WHITE PAPER

Technical Specifications in the Public Procurement of Computers

November 2007

Achieve 'Best Value For Money' and stay legal

The goal of this white paper is to help contracting authorities draft technical specifications in the tender documentation for public contracts for the supply of computer systems (desktops, notebooks and servers).

Public procurement legislation aims to enable contracting authorities to achieve 'Best Value for Money' in their purchasing activities by removing unjustified obstacles to the opening up of public procurement to competition. The way in which technical specifications are defined in the contract documents is critical to this purpose.

Whilst public procurement laws and regulations vary from country to country around the world, the main principles governing the drafting of technical specifications in the contract documentation of public tenders are substantially similar across jurisdictions:

- *transparency and openness of tenders to competition*, which requires that technical specifications be drafted so as to afford equal access for tenderers;
- *equal treatment of tenderers and non-discrimination*, which means that technical specifications should not discriminate, either directly or indirectly, against tenderers.

In particular, many countries prescribe not to use brands in technical specifications, unless objectively justified by the subject-matter of the contract or the need to describe the contracted goods, if these goods cannot be adequately described otherwise.

When purchasing computers, contracting authorities should identify the features of the CPU and the computer system they would like to purchase, taking primarily into account their intended use of the product and the available budget. To this effect, they should carefully consider all the features brought to the market by research and technological innovation in the manufacture of computers, with a view to identifying those features that best meet their requirements and help them achieve their goals.

While performance is important, other CPU and computer system features unrelated to performance are equally important and sometimes critical in relation to the contracting authority's intended use of the product. Sections 1 and 2 of the White Paper provide practical guidance on how to describe, respectively, performance and other CPU and computer system features unrelated to performance in the technical specifications.

This White Paper will therefore enable contracting authorities to *effectively achieve their 'Best Value for Money' objectives, while remaining compliant*.

TABLE OF CONTENTS

1. Performance 3

1.1. Using benchmarks to describe performance 4

1.1.1. What benchmark should I use? 4

1.1.2. How to use performance benchmarks 7

1.2. Other methods to describe performance 8

2. Features unrelated to performance 10

2.1. Energy efficiency and energy-efficient performance 11

2.2. More examples of features unrelated to performance 12

3. Sample technical specifications 15

This white paper is published on http://www.intel.com/go/tenders.

Updates of the white paper will be made available on this web page from time to time.

1. Performance

Amongst all CPU and computer system features, performance is certainly the first that comes to mind.

The 'Central Processing Unit' (CPU), also known as 'processor' or 'central processor', is the brain of a computer system, being the component that actually executes instructions. In personal computers (desktops and notebooks) the CPU is contained on a single silicon chip called 'microprocessor', and so in the PC world the terms CPU and microprocessor are used interchangeably. In large machines (e.g., servers), CPUs are contained on one or more printed circuit boards.

Referred to the CPU, performance indicates how fast a CPU does its main job - processing instructions. CPU performance is the result of various factors, some of which have to do with how the CPU is built, while others depend on the configuration of the computer system in which the CPU is integrated.

The following CPU features directly determine the CPU performance:

- internal architecture;
- clock speed, also known as frequency, which is measured in MHz or GHz (e.g., 1.86 GHz, 3.20 GHz);
- width of instructions/data that the CPU is able to process, which is measured in bit (e.g., 16-bit, 32-bit or 64-bit);
- number of execution cores per CPU (e.g., single core, dual-core, quad-core);
- size of the local cache memory of the CPU (e.g., 1 MB, 4 MB);
- internal architecture, frequency (measured in MHz: e.g., 800 MHz, 1066 MHz, 1333 MHz, etc.) and width (measured in bit: e.g., 16-bit, 32-bit, 64-bit) of the system bus; and
- multithreading and multitasking enabling technology.

The other hardware and software components of the computer system in which the CPU is integrated and their settings have also an impact, at varying degrees, on the CPU and computer system performance. Amongst others:

- chipset;
- motherboard;
- system (or main) memory;
- hard disk;
- graphic card;
- operating system;
- software applications loaded on the machine.

