre with a solid n+1 redundancy concept.

The following table shows the magnitude of misinvestments that can occur for various scenarios in terms of quantity, power assumption and actual power requirement per rack.

Equipment	Size	Projected performance data		Actual demand		Unnecessary
Racks	Whitespace	DC design	Total perfor-mance	Demand	Total perfor-mance	investment (*1)
8 racks	49 m²	8 kW per rack	64 kW	5 kW per rack	40 kW	€ 192 – 240 K
8 racks	49 m²	10 kW per rack	80 kW	6 kW per rack	48 kW	€ 256 – 320 K
16 racks	86 m²	12 kW per rack	192 kW	8 kW per rack	128 kW	€ 512 – 640 K
16 racks	86 m²	15 kW per rack	240 kW	10 kW per rack	160 kW	€ 640 – 800 K
16 racks	86 m²	20 kW per rack	320 kW	12 kW per rack	192 kW	€ 1.02 - 1.28 million.
32 racks	149 m²	25 kW per rack	800 kW	15 kW per rack	480 kW	€ 2.56 - 3.2 million.
64 racks	275 m²	15 kW per rack	960 kW	10 kWper rack	640 kW	€ 2.56 - 3.2 million
64 racks	275 m²	20 kW per rack	1.280 kW	12 kW per rack	768 kW	€ 4.1 - 5.12 million.
64 racks	275 m²	25 kW per rack	1.600 kW	15 kW per rack	960 kW	€ 5.12 - 6.4 million
96 racks	393 m²	20 kW per rack	1.920 kW	12 kW per rack	1.152 kW	€ 6.14 - 7.68 million.
128 racks	510 m²	15 kW per rack	1.920 kW	10 kW per rack	1.280 kW	€ 5.12 - 6.4 million
128 racks	510 m²	20 kW per rack	2.560 kW	12 kW per rack	1.536 kW	€ 8.19 - 10.24 million.
128 racks	510 m²	25 kW per rack	3.200 kW	15 kW per rack	1.920 kW	€ 10.24 - 12.8 million

Table 1: Effects of miscalculation taking into account the number of racks and their high-performance
(*1 related to building, technical building equipment [DIN 276 cost groups 300 and 400])

A strategical approach

With regard to total investments in DC buildings, rooms and technical equipment for DC operation, any wrong investments can quickly amount to significantly more than 30%, in rare cases even up to three times the total DC investment costs.

The effects on operability and operating costs must be determined individually in each case, but will lead to higher costs and lower efficiency. This may jeopardise the legal requirements (e.g. in Germany, the German Energy Efficiency Act) and certification in terms of efficiency.

These examples clearly show that the effort required to qualify the planning assumptions represents an important and correct investment for the planning of a data centre.

A continuous dialogue between management, IT, and DC planning is essential to avoid significant discrepancies between planning and actual requirements, which can lead to bad investments and risks for the efficient operation of the data centre.

This is the only way to account for changes in the company and IT strategy, including changing company portfolios and business models that directly influence the IT strategy.

Technical advances in IT technology and changes in service and sourcing models can also be incorporated into long-term DC planning.

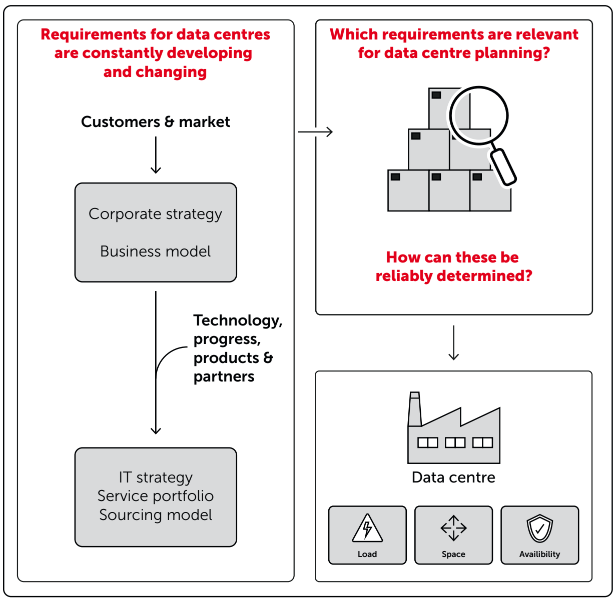

Figure 3: Getting started with the DC planning process

It typically takes between 24 and 48 months from the planning process to the completion of a data centre.

Therefore, it is advisable to draw up a comprehensive requirements specification at the start of the planning process, taking particular account of company and IT requirements.

This should be validated in a planning iteration before the start of construction to incorporate any recognisable changes into the planning.

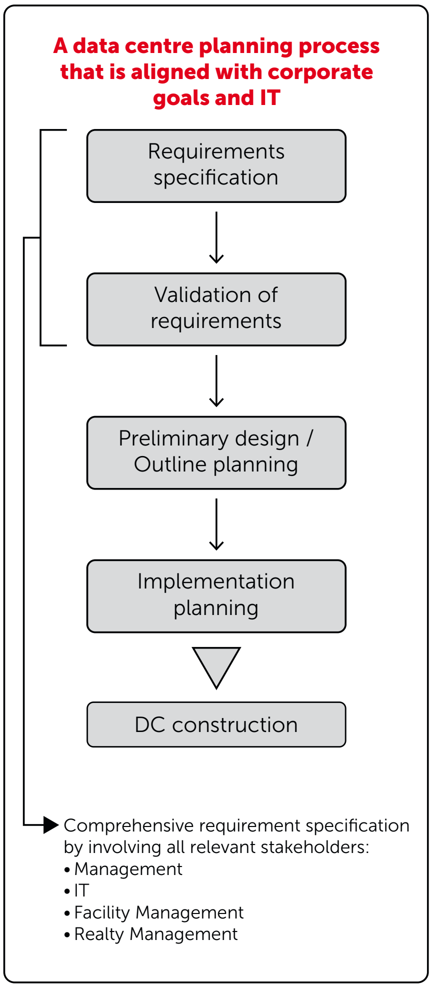

Figure 5: Data centre planning process

It is crucial to integrate changes to regulatory and compliance requirements (e.g. KRITIS, GDPR, NIS-2, IT Security Act), as well as new requirements relating to sustainability and energy efficiency, which are becoming increasingly important, into the planning process.

This approach ensures a comprehensive 360-degree analysis of the company, IT, technical and facility management requirements as part of data centre planning, which leads to valid and needs-oriented planning of the data centre project.

It is recommended that a qualified IT consultant be involved in the planning process in addition to the specialist planner for data centres and technical building services systems. This recording of requirements establishes a strong foundation for evaluating capacity and space needs as a planning impulse according to the following explanations.

