5 Myths in Systems Performance

Elisabeth Stahl

Have you heard? There is so much overhead with virtualization. A great performing system certainly can't be green. Private clouds do not perform . . . This paper will highlight 5 myths that revolve around the latest topics in systems performance and debunk them using real data and benchmarks. The paper will conclude with recommendations on how to be your own "MythBuster" in the systems performance arena.

Introduction

A "MythBuster" uses elements of the scientific method to test the validity of rumors, myths, adages, and news stories. This paper will highlight five popular myths in systems performance that focus on virtualization, energy efficiency, cloud performance, software impact, and benchmark value. The myths will be deconstructed using data and detailed analysis. The paper will conclude with recommendations on how to test the validity of myths in the systems performance arena.

Myth #1: Virtualization creates notable overhead.

Virtualization technologies allow IT organizations to consolidate workloads running on multiple operating systems and software stacks and allocate platform resources dynamically to meet specific business and application requirements. Leadership virtualization has become the key technology to efficiently deploy servers in enterprise data centers to drive down costs and become the foundation for server pools and cloud computing technology. Therefore, the performance of this foundation technology is critical for the success of server pools and cloud computing.

Virtualization may be employed to:

- Consolidate multiple environments, including underutilized servers and systems with varied and dynamic resource requirements
- Grow and shrink resources dynamically, derive energy efficiency, save space, and optimize resource utilization
- Deploy new workloads through provisioning virtual machines or new systems rapidly to meet changing business demands
- Develop and test applications in secure, independent domains while production can be isolated to its own domain on the same system
- Transfer live workloads to support server migrations, balancing system load, or to avoid planned downtime
- Control server sprawl and thereby reduce system management costs

To compare the overhead of virtualization technologies, it is useful to highlight industry standard and ISV benchmark results. An initial comparison could be made between a "single system image in native mode without virtualization" results vs. results using x86 virtualization. Note that in these cases, the "virtualized" results achieve fewer users per core. For these configurations, a price is indeed paid in capacity for the ability to participate in x86 virtualization technologies. But this virtualization overhead does not have to be a given.

The SAP Sales and Distribution (SD) Benchmark covers a sell-from-stock scenario, which includes the creation of a customer order with five line items and the corresponding delivery with subsequent goods movement and invoicing.

Figure 1 compares a "Full Resource Partition" SAP SD 2-tier result with a result that is noted as using "virtual CPUs." It might be expected that the virtual result would not achieve the performance of the non-virtual result because of overhead. Note that the users per core is approximately the same for each result. Both configurations contain a hypervisor and have in a sense implemented virtualization so the performance is similar and no degradation is seen from the "virtualized" result. The virtualization is built directly into the firmware of these systems as opposed to x86-based virtualization add-ons.

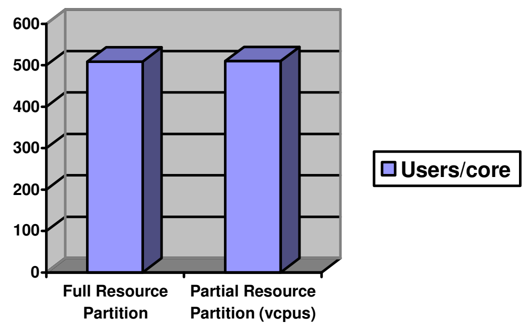

Figure 1.IBM SAP SD 2-tier "Full Resource Partition" result vs. "Partial Resource Partition (vcpus)" result

	System	Proc chips/cores/threads	Cert #	Users	Users/core	OS	ERP release
SAP SD 2-tier result	IBM Power 570	2/4/8	2007037	2035	508.75	AIX 5.3	6.0
SAP SD 2-tier "virtual"	IBM Power 570	2/4/8 using 2 virtual cpus	2008080	1020	510	AIX 6.1	6.0

Source: http://www.sap.com/benchmark   Results current as of 6/8/10.

Rating: Busted

Virtualization does not always have to contribute to performance overhead.

Myth #2: Energy efficient systems can not have great performance.

The power and cooling crisis is getting pervasive market attention. Macroeconomic factors abound: rising electric rates, higher oil prices, technology trends in higher power densities, more and more servers. Data center performance, space considerations and operating costs are issues that IT managers now confront on a daily basis. The current unpredictability of energy costs, the need to tightly manage data center budgets while striving to increase data center throughput, and constraints on capital spending for infrastructure upgrades create ongoing challenges for data center managers.

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. Numerous tools exist now as in Figure 2 to contribute to the integration of power efficiency and performance.

