Strategies for increasing data centre power efficiency

An overview of CERN's approach to energy efficient computing

Dr. Andreas Hirstius,
Sverre Jarp,
Andrzej Nowak

CERN openlab
19 February 2008

Table of Contents

1. Executive summary

Growing workloads and rising energy costs have made data centre energy efficiency a critical issue for many organizations, especially for those with older facilities that lack the power and thermal capacity needed to support growing needs. This is a challenge that will not go away anytime soon. According to a report published by the European Commission in mid-2007, the amount of electricity consumed by western European data centres can be expected to increase two-fold from 2006 to 2020.[1]

The cost of energy is also increasing. A recent report from Eurostat[2] shows that growing demand has lead to rising energy prices. Some countries, such as the United Kingdom and Norway, have added as much as 25% to the prices paid by household

[1] "European Code of conduct on data centres" – Paolo Bertoldi, Robin Murray; http://sunbird.jrc.it/energyefficiency/pdf/Data%20Centers%203%20july2007/Murray%20AIMS_Paris_%202%20July%202007.pdf

[2] "Statistics in Focus – Environment and Energy"; http://epp.eurostat.ec.europa.eu/cache/ITY_OFFPUB/KS-SF-07-080/EN/KS-SF-07-080-EN.PDF

consumers during 2006. On average in Europe, the cost of 1KWh increased by 9%. As this trend continues, a well-designed strategy for reducing electricity usage in data centres will provide organizations with significant and increasing financial benefits. It may also help them satisfy local or governmental environmental policies and improve their brand reputation by demonstrating support for eco-friendly activities.

Few organizations have stronger motivations for addressing these issues than CERN, CERN is the largest particle physics laboratory in the world, and is currently finalizing the construction of a new particle accelerator, the Large Hadron Collider (LHC), which will be used to extend our understanding of the fundamental structure of the universe.[3] A colossal computing infrastructure will be required to process the 10-15 petabyes of data that will be generated annually by experiments on this new system.

CERN's computing model is based on a high throughput implementation where independent jobs can efficiently be executed in parallel in a computer with multiple processors, provided there is sufficient memory. Because of this model, it is relatively easy to expand total capacity by adding more computers. Unfortunately, thermal issues in today's computers limit the capacity of CERN's current data centre. This represents a real challenge to the full deployment of the required resources.

This paper outlines the hardware and software strategies CERN is using to increase energy-efficiency and extend the life of its data centre, so it can meet these immense computing requirements. It also discusses the energy-efficiency benefits CERN has realized from the computing industry's move from single-core to multi-core processors. Intel's new and more power-efficient microarchitecture is also discussed. The authors hope this information will prove helpful for other organizations seeking to increase energy-efficiency and reduce costs as they grow their computing infrastructures.

2. The need for massive new computing resources

CERN, the European Laboratory for Nuclear Research, is in the finals stages of construction of a new particle accelerator, the Large Hadron Collider (LHC). At 27 km in circumference, the LHC will be the largest particle accelerator ever built, and will enable physicists to recreate the conditions that existed just a fraction of a second after the Big Bang.

Four detectors (ATLAS, CMS, ALICE, and LHCb)[4] are being built to record the results of these collisions, which will occur at a frequency of 40 MHz. The resulting data will be filtered in real-time, and the most relevant information will be transmitted to the CERN computing centre at continuous rates between 300 and 1200 MB/s for

---

[3] For extensive information about the Large Hadron Collider , including progress on its construction, visit the CERN website at: http://lhc.web.cern.ch/lhc/

[4] For information about each of the LHC experiments, visit the CERN website at: http://lhc.web.cern.ch/lhc/LHC_Experiments.htm

storage and analysis. The overall computing task is truly gigantic, since about 15 PB of data, including simulations, will be accumulated per year.

To prepare for this, the CERN Computing Centre is being equipped with state-of-the-art compute and storage servers interconnected by high-speed networks. Additionally, the High Energy Physics community has agreed to interconnect nearly 200 computing centres world-wide in a so-called "Grid", so that the LHC computing tasks can flow seamlessly to the most suitable site. The Grid could also become very important in optimizing power, since it can balance the load both within and among the individual sites. The "World-wide LHC Computing Grid" (WLCG) Collaboration has been established to manage this global effort.