Starting point for step-by-step planning

CPU technology

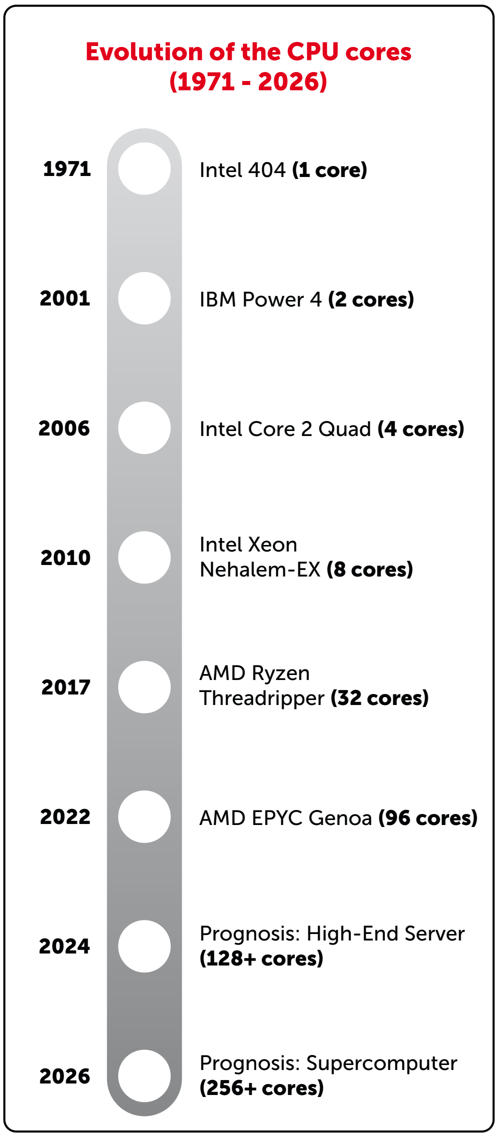

Figure 6: Evolution of the CPU

The development of central processing unit (CPU) technology from 1971 to the present day has brought about a significant change in server systems, characterised by several key advances.

The first microprocessors in the 1970s led to basic computing capabilities, while the importer of 16-bit and 32-bit architectures in the 1980s laid the foundation for more efficient data processing.

With the establishment of the 64-bit architecture in the 1990s, larger memory areas could be addressed, which was particularly relevant for server systems. At the same time, the first multi-core CPUs began to appear, which enabled several processes to be processed in parallel.

From the 2000s onwards, multi-core architecture became increasingly popular. The increasing number of cores in CPUs, combined with technologies such as Simultaneous Multi-Threading (SMT), led to a dramatic increase in parallel computing power in server systems. This development enabled better resource utilisation and optimised the processing of large volumes of parallel tasks, such as those found in databases, virtual machines and cloud services.

Since 2015, the many-core architecture has increased server performance exponentially by utilising processors with 64 or more cores. In addition, heterogeneous architectures have been established that integrate specialised computing units such as GPUs or AI accelerators to process workloads more efficiently. These advances are crucial for modern data centres and cloud applications that require high parallelism and specialised processing.

Overall, the increasing number of cores and parallel processing have significantly increased the efficiency and performance of server systems, which has enabled the development and scaling of modern applications and infrastructures.

GPU technology

A graphics processing unit (GPU) was originally designed to display graphics and relieve the CPU by rendering 2D and 3D graphics. However, due to its architecture with many parallel computing units, the GPU developed into an accelerator unit for a variety of computationally intensive tasks. Today, it is used in areas such as AI training, simulations, scientific calculations, data analysis and crypto mining. Due to its ability to perform large amounts of calculations in parallel, the GPU is no longer just essential for graphics, but for general computing-intensive applications.

These are mostly used in the field of High-Performance Computing (HPC).

Trends in CPU energy demand 2017 – 2026

To compare the energy requirements over several years, a specified amount of IT load, calculated in 2017 on a high-performance CPU at that time, was set as 100% power consumption. The curve falls sharply over the years under consideration. If the same IT load from 2017 were to be calculated on a modern system with a CPU with 256 cores in 2026, the computing operation would require approx. 70% less electrical energy compared to 2017.

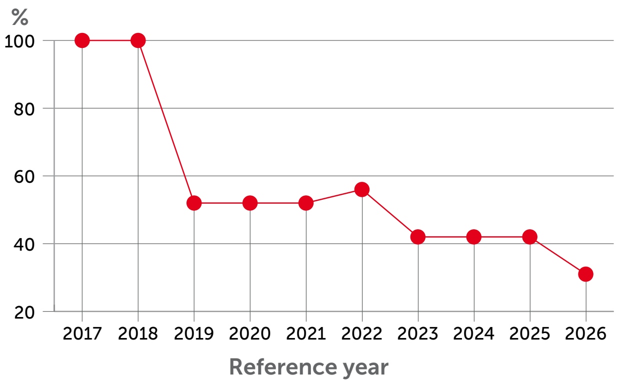

Figure 7: CPU energy demand 2017 – 2026 (© dc-ce)

The CPU server component was selected because, alongside the GPU, it is the component with the highest energy requirement in the server system. This analysis considers server systems whose GPU is only intended for displaying graphics output.

This example of the CPU could also be transferred to GPUs and memory chips such as RAM and SSD in the server in a similar form.

Trends in white space requirements 2017 – 2026

Here is the same example illustrated in terms of the white space required

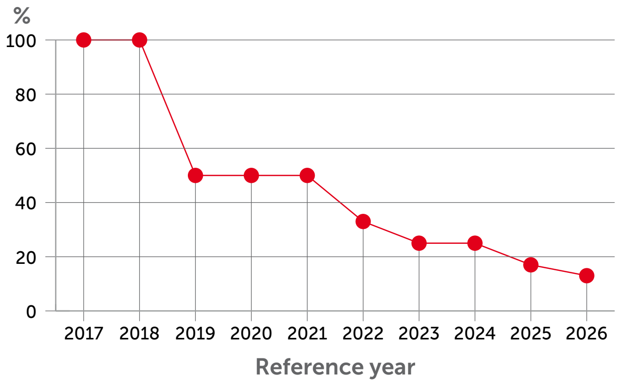

Figure 8: White space requirements 2017-2026 (© dc-ce)

If the space requirement for a specified IT load in 2017 is set at 100% and the same IT load is calculated on modern IT systems in 2026, the systems required for this will only take up just under 15% of the original space in 2017.

It is precisely this rapid change and further development, here only using the example of electronic components, that makes it clear that tried-and-true methods can only be used to a very limited extent in modern DC planning.

Added to this are the changes in the use of hybrid systems or the complete outsourcing of IT services; factors that can be attributed to strategic changes in usage behaviour and the way IT is handled.

Trends in future CPU performance development

Current observations and trends in microprocessor research and development show that the increase in clock rates is becoming more and more stagnant, but that the increase in performance is increasingly being achieved through optimisation in parallelisation, efficiency and flexibility. Technological innovations such as chiplet designs, heterogeneous architectures and specialised accelerators, together with advanced materials and 3D stacking technologies, are laying the foundations for further increases in CPU performance.

