

Virtualization of IT environments as ecological alternative in the data center

René Bersier
IBM Switzerland
Vulkanstrasse 106
CH-8010 Zurich
rene_bersier@ch.ibm.com

Keywords: IT server virtualization data center

ABSTRACT

Modern data centers are faced with requirements to deploy dozens or even hundreds of additional servers each year. The number of installed systems is reaching hundreds up to several thousands in a single data center. This is creating an ever increasing load on electricity, cooling power and is accelerating the required replacement of outdated equipment. Virtualization of server and storage systems allows reduction of physically needed assets with a striking reduction of needed power and cooling requirements for data centers and large IT installations.

Accelerated virtualization technology exploitation in the IT industry will allow not only a reduction of environmental stress by building 'green data centers' but also a reduction of operating costs, better usability of assets already in place and creation of a more flexible IT infrastructure with lower resistance to change. This poster presentation will focus on the current status of data centers, an outlook where we are heading without a turnaround and the virtualization technology available today to build and optimize new virtualized and power efficient IT infrastructures.

Current IT data center status – some facts

- Data centers typically consume 15 times more energy per square foot than a typical office building and, in some cases, may be 100 times more energy intensive^[1].
- According the Alliance to Save Energy, if the current rates of growth continued and data center efficiencies remained unchanged, data center electric bills and power requirements would double in less than ten years, and data center electricity bills would increase by an additional \$200 to \$300 million each year.
- An overwhelming majority of consumers – 92 percent – agree that business, government, and consumers have an equal responsibility to reduce energy use, according to the Alliance to Save Energy.
- A February 2007 study by the Energy Department's Lawrence Berkeley National Laboratory found that power consumption for servers and data center-related equipment in the United States was equivalent to about five 1,000-megawatt power plants while worldwide electricity consumption for servers was equivalent to 14 power plants. The total electric bill to operate those servers and related infrastructure equipment was \$2.7 billion in the United States and \$7.2 billion worldwide.

- According to IDC, by 2010 for every \$1 spent on hardware, 70 cents will be spent on power and cooling and by 2012 for every \$1 spent on hardware, \$1 will be spent on power and cooling.
- IDC found that power and cooling spending in the data center is growing at eight times the rate of hardware spending and 46.8% of data center managers do not know how many watts per square foot their data centers can or do support. The other 50%+ of respondents were thought to be largely “guessing” when they said they did know.
- The US Energy Information Administration (EIA) approximates that the 6,500 or so data centers in the United States consume as much power as the entire state of Utah in the average year, equal to the total electricity generated by 30 large power plants annually. More than half of the enterprise data centers in the United States are in major metropolitan areas where real estate and energy are at a premium cost.
- If data center efficiencies remain unchanged, data center electric bills and power requirements would double in less than ten years, according to the US EIA.
- The six million servers in America’s data centers, when factoring in the energy needed to run and cool the servers, consume more energy than the 300+ million televisions in the United States, according to IDC.
- According to Flowmetrics, data centers waste 875,000,000kWH of energy a year, or 436,000,000 tonnes of CO₂
- By 2008 50% of all data centers are power constrained and energy will be second highest cost in 70% of data centers by 2009, according to Gartner
- Green computing is more than operational issues: manufacture, managing suppliers, logistics and disposal
- Optimisation of Assets will become key: Reducing CPU utilisation by application, more efficient SW coding will be demanded and increasing utilisation by asset

Virtualization Introduction

A virtualized infrastructure is a logical representation of resources not constrained by physical limitations. Resources can be data, computing power, storage capacity and others. This allows to

- create many virtual resources within a single physical device - sharing the physical device
- see and manage many virtual resources as one – reach beyond the box
- dynamically change and adjust across the infrastructure – adapt to business requirements on demand

For a long time computer systems were built based on their physical and geographical constraints rather than business requirements. With virtualization we allow actually the business to dictate what it needs with a sub-second response time to change the infrastructure, versus the service dictate what they can support.

By virtualization of IT infrastructures we save many times:

- Typically a single not virtualized system will be dedicated for one particular application and will be bought with enough spare capacity to be able to handle the highest peak workload this system may experience. Sometimes there are upgrades planned over the systems life-cycle, sometimes a system just gets installed and used until replaced. Since the planning of a systems future utilization rate can be a more complex process, many (especially smaller and concerning their size limited) systems will include more spare capacity than actually needed (“just in case”, change due to external effect not known at the

beginning, demand is changing, even the smallest system is too large for a particular workload etc.).

This fact alone creates - summarized over a single data center - huge unused capacities which have to be produced, installed, run and cooled and finally disposed. This is not necessary in a virtualized environment since one virtualized resource is able to utilize the capacity not used by other virtual resources sharing the same physical hardware.