From a performance assessment perspective, therefore, *it is important to consider all the aforementioned features together, instead of focusing only one or two of them*. For example, although the clock speed of a CPU is one of the factors impacting performance, it cannot generally be regarded as a reliable proxy for the performance of the CPU and, even less, for the performance of the computer system in which that CPU is integrated.

1.1. Using benchmarks to describe performance

Benchmarks are an established and widely accepted method to gauge performance, and are commonly used by contracting authorities in the technical specifications to describe the required level of performance for the computer systems they would like to purchase.

The performance assessment is achieved by running a special computer program called 'benchmarking program' on the system under evaluation. The benchmarking program executes a series of standard tests and trials simulating a particular type of workload on the system, and provides a final performance score based on the metric specific to each particular benchmark.

Provided the correct methodology is followed (See Section 1.1.2.), performance benchmarks also enable comparison of the performance of two or more computer systems. It is exactly this comparison capability of performance benchmarks that makes them useful tools in the context of public tenders.

It is extremely important to say from the outset that benchmarks gauge either the *overall performance of a computer system* (this is true of 'system-level benchmarks' such as those developed by BAPCo*) or, as a minimum, *the combined performance of CPU, memory and compiler* (this is the case of 'compute intensive' benchmarks such as those developed by SPEC*). In this light, benchmarks may provide only *an indirect and not entirely accurate measurement of the performance of the CPU* integrated into the computer system benchmarked.

While this may perhaps seem obvious, for the sake of clarity it should also be noted that *performance benchmarks tell nothing about other important CPU or computer system features unrelated to performance*. Energy efficiency and features such as those enabling greater security, reliability or manageability, therefore, should always be described separately in the technical specifications. (See Section 2)

1.1.1. What benchmark should I use?

There are many different types of performance benchmarks around, and for a contracting authority to choose the most appropriate benchmark(s) for a specific tender is not always easy.

The wrong benchmark may cause contracting authorities to buy a computer system which is different from what they really intended to buy, and in some cases may even lead to discriminate against specific vendors and their products.

At a high level, a good performance benchmark should always have the *minimum attributes* described below:

- **Relevant:** Contracting authorities should choose a benchmark or a combination of benchmarks, which measure performance via tests representative of the actual everyday use that contracting authorities intend to make of the system.

  If the benchmark is not relevant, contracting authorities would risk purchasing a product which is different from what they need.

- **Regulated and recognised:** Contracting authorities should prefer performance benchmarks developed and regulated by well-recognised

industry consortia, i.e. independent non-profit standardization bodies with wide industry representation. Being the result of open and healthy debate and having well-defined methodologies, these benchmarks are generally *objective, impartial, reliable, reproducible and widely accepted*.[1]

Privately developed benchmarks may be instead influenced by individual computer companies, which would make them *insufficiently representative and sometimes even biased*. Furthermore, the general lack of a well-defined methodology in these benchmarks makes them *unsuited to compare the performance of different computer systems*, especially in the context of public tenders, as variations in the methodology influence the scores and make them not-comparable.

This does not necessarily mean that all privately developed benchmarks should not be used. The performance benchmarks of Futuremark Corporation*, for example, are developed following a neutral and transparent development process, build on Futuremark's extensive benchmark development experience and are widely recognised in the industry.[2]

- **Up-to-date:** Contracting authorities should always use the most recent version available of any given benchmark.

  Performance benchmarks are continuously updated and new benchmarks are regularly introduced to keep pace with relentless development and innovation in the computer industry. A benchmark that is not up-to-date loses its capability to effectively measure performance, as it will not account for the different way in which computer systems would work following the implementation of new features (e.g., multithreading and multitasking capabilities).

  Importantly, an outdated benchmark would also lose its capability to provide a comparison between two or more computer systems that effectively and genuinely reflects their real performance. If the innovative features of a better performing computer system are not accounted for by the benchmark, the performance benefits of that computer system would not be recognised vis-à-vis another computer system that, while inferior to the former in terms of real performance, could well achieve a higher score under the outdated benchmark. In short, outdated benchmarks lead to discrimination against the most recent, innovative and better performing products.