Step-by-step planning

For a data centre project, whether it is a conversion or a new build, a sound planning basis is required that plausibly and reliably describes the IT technology and service requirements over the entire planned operating life of the new data centre. It is inevitable that various assumptions will have to be made, which means that a certain amount of uncertainty will be incorporated into the planning. In this context, a distinction is made in the planning phase between calculable and incalculable factors. A goal is achieved when the final result has a plausible and reliable ratio of approximately 80/20. This provides a sufficient basis for the initiation of further project phases.

In addition, the need for migration and growth areas must be taken into account to implement subsequent adjustments efficiently. Additional security requirements for the IT capacities that influence the dimensioning of the essential design elements in the "white space" should only be recorded once, if at all. Modifications in later project phases should only be made in close consultation with the responsible technical architects.

If errors are detected or unforeseen changes occur during planning, it is essential to immediately check the impact on the data centre and its operational integrity. It must be evaluated whether the changes justify corrections and what adjustments are required. The earlier planning errors are identified, the easier and more cost-effective it is to correct them. In advanced project phases, the costs for corrections increase considerably, as project phases that have already been completed may have to be reworked.

If corrections are no longer possible, once an error has crept in, it can negatively impact the entire service life of the data centre and significantly impair the targeted profitability of the operation. This project phase should be regarded as particularly critical, as it lays the foundation for future data centre operations. This is important both from an economic perspective and with regard to legal requirements, which clearly define the operation of data centres through new regulations.

Classification of systems and racks

The electrical power consumption of server systems varies increasingly, which places specific requirements on a data centre's white space. In this context, server systems have been categorised into five performance classes as they are currently in use. This classification takes into account different performance profiles and their associated requirements for the data centre infrastructure.

When planning a new data centre, the focus is usually placed on classic and modern IT architectures. These systems typically include standard servers, storage units and network components that can be easily integrated into modern data centres. Specialised systems and legacy systems (older systems that are still in operation), on the other hand, require an individual requirements specification. A thorough validation of their requirements in terms of air conditioning, space requirements and infrastructure compatibility should, therefore, be carried out and taken into account when planning a data centre.

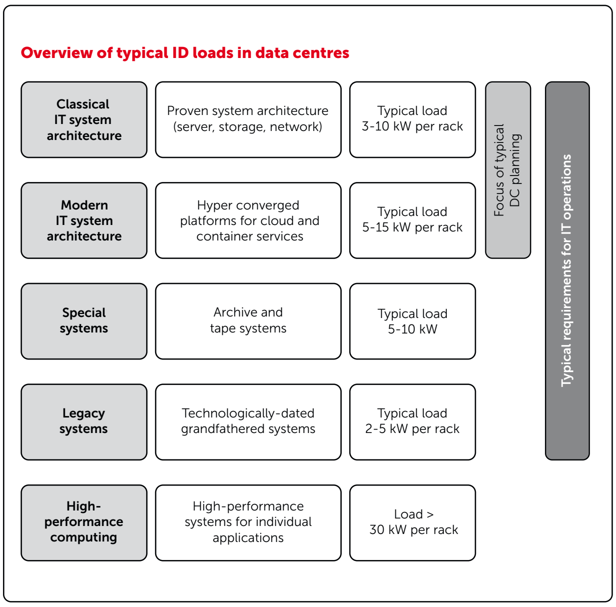

Figure 9: Typical IT loads in data centres

Classic IT system architectures

Classic system architectures in data centres include server systems, storage units and network components. These systems are typically implemented using standard 19-inch technology, which facilitates their integration into modern data centres. However, a key challenge of these architectures lies in the design of the air routing, as the supply and removal of air in the cold and hot aisles often has to be implemented using customised solutions in order to ensure efficient cooling.

The power requirement per rack varies depending on the hardware configuration used and can now range from 2 KW to 10 KW and, in certain cases, even more. These power requirements are largely dependent on the server density and the cooling solutions installed in the respective racks. Therefore, precise planning of air routing and cooling is crucial to avoid overheating and ensure the data centre's operational reliability.

Modern system architectures

The introduction of hyperconvergent systems has established a new class of hardware architectures characterised by the integration of computing power, storage and network resources. These systems are preferred in modern cloud and container services and can be used effectively in cold/warm aisle models in data centres. This architectural flexibility enables better adaptation to the requirements of dynamic IT environments.

The power requirement per rack varies depending on the specific expansion stage and the underlying platform architecture. Typically, this requirement is in the range of 5 kW to 15 kW but can also exceed this for more power-intensive applications or special configurations.

For optimised cooling and maximum energy efficiency, it is important to coordinate the data centre air conditioning and the IT systems through appropriate BIOS settings, particularly to adjust the fine-tuning of the air volume flows. This helps ensure efficient temperature regulation and increases the systems' overall performance and reliability.

Special systems

Special systems, such as solutions for backup and recovery, pose a particular challenge, especially in data centres that are dominated by SMEs. Classic tape systems require specific climatic conditions in terms of temperature and humidity, which often cannot simply be integrated into standardised cold or hot aisle models. These systems generally have a low power requirement in relation to the space required. The integration of such special systems in modern data centres, therefore, requires an adapted air conditioning strategy to ensure operational reliability and efficiency.

Legacy systems

The term 'legacy systems' usually refers to outdated systems that often run older legacy applications and are still integrated into the overall architecture of a data centre. These systems typically have a low power requirement of around 2 to 5 KW per rack. Due to their often non-standardised design, they frequently require individual housing, as they are not always 19-inch rack-mounted systems or may require special air routing for cooling. This poses a particular challenge for modern data centres that operate with cold/hot aisle systems, as the integration of legacy systems into these air conditioning models is often problematic. An appropriate solution, therefore, requires a customised air conditioning strategy and specific infrastructure measures to ensure operational stability and efficiency even when integrating older systems.

High-Performance Computing

The systems from the HPC sector represent a specific group, as both the performance requirements and the installation in modern server or system rooms generally pose a particular challenge. The power density of such systems, which are designed for specialised applications, can reach over 100 KW per system rack. A typical example of such systems are high-performance computing (HPC) systems, which are often used in areas such as research and science.

Due to the high heat development, systems in this performance class can only be operated with liquid-cooled solutions. This requires a special infrastructure in the data centre to ensure effective cooling and a safe environment. Due to their specific requirements, such systems are not discussed further in this white paper.

The integration and operation of these systems require precise planning and extensive adjustments to the infrastructure in order to ensure operational reliability and efficiency.

Flow chart for area and load determination

The flowchart below provides a guide to the main processes. The individual planning steps of the respective sub-processes are described in detail in the following chapters.

