A Power Benchmarking Study on the IBM System z9: Applying Energy Efficiency Metrics to Performance

October 4, 2007

Elisabeth Stahl
Michael Buechele
John P. Rankin
IBM Systems and Technology Group

Executive Overview

Over the past few years energy efficiency in the data center has become one of the top concerns for IT managers. As the cost of power grows significantly, the application of energy efficiency to systems performance becomes a metric that can not be ignored. IBM has long been a champion of energy efficiency and the IBM System z9 is an outstanding platform for the integration of power efficiency and performance.

This paper will demonstrate the role of the IBM System z9 as an excellent platform for power efficiency and performance, discuss tools used for measuring z9 power, highlight a Java power benchmarking study on this platform, and analyze client server consolidation as an energy management strategy. The paper will conclude with next steps in power benchmarking for the IBM System z9.

IBM System z9 - An Outstanding Platform for Power Efficiency

The IBM System z9 mainframe is the world's most sophisticated business server. "Specialty" processors in the IBM System z9 are designated for processing Linux, Java and data workloads as well as encrypting and decrypting certain data. The System z9's Hipersockets technology provides fast communication among virtual servers contained in a single machine.

System z mainframe attributes, such as outstanding security capabilities, are vitally important to business applications. In security certification (Common Criteria's Evaluation Assurance Level), the IBM mainframe achieved one of the highest levels of certification for logical partitioning, IBM's premier virtualization technology.

The IBM System z9 has the ability to handle massive workloads. The mainframe recently achieved the world's largest core banking benchmark result, delivering a record 9,445 business transactions per second (tps) in real-time based on more than 380 million accounts with three billion transaction histories. (http://www.ibm.com/press/us/en/pressrelease/21044.wss)

For comparable levels of computation, mainframe systems take up dramatically less space. Mainframe computers are proving to be dramatically efficient in cooling and power consumption. This paper will discuss power measurement tools for the IBM System z9, highlight a power efficiency and performance benchmark study on the IBM System z9, and review server consolidation as an energy management strategy.

Power Measurement Tools for the IBM System z9

In order to analyze energy efficiency on any system, it is important to understand the different tools available for measurement. This paper will highlight power estimators, nameplate power, external power meters and internal power monitors for the IBM System z9.

Power Estimators

Power estimator tools are available for the IBM System z9 to assist in environmental planning and were used in this study. The zPower power consumption tool for System z9 is available through the IBM Resource Link Web site. (http://www.ibm.com/servers/resourcelink) This tool provides an estimate of the anticipated power consumption of a particular machine model and its associated configuration. A user will input the machine model, memory size, number of I/O cages, and quantity of each type of I/O feature card. The tool will produce an estimate of the power requirements needed for this system. This tool is designed to assist in power and cooling planning for new or currently installed IBM System z9 servers.

Nameplate Power

Nameplate power is the potential maximum power required by a system that is specified on the physical system and in the product planning guides for electrical specifications for that system. It is dependent on the server configuration in terms of the number of books and the number of I/O cages installed. The nameplate power figures listed below assume the maximum configuration:

Model	One I/O cage	Two I/O cages	Three I/O cages
z9 EC Model S08	6.3 kW	9.2 kW	12.1 kW
z9 EC Model S18	8.8 kW	11.8 kW	14.7 kW
z9 EC Model S28	10.9 kW	13.9 kW	16.9 kW
z9 EC Model S38/S54	12.8 kW	15.7 kW	18.3 kW

Input power in kVA is equal to the output power in kW. Heat output expressed in kBTU per hour is derived by multiplying the table entries by a factor of 3.4. The maximum allowed circuit breaker rating is 60 Amps, which is to be used for both power feeds where 200-240V is applicable. For 380-4145 Volts, 32 Amps, and for 480 Volts, 30 Amps are recommended for both power feeds.