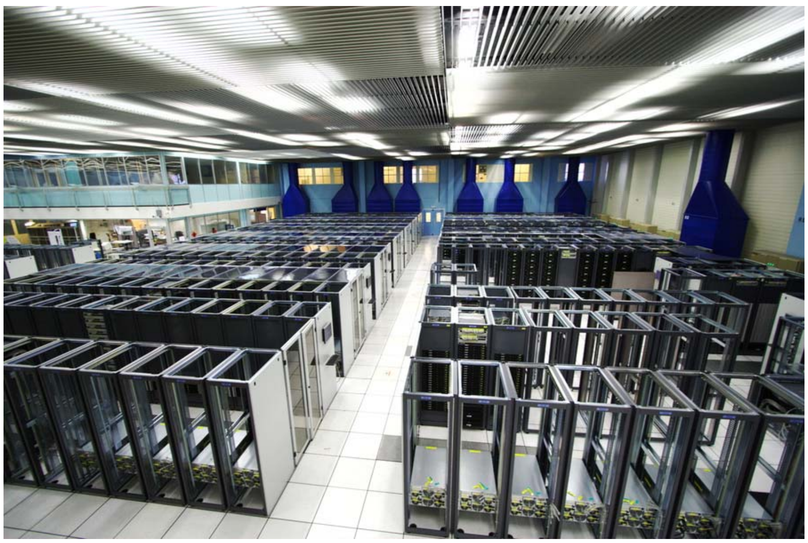

One of the biggest challenges CERN faces in successfully deploying the required resources is the power and thermal limitations of CERN's computing facilities. With the increased density of today's computing systems, space is not a key issue; but providing sufficient power and cooling for the massive number of new systems is a major challenge. For the first time in the history of CMOS processor technology, power and thermal factors have become a real threat to growth. This paper reviews the issues and the solutions CERN has implemented to tackle them.

3. Motivation for power savings

3.1. The current situation and historical limitations

The main CERN computing centre is located in Geneva, which has a temperate climate. The building which hosts the computing facilities was constructed in 1972, and was based on a design from the 1960s that was appropriate for hosting large mainframe systems. As with all computer centre buildings, it has a thermal limit. Using modern cooling technologies, the maximum amount of generated heat that can be removed by the cooling and ventilation system is 2.5 MW.

The old design inhibits certain thermal and efficiency optimizations, but others are being implemented. Quite recently, most of the cold aisles have been sealed off to avoid mixing hot and cold air in front of the racks.[5] Because of the 2.5 MW limit and the expected growth, the current facility is expected to reach its cooling limits by 2010. CERN is already considering building an additional 5 MW computing centre. To support research goals and contain costs, it will be essential to utilize both existing and future capacity as efficiently as possible.

There are currently about 2400 compute nodes used for physics applications in CERN's computing centre: about 400 single-core systems, about 700 dual-core systems, and about 1300 quad-core systems. In addition to the main computing centre, each of the four experiments associated with the LHC has its own data acquisition and processing centre, which transforms signals from the detector hardware into storable data and transmits the result to the main computing centre. These facilities incorporate between 500 and 2,000 machines each, and have even more restrictive requirements than the main centre. For example, one of those facilities is in a manmade cavern, 100 meters underground. In this situation, common cooling strategies do not apply. In addition, numerous stringent safety regulations have to be taken into account, which specify such details as the allowed materials from which computer components are made.

3.2. Projections for the upcoming years

Thanks to multi-core and many-core technologies developed by Intel and other manufacturers, computing resources have become more abundant, even though the power envelope per processor has remained more or less constant. Since the installation of a large number of dual-core servers, CERN's main computing resource pool has been running at around 1.5 MW. With the introduction of nearly 1300 new servers with quad-core processors between December 2007 and February 2008, the total dissipated power will rise by about 0.5 MW. Given the need for annual increases of a similar level, it is foreseen that the 2.5 MW limit will be reached within a year or two. At that time, the Laboratory could find itself unable to increase its computing capacity due to power and thermal constraints. As a consequence, new optimization measures have to be considered.

3.3. Power Consumption and Cooling Issues

For CERN's refurbished computer centre, one Watt saved on server power consumption has the added benefit of saving an additional Watt on power requirements for the cooling system. In other words, one Watt saved on power consumption translates into a total saving of two Watts. For many organizations, achieving a similar reduction would translate into direct savings. For CERN, the benefits are indirect. Since the computing requirements for the physics applications are basically without limits, any power-efficiency gains will simply be used to install additional computers.

---

[5] The mixing of hot and cold air is a common source of energy inefficiency in data centres. In an ideal scenario, all the cold air from a facility's chillers would be used to cool the computing equipment, and all the exhausted hot air would be returned directly to the cooling systems. Mixing indicates that some portion of the cold air is either not reaching the systems in need of cooling, or is reaching them only after being warmed to some extent by the hot air.

4. Measures taken to ensure power usage optimization at CERN

In order to tackle current and future power consumption issues, a number of software and hardware related solutions have been introduced. Several important approaches are discussed in this section, along with CERN's comprehensive hardware acquisition process, which aims at optimizing the current computing centre in terms of utilized power.