- By virtualization and sharing of physical assets, one can afford easily a more luxurious environment for each single application than one would have been able to provide for only a single particular application alone – and saving energy, cost and be more ‘green’ at the same time.
- A single not virtualized system will reach its practical end of life rather earlier than a virtualized system due to the fact that a non virtualized system may reach its typically smaller capacity limit earlier, insufficient planning or changed requirements ask for a configuration the system is no longer able to handle which requests for a replacement or for many other reasons.
With a virtualized system along with the decoupling of actual hardware from the logical representation of the needed systems, the boundaries of a single system are just simple definitions. These definitions can one (or even the virtualizing part itself) change on the fly to adapt for new requirements. With virtualized environments just a small number of larger systems will be able to cover for most or even all requirements a single virtual system may ever ask for in a data center.

IBM pioneered resource virtualization more than four decades ago and continues to deliver industry leading server, network, and storage virtualization products that enable customers to get more from their existing resources. Resource virtualization allows multiple users to share the same physical resources, use pooled resources as a logical whole, and emulate function. The benefits typically include increased utilization of assets, higher availability and scalability, greater flexibility, lower management and IT costs, and better security and provisioning.

By introducing virtualization functions on Unix, midrange and storage systems beginning of this decade, IBM provides new options for even broader virtualization in the data center. Early adopters on a worldwide basis are creating more energy efficient data centers than was ever possible before.

Find below some examples that this is a real, existent and valuable option to enhance current IT data center energy efficiencies.

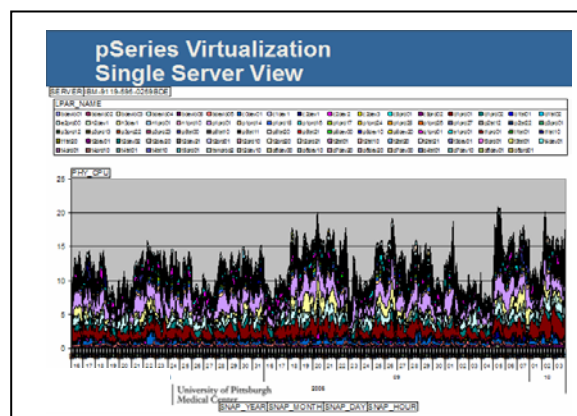
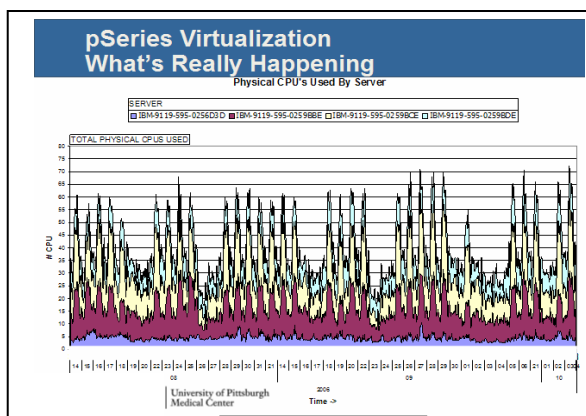
Example 1: PG&E Collaborates with IBM to Reduce Energy Consumption in the Utility's California Data Centers by 80 Percent ^[2]

- Companies Team to Design and Test New Energy Efficient Data Center Technology. Armonk, NY and San Francisco, CA - 10 May 2007: IBM (NYSE: IBM) and Pacific Gas and Electric Company (PG&E) today announced a collaborative effort to optimize energy efficiency opportunities in PG&E's California IT operations through a server consolidation program and by jointly developing a new way to measure and reduce heat in data centers. IBM also announced today its participation in PG&E's Energy Efficiency Incentive Program.
- A world leader in energy efficiency, PG&E teamed with IBM to deploy a server consolidation and virtualization initiative with the explicit goal of reducing energy consumption in its own data center facilities spanning over 40,000 square feet of raised floor in San Francisco, Fairfield and Diablo Canyon, California.
- PG&E will consolidate nearly 300 Unix servers onto 6 IBM System p servers, helping to reduce 80 percent of its energy and facilities consumption, and will use IBM virtualization technologies to boost utilization of the systems from 10 percent capacity to over 80 percent. In addition, PG&E will deploy IBM Rear Door Heat eXchanger water cooling technology on the System p servers to reduce heat in the data center by up to 60 percent.

