



Fortuitous Technologies

<http://fortuitous.com>

512-351-7783

Green Computing Guide

Philip Carinhas, Ph.D.

September 16, 2009

Contents

1	Why Go Green?	2
2	Performance Tuning	3
3	Capacity Planning	3
4	Reliability Considerations	4
5	Cooling Consideration	4
6	Maximize What You Have	5
7	Design to Fit	5
8	Deploy in (the Right) Time	6
9	Designing for Low-Power Systems	7
10	Virtualization	8
11	Storage Area Networks	10



1 Why Go Green?

Green computing is a very hot topic these days, not only because of rising energy costs and potential savings, but also due to the impact on the environment. Energy to manufacture, store, operate, and cool computing systems has grown significantly in the recent years, primarily due to the volume of systems and computing that companies now heavily rely upon.

Computing power consumption of companies has reached a critical point. For example, an E-commerce business with 100,000 servers can easily spend up to \$20 million a year on server power. Add another \$10 million for a/c cooling and it tops \$30 million a year in power alone. Clearly there is a huge potential for savings in their infrastructure.

Despite the huge surge in computing power demands, there are many existing technologies and methods by which significant savings can be made. This series is dedicated to the ways a typical organization can reduce their energy footprint while maintaining required levels of computing performance.

So why should a company promote green, or energy efficient computing?

- ★ **Climate Change:** First and foremost, conclusive research shows that CO₂ and other emissions are causing global climate and environmental damage. Preserving the planet is a valid goal because it aims to preserve life. Planets like ours, that supports life, are very rare. None of the planets in our solar system, or in nearby star systems have m-class planets as we know them.
- ★ **Savings:** Green computing can lead to serious cost savings over time. Reductions in energy costs from servers, cooling, and lighting are generating serious savings for many corporations.
- ★ **Reliability of Power:** As energy demands in the world go up, energy supply is declining or flat. Energy efficient systems helps ensure healthy power systems. Also, more companies are generating more of their own electricity, which further motivates them to keep power consumption low.
- ★ **Computing Power Consumption has Reached a Critical Point:** Data centers have run out of usable power and cooling due to high densities.



2 Performance Tuning

Performance tuning is the process of adjusting a computer so that it will perform to the best of its ability, given its current or aggregate workload. By doing so you can squeeze out every drop of computing goodness that you current systems have, which is often a lot. This will reduce energy footprint by avoiding unnecessary hardware upgrades, cooling, and associated e-waste.

Performance tuning and management can also lower the overall energy a given system uses because there are less resource allocated internally for a given software process. Disk I/O, CPU, memory reduction can lead to measurable energy savings.

3 Capacity Planning

Capacity Planning allows an organization to grow in a coherent and organized way while minimizing the amount of hardware needed to perform all required computing. Resource utilization is measured over time and analyzed to determine the precise computing infrastructure requirements.

Organized growth is important because it avoids having heterogeneous and ill-performing systems that drive up costs and excess energy need. It also avoids purchasing of overpowered or under-powered equipment that must be later disposed or would otherwise sit and suck power.

Because the success of any IT operations center hinges on the performance requirements of its customers, it is vital to know what the current and future capacity needs are in advance so that designs will accurately reflect future needs. Server consolidation, virtualization, and performance all rely on accurate forecasts of future needs which are provided by capacity plans and forecast models.



4 Reliability Considerations

Reliability of hardware is an important aspect of green computing because it reduces the overall costs of energy associated with system failures and E-waste. Performing coherent and routine reliability analysis of the IT infrastructure can identify unreliable systems and designs. In addition, companies benefit by realizing a more reliable IT and e-commerce framework.

System failures cost significant amounts of energy expenditure due to human intervention in cooled environments, travel to and from data centers, and other related activities.

E-waste, or scrapped computer hardware, is a very serious problem because of the energy required to dispose of the materials, and the hazardous chemicals embedded in the waste materials.

Reliability is directly related to manufacturing, repair, and ROI costs of systems. Increasing the reliability of the IT infrastructure leads to significant energy savings while minimizing hazardous waste materials and their disposal.

5 Cooling Consideration

Cooling the tremendous amounts of heat generated by computers is essential in order to keep the equipment and the software stable. This cooling in turn requires a significant amount of energy which can be estimated at between 30% to 50% of the energy used to power the hardware. This means that for each 100 kilowatt-hours used to power a system, 50 kilowatt-hours could be required to cool it.

You can reduce energy consumption by concentrating on the initial source of the energy burned in the computer. Energy efficient fans, CPUs, power supplies, disks, and related components can make a big difference that pays huge dividends. Other ways to effect efficient cooling [[9],[10],[11],[12]]:

1. System design, and load matching
2. Use of efficient power equipment which exist today
3. Floor layout and vent layout design
4. Coupled cooling systems that directly cool and exhaust



6 Maximize What You Have

One of the first things that smart companies do when they run into performance problems with their IT is to optimize and tune their infrastructure. This can lead to tens of percentage points worth of extra performance. Although we discuss *Performance Tuning* at length in 2, its worthwhile to realize that many companies often have the computing power they need right in their back yard. Collective computing approaches like *grids* and *virtualization* will be discussed in detail in sections 10.2 and 10

Optimizing IT systems is a very effective way to get better performance without initiating purchase and design expense. Make sure you have a comprehensive plan that implements periodic performance tuning.