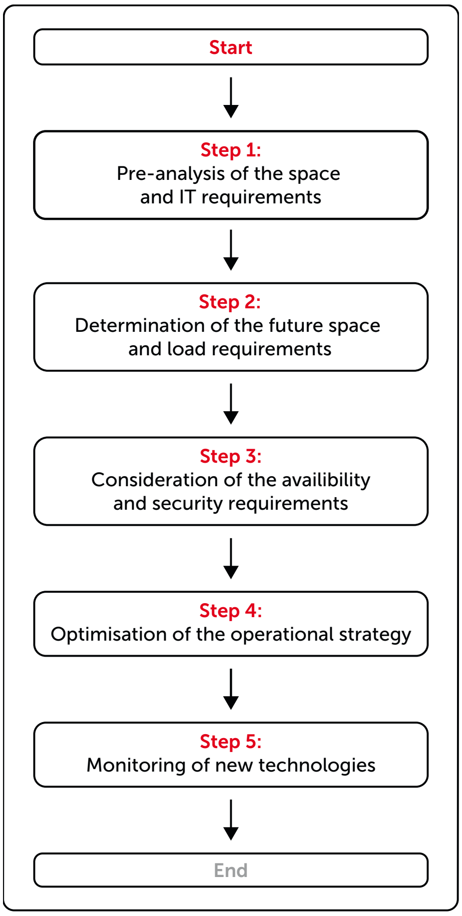

Figure 10: Flow chart for area and load determination

Step 1: Preliminary analysis of IT requirements

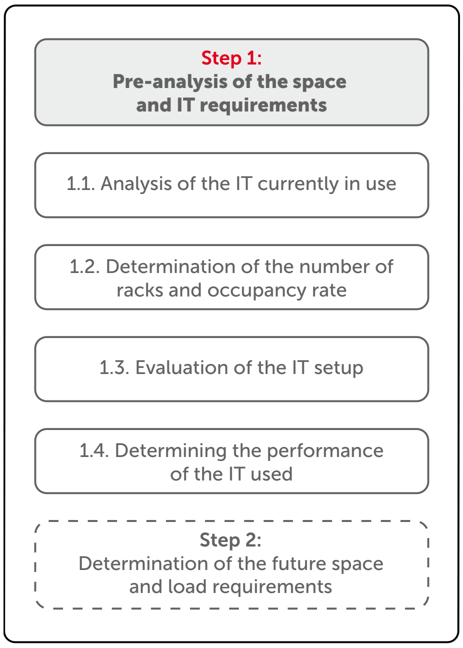

Figure 11: Step 1: Preliminary analysis of IT requirements

Analyse the IT currently in use

The detailed analysis of the IT infrastructure currently in use includes servers, switches, storage systems and other specialised systems. All relevant IT locations that are to be integrated into the new data centre must be included in the analysis.

Supplementing the classification of these systems according to criteria such as age, energy efficiency and utilisation provides additional insights to better assess the future requirements of these components and their potential replacement or decommissioning. This differentiated view enables precise planning and supports the optimisation of the IT infrastructure in terms of efficiency and future-proofing.

Determination of the number of racks and occupancy rate

When taking stock of the IT infrastructure, it is necessary to record the number of existing racks and determine their utilisation. This step involves determining the rack units (U) used in the racks and analysing their actual utilisation.

In addition, the racks should be evaluated for future-proofing to determine whether they are suitable for higher power density and the integration of future technologies. This assessment enables an informed decision about whether the existing racks can meet the requirements of a changing IT infrastructure or whether adaptations are required. This analysis thus contributes to long-term planning and investment security.

Evaluation of the IT setup

With regard to the timing of putting the new data centre into service, it is crucial to carry out a detailed analysis of which IT systems should definitely be transferred to the new data centre. Depending on the date of relocation, evaluation of which systems will be replaced by new technologies, and which will require renewal within the next one to five years should be carried out.

In addition, tasks and applications that can be moved to the cloud must be identified in order to enable the corresponding systems to be decommissioned on site. This strategic analysis helps to increase the efficiency of the IT infrastructure and ensure the long-term flexibility of the data centre. In doing so, consideration is given not only to whether the technology is up to date, but also to the potential for optimising resources through the use of cloud services.

Determining the performance of the IT used

The historical development of the power consumption of the existing IT infrastructure should be analysed using the data provided by the power distribution units (PDUs) or the uninterruptible power supply (USV). This analysis makes it possible to recognise long-term trends and evaluate the consumption patterns of the various IT components. The power densification of IT infrastructures following generational changes must be taken into account when analysing the data.

Based on this data, the average consumption per rack unit (U) should be calculated. For better differentiation, it is advisable to divide the results into three to five hardware categories. This categorisation is based on the age and type of systems and helps to evaluate the efficiency of the different hardware generations. This differentiation supports planning and makes it possible to make informed decisions about optimisations, renewals or upgrades in the IT infrastructure.

Step 2: Determination of future space requirements

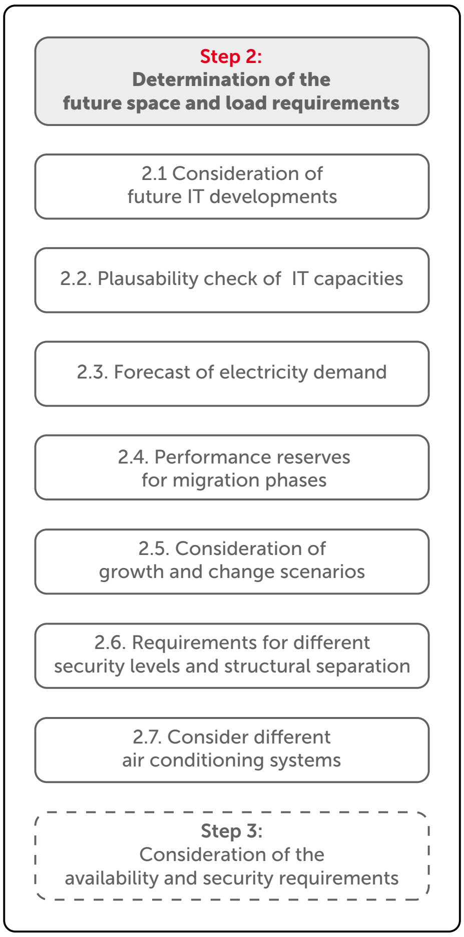

Figure 12: Step 2: Determination of future space requirements

Consideration of future IT developments

The evaluation of historical data on the development of space and energy requirements makes it easier to estimate future IT requirements until the new data centre is put into service. This data analysis helps to identify trends in IT utilisation and to make a well-founded forecast of future resource requirements. In addition, potential workloads that can be moved to the cloud should be identified in order to relieve the burden on local infrastructures and increase flexibility.

In addition, it is crucial to plan a future-proof IT architecture that not only meets current requirements but also considers future technological developments. To ensure this, regular scenario analyses should be carried out to evaluate the potential development of new IT business areas. These scenarios provide a basis for flexible planning and the ability to respond adequately to unknown or future requirements. Proactive, adaptable planning thus increases the future-proofing of the data centre and enables the effective use of resources.

Plausibility check of IT capacities

It must be ensured that the planned capacities for computing power, storage and network fully cover the forecast requirements. This requires calculating the resulting rack units (U), which should then be converted into a corresponding number of racks.