4.1. Hardware technologies

4.1.1. Processor and memory technologies

4.1.1.1 The value of multi-core processors

First and foremost, it should be mentioned that the multi-core evolution, which enabled a significant increase in computing power within a constant processor power envelope, has been extremely beneficial. For years, the heart of CERN's computing resources consisted of hundreds of ordinary PC workstations, each equipped with two Intel Pentium III or Pentium 4 processors. Although frequency increases in successive processor generations delivered needed performance gains, the corresponding rise in the amount of dissipated power eventually became a serious problem.

Fortunately, the move to multi-core processors enabled ongoing improvements in overall performance, without a corresponding increase in processor power consumption. Although the amount of memory had to be kept constant per core, the power savings compared to a non-multicore scenario have been impressive. It is questionable whether CERN would have been able to achieve its current computing capacity without multi-core processors.

Consider a scenario in which multi-core technologies were not available. In order to determine the performance and attractiveness of a computer system, CERN (like many computing centres) uses a benchmark called "SPECint," developed by the Standard Performance Evaluation Corporation (SPEC).[6] SPEC benchmarks are an industry standard for performance evaluation and are crafted to generate workloads very close to those running in typical computer centres. SPECint was chosen as the default benchmark, because changes in reported values match very well the changes in CERN's application performance. Dividing a computing system's SPECint result by the amount of power it consumes provides a measure of "SPECint per Watt". This result is quite important, since it indicates how much processing power is provided for each watt consumed by a running system.

[6] For detailed information about the Standard Performance Evaluation Corporation (SPEC) and the SPECint benchmark, visit the SPEC website, at: http://www.spec.org/

Around 2004, CERN had been using mostly enterprise class Intel processors, namely Pentium 4 Xeon processors, based on the Intel Netburst microarchitecture. The systems using these processors, which made up the backbone of CERN's computing centre at the time, had a SPECint/Watt ratio of about 6.

With the introduction of dual-core systems based on the Intel® Core™ microarchitecture (first used in Intel's "Woodcrest" processors), the SPECint/Watt ratio rose to values between 20 and 30. Today, Intel's quad-core systems ("Clovertown") are beginning to dominate in CERN's computing centre, since they provide a similar SPECint/Watt ratio to dual-core, while occupying less physical space. Intel's next generation quad-core Xeon processors, called "Harpertown", which are now available, improve the SPECint per Watt ratio even more. Had the new systems not been made available, CERN would have faced an impossible challenge of upgrading with computers using the Netburst microarchitecture. In order to maintain growth in terms of processing power, clock speeds would have had to exceed 10 GHz.

The following graph (Figure 1) shows a projection of CERN's computing needs through 2012 (in kilo-SPECint), along with the number of servers required to deliver that performance if they were based on Netburst processors (middle row) and if they were based on quad-core processors using the Intel Core microarchitecture (front row).

