Vendor-Neutral Tendering of x86 Servers

Guideline for Public Procurement
As of: December 2023 | Version 1.0

Guideline

bitkom

Published by
Bitkom e. V.
Albrechtstraße 10
10117 Berlin
Phone: +49 (0)30 27576-0
bitkom@bitkom.org
www.bitkom.org

Contact
Esther Steverding | Bitkom
Phone: +49 (0)30 27576-216 | e.steverding@bitkom.org

Responsible Bitkom committee
Expert Committee Vendor-Neutral Tendering

Project management
Esther Steverding | Bitkom

Layout
Lea Joisten | Bitkom

Title image
© Evgen3d – stock.adobe.com

Content

Acknowledgements

The present guidelines is the result of intensive collaboration between experts of public administration and representatives of partner companies of Bitkom. It owes its existence to the comprehensive groundwork of the Project Group »Product-Neutral Performance Description x86 Servers«. We would like to express our particular gratitude to:

- Peter Dümig, Dell GmbH
- Volker Kwasnicki, Lenovo
- Dr. Heiner Genzken, Intel Deutschland GmbH
- Jan Gütter, AMD GmbH
- Stefan Kreger, Fujitsu Technology Solutions GmbH
- Patrick Mitternacht, Bechtle AG
- Thomas Zapala, German Federal Office of Bundeswehr Equipment, Information Technology and In-Service Support [Bundesamt für Ausrüstung, Informationstechnik und Nutzung der Bundeswehr, BAAINBw]
- Esther Steverding, Bitkom e. V.

1 Introduction

1.1 Using this guideline

This guideline provides an overview of the foundations and criteria for the procurement of servers by contracting entities. It is the product of a working group under the chair of the German Ministry of the Interior's Procurement Office and the Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e. V. (Bitkom). This document aims to provide contracting entities of Germany's federal, state, and local governments with a dependable tool – that is easy to understand – to help them formulate their tenders for the procurement of servers in a vendor-neutral manner free of discrimination, i. e. without the use of trademarked names and without mentioning individual manufacturers, while taking into consideration current technological standards.

At the heart of this guideline stands the list of technical criteria, which can be used to describe and compare the devices themselves as well as requirements placed on their operational environment together with other properties. Besides the technical criteria, compliance with which guarantees proper device functioning for the reason they were procured, the guideline also contains references to environmental protection, energy efficiency, and IT security. Even if statutory requirements only partly obligate procurers to observe these interests, they are of key importance to public administration.

Nevertheless, it is pointed out that the listed technical criteria and requirements are subject to continuous change, and are to be weighted differently on the basis of the intended application of the devices to be procured. The more discerning the requirements placed on the product, the higher the offer price will generally be. Moreover, there will be fewer suitable products available on the market, which in turn might restrict competition and result in higher prices. This guideline cannot replace independent and reasonable assignment and prioritisation of the relevant criteria by the procuring entity.

The authors of this guideline also seek to support contracting entities by making special reference to sensitive criteria and requirements – i. e. those that might lead to market restrictions –, as well as decisions that are relevant to the costs. The symbols defined in the following are used:

Symbol	Bedeutung
	The requirement of criteria with this symbol can lead to cost increases and / or market restrictions.
	This symbol indicates the correction of a widespread error or marks particularly important statements in the text.
	This symbol indicates whether criteria can be verified with certificates (in German: Zertifikate).

1.2 Scope of the guideline

The present guideline exclusively concerns server hardware and software compatible with the x86 instruction set architecture and its extensions. Express reference is made that there are additional system architectures besides the x86 server architecture. Even with some aspects of their application cases overlapping, there are significant differences that would make a common classification into a uniform guideline significantly more difficult. To avoid complicating the use and handling of this guideline in the creation of tender documents, a conscious decision has been made to restrict its scope to the x86 server architecture.

The following server architectures are **not** compatible with the x86 instruction set architecture and therefore do **not** fall under the scope of this guideline.

- IBM POWER
- ARM
- MIPS
- SPARC

Furthermore, this guideline does not cover any HPC (high-performance computing) servers or services procured over the cloud.

1.3 Vendor-neutrality as a legal requirement

Procurement law stipulates an obligation to equal treatment, non-discrimination, and transparency of providers and products offered in the tender procedure. The aim is to ensure fair competition between economic players. German procurement law is an expression of the European Procurement Directive (2014/24/EU) and is (as concerns the scope of this guideline) implemented in German law through the Act against Restraints of Competition (GWB) and the Ordinance on the Award of Public Contracts (VgV). The regulations of the Procurement and Contract Procedures for Services (VOL), the Regulation on Awarding Contracts in the Areas of Transport, Drinking Water, and Energy Supply (SektVO), the Regulation on Concessions (KonzVgV), the Regulation on the Awarding of Public Contracts in Defence and Security (VSVgV), as well as the Regulation on Sub-threshold Procurement (UVgO) also apply.

In accordance with these legal foundations, the subject of the procurement is to be described using professional criteria free of discrimination, i.e. in a product-neutral manner (cf. § 97 Act against Restraints of Competition (GWB) and § 31 para. 6 Ordinance on the Award of Public Contracts (VgV) for tender procedures across the EU and § 2 para. 2 Regulation on Sub-threshold Procurement (UVgO) for sub-threshold procurement). Certain product descriptions or brand names may only be used in tenders in duly justified exceptions, if a sufficiently exact description using common descriptions or general criteria is not possible.

This guideline specifically addresses this problem, by providing a compact tool supporting compliance with the legal requirements, thus ensuring fair competition. The guideline specifies and explains current technical standards to describe servers using general and pertinent characteristics. The guideline will be updated regularly, taking into consideration new developments in technology while aligning the proposed criteria and requirements with the current state-of-the-art.

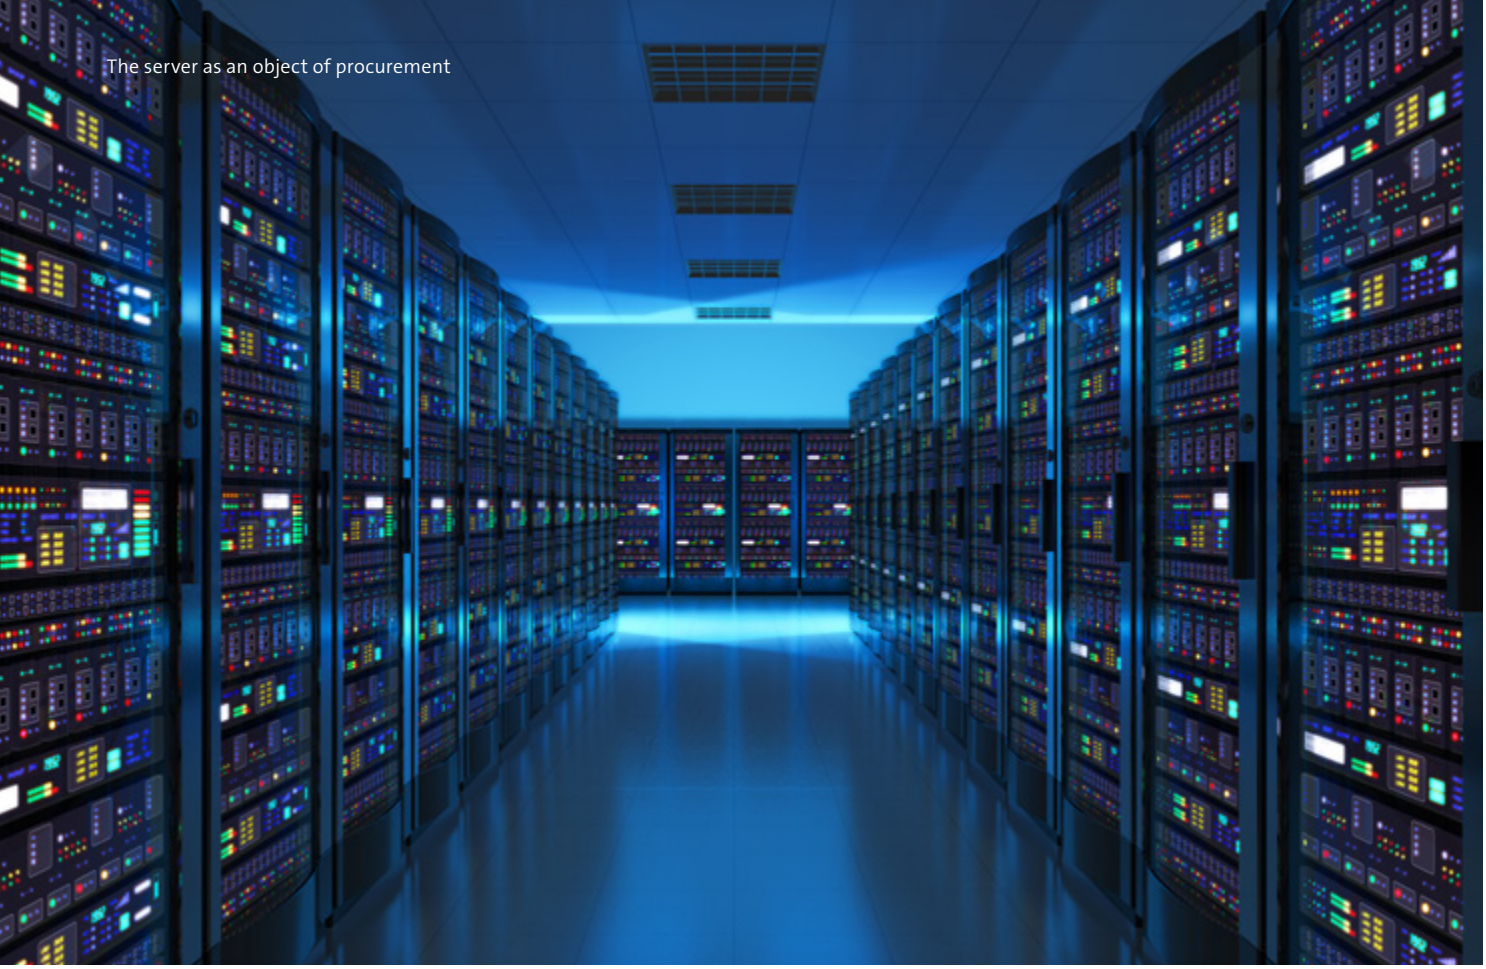

2 The server as an object of procurement

2.1 Definition of a server

A server is a data processing device that provides services and manages network resources for client devices such as desktop computers, notebook computers, thin clients, Internet protocol phones, smartphones, tablets, telecommunications, automated systems or other servers. Access to a server is mainly via network connections and not directly via user input devices such as a keyboard or mouse.