In this planning, it is also important to take into account temporary space requirements for possible migration processes. This additional space is necessary to ensure a smooth transition from old to new systems without disrupting ongoing operational processes. Careful planning of physical capacities enables optimised space utilisation and ensures the flexibility of the data centre during the migration phases.

Forecast of electricity demand

The use of configuration programmes and forecasting tools helps to precisely determine the power requirements of the newly planned systems. It is crucial to bear in mind that power consumption is becoming increasingly dynamic and that it is possible to switch off entire systems if necessary (load balancing at hardware and virtual level) in order to achieve energy savings.

These plans should be supplemented by load management strategies that incorporate seasonal and time-of-day power fluctuations. Such fluctuations can be balanced out with flexible electricity and air conditioning supply systems. These approaches allow the supply to adapt to changing power requirements in line with demand, thereby increasing operating efficiency and minimising energy wastage. Proactive load management thus contributes to the sustainable and operational reliability of the data centre.

Performance reserves for migration phases:

When planning transition phases in which systems have to be operated in parallel, sufficient reserve capacity should be planned to ensure a smooth migration and avoid bottlenecks. These reserve capacities are crucial to ensure stable and uninterrupted operation during the parallel operating phases.

It also makes sense to examine whether temporary power supply solutions could be a more cost-efficient alternative. Such temporary solutions offer the opportunity to avoid the high costs of maintaining oversized capacities in the long term.

Consideration of growth and change scenarios:

To enable flexible expansions in the data centre, it is essential to differentiate between secured and uncertain space growth. This differentiation supports sound planning and helps to adapt the infrastructure to realistic requirements. It should be borne in mind that IT performance often increases faster than the associated additional capacity requirements.

One promising approach could be the use of modular systems that allow IT capacities to be expanded at short notice without the need for major infrastructural changes. These modular concepts offer the necessary flexibility to react quickly and efficiently to changing requirements and help to optimise operational efficiency. The implementation of such flexible expansion options also ensures the future-proof nature of the data centre.

Requirements for different security levels and structural separation

It should be checked whether a structural separation of certain IT components is necessary, based on the different security requirements. This applies in particular to special systems such as conveyor robots or other critical IT components that may have higher requirements in terms of safety and protection.

Different levels of security within the data centre may require physical separation into separate areas. Such separation enables compliance with specific security requirements and minimises risk by physically isolating sensitive systems. These measures help to ensure data security and protect critical infrastructure.

Consider different air conditioning systems

When planning air conditioning systems, it is crucial to take into account the different temperature and humidity requirements of the future IT systems. These requirements can vary depending on the system and influence the choice of suitable cooling solutions.

It should also be evaluated whether air-cooled or water-cooled servers are a better solution, depending on the expected power density of the IT components. The decision for one of the two technologies should therefore be based on a detailed analysis of the expected heat development and long-term operating costs in order to maximise the efficiency and sustainability of the data centre infrastructure.

Step 3: Consideration of availability and security requirements

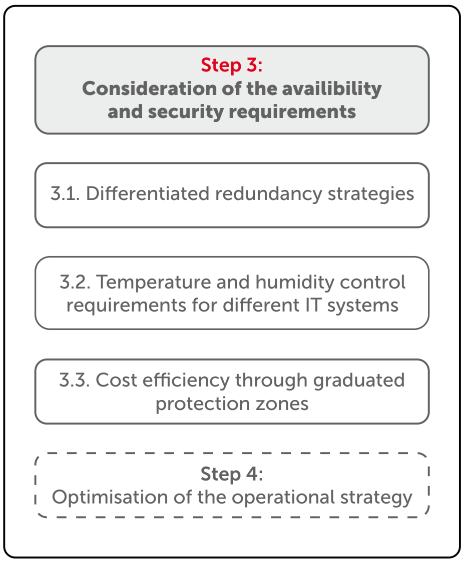

Figure 13: Step 3: Consideration of availability and security requirements

Differentiated redundancy strategies

The redundancy requirements from the security concept (DC policy) of the data centre must be taken into account. These requirements define the necessary measures to ensure operational reliability even if critical components fail.

As part of the planning process, it should be checked whether different redundancy concepts, such as n-redundancy for non-critical systems (e.g. systems for artificial intelligence or high-performance computing), are sufficient. With n-redundancy, the infrastructure is operated without additional redundancy, which may be acceptable for systems with less critical requirements.

For business-critical applications, on the other hand, proven redundancy concepts such as n+1 or n+n should be used. These concepts offer an additional level of security by providing either one component (n+1) or complete duplication (n+n) to maintain operation even if one component fails. The choice of the appropriate redundancy concept depends on the criticality of the applications and the security requirements and helps to minimise risk and maintain operations.

Temperature and humidity control requirements for different IT systems

A comprehensive review is required for different IT systems, such as HPC or AI, which may have different temperature and humidity requirements. In particular, it should be checked whether the use of water-cooled IT hardware makes sense in the future to ensure more efficient cooling. This review is necessary to reduce the operating costs of air conditioning through improved energy efficiency.

In addition, it is advisable to consider the use of efficient cooling techniques, such as direct or indirect free cooling or the use of water-cooled servers in certain areas of the data centre. These techniques can make a significant contribution to reducing energy consumption. A differentiated selection and planning of these cooling systems, depending on the respective system requirements, enables the optimisation of air conditioning costs and supports the sustainable design of data centre operations.

Cost efficiency through graduated protection zones

The safety and protection requirements from the safety concept (DC policy) of the data centre must be taken into account. These requirements define the necessary measures to ensure physical and IT protection.

In addition, the possible economic advantages that can result from the establishment of tiered safety zones should be considered. Applications with supposedly lower protection requirements, such as HPC or AI, could be housed in less heavily secured areas of the data centre. This would make the security infrastructure more efficient and reduce costs.

Differentiated security planning enables needs-based protection and optimises the use of resources without jeopardising the safety of critical applications.

Step 4: Optimisation of the operating strategy

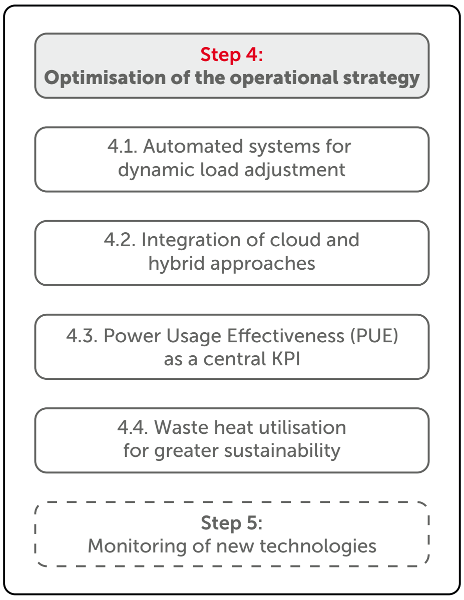

Figure 14: Step 4: Optimisation of the operating strategy

Automated systems for dynamic load adjustment:

Implementing automated control systems that are able to react dynamically to load changes in the data centre is recommended. These systems should be designed in such a way that they can automatically switch capacity on or off as required in order to efficiently optimise both air conditioning and energy supply. The automation of these processes enables demand-based control, which avoids overcapacity and minimises energy consumption.