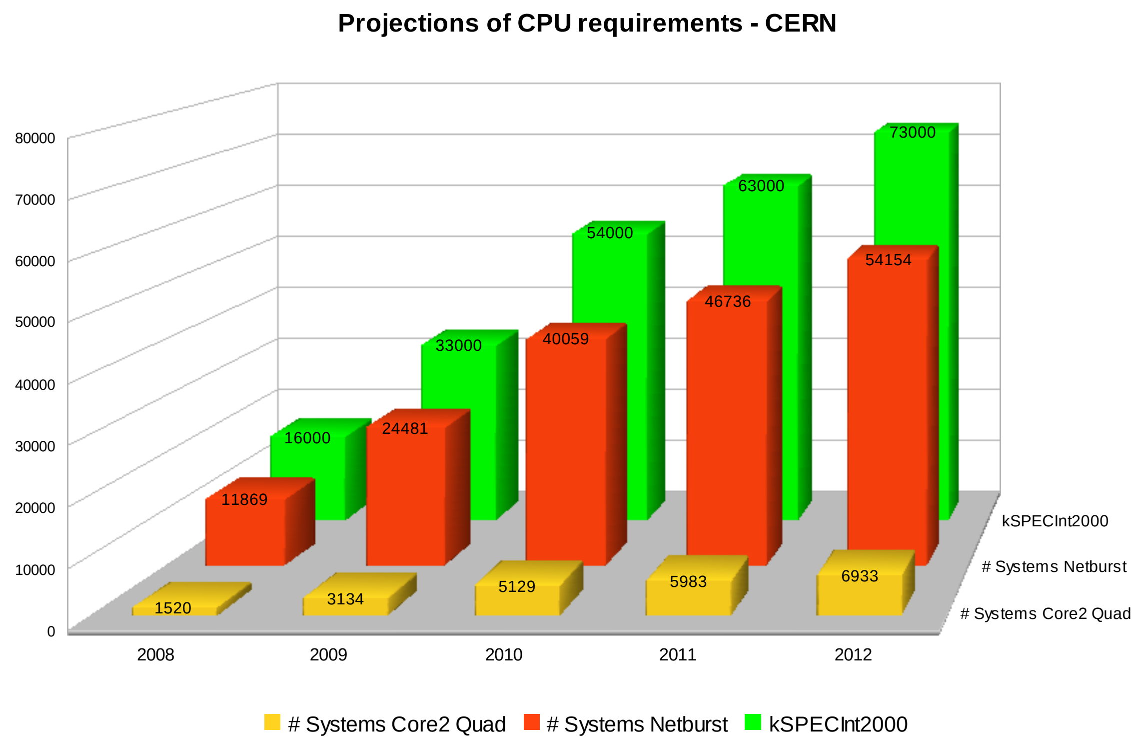

Figure 1. The number of servers required to meet CERN's projected computing requirements through 2012. The back row shows CERN's projected computing requirements. The middle row shows the number of Netburst based servers required to meet those requirements. The front row shows the

number of quad-core servers (based on the Intel Core microarchitecture) required to deliver the same level of performance.

One can also project the corresponding power consumption by comparing a hypothetical computing centre that would use only Netburst-based processors to a computing centre using the latest quad-core processors based on the Core microarchitecture. The difference in projected power consumption is striking, with the quad-core processors reducing total power consumption by nearly a factor of five (Figure 2).

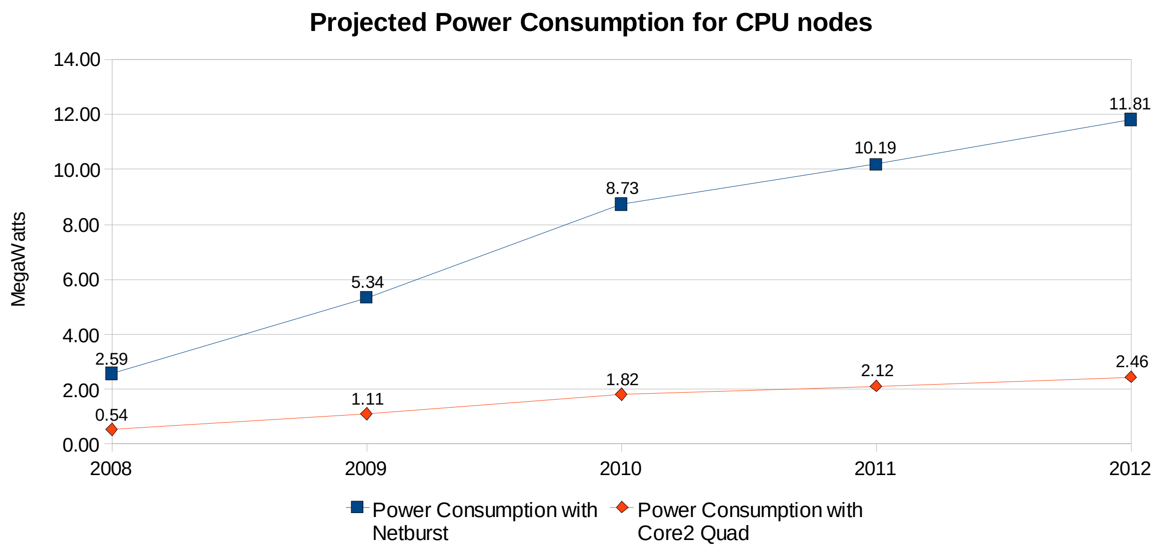

Figure 2. Projected power consumption for the main CERN data centre through 2012.
If CERN were using only quad-core processors based on the Intel Core microarchitecture rather than single-core processors based on the older Netburst microarchitecture, it would reduce power consumption by nearly a factor of five, as shown here.

4.1.1.2 Constraints due to memory power consumption

Forthcoming technologies in the CPU world will enable even more efficient processing with more cores per processor. Unfortunately for CERN, adding additional processing resources also means adding additional memory (RAM), as most physics applications have a requirement of 2 GB of memory per process (per processing core). As a result, memory has become a significant factor in power calculations. Depending on the technology, memory modules consume between 5 - 10 Watts per GB. Currently, modules with 2 GB are the most cost efficient solutions, and, depending on the technology details, consume between 10 - 15 Watts.

Using quad-core processors it is possible to build very cost-effective dual-socket computers with 8 cores and 16 GB of total memory. In this scenario, the memory alone consumes about 80 - 120 Watts. To help with this issue, Intel is planning to shift its focus from FB-DIMM memory modules to other, more energy-efficient, technologies in the near future, which will enable power savings in this area, as well.