System z9 Power Requirements and External Power Meters

The IBM System z9 EC requires at least two power feeds and uses two redundant three-phase line cords, allowing the system to survive the loss of power to either one. In case of a power failure of one of the line cords, the other one is able to take over the entire load to keep the system operating without interruption. The z9 EC is installed with three-phase wiring and operates with 50/60Hz AC power, and voltages ranging from 200V to 480V. For additional equipment (such as the Hardware Management Console, its display, and the modem), additional single-phase outlets are required.

For the power benchmarking study described in this paper, the Square D PowerLogic Energy Meter (http://www.powerlogic.com/products.cfm?id=19) was used as the external power meter to obtain actual results as the benchmarks were run. The external meter was also correlated with the internal monitor described below.

IBM System z9 Internal Power Monitor

The IBM System z9 provides the internal capability to monitor the actual power consumption and temperature of the system. Power Monitor is an additional function for the System Activity Display on the z9 EC Hardware Management Console (HMC). It displays Watts and BTUs per hour as well as cooling air input temperature. It is designed to help verify power consumption for currently installed System z9 EC servers and for this study was correlated with the external meter. Actual correlation data is highlighted in the next section of the paper.

Power Benchmarking Study: System z9 Java Measurements

The IBM System z Performance Lab in Poughkeepsie, New York, along with an extended team in Bangalore, India conducted a series of System z9 Java measurements in 2007.

The z9 measurement systems used consisted of an S08 (1 book) and S18 (2 book) models. Both of these machines were configured with a single cargo I/O cage. The number of processor cores used during measurements varied by a factor of 4. Using a Java workload on z/OS 1.8, a workload was executed which drove the CPU above 99% utilization.

Power used by the System z9 GA3 system was viewed on the Power Consumption Display on the z9 Support Element (SE) System Activity Display (SAD). Two power display points were taken on every measurement, one prior to execution and the other while driving the CPU above 99% utilization.

The System z9 SE SAD measurements were validated using a Square D PowerLogic energy meter installed for this study. The accuracy of the external PowerLogic energy meter was rated plus or minus 1% accuracy. Observed readings from the power meter and the z9 SE SAD were within 100 watts.

Study results are summarized in Table 1.

The first observation is that running the workload has a very small impact on the amount of power used. In fact, the difference in power displays from low CPU utilization to high CPU utilization was within the cost of an additional light bulb or two, as shown in Figure 1.

This observation is also dramatically shown in Figure 3, transactions per marginal or incremental KW. In this chart, throughput is scaled by the extra power used to run the workload over that required by an idling system. This figure shows how little power is actually required to run work once the machine is up and running.

The next observation is that increasing the workload from 4 to 8 processor cores or 12 to 16 processor cores has very little difference on power consumed. Both 4 and 8 core processors are model S08 systems; the 12 and 16 core processors are S18 systems. The study demonstrates that increasing the number of cores used within a model does not significantly increase the power used. This observation is illustrated in Figure 1.

Going from an 8 processor core system to a 12 processor core system moves from an S08 system to a S18 system. The study showed that this change results in approximately a 10% decrease in transactions / busy KW. But, as previously noted, the additional growth potential going to a 16 processor core system requires very little extra power. According to test results, the 16 processor core system has the best transaction / busy KW ratio measured here. These results are illustrated in Figure 2.

Processor	Processor Cores	Thruput	Idle KW	Busy KW	Thruput / Busy KW	Thruput / Marginal KW	Idle BTU / Hour	Busy BTU / Hour
2094-704	4	5.0	3.752	3.916	1.3	30	12,802	13,361
2094-708	8	9.0	3.919	3.987	2.3	133	13,372	13,604
2094-712	12	12.4	5.678	5.973	2.1	42	19,357	20,380
2094-716	16	14.9	5.727	5.920	2.5	77	19,541	20,199