The performance benchmarks recommended by Intel for desktops, notebooks and servers, at the time of publication of this white paper, are listed below:

Recommended benchmarks for desktops

With regard to desktops, Intel recommends the use of **BAPCo SYSmark 2007 Preview*** or later versions. This benchmark measures the performance of a computer system, in a Microsoft Windows XP* or Microsoft Windows Vista* operating system environment, by running relevant real-world computer programs typically used by business users.

SYSmark 2007 Preview features four test modules focusing on various scenarios (both a scenario rating and an overall rating are possible):

[1] Business Applications Performance Corporation (BAPCo), Standard Performance Evaluation Corporation (SPEC) and Transaction Processing Performance Council (TPC)* are all well-recognised industry consortia. For more information see, respectively, http://www.bapco.com, http://www.spec.org and http://www.tpc.org.

[2] For more information see http://www.futuremark.com.

- *Office Productivity:* performance common office application tasks such as communication, data processing, project management and file manipulation;
- *3D Modeling:* creation of an architectural presentation for a construction project with a flyby animation and a photorealistic image of the completed building;
- *E-learning:* creation of online material by integrating images, video and audio content in a web page for training; and/or
- *Video Creation:* creation of a video with special effects and images from multiple sources (the video is rendered in two formats – a low-bit rate for online streaming and a high-bit rate for downloading).

The following performance benchmarks could also be used as an alternative to BAPCo SYSmark 2007 Preview:

- **SPEC CPU2006***, which provides performance measurements that can be used to compare compute-intensive workloads on different computer systems (SPEC CPU2006 contains two benchmark suites: *CINT2006* for measuring and comparing compute-intensive integer performance, and *CFP2006* for measuring and comparing compute-intensive floating point performance); or
- **PCMark05***, which gauges both system level and component level performance.

Recommended benchmarks for notebooks

With regard to notebooks, Intel recommends the use of **BAPCo MobileMark 2007*** or later versions. This benchmark measures the battery life of notebook computer systems, in a Microsoft Windows XP or Microsoft Windows Vista operating system environment, while providing also a performance evaluation that reflects their typical day-to-day use by business users. The performance is measured *while the notebook is not plugged-in to A/C mains power*, which makes it ideal for the 'mobile worker' usage model.

MobileMark 2007 features four test modules focusing on various scenarios:

- *Productivity 2007:* it measures the performance and battery life of a notebook when carrying out typical office productivity activities such as browsing and creation of text documents, spreadsheets and presentations;
- *Reader 2007:* it measures the battery life of a notebook when reading e-books on it;
- *DVD 2007:* it measures the DVD playback battery life of a notebook; and/or
- *Wireless 2007:* it measures the web browsing battery life of a notebook in a wireless environment.

MobileMark 2007 does not return one overall score, but provides distinct ratings for each of the four aforementioned scenarios. With regard to the Productivity 2007 module, both a performance rating and a battery life rating are possible.

For a performance rating more reflective of the computer system's real capacity for handling performance sensitive application, *when a notebook is plugged-in to A/C mains power*, a performance-specific benchmark such as **BAPCo SYSmark 2007 Preview** (described in the Desktops section above) is recommended.

In this case, MobileMark 2007 would be used only to measure the battery life of the notebook.

Recommended benchmarks for servers

With regard to servers, Intel recommends the performance benchmarks listed in the table below and classified in consideration of the workload they use for the measurement:

WORKLOAD	PERFORMANCE BENCHMARK
Enterprise Resource Planning Customer Relationship Management	– **SAP Standard Application Benchmarks*** – **Oracle Applications Standard Benchmark*** – **SPECjbb2005*** – **SPECjAppServer2004*** – **TPC-C*** – **TPC-E*** – **Siebel PSSP***
Virtualised environment	– **vConsolidate***
Online Transaction Processing	– **TPC-C** – **TPC-E**
Data Warehousing / Data Mart Data Analysis / Data Mining	– **TPC-H***
Email	– **MAPI Messaging Benchmark 3 (MMB3)*** – **Lotus NotesBench***
Streaming Media	– **Windows Media Load Simulator (WMLS)***
Web Serving	– **SPECweb2005*** – **WebBench***
Application Development Scientific / Engineering	– **SPEC CPU2006**

1.1.2. How to use performance benchmarks

Following the right methodology is extremely important when using performance benchmarks, as *variations in the way a benchmark is run may lead to the results being unreliable and not comparable, which would affect the legality of the tender* and put contracting authorities at risk of being challenged over the award of the contract.