Moreover, servers are defined by the following properties:

- They are designed to support server operating systems and / or hypervisors (for virtualisation applications) and run user-installed enterprise applications;
- They have equipment to protect against hardware defects and failures or faults, such as: ECC memory, redundant PSUs and / or fans, RAID controllers;
- They are designed for continuous operation (24 hours x 7 days x 52 weeks);
- They may have a central shared memory;
- They have comprehensive system management tools for easy administration and maintenance, optionally also within the used operating system or virtualisation software.

2.2 Applications

x86 servers are primarily used in companies and public authorities for the following applications or functions:

Virtualisation servers

Virtualisation servers are used when several, usually older servers are replaced by new systems, with the old servers represented on the new hardware platform in the form of »virtual machines«. Existing applications can then still run on their own operating system (or different versions of an operating system). Virtualisation has become the standard and is used comprehensively. A fundamental difference is made between two types:

a. classic virtualisation, with a central storage system on which virtual machines and files are stored;
b. converged or hyper-converged. Hereby, virtual machines and data are stored on the relevant servers and no central storage is required. Each server provides computing power, storage capacity, and network connections, and is manged through virtualisation software.

We recommend asking the support of a consulting company that can help in choosing and configuring the virtualisation servers.

File servers

File servers function as central file storage; therefore, it should be ensured that the media have sufficient capacity and that files are backed up adequately with RAID.

Print servers

Print servers provide central printing services. This function is usually combined with others, such as those of a file server, because this does not require a lot of performance.

Mail servers

Mail servers are the central server instance used to manage electronic mail traffic. With common mail server programmes, it can be freely configured whether e-mails should also be stored locally on the workstation computers of the user or whether the e-mails should be stored centrally, on the mail server. Larger hard disk capacities might be required, depending on this configuration. If it is necessary, for legal reasons, to store e-mails for a longer period or to archive them, a professional consulting company should be involved to include the legally mandated requirements in the tender.

Databases

These servers can use a range of different database programmes for the storage and processing of a wide variety of data.

Application servers

This server type has been included in this list for the sake of completeness. Certain applications use this server type to calculate data, whereby the actual data are read from another database server. The application servers of the standard software SAP are a good example. Whereas the application data are exported, e. g. through a web server, the application server takes care of actual data processing, and the actual data are stored on a database server in the background (3-tier architecture). Such multi-tier applications are only used for larger applications (for example, at the tax authorities or with SAP installations).

Terminal servers

On a terminal server, the application is executed from a central computer, with the graphic output provided through a network connection. Various workspace computers can simultaneously access the terminal server and use the application. The generated data are stored centrally.

2.3 Dependencies in selecting a server

Prior to procurement, the scalability and licensing of software should certainly be taken into consideration. Increases in server performance are no longer exclusively achieved through higher clock rates, but rather also through a higher number of processor cores that run in parallel. This has changed the licencing models of many software providers. For many applications, the licencing costs are based on physical sockets for certain CPU types (e. g. operating systems), or available cores in the overall system (e. g. databases or middleware). By finding the right balance between the licencing model and the number of sockets / processor cores, significant costs savings can be realised. Processor manufacturers offer various processor types for such cases, with the number of cores reduced accordingly and a higher clock rate, or conversely CPUs with more cores.

3 Sizing

The most important question to be asked before procurement relates to the required server configuration relative to the specific application, e.g. number of users, software used, etc. This is called »sizing« (= design of CPU / Memory / HD8 / IO). Software and hardware manufacturers provide specific recommendations or special planning tools. Because of the number of potential parameters, no conclusive recommendations which would allow the procurer to define the dimensioning of a server can be made in this document. Some parameters are listed using examples

Example: If a public entity seeks to procure a server for virtualisation, the following questions must be answered beforehand:

- How many VMs should the system be able to service?
- What virtual CPU and memory performance is required by the individual systems?
- Should the data be stored on data carriers in the server or cluster, or is a central storage system used? What performance, scalability, and availability requirements do I place on my storage, relative to the number of input / output operations per second, (IOPS), RAID support of the hardware controller, cache size, write stability of flash storage media (NVMe and SSD)?
- How is a central storage connected (fibre channel, iSCSI, ethernet)?
- What type of network connection is required for my server (in gigabits per second)? Should the server be connected using optical connections (fibre optic cables)?
- What is the power draw / heat development (commonly specified in watts and BTU) that I can provide / tolerate at the server installation site?
- What power connections do I have available at my data centre, rack, UPS, socket?
- Can I realise very high server densities in my rack / data centre? (Highly compact servers can currently achieve power densities of around 1 kilowatts per height unit. These servers might require rack water cooling.)
- What are the redundancy requirements I place on power supply, data storage, network connections?
- How will server or data centre maintenance work affect operations?
- What is the optimal number of server systems (a very low number of 4-way servers (with 4 CPU sockets), a higher number of 2-way systems, or many 1-way servers)? This influences the amounts of space and power one will need, thermal output, licencing, etc.
- How much space do I have to lay cables?

These and additional parameters should be identified together with specialist departments.

4 Construction types

There are currently four types of housing available on the server market:

4.1 Tower systems (floor-stand systems)

This type of housing is primarily used outside of data centres, or in situations where rack installation is not possible. A typical property of tower systems is that the number of hard drives and I/O connectors can be expanded. With some tower systems, parts are available in rack-mountable form.

Tower systems should be fitted with a lockable screen to prevent data theft and inadvertent shut-down.

4.2 Rack systems

The unique property of rack systems is that they can be installed into 19" racks to take up only a few height units (HU, or rack units=RU). The compact design might restrict internal expandability of the hard drives and I/O connectors.

The racks are available in various sizes and depths, with a range of RUs, depending on the manufacturer. To guarantee seamless integration into the computer landscapes, rack dimensions should correspond to one of the following standards: EIA 310D, IEC 60297, or DIN 41949. Structural conditions (e.g. height of doors, lifts, (raised) floors, load-bearing capacities) should also be kept in mind. The number and capacities of available power connections, as well as the thermal capacity, should be taken into account.

On general principle, racks for rack systems should be fitted with a lockable front and back door as well as lockable side walls and covers.

4.3 Blade systems

Blade systems are even more compact than rack systems. A blade server, server blade, or just blade, is a component that is inserted into a blade server chassis together with similar blades. Manufacturers use unique chassis designs. The fitted infrastructure – such as PSUs, fans, backplanes – is shared with other installed blade servers. There are various expansion options available – with I/O modules (LAN and SAN), storage blades, as well as management modules – that can be fitted to match specific requirements.

Blades come with the benefit of a compact design, scalability and flexibility, as well as easier cabling using less cable. Blade servers and their components are specific to a manufacturer, with each manufacturer offering different HUs and expansion options.

4.4 Multi-node server systems

Multi-node server nodes are so-called scale-out systems that are mostly used in data centres. They can be expanded intelligently by bundling several independent server nodes and local flash drives into a common, compact rack housing. Moreover, with this technology, complete server nodes can be added or removed to match operational needs.

The modular housing with modes (servers) allows for higher densities than classic rack systems. This helps in creating a scalable IT infrastructure to quickly and flexibly respond to the rapidly changing demands placed on IT. They provide a massive scale-out x86 server performance for large data centres and are the most economical option in terms of server density, energy consumption, thermal output optimisation, and lower overall operating costs.

They offer an optimal platform for scale-out workloads such as virtualisation, all types of business applications, high-performance computing, deep learning and data analysis.

4.5 Edge servers

Edge or rugged servers are a variant of servers specifically designed for use outside of data centres or offices. They offer similar performance data to standard rack-, tower- or multi-node server systems, but allow, for example, higher temperatures, stronger vibrations or special filter panels against dirt. In some cases, they are also more compact, meaning that they can also be used in confined spaces.