4.1.1.3 Balancing processor frequency versus energy efficiency

In choosing the most energy-efficient processor, it is also important to consider the relationship between so called "Thermal Design Power (TDP)" and processor frequency. For a specific processor family, the processor with the highest frequency and cache will typically provide the highest performance. However, it will also have the highest TDP. The processor with the second highest frequency typically provides somewhat slower performance but has a much lower TDP.

For example, decreasing the frequency from 3.2 to 3.0 GHz might reduce computing performance by less than 10%, but decrease power consumption by as much as 30-40% (from 150W to 95W). Though individual jobs will run somewhat slower due to the lower processor frequency, total throughput per energy consumed may be much higher. Since CERN's primary goal is to get the best possible throughput within a given power envelope, choosing processors with a somewhat lower frequency is often the best approach.

4.1.2. Power supplies

CERN is interested in total power consumption for a given computing system, which includes the power consumption of all the internal components, including the system's power supply. The power supply is characterized by several properties. One of these is efficiency, which describes how efficiently the power supply can convert the incoming AC voltage into the different voltages that are needed inside a computer. With high quality power supplies, this efficiency can be on the order of 85 - 90%.

Another property of a power supply is the so called *power factor*. This is the ratio of the "*real*" or "*active*" power, which is the ability of a circuit to perform work, and the "*apparent*" power, which is the product of Voltage and Current in the circuit. The power factor of a computer power supply is usually load-dependent. At a low load, when the computer is idle, the power factor is smaller than under full load.

High quality power supplies tend to have a fairly high power factor under low load conditions and an even better power factor, close to 1.0, under full load. An overall low power factor increases losses in the power distribution system. This leads to higher energy costs, because most, if not all, major industrial users of electrical power are charged according to their usage of *apparent* power. Regulations for power factor corrected power supply already exist, but their reach is somewhat limited. Since power supplies are ubiquitous in every IT environment, enforcing stricter regulations would have a tremendous impact.

The industry has already developed a number of strategies that can help organizations address the power supply issue. For example, in blade systems a small number of large and mostly redundant power supplies are used to provide power for a high number of compute nodes. Since large power supplies tend to be of much higher quality than smaller ones, this can substantially improve overall energy-efficiency. There are also several approaches that involve performing the AC to DC conversion for the entire data centre, so only DC to DC conversion is required in the individual computers. However, though this strategy is generally more efficient, it may still be inefficient in certain cases.

4.1.3. Thermal issues and cooling

Since basically all the power that goes into a computer is converted into heat, extracting heat is a vital issue for a data centre. The process starts inside the individual computers, where there are two main approaches: air cooling and liquid cooling.

Direct liquid cooling is, in general, more efficient for heat exchange. However, it poses problems for two primary reasons. First, there are at least two main sources of heat in a computing system, processors and memory, and it is very difficult to properly cool memory using liquid cooling. Second, providing liquid cooling to a large number of computers would be very challenging. For these reasons, CERN is not using direct liquid cooling at this time.

Air cooling is still the standard cooling solution in most data centres and there are a number of factors to be taken into account. Almost all computers use fans to move air through the enclosure. Commonly used servers are one height unit (1U) high, and use a relatively large number of small fans that spin very fast, sometimes at more than 10,000 rpm. Such fans are quite inefficient compared to the much larger fans that can be used in computers with taller form factors, such as 2U, 4U and blade systems. The larger fans can spin more slowly and are typically much more efficient, which leads to a significant improvement in overall power efficiency.

A third cooling solution is based on water cooled heat exchangers mounted in the backs of the computer racks. The computers themselves are air cooled, but their hot exhaust air is cooled in the heat exchangers. Such solutions are being used in the online compute farms for the LHC detectors, but only because they are located in constrained environments with limited air cooling capabilities.

4.1.4. Computing centre layout

The CERN computing centre was built for large mainframe computers which had different cooling and power supply requirements than modern rack-mount servers. The computing centre was later refurbished to accommodate as many servers as possible. The major change introduced was a redirection of the airflow. Instead of blowing cool air in from the ceiling, it is now pushed through perforated tiles in the raised floor. A few compromises had to be made during the refurbishment, since the computing centre was in full production and only very short service outages could be tolerated.