7 Design to Fit

The number one fallacy that many companies have is that throwing more hardware at a problem will solve every computing problem known to man. Adding more hardware does not always imply better performance. Complex systems react very differently to the numbers of servers you have. You might be overloading the network, database, or other resources.

Design to Fit (or sizing) means that the computing system is designed and built to accommodate only what the computing requirements plus some "reasonable" surplus capacity. This is one of the basic tenets of *Capacity Planning* to be discussed in section 3.

Design to Fit is a very effective way to insure performance without overspending on energy and hardware. Make sure you spend sufficient resources to determine your overall needs, and design that into your systems and infrastructure.



8 Deploy in (the Right) Time

Another mistake you see floating around is that you need to have all the capacity you'll ever need right now. This can lead again to over-built systems that sit idle for years until they see some action or the scrap heap. A more sensible plan is to use *capacity planning* techniques to limit and build out your systems as you need them. After all, newer systems use less energy and give you better performance.

Have capacity planning management in place allows you to grow your environment as you need it, just like the *just in time* (JIT) processes used in manufacturing. It allows you to grow your IT as needed, save energy with fewer components, and purchase more power efficient systems later in the purchase cycle.

Just In Time (JIT) is an inventory strategy implemented in manufacturing proven to save millions by reducing in-process inventory and its associated storage and energy costs. The same techniques can be used by IT.

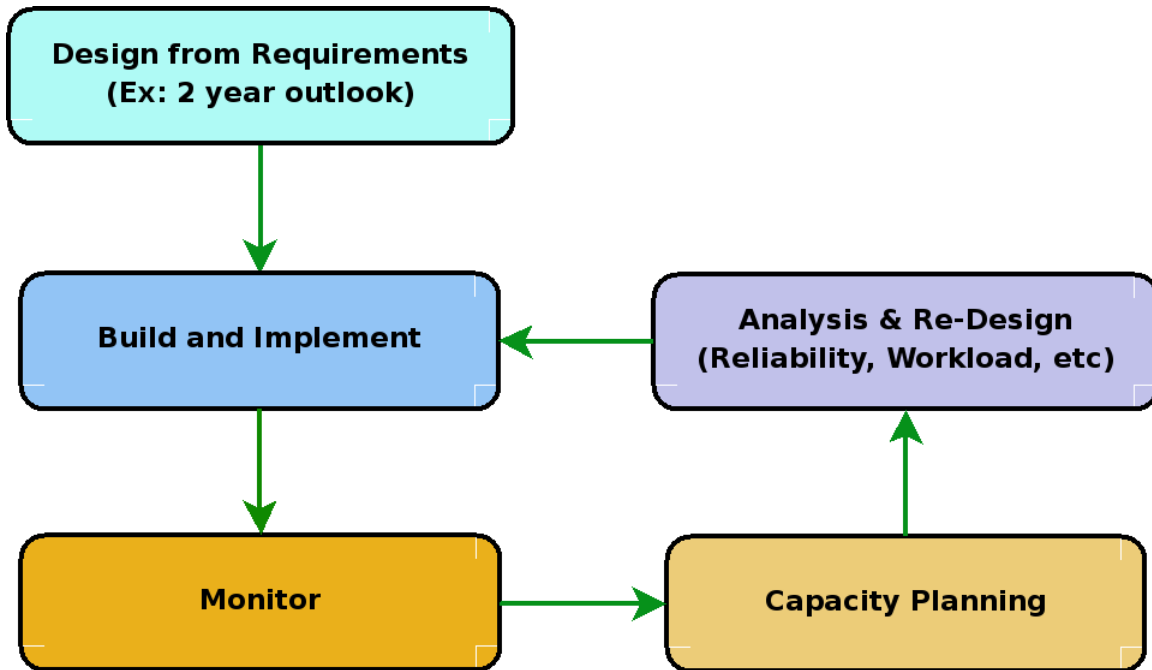


Figure 1: The DIT Process



9 Designing for Low-Power Systems

Using low-powered systems are an excellent way to reduce overall energy needs. The trick to doing this successfully is to properly design your system with performance metrics in mind. If you need a special purpose computing system, why not design it with the least energy in mind? Investing in low-power systems will save you many times the cost of design in terms of energy savings.

9.1 Energy Efficient Processors

New low energy processors can really save a lot of energy. That is because processor power is one of the largest consumers of a computer's power. Many new energy efficient cpu models enter the market each year by Via, AMD, Intel, and others. Design around the best performance/watt system you can.

9.2 Multi-Core Processors

Multi-core processors can enhance computing without adding significant extra power requirements to the overall systems. Most PC CPUs are multi-core these days. Many embedded systems now take advantage of low-power multi-core processors as well.

Sun's multicore chips, for example, are leading power efficiency: its 32-thread Niagara 1 chip, UltraSPARC 1, consumes about 60 watts, while the Niagara 2 chips have 64 threads and run at about 80 watts. Furthermore, Sun has integrated most of the peripheral components into the chips as well, further reducing power consumption and e-waste.