If performance benchmarks are to be used in a public tender, the legal principles of transparency and non-discrimination would mandate that the contracting authority defines in the contract documents the detailed methodology for running the benchmarks, reporting the results and reviewing the submitted results.

As a minimum, *in the contract documents* the contracting authority should expressly:

- identify the benchmark or combination of benchmarks to be used, as well as setting a minimum score under each benchmark, taking into account their intended use of the system and the available budget (when a combination of benchmarks is used, scores returned by different benchmarks should not be averaged as different benchmarks measure different things and use different metrics);
- require that tenderers should strictly follow the methodology set out by the benchmark's developer for running the benchmark (so-called 'run rules');
- require that tenderers run the benchmarks on the *very same* computer system that they will actually supply in case they win the tender (specific provisions should be made to the effect that any software to be supplied as part of the contract be loaded on the computer system before this gets benchmarked, and that the hardware and software settings of the computer system benchmarked be exactly the same as the settings of the computer system that tenderers will actually supply in case they win the tender);
- require that tenderers submit, *together with their offer*, adequate and complete system configuration information (both on the hardware and software components, including their settings), as well as benchmark documentation for the tested system, so as to enable the subsequent verification of the accuracy of the results submitted;
- require that, *after the award of the contract*, the winner of the tender should submit adequate and complete system and benchmark documentation for the tested system ('Full Disclosure Report') to the industry consortium or the organisation regulating the benchmark for its publication, so as to allow other tenderers and any other interested party to check the results, in accordance with the principle of transparency of the tendering process.

Full compliance with the principles of transparency and non-discrimination would also mandate that, before the contract is awarded, contracting authorities should verify in their laboratory, or have verified in a third party independent laboratory, the accuracy of the results submitted by tenderers, at least on the short-listed systems. This measure is obviously needed to prevent any potential fraud on the part of some tenderers, which could affect both the contracting authority and other tenderers.

1.2. Other methods to describe performance

Whilst performance benchmarks are an established method to measure the performance of a computer system, they do have their limitations and do not fit perfectly all situations.

Benchmarks require technical knowledge, take a long time to run and, importantly, represent a significant cost for both tenderers and contracting authorities.

Each tenderer has to run a benchmark on multiple computer system configurations to identify the right configuration that, in his opinion, maximises his chance of winning. Furthermore, to have more reliable results that are audit-proof, each tenderer may need to run the benchmarking application on the same configuration more than once (in which case the final score submitted at the time of the offer represents an average

of the various trials). If more than one benchmark is required by the contracting authority, the two steps described above have to be repeated for all the benchmarks.

The time it takes to run benchmarks and the need to set up a dedicated benchmarking laboratory with dedicated, skilled personnel represent a significant cost for tenderers. Outsourcing the benchmarking activity is also very expensive.

Importantly, these costs are generally passed on contracting authorities by way of higher prices. Furthermore, in addition to these indirect costs, contracting authorities will also incur into direct costs to verify in their laboratory, or have verified in a third party independent laboratory, the accuracy of the results submitted by tenderers.

In conclusion, contracting authorities should carefully consider the burden and cost of using benchmarks to describe performance in the technical specifications. There may indeed be circumstances in which the use of benchmarks could be inappropriate, uneconomic for both tenderers and contracting authorities, discourage wide participation of tenderers and lead to higher prices.