Table 1

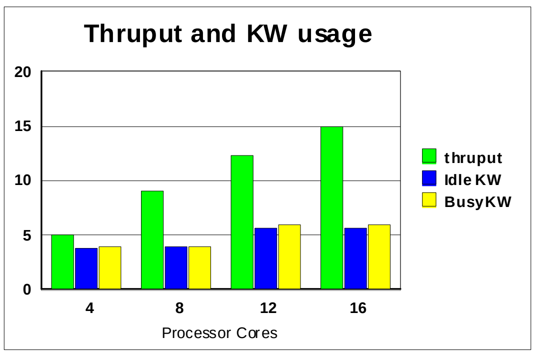

Figure 1

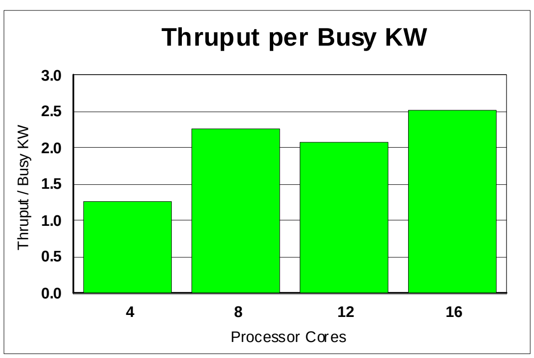

Figure 2

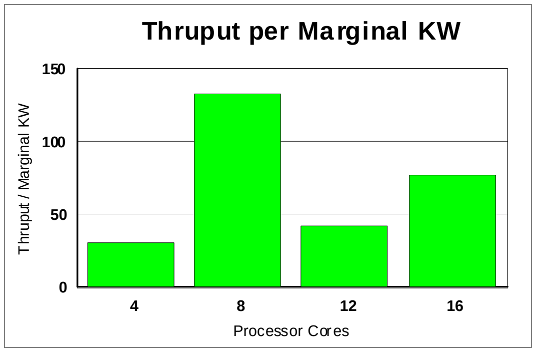

Figure 3

This study demonstrates the scalability of the system while power consumption remains relatively fixed; the marginal or incremental power requirements to run extra work are very small. IBM's System z9 mainframe offers a high utilization rate with systems typically operating at 80 to 100 percent capacity. (http://www.ibm.com/press/us/en/pressrelease/21271.wss) Based on this analysis, server consolidation onto a large server can prove to be an energy management strategy as outlined in the next section of this paper.

IBM Server Consolidation as an Energy Management Strategy

Server consolidation, virtualization, and total cost of ownership are energy management strategies which every IT manager should be familiar with. These energy management policies can conserve energy and reduce costs. Employing server consolidation with the IBM System z9 platform has increased power efficiency and performance for many clients.

Server consolidation is the process of centralizing business computing workloads to reduce cost, complexity, network traffic, management overhead and to optimize and simplify the existing IT infrastructure and provide a foundation for new solution investment and implementation. Server consolidation can be categorized into the following types:

- Centralization
- Physical consolidation
- Data integration
- Application integration

Increasing emphasis on the cost of delivering IT services has caused an unprecedented interest in server consolidation. Server consolidation is the simplification and optimization of IT environments by a reduction of the number of discrete components of infrastructure. Consolidation can be performed on application environments, such as application servers and databases, and on physical hardware, such as servers, routers, and storage.

Server consolidation offers possible advantages including improving availability, reducing costs and complexity, and simplifying operations. Server consolidation can help maximize an organization's return on investment in the data center. Consolidation also promotes high levels of security and data integrity that are difficult to achieve in a distributed environment.

Controlling IT costs is the main reason most businesses give for consolidating servers. There are several dimensions that make up the total IT cost picture:

- Hardware costs: Reducing the number of server, storage devices, peripherals
- Software costs: Lowering software licensing and renewal costs, operating systems, and supporting software
- Support staff: Reducing the number of human resources needed to administer or manage the environment, as well as ongoing training costs
- Operational costs: Reducing costs of floor space, power consumption, cooling systems
- "Hidden costs": Costs associated with inconsistent or incompatible hardware and supporting software, system failures, data loss, security exposures

In the past, one of the greatest advantages that server consolidation provided was savings on floor space. By applying server consolidation, IT managers now talk about the power they have saved.