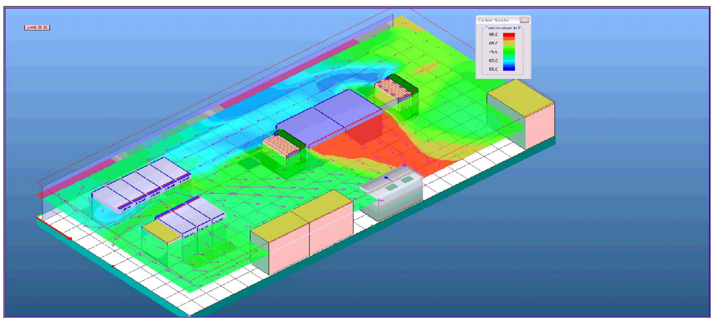

Figure 2.Computational Fluid Dynamics Tool

There are now industry standard metrics that measure energy efficiency and performance of systems. Vendors must be aware in the design of new systems how a system will perform but also how energy efficient it is.

- SPECpower_ssj2008 (http://www.spec.org/power_ssj2008)

SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of volume server class computers. A general methodology for power and performance has been developed (http://www.spec.org/power_ssj2008/docs/SPECpower-Power_and_Performance_Methodology.pdf).

- TPC-Energy (http://www.tpc.org/tpc_energy)

TPC-Energy is a new TPC specification which augments the existing TPC Benchmarks with Energy Metrics developed by the TPC.

- Green500 (http://www.green500.org)

The Green500 provides rankings of the most energy-efficient supercomputers in the world. The Green500 raises awareness about power consumption, promotes alternative total cost of ownership performance metrics, and ensures that supercomputers only simulate climate change and not create it.

With this focus on performance-per-watt type metrics, current systems are being evaluated and new systems are being designed with both performance and energy efficiency as a priority.

Rating: Busted

Systems can perform and be energy efficient.

Myth #3: Private clouds don't perform.

Cloud computing, which has already proven useful for numerous applications, is a flexible and cost-effective delivery platform for providing IT services over the Internet. Cloud resources can be rapidly deployed and easily scaled, with all business processes, applications, and services provisioned on demand, regardless of the user location or device. Cloud computing provides organizations the opportunity to increase service delivery efficiencies, streamline IT management, and better align IT services with dynamic business requirements. It provides support for core business functions along with the capacity to develop and deploy new and innovative services, making them more accessible.

The cloud computing model offers many advantages to IT organizations deploying infrastructure architecture. It has the potential to expand and automate resource virtualization, datacenter resource billing and metering, and self-service catalogues and requests. Cloud computing enables workload-optimized solutions with efficiencies and innovations across the business in areas such as development, test, analytics, infrastructure, and storage that can be quantified as:

- Reduced capital expenditures and labor costs.
- Rapid provisioning and de-provisioning of services.
- Enhanced resource pooling as computing resources are pooled to provide multiple capabilities to consumers. Resource pooling allows virtual and physical resources to be dynamically configured and assigned based on service level agreements as well as demand.
- Superior service management with visibility, control and automation across IT and business services.
- New deployment choices over the cloud, behind the firewall or as an integrated service delivery platform.

Different service layers can be distinguished in the cloud arena, depending on the level in the solution stack at which the service is being delivered:

- The top layer, corresponding to the business application viewpoint, hosts Software as a Service (SaaS)
- The middle layer, corresponding to the middleware services viewpoint, hosts Platform as a Service (PaaS)
- The lowest layer, corresponding to the physical environment viewpoint, hosts Infrastructure as a Service (IaaS)

A number of challenges need to be overcome before the promises of cloud computing can be realized, including performance. Performance concerns for cloud computing may focus on unpredictability -- performance risks caused by such areas as inefficiencies in I/O sharing or high performance computing. The degree to which cloud services can meet agreed service level requirements for availability, performance and scalability can be estimated using performance modeling techniques. Potential performance anti-patterns can be detected before they happen. During the assembly of the cloud solution, estimation, testing and measurement activities will be required on the part of the Integrator and/or Provider to finalize cloud SLAs, and ongoing monitoring will be required to ensure that the cloud SLAs continue to be met after deployment. Response times, throughputs, and capacity metrics for SLAs need to be established and "round trip" considerations made.

The cloud delivery model determines what performance indicators must be covered in an SLA.

- If the delivery model is Business Process as a Service, the service consumer is likely to have requirements regarding response times, throughput, availability, reliability and scalability of business processes. They will leave the derived capacity requirements to the provider.
- If the delivery model is Software as a Service (SaaS) or Platform as a Service (PaaS), the service consumer is likely to have requirements regarding response times, throughput, availability, reliability and scalability of transactions, supported by the software. They will leave the derived capacity requirements to the provider.
- If the delivery model is Infrastructure as a Service (IaaS), however, the service consumer is likely to have requirements regarding throughput, capacity, availability, reliability and scalability of

infrastructure components. In this case the services to be provided on top of the cloud infrastructure are not in the scope of the contract. Therefore, service response times or elapsed times cannot be addressed in the SLA.

The cloud deployment model determines how the SLA is offered to the consumer. If the cloud deployment model is private, the cloud consumer and provider may choose to agree upon a custom SLA.