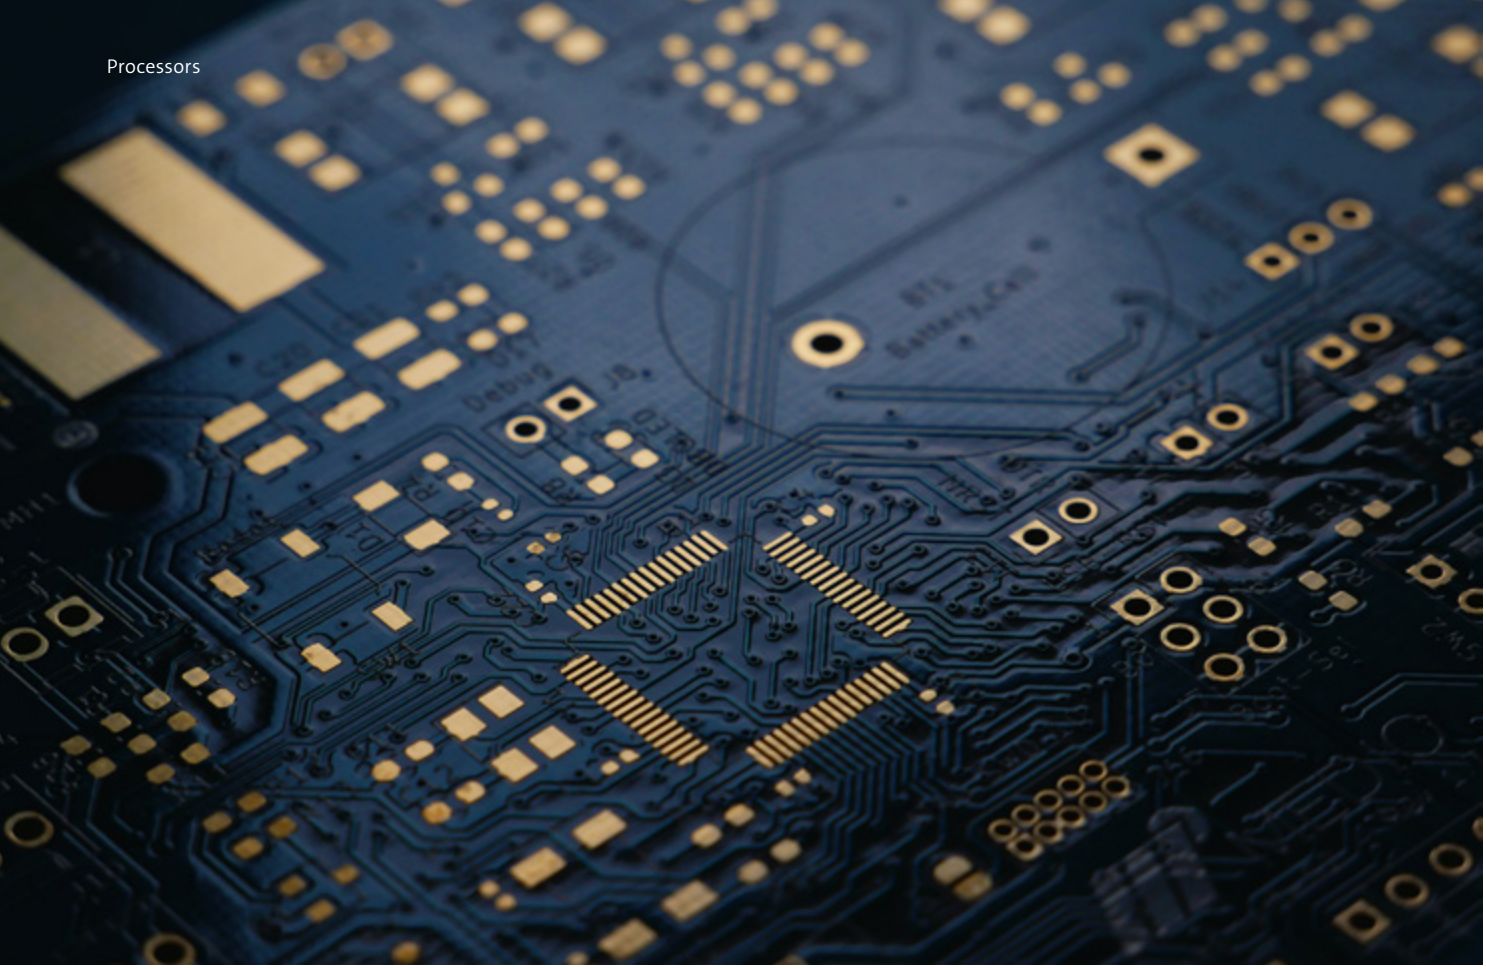

5 Processors

5.1 Server processors

A server processor is a processor with a design geared towards server requirements, which differ in many ways from processors used in desktop PCs and notebooks. This includes:

- lifetime / reliability / fail safety
- Support for secure memory technologies: with error correction code (ECC) and additional registers (Registered ECC)
- The number of processors that can be connected to each other on a mainboard (e. g. 1, 2, 4, and 8-socket servers)
- The number of cores per CPU (typically 8, 16, 32, 64, 128 and more)
- Far higher number of, and speed on, I/O interfaces for faster data transfers than desktop and notebook CPUs; in particular PCIe and storage interfaces
- Support for external hardware accelerators for specific tasks (e. g. GPU, FPGA, or DPU)
- Support for special virtualisation and security features

Therefore, a »server CPU« or »server processor« should be explicitly requested in the technical specifications, to benefit from the special properties to improve virtualisation task performance and boost security in server environments. However, besides the CPU speed, a well-balanced overall system influences server performance as well (see Section 5.2 Benchmark)

Under certain circumstances, an exception can be made and servers can be fitted with desktop or notebook processors (e. g. microservers). In this case, not all performance features might be available.

A new trend in server processors is processor-integrated accelerators, which have been specially developed for certain fields of application and then accelerate these applications. The idea is not new in itself, as graphics accelerators or accelerators for encrypting and decrypting data have long been integrated into processors.

In particular, data protection through encryption can be handled significantly more efficiently with accelerators. Modern processors have integrated accelerators for the native encryption of stored data or data to be transferred (AES-NI, SHA-NI, Crypto-NI), which are used automatically by suitable software. This is another example of why servers should only be operated with up-to-date software.

Examples of new fields of application for accelerators include artificial intelligence, security, analytics, storage, networking and high-performance computing (HPC). The higher computing power for certain applications may result in lower overall costs, as less hardware and possibly no special hardware is required to run certain applications. Energy efficiency is also improved when applications are processed faster. In individual cases, however, it must also be checked whether the software used capitalises on the advantages of the accelerators.

5.2 Benchmarks

Public tenders must be formulated in a vendor-neutral way, i. e. without reference to a certain product of a manufacturer, which is why server tenders may not include any specifications for certain processor manufacturers or certain processor models. Servers have a wide range of applications, with the most frequent ones summarised in Chapter 2.2 together with associated preferences for hardware equipment. The benchmarks commonly used for PCs or notebooks (e. g. BAPCo Sysmark) cannot be applied to servers.

Benchmarks have the benefit of providing a specific, comparable, and reproducible method to measure system performance; furthermore, they are relevant, current, objective, and credible. Especially with smaller procurement volumes and individual purchases, both the procurer and the bidder might find it difficult to carry out their own benchmarks. These are expensive and require great effort. In these cases, standardised and generally recognised benchmarks for the application case in question

should be used. These application benchmarks are based on stringent test methods, that were developed by either independent industry consortiums or by software developers, and are recognised and supported by server manufacturers.

The criteria to assess the quality of benchmarks above all include detailed documentation on the benchmark procedure as well as on the benchmark results. Both aspects ensure that users can, on the one hand, assess whether a benchmark objectively illustrates the performance to be expected, without any bias, regardless of the platform used, and, on the other hand, that another party can reproduce the results.

The revision or version of the benchmark is also important. If implementations change as a result of technological developments or new models of use, the usage patterns change accordingly, even if the application itself remains the same. This is why credible benchmarks are updated regularly. To depict the requirements of current applications in the most precise way possible, it is therefore important to always use the latest and most current version of a benchmark. This procedure recognises the effort potential bidders make to implement the latest technologies and improve system performance. Tenders that specify old benchmarks or old benchmark versions, never reflect the state-of-the-art. The procurer runs the risk of purchasing devices that will not deliver the expected performance during the planned period of use.

Server applications scenarios vary widely. Therefore, this guideline cannot recommend any specific guideline or even define specific targets for benchmark values to be achieved. The procurer should rather check precisely which type of application will be used on the servers to be procured and, based on this, choose a suitable benchmark.

The best-known consortium in the field of benchmarks is the SPEC (Standard Performance Evaluation Corporation). SPEC Benchmarks are the industry standard for comparing server performance. A range of published benchmark results for a wide variety of server configurations can be found on ↗ https://spec.org. These can be used as a guideline for determining tender specifications (details on the various partial benchmarks of the SPEC Benchmark Suite can be looked up in the Glossary).

As is the case with all published benchmark results, the user should keep in mind that these were only carried out for certain configurations, which do not necessarily match the servers to be procured. If benchmark values are used to compare systems, it is therefore recommended to also study the so-called »Full Disclaimer« on the websites of the benchmark providers. These specify the exact configuration of the measured system.

In addition to the SPEC Benchmarks, which are designed for a wide range of uses, other benchmarks are application specific. These should be relevant for the intended use of the systems to be procured. The requested benchmark should be appropriate to the intended workload.