Two examples of benchmarks not perfectly fitting the situation are small tenders and tenders for servers. In small tenders the use of benchmarks may discourage wide participation of tenderers and be uneconomic. If the number of units tendered is limited, the administrative burden and costs that both tenderers and contracting authorities would incur with benchmarking may not be justified. In particular, as benchmarking costs cannot be spread out across thousands of units, as it happens in larger tenders, such costs will become substantial and result in higher prices. As far as servers are concerned, performance benchmarks generally take very long and are very expensive to run, which makes it challenging to comply with the methodology necessary to guarantee the transparency of the tendering process (as described in Section 1.1.2. above).

In these cases, alternative and more practical solutions should be explored and adopted by contracting authorities to describe the performance of a computer system in the tender documentation.

A possible solution would be to use a combination of product features related to performance. When adopting this approach, contracting authorities should pay attention in selecting a combination of product features that do not discriminate, directly or indirectly, against any tenderers or products. With this proviso, a combination of product features carefully chosen can provide a good indication of the performance of a computer system, while avoiding the administrative burden and costs of benchmarks.

Another possible solution would be to describe computer system performance in terms of functional requirements. In this case, such parameters must be described sufficiently precisely in the contract documents to allow tenderers to determine the subject-matter of the contract and to allow contracting authorities to award the contract. In his offer, each the tenderer must prove to the satisfaction of the contracting authority and by any appropriate means that the computer system satisfies the functional requirements of the contracting authority, as set out in the contract documents.

2. Features unrelated to performance

Whilst performance is important, other CPU and computer system features unrelated to performance are equally important and sometimes critical, depending on the contracting authority's requirements and intended use of the product.

These features may help contracting authorities reduce their electricity cost, improve system and network security, reliability and availability, minimize the need for IT support and achieve better manageability of their computer assets and networks. This would result in reduced maintenance and running costs as well as increased user productivity, which in turn would significantly contribute to lowering the so-called *'Total Cost of Ownership (TCO)'*.

TCO analysis was first developed and applied by leading IT consultant firm Gartner[3] in 1987 and has been a consistent, industry-standard way of assessing IT costs ever since. TCO enables purchasers of a computer system to evaluate the economic value of their investment by providing them with a 360 degree cost assessment that is not limited simply to the upfront cost of purchase of the system, but also covers other categories of costs arising from the use of that system during its entire life cycle (electricity costs; administration and asset management costs; costs associated with the integration of new systems into the existing IT environment; upgrade costs; training costs; technical support and maintenance costs; costs associated with security breaches or system failure; loss of end-user productivity; decommissioning costs; etc.).

Importantly, *contracting authorities will be able to truly achieve their 'Best Value for Money' objectives only if they go beyond the upfront cost of purchase of a computer system and focus instead on the TCO of that system over its life time*.

To this purpose, contracting authorities should identify all the CPU and computer system unrelated to performance that may contribute to reduce the TCO of that system, taking into account their budget, requirements and intended use of the product. Once all such features have been identified, contracting authorities should describe them in the tender documents *precisely enough* to allow tenderers to determine the subject-matter of the contract and to allow the contracting authority to award the contract to the winner. Furthermore, the description of these features in the contract documents must also be *objective and non-discriminatory*.

It is important to clarify that *contracting authorities would not infringe the principle of non-discrimination by requiring product features unrelated to performance specific only to some tenderers or product manufactures, as long as contracting authorities can objectively justify their request* based on their requirements, their intended usage of the computer system and the subject-matter of the contract.

As a way of illustration, if only one product manufacturer in the market provides a certain security feature in a product to be procured by a contracting authority through a public tender, the contracting authority could still legitimately identify that feature in the technical specifications, if this is justified, for example, on the basis that the feature would significantly increase the security of computer systems used to store sensitive data.

[3] For more information see http://www.gartner.com.

2.1. Energy efficiency and energy-efficient performance

Amongst all features unrelated to performance, energy efficiency deserves a special mention. Energy efficiency has been important for servers for some time, and now desktops and notebooks are starting to receive the same degree of scrutiny.

Whilst computer assets in every organisation continue to increase to satisfy ever demanding requirements, energy consumption also increases. Greater electricity cost, limited energy availability and concerns on the environmental impact of producing electricity and energy waste are leading contracting authorities to pay greater attention to energy efficiency as an important element in their purchasing decisions.