Performance Considerations

- Consider service level agreements and workload optimization in order to manage capacity and prevent bottlenecks from occurring within the private cloud.
- Organize work loads so that they can scale to support anticipated and unanticipated growth using techniques that will optimize use of storage such as virtualization.
- Consider the skill-set of teams as well as the systems management capabilities to help drive service level performance and implement tools such as Tivoli Service Automation Manager (TSAM).
- Consider security implications.

Rating: Busted

Performance considerations can be addressed throughout the lifecycle of the cloud solution to deliver service levels that are typical of well constructed enterprise applications.

Myth #4: You always need the latest and greatest operating system software to get great performance.

What kind of performance degradation is really seen running older software on a new system?

On Power System POWER7 servers, partitions can be configured to run in several modes, including:

- POWER6 compatibility mode - This execution mode is compatible with v2.05 of the Power Instruction Set Architecture (ISA).
- POWER6+ compatibility mode - This mode is similar to POWER6, with 8 additional Storage Protection Keys.
- POWER7 mode - This is the native mode for POWER7 processors, implementing the v2.06 of the Power Instruction Set Architecture. This mode takes advantage of the 4-thread SMT throughput.

The selection of the mode is made on a per partition basis, from the HMC, by editing the partition profile.

A partition using AIX Version 5.3 executes in POWER6 or POWER6+ compatibility mode. A partition that uses AIX 6.1 with TL2, TL3, or TL4 up to SP2 executes in POWER6 or POWER6+ compatibility mode. Starting from TL4 SP3, AIX 6.1 fully supports POWER7 mode.

The advantages of these modes is that older operating system versions can run on the new hardware, and live partition mobility can occur between POWER6 servers and the new POWER7 machines. The binary compatibility between POWER6 and POWER7 allows migration of partitions between current servers and new ones. LPARs can be moved to and from POWER6 and POWER7 hardware based on business needs.

What kind of performance is there with POWER6 compatibility mode on a POWER7 system? Measurements have shown a relatively small decrease with POWER6 mode as seen in Figure 3, primarily due to simultaneous multithreading capabilities.

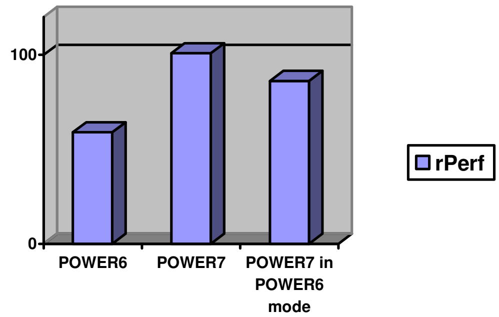

Figure 3.POWER6 relative performance (rPerf) vs. POWER7 system (rPerf) vs. POWER7 in POWER6 mode (testing)

	System		Cores	GHz	L1 Cache (KB)	L2/L3 Cache (MB)	rPerf
rPerf	IBM Power 570	POWER6	8	3.5	64/64	32/128	58.95
	IBM Power 780	POWER7	8	3.86	32/32	2/32	100.75

Source: http://www.ibm.com/systems/power/hardware/reports/system_perf.html

Rating: Busted

Running older levels of software on new technologies does not have to significantly impact performance.

Myth #5: Benchmarks don't help evaluate systems.

Industry standard and ISV benchmarks are certainly not perfect indicators of performance. But having some good data is better than having no data.

Rating: Busted

1 > 0

Recommendations on How to Be Your Own Systems Performance "MythBuster"

- Understand that a systems performance "fact" may in reality be a rumor or myth.
- Perform thorough research using varied sources.
- Obtain real data and benchmark results.
- Analyze over a period of time.
- Rate the systems performance myth as Busted !

References

[COMPA09] Comparing Costs of IBM PowerVM and x86 VMware for Enterprise Server Virtualization (2009)
ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/pol03023usen/POL03023USEN.PDF
IBM PowerVM
http://www.ibm.com/systems/power/software/virtualization/index.html
IBM PowerVM Virtualization
ftp://ftp.software.ibm.com/common/ssi/pm/sp/n/pod03015usen/POD03015USEN.PDF
What's New in VMware vSphere™ 4
http://www.vmware.com/files/pdf/vsphere_performance_wp.pdf
SAP Standard Application Benchmarks
http://www.sap.com/solutions/benchmark/index.epx
A Comparison of PowerVM and x86-Based Virtualization Performance
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101574
Cloud Security Guidance - IBM Recommendations for the Implementation of Cloud Security
http://www.redbooks.ibm.com/redpapers/pdfs/redp4614.pdf
Cloud Computing: Save Time, Money, and Resources with a Private Test Cloud
http://www.redbooks.ibm.com/redpapers/pdfs/redp4553.pdf
Cloud Computing
http://www.ibm.com/ibm/cloud
Performance Implications of Cloud Computing
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101684
SPECpower
http://www.spec.org
TPC-Energy
http://www.tpc.org
Green500
http://www.green500.org
IBM Power 750 and 755 Technical Overview and Introduction
http://www.redbooks.ibm.com/redpapers/pdfs/redp4638.pdf

TRADEMARKS