Benchmarks often simulate certain application constellations and vary according to the operating system and application software used. Not in all cases can the results be taken over completely for the specific intended use. To ensure that a server tender is processor-vendor-neutral, while guaranteeing some sort of comparability across the offers, an objective classification of performance in the form of a benchmark is still recommended (e. g. SPEC).

In short, benchmarks should have the following characteristics: **comparability, reproducibility, relevance, currentness, objectivity, and credibility**. SPEC is the industry standard for a wide range of applications. In specific applications, usually whenever the system to be procured is intended to only execute a single application, application-specific benchmarks are relevant (e. g. those of manufacturers such as VMWare and SAP).

Refer to the Glossary at the end of this guideline for details on all relevant server-specific benchmarks.

6 Server memory and storage

Various memory and storage technologies are used in servers. On general principle, only media that were approved specifically for servers by the manufacturer should be used.

These media, their characteristics, as well as their benefits and drawbacks are briefly introduced here, to help you make the best procurement decision. In principle, there are two different use scenarios for servers: main and mass memory. Currently, the following memory and storage technologies are available: RAM, persistent memory, as well as flash drives and hard disks.

6.1 Main memory

The main memory, also called working memory, is essential to server performance. The read and write speeds as well as memory capacity are decisive. Two technologies are currently used in main memory: RAM (Random Access Memory) and persistent memory.

Compute Express Link (CXL) is an open standard for the high-performance connection of the processor to the memory or other components in the server. CXL is based on the serial physical and electrical interface PCI Express (PCIe). The serial communication and pooling capabilities of CXL can overcome the performance limitations of conventional DIMM memory at high memory capacities, and can optimise communication between other components in the system, such as accelerators.[1]

RAM (Random Access Memory)

RAM is very fast memory with very high data transmission rates. RAM is faster, but also more expensive, than flash. Contrary to flash, RAM is volatile memory. After an interruption to the power supply, all stored data are lost.

Persistent memory

A special variant of memory is so-called persistent memory, which, like RAM, is available in the form of DIMMs but also in the form of SSDs. Persistent memory (PM) is non-volatile, similar to flash, bridging the gap between quick and expensive classic RAM and affordable but slower flash memory. In terms of the procurement price, capacity, and flexibility, PM has some benefits that might be attractive depending on the scope of use.

DIMM modules based on persistent memory are available with higher memory capacities than RAM. This makes them a suitable alternative, as an extension of the main memory or as non-volatile memory with low latency situated close to the CPU. Various requirements must be observed when selecting the CPU as well as the operating system and the applications.

6.2 Mass memory

1 Wikipedia contributors. (2023, November 23). Compute Express Link. In Wikipedia, The Free Encyclopedia. Retrieved 12:50, December 18, 2023, from ↗ https://en.wikipedia.org/w/index.php?title=Compute_Express_Link&oldid=1186422698

Flash

Flash memory has:

- very fast access times;
- very high I/O performance;
- low energy consumption;
- no moving parts, therefore relatively immune to vibrations and shocks;
- higher service temperatures than hard disks;
- depending on the form, various levels of quality in rewritability (endurance).

Endurance is often specified in »Drive Writes Per Day« (DWPD); which the manufacturer guarantees. DWPD indicates how often an SSD can be rewritten completely per day over a certain period, typically 5 years. A rough distinction is made into categories, with prices varying accordingly:

Entry (relatively low write rates)	DWPD lower than 2
Mid-range (for medium-range write rates)	DWPD between roughly 2 and roughly 10
Performance (optimised for high write performance)	DWPD of roughly 10 and higher

SSDs

Solid state drives, or SSDs, are usually also based on flash. They are available for the SATA, SAS, and M.2 interfaces. SSDs are available with the 2.5 and 3.5 form factors, as well as in modules for the M.2 interface. SSDs offer several 1,000s of I/Os per second, far outperforming hard disks. SSDs come with various endurances. They are always connected using a suitable controller – see Chapter 6.5. By now, SSDs are increasingly used as boot devices in servers, because SSDs are often more affordable than hard disks.

NVMe

Non-Volatile Memory express, or NVMe, is connected directly to the PCIe interface. NVMe devices are available both as PCI adapters and with a U.2 / U.3 interface (form factor of a 2.5 or 3.5-inch hard disk or as EDSFF). NVMe devices offer much higher throughput rates than SSDs, with lower latency.

NVMe's based on PM and / or with a connection through PCIe 4.0 help solve memory bottlenecks in servers and help process higher data quantities quicker.

6.3 Hard disks

The following generally applies to hard disks: mature technology, have been available on the market for a long time, and sensitive to strong shocks during operation.

Hard disks come with the S-ATA and SAS interfaces. The SAS interface has been specially developed for servers, with higher throughputs as well as lower access times than the S-ATA interface. For this reason, SAS hard disks are used whenever there is a need for high performance (I/O per second). SAS hard disks additionally offer greater fail safety. In comparison, S-ATA hard disks have lower performance and fail safety, and are suitable for applications which mainly involve sequential data access (e. g. image processing and archive data).

When it comes to server hard disks, a difference should be made between various criteria:

- Hard disk type: S-ATA or SAS
- Hard disk rotations per minute (RPM): 7,200 (S-ATA) or 10,000 (SAS) or 15,000 (SAS)
- Data transfer rate: SATA at 6 GBit/s compared to SAS at 12 GBit/s or 24 GBit/s
- SAS: roughly 160–220 I/Os per second
- S-ATA: roughly 90–100 I/Os per second

Hard disks with a higher capacity are by now only offered with a 3.5-inch form factor. 2.5-inch disks with lower capacities are increasingly replaced by SSDs. By now, hard disks are rarely used whenever high I/O performance is required, e. g. for databases.

7 UEFI BIOS / Classic BIOS, Driver, Operating System

UEFI BIOS

The Unified Extensible Firmware Interface, hereinafter referred to as UEFI in short, is currently the central interface between the firmware, individual server components, and the operating system. It handles the management between the software and hardware during the BOOT process. From a logical perspective, it sits beneath the operating system. Current versions of UEFI have a secure boot component. This restricts booting to pre-signed bootloaders to prevent the booting of malware and other unwanted programmes. Its most important features are:

- embedded BIOS emulation to guarantee compatibility with the BIOS;
- high-resolution graphic cards that are already supported during computer start-up;
- various operating systems installed on the system can be selected and launched through the UEFI, without the need for special bootloaders;
- whereas a master boot record was usually used with BIOS, a GUID partition table is used for booting, which allows for the booting of hard disks with capacities over 2 TB;
- a universal network boot system (preboot execution environment) is an alternative to launch local media;
- the secure boot mechanism required by Microsoft – which restricts booting to predefined and correspondingly signed bootloaders – is supported. This makes system launch safer, preventing malware from already damaging the system before system launch.
- A Root of Trust (RoT) is the foundation of security. The term »Roots of Trust« refers to several features in the trusted computing model that the computer's operating system always trusts. RoT is similar to a standalone engine that controls the cryptographic processor, which is the core of the Trusted Platform Module (TPM).

Classic BIOS

Classic BIOS was the predecessor of the Unified Extensible Firmware Interface (UEFI), today also called legacy BIOS. With the BIOS (Basic Input Output System), the functionality of all system components is tested during the so-called POST (power-on self-test). System functions and system hardware configurations, such as security and energy-saving functions, server management, boot sequence, etc. are configured in the BIOS setup. The hardware configuration options vary between manufacturers; therefore, it should be possible to query the specifications of the BIOS manufacturer. Classic BIOS has by now been largely replaced by the Unified Extensible Firmware Interface (or UEFI in short, previously EFI).

Classic BIOS is a discontinued model and should no longer be used. All current servers and operating systems have supported UEFI for some time now. Therefore, BIOS or Compatibility mode should no longer be used today.

Drivers

All system drivers of the most-important operating system versions should be up-to-date and allow for seamless operation. Driver updates of individual standard components should not lead to system conflicts. These updates should be available offline or online, depending on the fitted components.

Operating system

The functions of common server operating systems will not be addressed in detail here. On general principle, the servers should be included in the Hardware Compatibility List of the operating system in question.

This above all applies to so-called converged and hyper-converged infrastructures. All components used in the server must be listed in the relevant Hardware Compatibilty List (HCL) of the software manufacturer in question, together with the server systems themselves.

The use of unapproved systems and components is strongly discouraged: this can result in undefinable disruptions to the entire system, up to irrevocable data loss

8 System management

Server management software is now an essential component for secure server operation which makes server management tools a fixture of server systems. They allow for server monitoring during operation and / or independent of the operating state.

Administrators or service technicians can access the server and carry out comprehensive monitoring tasks, also at decentralised locations, to efficiently carry out routine tasks and maintenance activities if server problems occur. These tasks are increasingly performed »out-of-band« (OoB) as well, via a baseboard management controller (BMC).

The baseboard management controller, also called microcontroller, is the centrepiece of the OoB architecture. The BMC has the following tasks:

- interface between the system management software and the managed hardware (through which the BMC is connected using IPMB and ICMB)
- independent monitoring
- independent event-logging
- recovery control.

Essential functions furthermore include update and profile management, deployment, and the capacity to integrate larger installation within the datacentre into enterprise management systems.

It can be operated via a web-based user interface; an additional command line interface can be appropriate for specific components. One of the ways in which this is carried out is via the REST interface (representational state transfer).