As computer systems are purchased for their ability to deliver a specified level of performance and capability, *contracting authorities should focus on the energy cost required to deliver the requested level of performance and capability ('energy-efficient performance')*.

Energy efficiency and energy-efficient performance, however, pose some unique measurement challenges.

Reference to Thermal Design Power (TDP) to assess the energy efficiency of a computer system would not be appropriate. TDP is a specification indicating the thermally-relevant power-draw of processor used to design computer platforms. TDP, however, conveys neither performance nor the energy cost at that performance level.

While there are standard practices that have wide industry acceptance for gauging performance through purpose-built performance benchmarks (See Section 1.1.), there is no standard methodology for evaluating energy-efficient performance. Until an industry consortium-developed solution becomes available, as an interim solution Intel would recommend contracting authorities to adopt the methodology described at http://www.intelcapabilitiesforum.net/EEP/.

The proposed methodology is described with reference to desktops and notebooks, but it can also apply to servers with some adjustments, as the core principles would be the same. The methodology is based on two main elements:

- definition of workloads representative of the actual everyday use that contracting authorities intend to make of the computer system, and choice of a relevant benchmark (See Section 1.1.1.);
- definition of a usage context that simulates how the computer system will be used over the course of an average day (e.g., time the system spends executing the defined workload, sitting idle, being in standby or off).

In order to calculate the energy cost, the computer system must be monitored with a wattmeter for actual energy consumption while the computer system is being benchmarked. In short, the proposed methodology leads to achieve *final results reporting the delivered performance of a computer system at a given level of energy cost*, communicating both performance and the energy cost of delivering that performance. Metrics that attempt to combine these measures are best avoided, as they can obscure both the delivered performance and energy cost. This is especially true of taking ratios, as there is not a guaranteed linear correlation of the value of performance or energy cost.

Please refer to the link provided above for more detailed information.

Importantly, a contracting authority wishing to have energy-efficient performance as one of the criteria for the award of the contract should describe in detail the methodology for gauging energy-efficient performance in the tender documents, and should request tenderers to fully comply with such methodology.

2.2. More examples of features unrelated to performance

Additional examples of CPU and computer system features unrelated to performance are provided in the table below:

PLATFORM	FEATURE & BENEFITS
Desktops Notebooks Servers	***Hardware-assisted*** **data execution disabling technology** When combined with a supporting operating system, the data execution disabling technology of some CPUs and chipsets can prevent certain classes of malicious buffer overflow attacks. In a typical attack, a malicious worm creates a flood of code that overwhelms the processor, allowing the worm to propagate itself to the network, and to other computers. Malicious buffer overflow attacks pose a significant security threat to businesses, increasing IT resource demands, and in some cases destroying digital assets. These attacks cost businesses precious productivity time, which can equal significant financial loss. Data execution disabling technology allows the processor to classify areas in memory by where application code can execute and where it cannot. When a malicious worm attempts to insert code in the buffer, the processor disables code execution, preventing damage and worm propagation. Hardware-assisted data execution disabling technology therefore contributes to significantly improve system and network security, reliability and availability. In addition, by reducing the need for software patches aimed at buffer overflow attacks and for virus-related repairs, it frees IT resources for other initiatives.
Desktops Notebooks Servers	***Hardware-assisted*** **virtualization technology** The virtualization technology of some CPUs and chipsets improves the robustness, reliability and performance of software-only virtualization solutions. Virtualization technology enhances security, improves manageability and increases system and network availability, limiting downtime and maintaining worker productivity. At its simplest, virtualization allows a single platform to run multiple operating systems and applications in independent partitions at once without incurring significant emulation costs. With virtualization, one computer system can function as multiple 'virtual' systems. Virtualization has emerged as a compelling technology for server platforms, leading to lower hardware acquisition costs as well as improved performance efficiency. In a data centre, for example, multiple servers could be combined into one system, running different applications and operating systems, providing advantages for IT tasks such as server consolidation, legacy migration and 'live' migration. Client platforms are also increasingly utilizing virtualization technology to enable secure partitions for system back and security features. For example, IT managers could isolate a portion of a managed PC to perform system upgrades and maintenance without interrupting the end-user. IT managers could also create one desktop PC build that can function independently as both a business and personal system, keeping software loads and virus attacks separate, or one that runs different operating systems and software for different or legacy tasks.