Control systems should be expanded to include machine learning-based algorithms that can analyse load trends and optimise them with foresight. These algorithms should use historical data to recognise patterns and make precise predictions about future load peaks or Modifications. This allows maximum efficiency of data centre operations to be achieved by dynamically adapting resources to the respective requirements. This combination of automation and predictive analysis promotes sustainable optimisation of energy consumption and operational processes.

Integration of cloud and hybrid approaches

The integration of cloud and hybrid infrastructures enables data centres to respond more flexibly to peak loads. If necessary, workloads that require high computing power or where availability is crucial can be outsourced to the cloud to relieve the local infrastructure. This reduces the load on on-site resources and ensures stable performance.

Strategies should be developed that enable the intelligent distribution of workloads between local data centres and the cloud. Such distribution helps to reduce costs, as resources can be utilised more efficiently and overcapacity can be avoided. At the same time, scalability is maintained so that the data centre can react flexibly to changing requirements. This combination of local and cloud infrastructure offers an optimal balance between flexibility, efficiency and cost control.

Power Usage Effectiveness (PUE) as a central KPI

The Power Usage Effectiveness (PUE) value describes the ratio between the total energy consumption of a data centre and the energy actually used by the IT systems. It serves as an indicator of energy efficiency and was defined as a reference value for larger data centres with a nominal output of over 300 KW as part of the German Energy Efficiency Act. Compliance with these requirements is legally binding and is intended to ensure efficiency standards in data centres.

The PUE value should be used as the primary KPI for evaluating the energy efficiency of a data centre. A lower PUE value indicates a more efficient use of energy, as in this case a higher proportion of the energy consumed is used directly for the IT systems. The aim when optimising energy consumption should, therefore, always be to reduce the PUE value as much as possible to increase operating efficiency and reduce the impact of the data centre on the environment.

Waste heat utilisation for greater sustainability

Finding ways to utilise the waste heat from the data centre for other applications, such as heating offices or other buildings. This measure helps to improve the energy efficiency of the entire operation and can significantly reduce operating costs. The German Energy Efficiency Act stipulates the use of waste heat as a reference value for larger data centres with a nominal output of over 300 KW. Compliance with these requirements is required by law in order to promote efficiency standards and reduce environmental impact.

In addition, the feasibility of coupling the data centre with district heating networks should be examined, or the use of waste heat for industrial processes in the environment. Such integration could further promote the sustainability of the data centre by efficiently feeding and utilising the waste heat into external systems. This strategic orientation contributes to the conservation of resources and strengthens energy efficiency at regional level.

Step 5: Monitoring new technologies

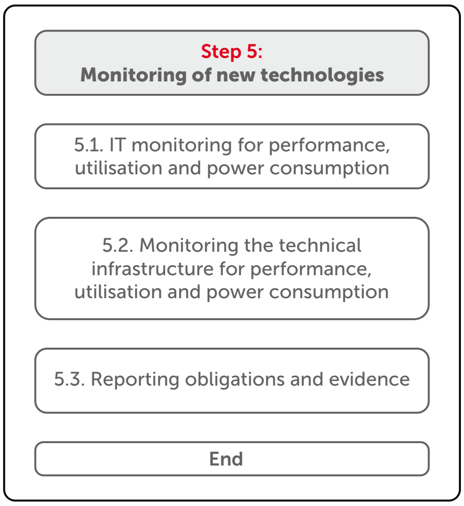

Figure 15: Step 5: Monitoring new technologies

IT monitoring for performance, utilisation and power consumption

Efficient operation of IT systems in data centres means achieving the highest possible utilisation of the individual systems without risking performance losses in data processing. To ensure this, comprehensive monitoring and control mechanisms are required that cover the various infrastructure levels of the data centre.

There are several approaches to this:

- Centralised fabric architecture: A centralised fabric architecture allows for the control of the computing, storage, and network infrastructure. This architecture provides an automated management level that enables dynamic adjustment and monitoring of resource utilisation in real time.
- Control at the hypervisor level: At the virtualisation level, the hypervisor takes over the management of IT resources. Virtualisation allows computing loads to be flexibly distributed and optimised in order to maximise resource utilisation while maintaining system performance.

- Dedicated monitoring tools: Integrated monitoring and control tools enable continuous analysis of system utilisation. These tools provide comprehensive insights into the IT infrastructure's performance, energy efficiency and security status and sometimes allow adjustments to be made in real time.

Monitoring the technical infrastructure for performance, utilisation and power consumption

Monitoring the technical infrastructure in data centres is necessary for several key reasons. The German Energy Efficiency Act stipulates a reporting obligation that requires comprehensive recording and documentation of energy consumption. In addition, continuous monitoring contributes to the data centre's operational safety. By recognising deviations from normal operation at an early stage, potential failures can be avoided and the availability of the systems can be ensured. Monitoring also makes it easier to identify and implement optimisation potential during operation, which leads to greater efficiency and cost savings.

Options for monitoring the DC infrastructure:

- Intelligent power distribution units (PDUs): Modern PDUs enable energy consumption to be monitored at the level of individual server racks or devices. Many of these PDUs are network-enabled and provide detailed data on power consumption, enabling precise control of energy flows.

- Data centre infrastructure management (DCIM) systems: DCIM solutions provide comprehensive management of the data centre's physical infrastructure. These systems monitor not only energy consumption but also temperature and humidity. They, therefore, contribute to the optimisation of the entire data centre infrastructure and help to control energy consumption efficiently.

- Building automation (BA) and building management system (BMS): Monitoring and analysing systems such as USV (uninterruptible power supply) systems, other electronic systems and HVAC systems (heating, ventilation, air conditioning) enables holistic monitoring of the infrastructure. By analysing the data from these systems, trends and potential weak points can be identified, enabling proactive maintenance and optimisation.

Reporting obligations and evidence

Reporting on energy consumption

Data centres with a nominal connected load of at least 300 kilowatts are required to document detailed data on their electricity consumption and its origin and to report this information to the relevant public authorities. From 2024, at least 50% of the electricity consumed must come from renewable energy sources and from 2027 this proportion must be increased to 100%.

Efficiency reports

Deployers of data centres are obliged to prepare regular efficiency reports on the energy consumption efficiency of their facilities. In particular, the key figures for Power Usage Effectiveness (PUE) must be documented and made available to the relevant public authorities. The PUE value serves as a benchmark for energy efficiency and enables an evaluation of the effectiveness with which energy is utilised for IT systems.

Waste heat data

Data centres that have an annual energy consumption of at least 2.5 GWh are additionally required to submit annual reports on the use and avoidance of waste heat. These reports are intended to provide the competent public authorities with information on the amount of waste heat generated and the measures taken to avoid or utilise waste heat.