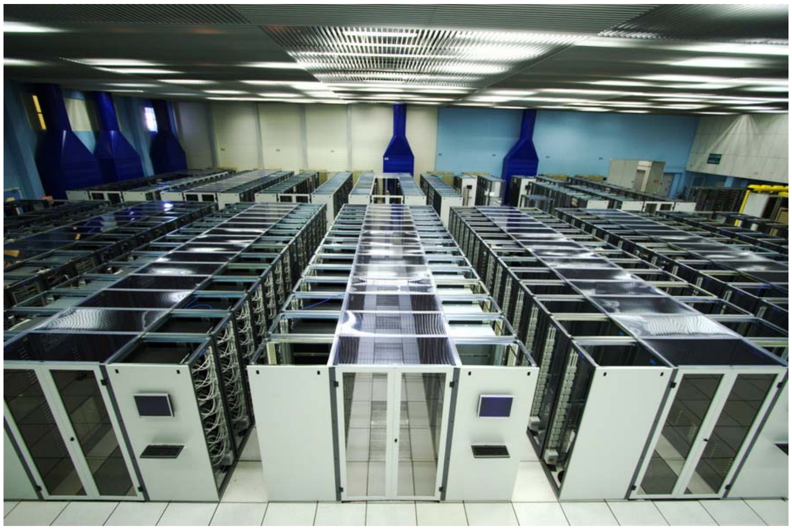

CERN uses a standard "hot aisle – cold aisle" approach in its data centre layout. In a cold aisle, two rows of racks are aligned with the front sides of the computers facing each other. In a hot aisle, the backs of the computers face each other. Cold air is blown up into the cold aisles, and the hot air is extracted from the hot aisles by convection. To increase efficiency, the cold aisles are being completely enclosed. This prevents hot air from "looping" into the cold aisles and keeps cold air from being blown off into the large air volume above the racks.

There are other possibilities for increasing cooling efficiency. Unfortunately, they can not always be used in the CERN computing centre. One option, for example, is to actively extract hot air from the hot aisles, a technique that is used in many new installations.

4.2. The hardware acquisition process at CERN

4.2.1. Including power efficiency in TCO projections

The acquisition process for installing computers in the CERN computing centre takes energy-efficiency into account. Each tender for compute nodes requests a certain capacity of computational performance, measured in SPECint2000 units. In order to ensure that the overall most cost-effective solutions are obtained, power consumption is measured under very well-defined conditions. The resulting power consumption is then multiplied by 6 CHF/W[7], which reflects the estimated electricity costs over a three year lifetime. By including this cost in the evaluation process, CERN not only purchases the overall most cost-effective systems, but also gives the suppliers an incentive to provide very efficient power solutions.

Computer systems are usually bought with a three-year warranty, and continue to be exploited until they reach the end of their warranty period. Until recently, energy consumption was not a driving force for decommissioning systems. This has changed now, and about 2100 single-core systems based on the Netburst processor microarchitecture are being retired to make way for more power efficient systems.

[7] Swiss francs per Watt

4.2.2. Measuring power

To measure power consumption, a power meter is installed in the primary AC circuit of the computer. Only the total power consumption of the entire computer is of interest, since it will include all consuming components and inefficiencies. Before testing, "Scientific Linux CERN" is installed on the system. This is the standard Linux version in use at CERN, and is based on Red Hat Enterprise Linux.

The first measurement is conducted while the computer is idle; meaning that nothing but the operating system is running. The second measurement is performed under two specified loads. The applications used were chosen because they generate workloads that closely model those of normal physics applications. The first application stresses only the processor; the second one also stresses the memory subsystem. The applications are replicated on all the available cores in the computer.

System utilization is also important in determining the most energy-efficient computers for a given environment. The average utilization of CERN's central computing facilities is currently about 70 – 80%, and will evolve towards 85 - 90% when the LHC experiments start taking data. To reflect these conditions, the power consumption value used in the adjudication process consists of 20% idle power consumption and 80% power consumption under load. (Most computing centres seem to see utilization figures below 50%, and should evaluate system power consumption accordingly.)

4.2.3. Power-sensitive benchmarks

SPEC has recently released a benchmark called "SPECpower,"[8] which measures power consumption under a certain workload. CERN participated in the beta testing of this benchmark. In time, this benchmark, combined with a suitable custom workload, could replace the current procedure for measuring power consumption in the tendering process for new computers.

4.2.4. Experience so far

Through its tendering process and earlier in-house power measurements, CERN has collected data that enables historical comparisons to be made on server power consumption. When comparing the measurements of a computer built with quad-core processors based on the latest Intel Core microarchitecture to a computer using processors based on the Netburst microarchitecture, the results are rather impressive. The best way to compare the two microarchitectures is to look at the performance per Watt divided by the CPU frequency. The system based on Netburst processors reached only 2.21 SPECint/(Watt*GHz) whereas the system based on the quad-core processors "Harpertown" reached 12.73 SPECint/(Watt*GHz). This implies that the quad-core based solution is more efficient per GHz by a factor of 5.7.