Desktops Notebooks Servers	***Hardware-assisted*** **remote hardware and software asset tracking and management, with out-of-band system access** Some processors and chipsets contain built-in capabilities that, combined with supporting software, can allow IT managers to remotely discover, access, heal and protect their networked computer systems, even those that lack a working operating system or hard drive or are turned off (provided the computer system is connected to a power source and an active LAN port). Amongst other things, these capabilities would allow IT departments to: - improve inventory tracking and achieve a better asset management; - improve manageability and system/network security, by making it easier to keep software and virus protection consistent and up-to-date across the organization; - minimize desk-side visits and reduce IT support costs; and - reduce downtime and maintain end-user productivity with less disruptive, remote 'lights-out' security patch and software application updates.
Desktops Notebooks Servers	**DDR2/DDR3 or FBDIMM system (or main) memory** Last generation system (or main) memories based on the standards DDR2, DDR3 or FBDIMM enable greater performance and reduced power consumption. In addition, a computer system (especially a server) with DDR2, DDR3 or FBDIMM main memory will be much easier to upgrade in future, as these are likely to be the new industry standards for multiple generations.
Desktops Notebooks Servers	**Warranty, after-sales services and support** Warranty, after-sale services and support, including training, should also be carefully considered. This may be particularly important for servers, as they generally need maintenance and upgrades over their life time.
Desktops Notebooks	**Guarantee of a stable image hardware platform for at least 12 months** As IT managers deploy client systems in the enterprise, unexpected changes to the previously qualified platform can force software image revisions and hardware re-qualifications. The increased complexity to the client environment can raise hardware support and image management costs. A stable image hardware platform is a standardized configuration for key hardware components of desktop and notebook PCs (processor, chipset and wired/wireless LAN) and reference drivers that IT departments can deploy into the organization for a set period of time, usually 12 months. During this time, IT managers can develop and maintain a standardised software image consisting of an operating system, hardware drivers and application software that can be loaded on any new PCs deployed into the environment without concerns about compatibility between the hardware and software. The guarantee of a stable image hardware platform enables greater software image stability and helps IT departments reduce the number of system qualifications and client configurations in the environment, resulting in lower hardware support costs and improved IT responsiveness.
Notebooks	**Battery life and mobility form factor** The usefulness of a mobile PC is directly related to its portability and the longevity of its batteries. So is the level of productivity of the mobile worker. Contracting authorities are therefore recommended to request that notebooks should not exceed 6 pounds in weight and 1.5 inches in thickness, while providing a minimum battery life of at least 3.5 hours under BAPCo MobileMark 2007.

Notebooks	**Integrated wireless Local Area Network (LAN)** Wireless connectivity transforms the workday by giving users new ways to stay productive and efficient as they move through fast-paced days, both in and out of the office. Notebooks with integrated wireless LAN and supporting software enable users to seamlessly connect to intranet or internet whenever a wireless LAN hotspot is available, thereby providing greater flexibility and freedom to work on the go.
Servers	**Automatic redundant processor booting, 'hot-spare' parts, 'hot-swap' parts and memory mirroring** One of the key measurements for a server is its 'percentage up time'. This determines the time the server is available for users to access and the amount of unplanned downtime/year that the server may experience. Unplanned downtime can result from a number of different factors such as hardware failure, software problems or operator errors. Automatic redundant processor booting, 'hot-spare' parts, 'hot-swap' parts and memory mirroring increase server reliability and contribute to minimizing the possibility of unplanned downtime. Automatic redundant processor booting is a capability of some multiprocessor servers that, in case of system boot failure by the primary boot processor, automatically enables the system to reboot from other processor(s) installed on the system, maintaining high availability for mission critical applications. Hot spare parts are redundant components within a server, which are powered on, but not actively functioning in the system. In case of failure, a defective component is automatically replaced by a hot-spare component with no disruption for the server or the application. Parts identified as hot-swap are capable to be changed ('swapped') without having to turn off the server. Contracting authorities should request that servers be equipped with as many hot-spare, or as a minimum hot-swap, elements within them as possible. The main hot-spare / hot-swap parts in a server are power source, fans, hard-drives and main memory. Mirroring is a technique used to protect the data stored in memory (hard-disk or main memory). Mirroring provides the ability to duplicate the contents (from main memory or hard-disk) in second memory or hard-disk array, so that if one array fails, data are not lost and are instantly available to the users preventing any server downtime.