The Redfish API (Redfish Scalable Platforms Management) specification for out-of-band server systems management should be supported. Redfish aims to replace the Intelligent Platform Management Interface (IPMI).

Server management must be secured through suitable access controls to prevent unauthorised access to management functions. Multi-Factor Authentication (MFA) should be considered.
A server management system lowers administration costs and increases server availability. Service costs are additionally lowered through preventive fault detection, as well as with various integrated diagnosis functions.

In virtualised environments and when using converged and hyper-converged infrastructures, it can be sensible to have server monitoring integrated into a management tool of the CI or HCI software manufacturer in question. This protects the investment made while also reducing administrative costs.

For some time now, management features have also been offered on a cloud basis. Here, telemetry data from the servers and the management interfaces are collected and transmitted to a secure cloud environment on the server manufacturer's servers, and clearly displayed in a dashboard. This environment can then be accessed by the administrator from anywhere in the world. If the transmission of telemetry data and management functions outside the client's premises is not desired, this should be noted in the specifications.

9 PSUs / power supply

General information

Servers are in many ways exposed systems, and great attention should be paid to their availability. For this reason, PSUs are often redundantly installed in servers, together with fans, hard disks, and I/O-adapter components. In addition to redundancy, a second criterion of PSUs is their capacity for hot plugging, which makes it possible to swap out defective components without interrupting operation. Redundant and hot plug-capable components are virtually the standard for large and important systems, i. e. systems in central functions, which are used simultaneously by many, such as databases and virtualisation or file servers. For smaller and affordable entry-level systems, which usually only have a single CPU socket, the manufacturer usually opts against redundancy and hot plug-capability for PSUs and fans, in order to cut costs.

PSUs

In tower and rack servers, redundancy usually implies that each system require at least two, or in rare cases more PSUs. Blades are provided with power by the blade chassis. In these systems, PSUs usually supply power to a significantly higher number of servers, so-called blades, for the most energy-efficient overall operation of systems. Redundancy should be required for the majority of use scenarios for blades.

With redundant PSUs, it should be ensured that the systems provide separate power connections, to also realise redundancy in terms of power supply. If all PSUs receive power from a single, shared supply, a disruption will inevitably lead to interruptions in operation.

The Ecodesign Regulation ErP Lot 9 for servers and data storage products has been in effect in the EU since March 2020. It contains relevant energy-efficiency requirements. Its scope covers all standard servers. With it come stricter requirements for PSUs. From January 1, 2024, the next stage will take effect, making titanium-quality power supply units mandatory for almost all server systems.

In practice, an upstream uninterruptible power supply (UPS) is commonly used for at least one power supply, which should protect the servers against supply interruptions as well as against frequently occurring voltage fluctuations and grid spikes.

Manufacturers commonly offer PSUs in a variety of performance classes, geared towards the requested configuration, to efficiently operate the PSUs within a defined load range. This must be taken into consideration when planning overall system expansions!

Note: During procurement, ensure that the overall server power consumption does not exceed climate-control capacities and the capacities of electrical connections. Many manufacturers offer suitable calculation programmes online.

Through their system management, some manufacturers also offer functions to restrict system power consumption to far lower levels than the performance values specified on the PSUs. This primarily serves to more efficiently use available rack and/or climate-control capacities.

Fans

Moving mechanical components such as fans are subject to wear, which is why they should be designed redundantly and monitored.

When procuring rack servers, there is no need to set noise restrictions, because these are operated in separate rooms. With tower servers, noise restrictions should be specified depending on the location where they are installed.

With rack, blade, and multi-node servers, fans are usually both redundant and hot-plug capable, which guarantees continuous operation. For cost reasons, cheaper, entry-level tower systems usually do not feature redundancy or hot-plug capability.

Power supply

In rack-optimised servers as well as blades, power is frequently supplied through power strips situated in the rack or PDUs (power distribution units). In practice, every rack usually has two separated supplies to ensure appropriate redundancy. Either USPs in the rack or external USPs are frequently used to supply these units with power.

These power strips usually have two types of power connectors for the servers. The most commonly used connection for servers in Germany is the IEC 320-C13, also known as the cold appliance connection, or the IEC 320-C15, also known as the warm appliance connection. Devices that use heavier loads, typically over 2,000 W, are usually fed with IEC 320-C19 or IEC 320-C21, a rectangular plug with three horizontal contacts arranged in a triangle.

The small power strips or PSUs with IEC 320-C13 connectors are in turn supplied with power through either a two-pin earthed plug or IEC 320-C19, e. g. through a UPS. Larger power distribution systems usually require the use of single-phase connections for 32 or 63 ampere. In Germany, a three-phase 16 or 32 ampere connection is most frequently used – in German: »Kraftstrom« – which also provides the highest performance.

When planning the system, the latter power connections must be realised by a specialised electrical company. With a sufficient number of connections of this type, racks can be supplied with wattages far exceeding 10 kW, with adequate climate control.

Uninterruptible power supply – UPS

Uninterruptible power supplies are available from approx. 300 VA up to the range of several 100 KVA. Uninterruptible power supplies intended for installation into racks usually offer several IEC 320-C13 connectors, depending on the total output, and from approx. 3,000 VA, systems are usually designed with one or more IEC 320-C19 connectors, which can be used to connect additional sockets or PDUs. The UPS must report a disruption in power to the connected servers, to

allow relevant precautions to be taken with automatic shut-down, ensuring the consistency of data to be processed.

When installing individual servers, signalisation generally takes place serially, through an RS232 or USB interface. In rack installation, the most frequently used signalisation medium is ethernet, with SNMP as the protocol. The UPS needs a suitable interface card for this. It must be ensured that the network components located between the UPS and the server are also secured by the UPS.

To allow for proper server shut down, UPS's must bridge a certain minimum interruption period. The periods required for the systems are relative to the type of applications; information should be requested from the specialist department. The Internet websites of manufacturers contain relevant tables to determine the number of additional batteries required, if necessary.

For planning purposes, it should be noted that the electrical connection of UPS above 3,000 VA requires a specialist electrical company.

10 Data backup and restore

Data backup refers to the copying of data on a computer system onto another storage medium or another computer system, either in part of in full. The Recovery Time Objective (RTO) and acceptable data loss as a result of disruptions (Recovery Point Objective (RPO)) significantly influence the selection of hardware and software.

Data secured on a storage media are called a backup. Recovery of the original data from a backup is also called data retrieval, repair, or restore.

Data backup is intended to protect data against hardware loss, e. g. due to theft, fire, surges in voltage, or against software errors, viruses, worms, as well as user errors, which occurs frequently. The backup should be physically separated from the system it backs up, to ensure data protection.

A difference is made between archiving, i. e. long-term and unchangeable data storage for auditing purposes, and simple data or system backups. The Federal Ministry of Finance published a summary of principles on archiving and auditing in its Principles of Data Access and Auditability of Digital Documents (GDPdU), to which companies have been subject since January 2002. Archiving will not be addressed in further detail here.

Tape drives are frequently used to back up small or individual systems. Current systems usually use the serial SCSI interface – or SAS.

If the sole objective is to restore system operations, this is frequently realised using tape drives running on the LTO (Linear Tape Open) standard. LTO has been available on the market for many years, with a range of available capacities.

Separate backup systems are frequently used to protect servers and data within company networks and in larger systems. These are usually connected to larger tape loaders (a single drive with several tape cartridges – for larger data quantities), libraries (tape robots with several tape drives, significantly larger number of tape cartridges – for larger data quantities and higher speeds), as well as larger separate storage systems with hard disks (so-called virtual tape libraries for online backup – quicker recovery, but might consume more power and no option to ensure data security in a safe). These components are usually also connected via SAS, with fibre channel being used as well. Tape drives and virtual tape libraries can also be used in tiers, if necessary.

Suitable backup software is always required for backups. In this respect, it is pointed out that there are many dependencies between the server operating system, the backup software, and the tape units used. It is therefore recommended to only use systems certified for backup use with each other. It is recommend to ask the support of a consulting company that can help in choosing and configuring the backup and restore concept.

It is also possible to back up to cloud environments, but the parameters discussed in paragraph 1 and the data protection situation should be examined in detail beforehand.

11 Other services

11.1 Establishing the power connections

From the perspective of the contractor, the customer is responsible for connecting the servers and / or racks to the on-site power supply. Standardised power lines and fuses must be installed. A contractor can inform the customer of the requirements (connected loads and consumption ratings). The customer is responsible for realising the required electrical installations. Only the operator or its building services are familiar with the exact infrastructure. The contractor ensures all server or rack requirements are met. When commissioning a rack / UPS, building services should then connect the power cable of the racks / UPS to the corresponding CEE plug or realise a fixed connection. This is not always necessary when installing an individual server. The customer can also commission these required services when placing an order for servers / racks. In this case, liability must be arranged contractually.

11.2 Support

If necessary, relevant support should be agreed, including details on response / recovery times.

Common offers vary according to:

- Contract duration
- Response times (period between reporting a disruption and receiving a first response from support)
- Recovery times (period between reporting a disruption and restoration of server functionality)
- Categorisation into disruption classes (serious, critical, and non-critical disruptions)
- Defining service transfer points
- Spare-parts logistics
- Additional technical services (supporting operations, etc.) based on expenditure (hourly rates, travel expenses)

Following demand, requirements might include:

- 3, 4, or 5 years of on-site service
- On-site service with a response time of x hours
- On-site service with a recovery time of x hours
- Hotline can be reached x hours y days a week
- Hotline in German (or English)
- Delivery of spare parts which do not require swapping by a service technician (for hot-plug components)
- Spare parts storage at the customer

Individual agreements can be made when procuring solutions with high availability or safety-relevant solutions. An assessment of the necessity of the requirements and the resulting costs must be made here.