---

[8] For more information about the SPECpower benchmark, visit the SPEC website, at: http://www.spec.org/power_ssj2008/

Interestingly, the move from systems with dual-core processors (“Woodcrest”) to systems with quad-core processors (“Harpertown”) does not show such an impressive jump in efficiency. This may seem surprising, since the quad-core processors basically deliver twice the compute power of a dual-core processor, but with the same TDP. However, as already discussed, systems with quad-core processors require twice as much memory in the CERN computing environment. Since this memory currently consists of FB-DIMM memory, doubling the memory size means that the memory is now consuming about as much power as the processors. The internal cooling fans also use more power to cool the memory and some additional power is needed by the chipset to exploit the additional memory. For these reasons, in CERN’s case, the quite significant improvement in power efficiency at the processor level does not translate into a significant improvement in power efficiency per server.

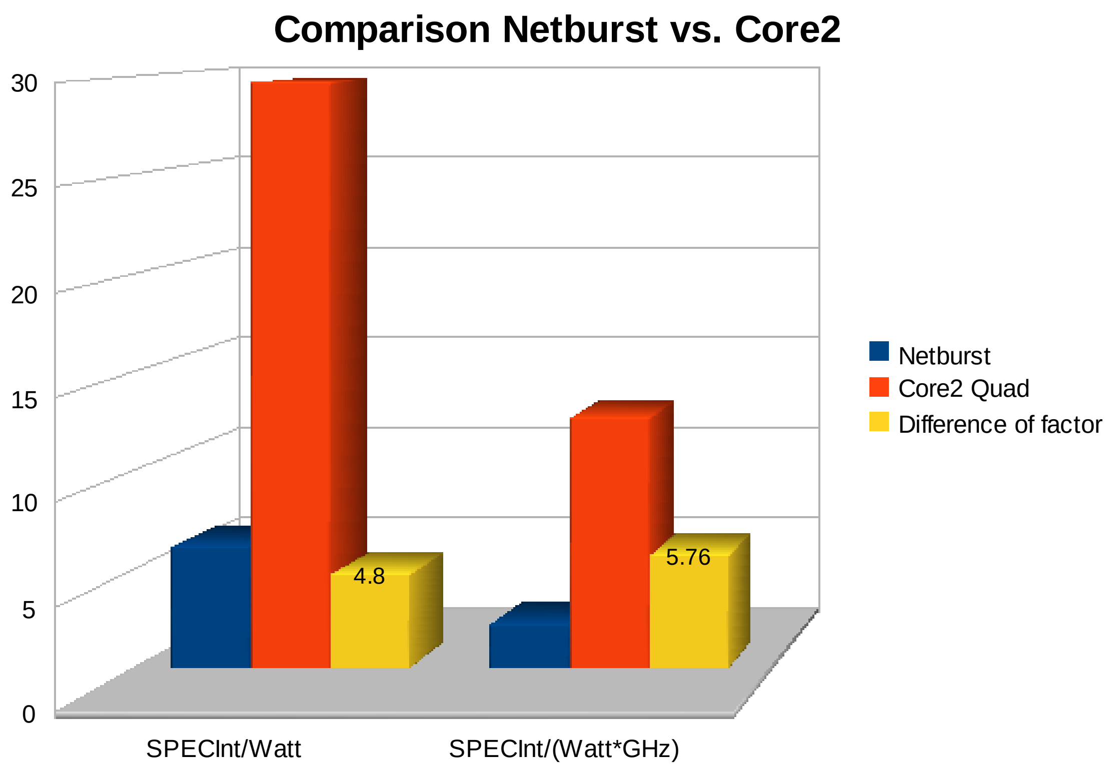

Figure 3. Performance as a function of energy consumption and processor frequency.
Quad-core processors based on the Intel Core microarchitecture deliver 4.8 times better performance per Watt than earlier single-core processors based on the Intel Netburst architecture, and 5.76 times better performance per Watt per gigahertz. This last figure shows the efficiency improvements in the architecture itself, independent of processor clock speed.

4.3. Software technologies

Other sources, including the Gartner Group, have indicated that substantial energy savings are possible through software modifications, as well as hardware optimization.[9] CERN's experience corroborates those claims, and indicates that selecting proper software technologies can be at least as important as choosing optimal hardware. The following sections discuss some of the most prominent software techniques used at CERN to reduce power usage and increase computing efficiency.

4.3.1. Virtualization

Virtualization is a software technique used to run multiple operating systems (each typically representing a "service") on a single hardware system. Though many services are only occasionally busy, the power consumed by the idle hardware is almost as much as the power required for active operation. Consolidating such services onto a smaller number of servers using virtualisation can provide significant energy savings.

In the CERN environment, only a relatively small number of services run at low CPU utilisation, and most of those services are already running in a virtualised environment. Potential energy-savings from virtualization may be much greater for organizations with lower average utilization rates.

4.3.2. Platform optimization

There are numerous ways to tune a system to provide optimized energy usage. One example is online processor frequency scaling. When the processor is idle, its frequency can be reduced automatically using control software. Another technique is to put the idle processor to sleep. This can be done by using external management interfaces, such as IPMI, which is offered by Intel (the underlying system software needs to be made aware of this process). These strategies are not extremely beneficial in CERN's environment, where processor utilization exceeds 80%, but might prove useful for other organizations.