Calculation example

This sample calculation considers a 2014 data centre with a total output of 285 KW. The data centre's white space has 148 racks, covering an area of around 600 m². Due to changes in the neighbourhood, operational security could no longer be guaranteed to the extent originally planned, so it was decided that a new data centre would be built at an alternative location.

The data centre is currently operating close to its capacity limits in terms of energy supply and air conditioning. This leads to the creation of hotspots in the white space. Most of the racks are distributed without enclosures, which makes uniform cooling difficult. The appearance of the white space is also characterised by the heterogeneous use of the racks. Some are unused, while others contain obsolete hardware that has already been switched off. This inconsistent arrangement and utilisation affect efficiency and temperature control in the data centre.

Initial situation: Capacity of the legacy white space

Rack units (U)	Racks	Electrical performance	White space area
2928	148	272 kW	591 m²

Table 2: Legacy white space capacities

Step 1: Preliminary analysis of IT requirements

Analyse the IT currently in use

Type	VM	Server HW	Net-work	Storage	Special System	Age of HW	Assign-ment (U)	Electrical power [kW]	System utilisation
Mail server	X								
Mgmt-Sys.	X								
SAN-01, -02				X		9	260	21	80
Databases 1		X				6	173	16	40
Core network			X			4	99	6	
Network			X			10	166	3	
SAN-03					X	4	80	12	65
Databases 2					X	2	280	28	20
Tape library					X	10	430	3	100
Appl.	X								
Various app servers		X				5-10	98	5	23
Specialised APP servers					X	Partial >10	191	10	11
...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...
Priv.Cloud-01		X				9	180	21	42
Priv.Cloud-02		X				1	108	15	54
Total							2928	272	

Table 3: Record of the hardware currently operated in the legacy data centre

First of all, the IT hardware currently in use was recorded and analysed as comprehensively as possible in various categories in the table shown above. The hardware components were systematically categorised in order to assess their technical condition, performance and energy consumption.

When moving IT systems, it is helpful to categorise the systems in question as to whether they can be mounted in standard 19-inch racks or whether they require manufacturer-specific special racks.

The server cabinets and racks in the white space of the data centre were also recorded and classified in a similar way. The aim of this survey is to create a basis for the harmonisation and optimisation of space utilisation, cooling strategies and capacity planning.

Rack type	Number (racks)	Number (U)	Assignment (U)	Utilisation (%)
42 U, 19" Std.	72	3024	1653	55
52 U, 19"	22	1144	522	46
32 U, 19"	15	480	126	26
...	...	...	...	...
Special racks	36	829	627	76
Total	148	5477	2928	53%

Table 4: Determining the number of racks and utilisation rate

Evaluation of the IT used

The analysis of the defined server categories yielded the following results:

Some of the classic IT systems are already more than five years old and are about to be replaced, which will take place at the latest during the move to the new data centre. In contrast, modern IT architectures, such as the Private Cloud, are already in place and planned for the move. A possible expansion of this architecture is being considered.

The SAN systems currently in use are outdated and will be replaced by high-performance and energy-efficient models. With regard to the database system, which is relatively modern and powerful, the analysis shows a utilisation rate of just 20%. To increase efficiency, this system is to take on additional workload and it is planned to gradually expand capacity. To this end, a forecast of future growth is planned and corresponding resources in the white space are provided in the form of free rack units (U) in the racks, electrical performance and cooling at the forecast times.

Special systems that are marked as critical are analysed separately. The ten-year-old tape library will be replaced by modern archiving systems. To ensure data integrity, a concept for the migration of tapes from the long-term archive must also be developed.

Determining the performance of the IT used

The power consumption of the respective IT systems could be largely determined via the individual PDUs. In cooperation with the IT departments, it was possible to record the IT utilisation for key systems. These were entered in the above 'Table 3: Record of the hardware currently operated in the legacy data centre'. These results are used for further decisions.

Conclusion Step 1

The results of the analysis show that the data centre's electricity consumption increased by 30% over the period of operation when historical data is taken into account. At the same time, the space required in the data centre decreased by over 25 % during the period under review. Almost 50 % of the rack units (U) in the racks are unused, which indicates a sub-optimal utilisation of resources. Full capacity utilisation was not achieved at any time.

These findings from the past highlight the need for action to optimise the use of resources and the need for strategic planning of the future IT infrastructure.

Step 2: Determination of future space requirements

Consideration of future IT developments

As part of the analysis, it was determined that some applications are no longer required or are to be replaced by new ones. Although these modifications are not part of the project, they indicate that the expected resource requirements remain essentially constant.

Growth in storage resources is forecast at 10% per year. At the same time, the mailboxes are to be outsourced to the public cloud, which will relieve the burden on the private cloud and reduce internal resource requirements.

In addition to the existing private cloud, which is intended for the move, a new system will be introduced which is intended for virtualisation as well as containerisation. The previous SAN systems will be completely replaced by new and high-performance models.

Applications were also identified that are particularly suitable for containerisation. In addition, a new application that requires a high level of resources and will be operated in a fully virtualised manner is planned. These expansions and modernisations aim to achieve a sustainable increase in the efficiency and flexibility of the IT infrastructure.

Conceptual planning of new IT systems

New IT systems are designed based on the current workloads, the forecast growth rate and the additional applications planned. The aim is to ensure that the future workload can be supported in all areas in combination with the systems planned for the move.

The new IT systems will be designed to meet both current requirements and the expected increases in capacity. This is ensured through the use of state-of-the-art hardware, which should maximise efficiency and performance. The aim of this optimisation is to make the operation of the systems as effective as possible while at the same time being prepared for future workload expansions.

The use of advanced technologies should both increase energy efficiency and ensure flexible scalability for future developments. This strategic approach supports the sustainable use of resources and minimises the operating costs and construction costs for the new data centre.

Plausibility check of IT capacities

As part of the audit, the systems that are to be transferred from the existing legacy data centre to the new data centre are identified. At the same time, a list of the new systems to be procured will be drawn up, which will have the status of the most modern technologies available on the market in 2025.

This review aims to create an efficient and sustainable IT infrastructure by replacing outdated systems with high-performance and highly efficient technologies wherever possible. The selection is based on the goal of equipping the new data centre with the most advanced systems available that meet the requirements in terms of performance, virtualisation, safety and scalability.

Calculation

The hardware systems that are to be migrated 1:1 to the new data centre include all existing IT components that must continue to be operated for operational reasons or that meet the technical requirements for the new DC. These systems were carefully evaluated.

Designation	Server HW	Special Systems	Age of HW (years)	Utilisation	Electrical power [kW]	System utilisation
Data bases 2		X	2	280	28	20
Priv.Cloud-02	X		1	108	15	54
...	...	...	...	...	...	...
Dedicated APP server		X	Partial >10	79	3	11
Total				515	46	

Table 5: Hardware that will be migrated 1:1 to the new data centre

The network is being completely redesigned as part of the data centre relocation to meet future requirements. This redesign includes all network components, including switches, routers, firewalls and other connectivity systems, with the aim of creating a modern, scalable and high-performance infrastructure.