The downside of multi-core systems is that the software needs to take advantage of the new multi-cpu and threading architecture. If you don't redesign the software, chances are good that you are only getting a fraction of the performance goodness out of your investment.

Determining whether multi-core cpu systems can help you is a matter of design requirements, performance potential, and cost.




Figure 2: Niagara 8-CPU Core



10 Virtualization

Instead of having one computer for each service or set of services, you can instead consolidate each server onto a larger *virtualized* system that uses its resources to the fullest, and has a much smaller energy footprint. This benefits in several ways:

1. It reduces the total amount of hardware used in your environment
2. Idle Virtual servers can be powered off
3. The virtualized server will have much less idle time and waste less
4. The total volume of space, air, and rent will be reduced. Data centers can use up to 100 times the energy per square foot of typical office space.
5. Some power companies pay rebates for conversion to virtualized systems.

There is a strong connection between virtualization, capacity planning, and performance management because of the extreme performance requirements that are placed on virtual servers.

Once in place, virtual systems have a unique power flexibility that allows for power consolidation, efficiency, and ability to power-off unused systems.

10.1 Virtualized Servers

In fact, virtual servers are the simplest instance of a virtualized system. You can take several systems and consolidate them on a single server that runs virtualization software such as:

- ★ VMware (<http://vmware.com>)
- ★ Xen (<http://xensource.com>)
- ★ Parallels
- ★ Linux Virtual Server

10.2 Grid Computing

Grid computing allows organizations to fully use the existing computing resources they have in a collective way. This allows for the connecting of disparate systems and clusters together to achieve a much larger computing capacity, thus avoiding extra purchase of computer systems.



Grids systems can allow a company to consolidate data, computing, and software systems in a cohesive way that saves overall energy consumption. In many cases, the computing power is already there, waiting to be used; it just needs to be tapped.

There are several companies that offer both grid implementation via standard *open source* tools, and implementation via drop-in grid-enabling software.

10.3 Data Center Automation and Provisioning

Physical servers can be provisioned as well as virtual ones. Sometimes data-centers prefer to provision servers in this way because of simplicity and security. There are several projects and tools out there that address this ([13, 14, 15])

The key challenge for data centers energy consumption is to know how much to allocate to a project. The key points to consider are:

- ★ How to quantify the workload that a proposed resource will use
- ★ Given that workload, how much physical resource to allocate
- ★ Is the software and system infrastructure modular and portable?
- ★ How to monitor and predict changing workloads for capacity planning
- ★ Is the Network/Storage infrastructure sufficient to meet aggregate workload?
- ★ Do I have enough capacity for the next 3 cycles?

10.4 Virtual Networking and Telecommuting

Keeping people from traveling to work can save a massive amount of energy and green house gas emission. Most people have ample computing power at their homes to connect to their office through a Virtual Private Network (VPN). This technology is very mature and ubiquitous. Some companies often have reservations about letting their people work from home, but there are many who do it efficiently and successfully. Key technologies that save energy:

- ★ Virtual Private Networks
- ★ Teleconferencing Technologies
- ★ Email, Blogs, Bulletin Boards, etc.
- ★ VoIP, Cell, Hand-helds, Convergence



11 Storage Area Networks

Consolidating Storage is an important way you can reduce the total number of disks in a datacenter. Storage Area Networks or SAN allows you to build an efficient storage network that consolidates all storage. Considerations:

- ★ Does SAN reduce the total number of disks?
- ★ SAN allows you to avoid data duplication of OS data.
- ★ SAN can contribute to smaller backup systems via *Snapshots*
- ★ SAN gives a higher reliability factor. (Replication)
- ★ SAN allows greater scalability, avoiding upgrade cycles.
- ★ SAN allows tiered storage: Solid-State, SATA, SAS.



References

★ General References

- [1] <http://fortuitous.com/en/resources/>
- [2] http://en.wikipedia.org/wiki/Green_computing
- [3] <http://www.energystar.gov>
- [4] http://ecenter.colorado.edu/energy/projects/green_computing.html
- [5] <http://thomas.loc.gov/cgi-bin/bdquery/z?d109:h.r.05646:>
- [6] <http://www.sun.com/aboutsun/environment/index.jsp>
- [7] <http://www.openeco.org>
- [8] <http://www.natlogic.com>

★ Data Center References

- [9] Guidelines for Energy Efficient Data Centers,
http://www.thegreengrid.org/gg_content/Green_Grid_Guidelines_WP.pdf
- [10] The Green Grid Opportunity,
http://www.thegreengrid.org/gg_content/Green_Grid_Position_WP.pdf
- [11] Green Metrics,
http://www.thegreengrid.org/gg_content/Green_Grid_Metrics_WP.pdf
- [12] Green Data and Storage Primer,
<http://www.greendatastorage.com/powerprimer.htm>

★ Data Center Tools References

- [13] <http://www.openqrm.org>
- [14] <http://qlusters.com>
- [15] <http://openxvm.org>