Proactive service packages are now available, similar to system management, in which the servers transmit component telemetry data to the server manufacturer, so that in the event of a hardware defect, a suitable spare part is already proactively shipped to the customer. If a so-called »Call Home« function is not desired, this must also be noted in the service specifications.

11.3 Logistics

The following logistical services can be agreed, if necessary:

- Specification of the max. delivery time
- Free delivery at the facility
- Free delivery at the place of use
- Delivery abroad
- Delivery to different locations
- Server installation
- Manufacture of power connections (see Chapter 9)
- Installation into available rack systems

- Connection to the LAN and SAN infrastructure
- Customer-specific server pre-configuration
- Taking over asset management
- Documenting the configuration

With rack deliveries, the customer must inspect, or provide information on, the transport route between the delivery vehicle and the installation site (e. g. paved roads, heights and widths of doors, permissible loads of lifts, number of steps, and much more).

12 Environmental and health protection

12.1 General legal requirements

All legal requirements must be complied with, in particular Regulation 2019/424 laying down ecodesign requirements for servers and data storage products which has been in effect from 7 1 March 2020. (see also Chapter 9)

The Ecodesign Regulation for servers and data storage products specifies legal minimum requirements for placing these product types on the EU market. This above all includes requirements on energy and material efficiency. The requirements placed on PSU and material efficiency apply to all server and data storage products; energy efficiency requirements, on the other hand, only apply to servers with one or two processor sockets. The criteria of the Ecodesign Regulation for servers and data storage products can be accessed here: ↗ https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1553786820621&uri=CELEX%3A32019R0424

Legal requirements apply equally to all x86 servers and do not have to be included in the service description.

12.2 Certifications and labels for verification purposes

A distinction must be made between legal requirements and voluntary certifications and labels, which highlight special product characteristics or which serve to verify compliance with special requirements in certain usage environments.

Contracting entities can demand presentation of such verifications to more readily determine that the offer complies with the characteristics demanded in the description of service.

If the procurer demands presentation of such a certificate, it must be usable within the meaning of public procurement legislation, i. e. in particular providing suitable verification of the characteristics demanded in the service description (§ 34 para. 2 of the Ordinance on the Award of Public Contracts (VgV)). Moreover, alternative certificates that place similar requirements on the service must be accepted as well. A distinction should be made between the certificate as potential verification and the actual requirements placed on the object to be procured. Requirements must be formulated in a call for tender in a binding manner. Certificates can verify compliance with these requirements. Declarations of manufacturers should be recognised as evidence if their credibility can be suitably asserted, e. g. with test and inspection reports, or if they meet international standards.

The certificates and their areas of application for x86 servers that are relevant for certain requirements are listed below. The procurer must decide which of these verifications is required for the scope of use in question on a case-by-case basis.

ENERGY STAR®: ENERGY STAR® is a voluntary programme run by the EPA (US Environmental Protection Agency). ENERGY STAR® products are certified by independent certification authorities and are listed in the ENERGY STAR® database (↗ https://www.energystar.gov/productfinder/). The EPA also demands that a product sample is tested. The currently valid requirements are ENERGY STAR® for Computer Server Version 3.0. The criteria for ENERGY STAR® servers can be accessed here: ↗ https://www.energystar.gov/products/spec/enterprise_servers_specification_version_3_0_pd. After the EU Energy Star® Programme expired in 2018, this specific label should no longer be demanded in EU tenders. Alternatively, the Energy Star® criteria can be used in the tender documents.

EPEAT: EPEAT is a globally leading environmental label for IT. The EPEAT programme offers independent verification of manufacturer specifications, and the EPEAT Online Register has a list of sustainable products of a wide range of manufacturers. The criteria for x86 servers, NSF 426 20019, can be accessed here: ↗ https://epeat.net/servers-search-result/page-1/size-25?productName=x86

TCO Certified 8 for servers: TCO Certified is a global sustainability certification for various product categories. Comprehensive criteria promote social and ecological sustainability over the entire IT product lifecycle. Compliance with each requirement is assessed separately, both before as well as after certification. The criteria for x86 servers can be accessed here: ↗ https://tcocertified.com/files/certification/tco-certified-generation-8-for-servers.pdf

The Blue Angel: The Blue Angel (Blauer Engel) for servers and data storage products (DE-UZ 213) is a voluntary certification for environmental aspects that strives to honour particularly environmentally friendly products. For all products that meet the label criteria, a request can be submitted to RAL gGmbH, after which permission can be granted to use the environmental label for the product in question on the basis of a label use agreement. The label covers all common server types and data storage products, extending beyond the scope of the Ecodesign Regulation. Additionally, the award criteria are much stricter than the requirements of the Ecodesign Regulation. Using these criteria in tender documents might result in severe market restrictions. Before the tender, a careful assessment should be made if there is a sufficiently large number of servers or data storage products available on the market with the environmental label, to allow sufficient competition. The Blue Angle award criteria for servers and data storage products can be accessed here: ↗ https://www.blauer-engel.de/en/get/productcategory/169

Comparability of eco-labels

Today, taking environmental issues into account is one of the basic requirements for all office devices. Mandatory requirements for environmental sustainability (e. g. environmentally friendly disposal of old devices, a ban on certain substances in products, electromagnetic compatibility) must be met by the manufacturers of electronic products by law. If a manufacturer does not

meet these basic legal environmental requirements, they are not permitted to place their products on the EU market at all. Requirements that go beyond the legal minimum standard are being increasingly stipulated, especially in the areas of energy consumption, service life and noise emissions. Some requirements (both mandatory by law and those that go beyond this) are collectively checked and evaluated by eco-labels. However, one should tread lightly when using eco-labels in calls for tender, because, depending on the choice of quality mark, certain devices or providers will be excluded from submitting a tender offer, resulting in the market being narrowed accordingly. Moreover, not all eco-labels check for the same criteria pursuant to the same standards. Thus, they can hardly be compared. Not least for this reason, this guideline recommends issuing specifications for the device criteria and requirements in calls for tender. Not only should eco-labels be allowed as proof of meeting these criteria, but also test protocols. When updating eco-labels, there may be delays between the application and the approval of the new quality mark. In this case, self-declarations that state compliance with the relevant requirements should also be accepted.

12.3 Social sustainability

Besides economic and ecological criteria, social aspects should be considered in tender procedures (§§ 97 para. 3 Act against Restraints of Competition (GWB), 31 para. 3 Public Tender Regulation (VgV) for above-threshold procurement, §§ 2 para. 3, 22 para. 2 Regulation on Sub-Threshold Procurement (UVgO) for sub-threshold procurements). Such social aspects include, in particular, the rights of employees, the prohibition of child labour and employee discrimination, and compliance with the working hours framework at both the tenderer and their suppliers.

To make sure these aspects are taken into consideration in the tender procedure for IT products and IT services, the awarding party can require each bidder in the tender procedure to submit a declaration of social sustainability for IT. This declaration, one of the so-called text components for contract design and elaborations on the scope of application, can be retrieved from the website of the German Ministry of the Interior's Procurement Office. More detailed information on the declaration of commitment to social sustainability for IT can be found here: ↗ http://www.nachhaltige-beschaffung.info/SharedDocs/DokumenteNB/Verpflichtungserklärung_ILO_BeschA_Bitkom_2019.html?nn=3631266

The website of the Procurement Office of the Federal Ministry of the Interior provides a summarised overview of additional aspects of sustainable IT product procurement: ↗ http://www.nachhaltige-beschaffung.info/SharedDocs/DokumenteNB/Produktblätter/Informationstechnik.pdf?__blob=publicationFile&v=10

Accessibility

Public entities in Germany are legally obliged to procure accessible hardware and software. General accessibility requirements are legally defined in § 4 of the Equality for Persons with Disabilities Act (*Behindertengleichstellungsgesetz*, BGG, see: ↗ https://www.gesetze-im-internet.de/bgg/BJNR146800002.html) (cf. Annex B to this guideline for legal foundations and more information on accessibility). More details are laid down in, for example, Part 1 of the German Information Technology Accessibility Act [Barrierefreie-Informationstechnik-Verordnung, BITV 2.0] (↗ https://www.gesetze-im-internet.de/bitv_2_0/BJNR184300011.html) of the German Equality for Persons with Disabilities Act [Behindertengleichstellungsgesetz, BGG]). Tenders should refer to these or equivalent requirements (cf. Annex B.2). The provider submits a self-declaration laying out which accessibility requirement are met by the offered product and which cannot be met. DIN EN 301 549:2020-02 Accessibility requirements for ICT products and services should be used for this purpose. Direct reference to this standard is made in Part 1 of the Information Technology Accessibility Act BITV 2.0 (↗ https://www.gesetze-im-internet.de/bitv_2_0/BJNR184300011.html) of the German Equality for Persons with Disabilities Act (BGG). As laid down in § 31 para. 2 no. 1 Ordinance on the Award of Public Contracts (VgV), reference can be made to DIN EN 301 549 in the performance specifications, in order to appropriately take the user needs of persons with disabilities into account. Clause 4 of the Technical Report CEN/CLC/ETSI TR 101 552 (2014-03, ↗ https://www.etsi.org/deliver/etsi_tr/101500_101599/101552/01.00.00_60/tr_101552v010000p.pdf) provides self-declaration templates.