4.3.3. Multi-threading

Multi-threading is a software technique that can be used to increase computational efficiency, and thus to increase power efficiency. It enables numerous software threads to run simultaneously and access common data in memory. This can help keep multiple processor cores busy inside a reduced memory footprint. It must be kept in mind, however, that writing multi-threaded applications often requires specialized programming skills.

Multi-threading has recently come into the spotlight, mainly because of the introduction of multi-core processors. CERN is currently investigating the potential value of multi-threading for its own applications. Regular workshops are being held on site (in cooperation with Intel), during which the participants are shown how to program scalable applications, with both current and future systems in mind. For

[9] "Eight Software Approaches Can Enable Energy-Efficient Computing" – Nick Jones, Gartner Group

High Energy Physics, the crucially important goal of these efforts is to improve overall performance per watt.

4.3.4. Compiler technologies

Research shows that the compilers used for creating binaries from source code may have a significant impact on the performance of a computing system, and therefore on its power and efficiency characteristics. By using optimizing compilers, such as the Intel C/C++ compiler (icc), it is often possible to speed up an application because the underlying microarchitecture is used more effectively. One of the studies conducted at CERN indicated gains of as much as 65%, in terms of SPECint per Watt, when using Intel's compiler with profile guided optimization.

5. Conclusion

CERN will need a very powerful computing infrastructure to handle the 10 - 15 Petabytes of data per year that will be generated by its LHC experiments. The CERN computing centre was built over 35 years ago and, even after refurbishment, is only capable of providing 2.5 MW of electrical power and 2.5 MW of cooling capacity. In order to provide the maximum computing resources within this limit, CERN is using a variety of techniques and tools to achieve the best possible power efficiency.

Most notably, CERN has benefitted from the move to multi-core processors with a more efficient microarchitecture. Multi-core processors based on the Intel Core microarchitecture deliver about five times more compute power per Watt than single-core processors based on the earlier Netburst microarchitecture. This is enabling CERN to provide five times more compute power within the same data centre power envelope. As the number of cores per die is expected to rise in the near future, many-core technologies (with two-digit core counts) seem to be a promising way to further reduce power consumption, provided the memory requirements can be kept under control.

CERN is also focused on maximizing data centre performance per watt in many additional ways, including optimizing its tendering process, its data centre layout, and its power and cooling strategies. Additional power savings are possible on the software side, and CERN is adopting a variety of strategies, and investigating others. It is hoped that CERN's experience in these areas will provide a foundation on which other organizations can build to improve power efficiency and reduce costs in their own data centres.

For a Summary of Power Saving Strategies, see the appendix.

Appendix: Summary of Power-Saving Strategies

The following strategies have proved valuable in the CERN computing environment, and may offer value for other organizations seeking to increase computing capacity, while reducing power consumption and total costs.

1. Make power efficiency part of the server tendering process, by including the impact of power consumption in estimating total lifecycle costs.
   - Measure performance using workloads that closely reflect production applications.
   - Measure power consumption in the primary AC circuit, to account for all components and internal inefficiencies.
   - Divide performance results by power consumption to evaluate performance/Watt.
   - Take power and cooling costs into account when purchasing new systems and when determining if and when to decommission older systems.

2. When configuring new servers.
   - Consider multi-core processors, which can provide much higher computing capacity per Watt than single-core processors for many workloads. (Be aware, however, that increased memory requirements may impact total energy savings per server.)
   - If throughput is more important than response time for your applications, consider selecting the processors within a family that do not have the highest frequency. They may deliver better total throughput per Watt.
   - Select high-quality power supplies that are efficient and have a high power factor at both high and low loads.
   - Consider blade servers. The large power supplies and fans of these systems tend to be more efficient than the smaller versions used in typical pedestal and rack-mount servers.

3. Optimize your data centre for power efficiency.
   - Reduce hot and cold air mixing, by aligning racks to create hot aisles and cold aisles.
   - Consider sealing cold aisles and extracting heat directly from hot aisles.

4. Evaluate software technologies for improving the overall energy efficiency of your systems and data centre.
   - Use virtualization to consolidate small workloads and infrequently used applications onto fewer servers.
   - Take advantage of optimizing compilers, such as the Intel C/C++ compiler, which can increase performance significantly without an appreciable increase in energy consumption.
   - Consider existing technologies that scale online processor frequency based on workloads, and can put an idle processor to sleep.

- Evaluate multi-threading as a way to increase software performance, but be aware that new developer skills may be required.
- Evaluate grid computing as a way to distribute workloads more efficiently across multiple systems and facilities.