Example 2: University of Pittsburgh Medical Center Boosts Efficiency and Cost Savings with IBM Virtualization Technology ^[3]

- Leading Health Care Enterprise Embraces Virtualization Technology in IT Transformation. ARMONK, NY and PITTSBURGH, PA - 09 Apr 2007: IBM (NYSE: IBM) and the University of Pittsburgh Medical Center (UPMC) today announced that the medical center's deployment of IBM virtualization technologies is saving significant IT infrastructure costs and improving productivity, while maintaining high service levels. Virtualization allows customers to increase efficiency by dividing a server into smaller parts or by pooling smaller, individual machines so they can perform as one larger system.
 - ...
- UPMC is a national leader in implementing electronic medical records. The large amount of information generated by these systems requires a powerful, flexible IT infrastructure and an enterprise-wide approach to effectively manage storage, computers, memory, network bandwidth and disk input/output. As a result, UPMC is in the process of replacing HP and Sun Microsystems technology with IBM to create an on-demand data center. UPMC's new data center includes IBM servers, storage and software that feature industry-leading virtualization technology to consolidate and simplify the management of its IT infrastructure.
 - ...
- UPMC is using IBM virtualized servers and storage, as well as IBM management tools to help simplify the management of the total infrastructure. The most recent virtualization deployment is on UPMC's UNIX systems. The UNIX applications were migrated to AIX and standardized on four IBM System p5 595 servers. To accommodate the new application for outpatient electronic health records and stay within budget constraints, UPMC used IBM's Advanced POWER Virtualization (APV) for the majority of its System p5 servers to more effectively utilize the available processors. APV allows micro-partitioning to share adjustable fractions of a CPU across partitions, as well as a dynamic load distribution across the CPUs and memory.

Actual production environment: 4 IBM System p5-595, 64 CPUs, 512 to 640GB, 222 LPARS, fully virtualized, reduced peak requirements by over 50%, provides over twice the performance capability, provision new virtual servers in hours, created 184 spare virtual CPU's for new workloads – besides IBM Mainframe, only System p with its advanced POWER virtualization can do this. Actual utilization graphs (left: summary of all 4 systems with a peak of 72 CPU cores, right a single system with a peak of 21 CPU cores):




Example 3: IBM Unveils Plan to Combat Data Center Energy Crisis; Allocates \$1 Billion to Advance "Green" Technology and Services ^[4]

- IBM to Double Computing Capacity in its Own Data Centers, Uses New Technologies to Avoid Five Billion Kilowatt Hours of New Energy Use. NEW YORK, NY - 10 May 2007: IBM (NYSE: IBM) today announced it is redirecting \$1 billion per year across its businesses, mobilizing the company's resources to dramatically increase the level of energy efficiency in IT. The plan includes new products and services for IBM and its clients to sharply reduce data center energy consumption, transforming the world's business and public technology infrastructures into "green" data centers.
- The savings are substantial -- for an average 25,000 square foot data center, clients should be able to achieve 42 percent energy savings. Based on the energy mix in the US, this savings equates to 7,439 tons of carbon emissions saved per year.
- Called "Project Big Green," IBM's initiative targets corporate data centers where energy constraints and costs can limit their ability to grow. The initiative includes a new global "green team" of more than 850 energy efficiency architects from across IBM.

Example 4: Actual customer energy consumption calculation based on different technologies used for a new project planned to be implemented:

- Blue: New project based on dedicated x86 blades (100 blades per year, 5 years)
- Green: New project based on adding IFL (Integrated Facility for Linux) to already installed IBM System z and running 500 systems virtualized on the IBM System z.



With virtualization and Linux for System z, this customer is able to save the energy used to power a complete Swiss Ski resort (Flims), each year 2'535'075 kWh, and that's for a single project only!

(Flims: approx 2'500 people, average energy consumption of a Swiss family with 4 or more people equals 3'750 kWh per year ^[5])

IBM is a long-standing leader in environmental protection, having taken early action to establish its environmental affairs policy in 1971. For more information about IBM's Energy Efficiency Initiative, access to video and audio interviews with IBM and industry leaders, please visit: <http://www.ibm.com/press/greendatacenter>.

Other company, product and service names may be trademarks or service marks of others

References

- [1] Gary Shanshoian, Michele Blazek, Phil Naughton, Robert S. Seese, Evan Mills, and William Tschudi, High-Tech Means High-Efficiency: The Business Case for Energy Management in High-Tech Industries
- [2] <http://www-03.ibm.com/press/us/en/pressrelease/21517.wss>
- [3] <http://www-03.ibm.com/press/us/en/pressrelease/21341.wss>
- [4] <http://www-03.ibm.com/press/us/en/pressrelease/21524.wss>
- [5] http://www.ekz.ch/internet/ekz/de/kundendienst/stromrechnung/mein_stromverbrauch.html