3. Sample technical specifications

A *very basic example* of technical specifications is outlined below:

EXAMPLE

Desktop computer system with the following product features:

Computer system performance Δ

- The computer system must achieve a minimum score of XXX under BAPCo SYSmark 2007 Preview.

Other features †

- Energy efficiency (energy-efficient performance): (…)
- Hardware-assisted virtualization technology: (…)
- Hardware-enabled remote hardware and software asset tracking and management, with out-of-band system access: (…)
- Guarantee of a stable image hardware platform for at least 12 months: (…)
- Main Memory: (…)
- Hard-drive: (…)
- Operating System: (…)
- Monitor: (…)
- Etc.

Δ The methodology for running the benchmark and reporting the results should be described in detail in the tender documents (See Section 1.1.2.)

† The description of these features in the contract documents must be sufficiently precise to allow tenderers to determine the subject-matter of the contract and to allow the contracting authority to award the contract to the winner. With specific regard to energy efficiency and energy efficient performance, the methodology to be used by tenderers to gauge it should be described in detail in the tender documents. (See Section 2.1.)

IMPORTANT! PLEASE READ

INFORMATION IN THIS DOCUMENT MAY BE SUBJECT TO CHANGE WITHOUT NOTICE. UPDATES MAY BE MADE AVAILABLE ON http://www.intel.com/go/tenders.

WHILST BEST EFFORTS HAVE BEEN MADE TO ENSURE THE ACCURACY AND COMPLETENESS OF THE INFORMATION PROVIDED IN THIS DOCUMENT, INTEL CORPORATION AND ITS WORLDWIDE SUBSIDIARIES ACCEPT NO RESPONSIBILITY IN CONNECTION THEREWITH AND DISCLAIM ANY LIABILITY FOR ANY TYPES OF DAMAGES (INCLUDING, WITHOUT LIMITATION, NOMINAL, COMPENSATORY, DIRECT, INDIRECT, SPECULATIVE, SPECIAL AND CONSEQUENTIAL DAMAGES) THAT CONTRACTING AUTHORITIES MAY EVENTUALLY SUFFER BECAUSE OF THEIR RELIANCE ON SUCH INFORMATION.

IT REMAINS CONTRACTING AUTHORITIES' EXCLUSIVE RESPONSIBILITY TO IDENTIFY THE REQUIRED FEATURES OF THE PRODUCTS AND SERVICES THAT WILL CONSTITUTE THE SUBJECT-MATTER OF THE CONTRACT, AND TO DRAFT THE TECHNICAL SPECIFICATIONS IN COMPLIANCE WITH APPLICABLE LAWS AND REGULATIONS.

NO EXPRESS OR IMPLIED LICENCE, WARRANTY OR INDEMNITY OF ANY TYPE (INCLUDING INDEMNITY FOR THE INFRINGEMENT OF THIRD PARTY IP RIGHTS) ARE PROVIDED BY INTEL CORPORATION AND ITS WORLDWIDE SUBSIDIARIES BECAUSE OF, OR IN CONNECTION WITH, THIS DOCUMENT.

Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at www.intel.com.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

SPEC, SPECint, SPECfp, SPECrate, SPECweb, SPECjbb are trademarks of the Standard Performance Evaluation Corporation.

**Other names and brands may be claimed as the property of others.*