Packaging

The German Packaging Act (Verpackungsgesetz, VerpackG) regulates acceptance of returned packaging. If the private end user is left with the packaging, the distributor or dual system commissioned by the distributor is obliged to take it back. Other entities equal to private end users are, among others, administrative bodies, barracks and hospitals (see § 3 VerpackG para. 11). The Central Agency Packaging Register has a detailed overview (↗ https://www.verpackungsregister.org/en). In principle, the return of packaging should always be free of charge. With no additional costs incurred besides logistics costs at the moment, the demand for exclusion criteria should be assessed.

13 IT security

The implementation of hardware-based security is growing increasingly important for use within data centres, with more and more attacks on server systems taking place below the operating system, at the hardware level. Signatures built into the hardware, which verify that firmware of individual components is issued by the original manufacturer, already improve security. There should be an easy option to regularly/automatically check the hardware for firmware changes. If any changes are identified, there should be either the option to quickly recover the desired state or to carry out a forensic analysis. in any event, a server with undesired changes should never remain in normal operation, Servers should have the option of supporting the security architectures provided by the CPU manufacturers.

Here, the hardware accelerators mentioned during the processor selection are important for typical security functions such as encryption: Embedded software should support the desired accelerators, either by default or in a configurable manner, in order to keep the computational effort for encryption low. Therefore, also for general IT security reasons, up-to-date versions of operating software and applications should be used.

The configuration of the server in the UEFI / BIOS and on the management processor are critical to security. It should be possible to back these up and verify them because they govern crucial security settings, such as port deactivations. Care should generally be taken to ensure that unused functions and interfaces are deactivated. In the ideal case, there should be a clear visualisation of whether all the »best-practice« recommendations have been implemented on the server.

It should be possible to implement the settings over a larger number of servers in the easiest possible way, e. g. through management tools or scripts.

To guarantee server security at the end of their life, all data on hard disks, SSDs, and all chips and caches within the server should be deleted. This way, the hardware can be used again at other locations, which is very desirable in terms of protecting the environment and the circular economy. This should be feasible in the easiest possible way and already taken into consideration during procurement.

14 Award criteria

14.1 Proof of performance

For many award criteria, bidders can be adequately assessed based on their own statements and no further information is usually required. This above all applies to benchmarks.

However, the fulfilment – and consequently evaluation – of other performance requirements can more sustainably be verified through measurements on the specific performance object. These assessment criteria include, for example, power consumption.

If need be, the information from the data sheet can be used. However, only maximum values are usually given. Noise emissions can be disregarded here, because servers are usually installed at datacentres or utility rooms for technical equipment, and not in office environments.

14.2 Assessment process

To identify the most economical offer, the prices and performances of the offers are assessed using the assessment matrix. This matrix is drawn up beforehand. All assessment criteria are individually assessed and a score is given. The scores are then multiplied by defined weighting points. The results are added together to give each offer a number of performance points.

After assessing the performance figures, the price-performance ratio is determined. This is the most crucial element to help choose the most economical offer. To determine the price-performance ratio (Z), the total number of performance points (L) is divided by the total price (P).

$Z = L / P$

The broader reference value method is recommended for more complex tenders (functional service description, numerous B-criteria). A criterion for decision can be defined for offers with a price-performance ratio figure that falls within a range of fluctuation to be defined (5 – 10 percent) below the figure of the leading offer, which then helps identify the most economical offer from all preselected offers.

15 Contractual provisions

15.1 EVB-IT

The provision of the tendered services or the delivery of the tendered products after successful completion of the award procedure is based on appropriate contracts. The Federal Ministry of the Interior and Bitkom have worked out various sets of agreements to be used for this purpose, to support the contracting authorities. The sets of agreements can be found on the website of the Federal Commissioner for Information Technology (↗ https://www.cio.bund.de/Web/DE/IT-Beschaffung/EVB-IT-und-BVB/Aktuelle_EVB-IT).

16 Practical suggestions for the tender procedure

16.1 Market research

Market research is a helpful tool when it comes to preparing for a tender procedure. If done correctly, the results can be very helpful in carrying out a needs analysis and formulating the requirements or specifications in a manner compliant with procurement law. If the public contracting entity is well informed on common market products and requirements, this might also increase the efficiency of legally compliant procurement tendering.

Market research is expressly permitted by law:

»Before launching a procurement procedure, the contracting entity may conduct market research to prepare for the procurement and to inform undertakings of its procurement plans and requirements«. (§ 28 para. 1 Ordinance on the Award of Public Contracts (VgV))

Section 28 of the Ordinance on the Award of Public Contract (VgV) does not specify in detail the way in which market research should be conducted. Consequently, this ensures compliance with general principles under procurement law such as equal treatment and transparency. This guideline provides an introduction to server market research.

Commercial procurement models

Servers can be procured by renting, buying or leasing them. Rental and leasing are different in that, at the end of the contractual operating life, the client usually receives a purchase option for the leased item. The approach selected by the procurer depends not least on whether it has a one-off budget or a budget covering several years. Generally, one of the above-mentioned procurement models must be chosen in advance of the procurement measure in the context of an economic feasibility study.

However, the purchase of servers is the most common form of server procurement. For this reason, renting and leasing are not discussed in depth in this guideline. For more basic information on leasing and renting ICT equipment, see the other published guidelines.

17 Annexes

17.1 Glossary

Benchmarks

SPEC benchmark

The SPEC Consortium has published a wide range of benchmarks. Benchmark descriptions can be accessed through the website ↗ http://www.spec.org. It is usually unrealistic for the procurer to independently carry out SPEC Benchmarks. The benchmark procedure is very complex and expensive with special compilers and parameter adjustments. Moreover, it might take several days to correctly design and carry out a benchmark test.

The SPEC CPU 2017 benchmarks comprise standardised, CPU-intensive tests that process demanding tasks to measure and compare processor, storage sub-system, and compiler performance.

SPEC Benchmarks illustrate performance in processing complex computation tasks over a wide range of possible hardware combinations, using workloads that are based on real applications. The benchmarks are provided as source code and require the use of compiler commands and other commands through the input console. SPEC CPU 2017 additionally contains optional energy consumption measurement.

The SPEC CPU 2017 Suite comprises 43 benchmarks, divided into four benchmark sets:

- SPECspeed® 2017 Integer and SPECspeed® 2017 Floating Point compare the time it takes the server to process single tasks.
- SPECrate® 2017 Integer and SPECrate® 2017 Floating Point measure the throughput or output per unit of time.

The results of the individual benchmarks are expressed in two different ways: »Rate« vs. »Speed« or »Base« vs. »Peak«. The »Rate« values are non-coherent, this means that they are neither affected by connections between individual processors nor between individual systems. The »Speed« values, on the other hand, do take these connections into consideration.

The differences between the »Base« and »Peak« values are primarily the result of different compiler and library optimisations. »Base« reflects the performance without optimisations, »Peak« takes these into consideration.

TPC

The Transaction Performance Processing Council provides benchmarks for database tests. The current benchmark variants can be accessed at ↗ tpc.org.

SAP SD Benchmark
SAP is currently one of the most widely used standard software in companies, and is also increasingly used within public authorities The SAP company has developed its own benchmark to compare system and database manufacturers. Details on the SAP SD Benchmark can be found at ↗ https://www.sap.com/about/benchmark/appbm.html

VMMark
VMMark was developed by the company VMWare, to measure the suitability of x86 server platforms for virtualisation. Refer to ↗ https://www.vmware.com/de/products/vmmark.html for details on the benchmark procedure

Processor technologies

Sockets
A socket is a slot for a processor (CPU) in the server. A dual-socket server indicates that a server has two processors; single-socket servers have only a single processor (e. g. used in simple file servers or print servers). Other, synonymous terms are one, two, four, and eight-way systems.

Important note on systems with four or more sockets:

Servers with four or more sockets / processors are suitable for special, powerful applications. They should be planned, designed, and configured by specialists, which is why this guideline will not cover these systems.

Cores
Every modern processor has two or more cores which run in parallel. With more cores, a processor can master more tasks simultaneously (»multithreading«). Contrary to many desktop and notebook applications, server applications usually make much more efficient use of the available cores.

Clocking, clock frequency
Each processor is driven by a clock. All modern processors dynamically and gradually adjust their operating frequency to performance demands, to lower power consumption with low loads. Several clock frequencies are often specified for modern server processors, which apply to various operating states and different loads of individual cores:

Base clock frequency: Clock frequency achieved by all cores without overclocking.

Maximum clock frequency: The maximum clock frequency which can be achieved by individual cores.