Designation	Utilisation (U)	Electrical power [kW]
Core network	56	6
Network	42	2
...	...	...
Total	132	9

Table 6: Planned network hardware for the new data centre

The new IT systems in the data centre are planned on the basis of future workloads and the increasing requirements for performance, safety and efficiency. These systems incorporate state-of-the-art technologies and are designed to optimise performance, scalability and energy efficiency.

Designation	19" installation	Utilisation (U)	Electrical power [kW]	Forecasted system utilisation (%)
Priv. Cloud-03	X	42	22	80
Storage, expandable	X	32	8	60
...	...	...	...	...
...	...	...	...	...
Archiving, expandable	X	16	4	60
Total		122	41	

Table 7: Planned server hardware for the new data centre

When planning the new data centre, it is crucial to plan reserve space for migration and future growth.

When planning reserve space and the future IT infrastructure in the data centre, the progressive development of performance densification in IT systems must be taken into account, especially with regard to the upcoming generational changes in hardware. .

Designation	19" installation	Utilisation (U)	Electrical power [kW]
Reserve area	X	84	25
Total		84	25

Table 8: Planned reserve and migration space for the new data centre

Calculation bases and requirements

The calculation carried out considers various availability and security requirements in order to ensure the stability and reliability of the new data centre. Redundancy strategies were analysed and specifically covered by the private cloud systems to ensure a high level of operational reliability.

Measures for compliance with the requirements

Enclosures for the old systems: Special enclosures were designed for the legacy systems still in use to meet their specific requirements and ensure reliable heat dissipation. This measure ensures stable operation despite the use of older hardware.

Integration of modern enclosures: Modern IT systems have been integrated into optimised enclosures tailored to current cooling and security requirements. These enclosures enable targeted and efficient air conditioning of the IT components.

Load management and automation: The new systems were analysed in terms of their load fluctuations and adapted accordingly to the typical load profiles. By using automated control systems, the systems can react flexibly to fluctuations and optimise their performance.

Monitoring and control centre: A central control centre was set up, which is connected to the monitoring system and ensures round-the-clock (7*24) monitoring of all critical components. This monitoring enables the early detection of faults and proactive measures to prevent failures.

This comprehensive approach and implementation ensure that the technical requirements of the new data centre are fully met while guaranteeing high availability and operational efficiency.

Calculation result

After adding up the partial results, the capacities for the new white space in the new data centre are as follows

Rack units (U)	Racks	Electrical power [kW]	White space area (m^2)
853	22	121	118

Table 9: Planned capacities for the new data centre

Conclusion Step 2

The takeover of the legacy systems on a larger scale limited the optimisation of the data centre and had a significant impact on the infrastructure. When selecting new IT systems, care was taken to ensure that they were state-of-the-art

If it were possible to completely replace or virtualise the legacy systems, the space and power requirements could be significantly reduced. From this perspective, the legacy systems require 46 KW of energy and occupy 515 rack units (U), which corresponds to a volume of around 15 racks in white space. Such a measure would enable significant optimisation of the data centre, both in terms of space utilisation and energy consumption.

List of abbreviations

- AI: Artificial intelligence
- APP server: Application server
- BA: Building automation
- BMS: Building management system
- CPU: Central Processing Unit (processor)
- CRM: Customer Relationship Management
- DC: Data centre
- DCIM: Data Centre Infrastructure Management
- ERP: Enterprise Resource Planning
- GDPR: General Data Protection Regulation
- GPU: Graphics Processing Unit (graphics processor)
- HPC: High-Performance Computing
- HVAC: Heating, ventilation, air conditioning
- ICT: Information and communication technology
- KPI: Key Performance Indicator
- KRITIS: Critical infrastructures
- NIS: Network and Information Security
- PUE: Power Usage Effectiveness
- TBE: Technical Building Equipment
- U: Rack unit
- USV: Uninterruptible Power Supply
- VM: Virtual Machine

About eco

eco shapes the Internet

With around 1,000 member companies, eco is the leading association of the Internet Industry in Europe. Since 1995, eco has been instrumental in shaping the development of the Internet. eco is committed to an efficient, reliable and trustworthy ecosystem of digital infrastructures and services and has stood for an Internet with responsibility for 30 years.

Internet with responsibility

Together with our members, we are committed to a free, technology-neutral, network-neutral, and high-performance Internet. We thereby want to support the security and the reliability of the Internet, as well as build trust in it. Our goal is to shape the digital transformation of society and the economy in the best possible way so that successful economic action can succeed on the basis of our democratic values.

As the voice of the Internet industry, we assume societal responsibility for ethically oriented digitalisation.

eco connects across industries

Digital transformation is permeating more and more areas of our lives. This is also expanding the spectrum of our association work – and not only in terms of content. It is essential that we shape progress and change together!

Companies in the information and telecommunications industry need to interact closely with traditional industries in order to create a sustainably functioning digital ecosystem. eco serves as a neutral platform which allows points of view, goals and concerns to be discussed on an equal footing. We bring our members and industry stakeholders into dialogue with academia, society and politics.

eco creates standards

In our expert and competence groups, you will find the ideal platform for exchange on current and future Internet topics and for the further development of these topics.

Through its close connection to DE-CIX, eco itself is a part of the industry and an active shaper of digital transformation, with strong technological core competencies in the fields of infrastructure and security.

Seals of quality developed by eco set standards and make the market more transparent for providers and users. These standards sustainably strengthen the Internet and digital industry as the engine of the overall economy. eco's advisory services for members and its services for Internet users offer support on legal issues, increase security and improve youth protection.

Data centres

The eco Data Center Expert Group and the Data Center Infrastructure and Data Center Efficiency Competence Groups handle all topics to do with data centres. A particular focus is applied to the further rollout of broadband connectivity, outsourcing of data, computation time in data centres, and topics such as IT performance and energy efficiency.

international.eco.de/datacenter

Datacenter Star Audit (DCSA)

Since 2005, eco-authorised auditors have been objectively assessing the infrastructure and services of data centers through means of the DCSA – with this being implemented increasingly throughout Europe. The DCSA is suitable for any company that operates a data center/server room or who avails of colocation.

Regardless of your business model or the size of your IT area, the DCSA offers you the opportunity to transparently and understandably present your redundancy concept, as well as your quality, security, and availability. The audit is performed in a timely manner at a manageable financial cost. The seal of quality increases your trustworthiness vis-à-vis customers, auditors, banks and insurance companies.

Further information and contact:
international.eco.de/topics/datacenter/dcaudit
info@dcaudit.de

Notes

Orientation Guide
to Power Density and
Load Determination of Servers,
Data Cabinets and Data Centers

eco – Association of the Internet Industry
Lichtstrasse 43 h
50825 Cologne, Germany
Phone: +49 (0) 221 7000 48 0
Email: info@eco.de
Web: international.eco.de

eco
ASSOCIATION OF THE
INTERNET INDUSTRY