Server consolidation specifically exploiting virtualization can be a very effective tool in helping to reduce energy costs. Virtualization creates virtual server and storage resources and maps them to physical resources. By pooling, managing and optimizing IT resources across servers, virtualized systems and storage can improve the economics and operations of under-utilized IT assets, with the potential to decrease energy needs and usage.

Server consolidation may also help reduce the Total Cost of Ownership (TCO). An effective server consolidation plan can optimize computing capacity and manage IT resources more efficiently, resulting in TCO reduction.

TCO is a key financial measure that provides IT managers with estimates of direct and indirect costs on technology investments. TCO is used by organizations to align IT investment with organizational strategy. TCO may be calculated over a three year period to achieve a cost that reflects implementation of a system. Managers may also look at TCO over a shorter period for budgets and planning.

TCO includes hardware, software, operational, networking, and support costs. Major operational costs include electricity, cooling, and floor space.   Power and energy efficiency have begun to play a larger part in TCO's calculation. A system with a low initial acquisition cost may have such high power costs that the acquisition cost becomes a trivial part of the equation.

As an example, with System z9 server consolidation, Nationwide Insurance is seeing an 80% reduction in floor space requirements and power consumption. Nationwide has avoided substantial costs associated with expanding a data center, which was projected to cost more than $10 million. (http://www.zjournal.com/index.cfm?section=article&aid=394)

Conclusion

Many organizations around the globe are looking to reduce power consumption and many are facing data center power challenges. With the IBM System z9, clients can successfully integrate power consumption with performance.

IBM has led the technology industry in energy-smart innovation for over 40 years and is committed to climate protection. It is IBM's goal to sustain leadership in energy conservation and management by continuing to deliver power-management and cooling technologies. With these technologies, systems use less power, generate less heat and use less energy to cool the system.

What are the next steps for energy efficiency applied to systems performance? Industry consortia and government agencies are currently working on standards for server power performance metrics. When purchasing a large appliance or a server, energy efficiency can be key criteria. Soon it will be possible to directly compare systems based on new power performance metrics.

This paper has described IBM System z9 power measurement tools, highlighted a power benchmarking study, and discussed server consolidation strategies for energy efficiency. As additional power benchmarking analysis is performed, the IBM System z9 will continue to be viewed as an outstanding platform for power efficiency.

References

IBM's Project Big Green Spurs Global Shift to Linux on Mainframe
http://www.ibm.com/press/us/en/pressrelease/21945.wss

IBM and Financial Network Services Deliver World's Largest Core Banking Benchmark
http://www.ibm.com/press/us/en/pressrelease/21044.wss

S/390 Server Consolidation - A Guide for IT Managers
http://www.redbooks.ibm.com/redbooks/pdfs/sg245600.pdf

EPA Announces New Computer Efficiency Requirements
http://yosemite.epa.gov/opa/admpress.nsf/a8f952395381d3968525701c005e65b5/113b0c0647fee41585257210006474f1!OpenDocument

Insider Insights: Linux on the Mainframe Provides Big Savings for Nationwide Insurance
http://www.zjournal.com/index.cfm?section=article&aid=394

Server Consolidation Using Advanced POWER Virtualization and Linux
http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=SA&subtype=WH&appname=STG_PS_USEN&htmlfid=PSW03005USEN&attachment=PSW03005USEN.PDF

IBM System z9 Enterprise Class Technical Guide
http://www.redbooks.ibm.com/redbooks/pdfs/sg247124.pdf

Case Study: Nationwide Uses Linux and High-Power Virtualization for Web Presence
http://mediaproducts.gartner.com/gc/reprints/ibm/external/volume2/article13/pdf/article13.pdf

IBM Mainframe and Linux Can Lower Customer Costs, Independent Analyst Writes
http://www.ibm.com/press/us/en/pressrelease/21271.wss

Thanks to Rathi Vakkayil, Development Software Testing, IBM Systems and Technology Group, Bangalore, India.

Thanks to reviewers Karl Huppler and Dave Anderson.