Memory controllers

So-called RAID controllers are often used in servers. RAID (»redundant array of independent disks«) refers to the bundling of several hard disks / SSDs into a single unit, to achieve a higher speed, lower risk of data loss, and higher availability. The most-frequent RAID applications are:

- **RAID 0:** Striping – acceleration without redundancy: The simply bundling of (e. g.) two hard disks into one.
  - Benefit: Great boost in performance thanks to parallel writing on both hard disks.
  - Drawback: If a hard disk fails, the data on both hard disks are lost.
- **RAID 1:** Mirroring: Storing identical data on two different hard disks.
  - Benefit: If a hard disk fails, the data are usually still available on the other hard disk.
  - Drawback: No higher speeds.
  - Please note: RAID 1 does not replace regular back-ups!
- **RAID 5:** Performance + parity: a combination of the versions explained above (similar to all other unmentioned RAID types, which are rarely used). The RAID version generally used in servers is RAID 5. At least three hard disks are required (or four with a hot-spare disk). Similar to RAID 1, the disks are bundled into a unit, with additional parity information stored. This information can be used to recover information on a failed disk after swapping it out or to recover the information onto a backup disk (a so-called hot-spare disk).
- **RAID 6:** Similar to RAID 5 with double parity, slightly reducing write performance.
- **RAID 10:** Combination of RAID 1 and RAID 0 for applications with particularly high writing activity.

Interfaces

VGA Video

Server systems usually have a video or VGA port at the back of the housing, or sometimes also at the front. The VGA connector is an analogue image transfer standard for connector and cable connections between server graphic cards and a console. A VGA connector has a 15-pin D-Sub-Mini plug.

USB

Universal Serial Bus (USB) is a serial bus system used to connect external devices. USB connections are very widely used with a broad variety of available products supporting this standard. It has replaced many older interfaces. USB is the standard plug for keyboards and mice.

KVM

Several computer systems are often connected to a central console using a KVM infrastructure. KVM is an umbrella term and the abbreviation for keyboard, video, and mouse. If a VGA connection is available at the front, it can be used to connect a console if required. Integrated remote management solutions optionally allow for a truly IP-based KVM solution.

Serial

The serial interface for data transmission between a computer and the periphery was developed in the 1980s. A 9-pin SUB-D connector (RS-232) or an ethernet cable with an RJ-45 connector is used to connect the serial interface. It has serial data transmission; serial interfaces are rapidly growing outdated and have by now been largely replaced by USB.

PCI

PCI Express is an extension standard to connect peripheral devices with the chipset of server processors. PCI Express is also abbreviated as PCIe and PCI-E.

Different generations of PCIe interfaces are now available. Most current servers support PCIe Generation 4 and PCIe Generation 5. The theoretically possible transfer rates double from generation to generation. For example, one data line supports up to 8 Gt/s (giga transfers per second) for PCIe Generation 3, up to 16 Gt/s for PCIe Generation 4 and up to 32 Gt/s for PCIe Generation 5.

However, when requesting PCIe generations, it should be noted that servers are not usually equipped exclusively with the latest generation, as there are still very few PCIe Generation 5 peripheral devices such as network cards, FC-HBAs etc., available. Only a few high-end GPUs already support PCIe Generation 5. Therefore, manufacturers rely on a mix of Generation 4 and Generation 5 slots, as the cost of Generation 4 PCIe is lower and too many Generation 5 PCIe slots can cause higher inter-line interference. The generations are compatible with each other (e. g. Gen 3 card in Gen 4 slot or vice versa) and agree on the best possible speed. In the server environment, PCIe slots are offered with 4, 8 or 16 data lines. For example, a third-generation PCIe slot with 4 data lines is abbreviated to PCIe Gen3 x4. Servers can be extended with PCIe slot cards, which are plugged into the PCIe slots. PCIe slot cards are available in various form factors. A difference is made between cards with a »full height and full length« form factor and smaller »low profile« form factors. Depending on their design and technology, servers have different PCIe slots, e. g. PCIe Gen3 x8 full height / full length or PCIe Gen3 x4 low profile. In individual cases, one must check whether the desired PCIe expansion cards can be installed in the server.

Compute Express Link

Compute Express Link (CXL) is an open standard for the high-performance connection of the processor to the memory or other components in the server. CXL is based on the serial physical and electrical interface PCI Express (PCIe). The serial communication and pooling capabilities of CXL can overcome the performance limitations of conventional DIMM memory at high memory capacities and optimise communication between other components in the system, such as accelerators.

Ethernet

Ethernet is a LAN standard for data transmissions between two terminals. Ethernet was developed very early, in the 1970s, after which the IEEE developed the technology further and standardised it in the 1980s. Today, 1 Gb/s Ethernet (1 GbE) or 10 Gb/s Ethernet (10 GbE), and 25 GbE are standard. Server performance density is on the rise and better hardware loads are realised through virtualisation or software defined storage (SDS), requiring a network connection with the highest-possible data throughput. In part, even faster connections with 40 (by now outdated), 50, or 100 GbE are used.

Servers generally already have several ethernet ports integrated into their mainboard. Depending on the speed and required cable length, the connections are realised using copper technology (Base-T with RJ-45) or fibre optics (e. g. SFP+ or SFP28). Suitable cards can provide additional ports.

Network card drivers within an operating system make it possible to bundle several GbE ports to improve data throughput. Here, the channels can be operated actively / passively or actively / actively, with simultaneous disruption protection.

Storage (SAS / SATA, Fibre Channel, iSCSI, NVMeoF)

Storage can be connected to a server system in many ways. A difference is generally made between local storage within the server system and external storage outside of the server system. Access to local storage is generally realised through SAS / SATA or NVMe. Access to external storage can be realised using SAS, FC, iSCSI, or NVMeoF. NAS solutions should not be considered here.

SAS / SATA

SAS is the abbreviation of Serial Attached SCSI. SAS has replaced parallel data transfer through a shared SCSI Bus. The SCSI protocol is still used, but transfer takes place serially, through a point-to-point connection between the SAS controller and SAS or SATA hard disks. Contrary to an SAS controller, a true SATA controller can only address SATA hard disks. The SAS / SATA controller is usually already integrated in the server systemboard; alternatively, it can be retrofitted as a PCI card. An SAS controller with external ports is required to connect an external SAS / SATA hard disk unit. In this case, the RAID groups are formed through the server SAS controller, and can now be used directly by this server. If several servers access a common SAS / SATA storage solution, the storage solution must have one or two storage controllers with RAID function. For this, the server requires a SAS hostbus adapter with external ports. The RAID groups are formed through the storage controller of the external SAS / SATA storage solution Only a small number of servers can be connected to the storage solution, because each server is connected directly to an SIS port of the storage controller. The servers and SAS / SATA storage solution cannot be located more than a few metres apart.

Fibre Channel

Fibre Channel (FC) is an optimised protocol for serial high-speed transfers of large data quantities. Today, transfer rates from 16 to 64 Gb/s are common. Fibre Channel can be transferred through both a copper as well as a fibre optic line. Copper lines are only used within the FC storage solution. Depending on the storage solution, Fibre Channel hard disks or SAS / SATA hard disks are used in the storage backend. Servers, FC storage solutions, and tape libraries are connected to each other with fibre optic cables through a fibre channel network (storage area network), allowing for shared use of all resources. The server must be supplemented with a fibre channel hostbus adapter (PCI card). The availability of data is paramount when it comes to fibre channel. Therefore, the data of the fibre channel storage should generally be accessible through at least two redundant paths. All components are to be designed redundantly, to make sure that a disruption to one of the components does not affect data availability. The fibre channel network is highly scalable and supports a large number of servers. The servers and FC storage can be located several kilometres apart.

iSCSI protocol

iSCSI (Internet SCSI) provides an option to connect external storage solutions. SCSI packets can be transmitted through an ethernet infrastructure. iSCSI solutions are affordable with data transfer rates of 1 Gb/s or 10 Gb/s, building on the existing knowledge of networks within the IT department. iSCSI has a similarly high overhead as fibre channel and the TCP / IP protocol is not optimised for block-based data access. Therefore, it is sensible to keep the iSCSI network separate from the rest of the LAN. The iSCSI network topology allows many servers to access shared iSCSI storage. With the help of an iSCSI driver (iSCSI initiator), servers can access the operating system through a standard network card. With new and powerful processors, the additional load on servers is very low. iSCSI offload and iSCSI boot functionality can be realised using an iSCSI host bus adapter or a multi-functional network card. A problematic feature of iSCSI is the low throughput. To counteract this problem, data should be divide over several volumes, made accessible through various iSCSI ports of the iSCSI storage solution. The servers and iSCSI storage can be located several kilometres apart, as long as latencies are low.

NVMeoF

NVMeoF (NVMe over Fabrics) is a relatively new protocol to connect quick memories using existing connections such as ethernet or fibre channel as a transport medium. This technology is still new and relatively rare, which is why expert advice is always required in advance.

Bitkom represents more than 2,200 companies from the digital economy. They generate an annual turnover of 200 billion euros in Germany and employ more than 2 million people. Among the members are 1,000 small and medium-sized businesses, over 500 start-ups and almost all global players. These companies provide services in software, IT, telecommunications or the internet, produce hardware and consumer electronics, work in digital media, create content, operate platforms or are in other ways affiliated with the digital economy. 82 percent of the members' headquarters are in Germany, 8 percent in the rest of the EU and 7 percent in the US. 3 percent are from other regions of the world. Bitkom promotes and drives the digital transformation of the German economy and advocates for citizens to participate in and benefit from digitalisation. At the heart of Bitkom's concerns are ensuring a strong European digital policy and a fully integrated digital single market, as well as making Germany a key driver of digital change in Europe and the world.

Bitkom e.V.

Albrechtstraße 10
10117 Berlin
T 030 27576-0
bitkom@bitkom.org

bitkom.org

bitkom