Dallas

Computer Measurement Group

VALUE COMPUTING

CMG Value Computing

- IT research/analysis firm
  - 2 people; 30 years IT experience (each) in systems/storage/mgt.
  - International focus (based in Dubai last year)
    - Last year: Philippines, Malaysia, Kuala Lumpur, India, Russia, Switzerland, Germany, Romania, Italy, Egypt, Qatar, Dubai, Canada, & all over the U.S.
    - Previous years: China, Brazil, Taiwan, South Africa …
  - Case Studies; Reports; Speeches/Seminars
  - Foci:
    - Microprocessors, servers, infrastructure, management, process flow

CMG Value Computing

- Agenda
  - Not here to provide a litany of #s
    - You're the numbers experts!
  - Talk about three problems — and seek your ideas
    - Test suites
      - Measuring the wrong stuff (optimized application performance vs. general workload processing)
        - » The impact of virtualization
        - » A call for an industry standard mettle test
      - The TPC-C — wrong focus — and to expensive to run…
    - Availability data
      - Having big trouble obtaining reliability data from vendors and IT buyers — why?
    - Cloud computing and the new server order
      - Servers are consolidating around three microprocessor architectures
        - » X86 multi-core, POWER, and z
        - » Sun SPARC is dead; Itanium is dying
  - Summary Observations

CMG Value Computing

- Value computing?
  - Maximize investment in information systems
    - A focus on acquisition and utilization (virtualization)
  - Reduce total-cost-of-ownership (TCO)

TCO: A Range of IT Cost Factors – Often Not Considered

- *Availability*
  - High availability
  - Hours of operation
- *Backup / Restore / Site Recovery*
  - Backup
  - Disaster Scenario
  - Restore
  - Effort for Complete Site Recovery
  - SAN effort
- *Infrastructure Cost*
  - Space
  - Power
  - Network Infrastructure
  - Storage Infrastructure
  - Initial Hardware Costs
  - Software Costs
  - Maintenance Costs
- *Additional development/implementation*
  - Investment for one platform – reproduction for others
- *Controlling and Accounting*
  - Analyzing the systems
  - Cost
- *Operations Effort*
  - Monitoring, Operating
  - Problem Determination
  - Server Management Tools
  - Integrated Server Management – Enterprise Wide
- *Security*
  - Authentication / Authorization
  - User Administration
  - Data Security
  - Server and OS Security
  - RACF vs. other solutions
- *Deployment and Support*
  - System Programming
    - Keeping consistent OS and SW Level
    - Database Effort
  - Middleware
    - SW Maintenance
    - SW Distribution (across firewall)
  - Application
    - Technology Upgrade
    - System Release change without interrupts
- *Operating Concept*
  - Development of an operating procedure
  - Feasibility of the developed procedure
  - Automation
- *Resource Utilization and Performance*
  - Mixed Workload / Batch
  - Resource Sharing
    - shared nothing vs. shared everything
  - Parallel Sysplex vs. Other Concepts
  - Response Time
  - Performance Management
  - Peak handling / scalability
- *Integration*
  - Integrated Functionality vs. Functionality to be implemented (possibly with 3rd party tools)
  - Balanced System
  - Integration of / into Standards
- *Further Availability Aspects*
  - Planned outages
  - Unplanned outages
  - Automated Take Over
  - Uninterrupted Take Over (especially for DB)
  - Workload Management across physical borders
  - Business continuity
  - Availability effects for other applications / projects
  - End User Service
  - End User Productivity
  - Virtualization
- *Skills and Resources*
  - Personnel Education
  - Availability of Resources

Routinely Assessed Cost Factors

CMG Value Computing

- Maximize investment in information systems
  - Performance characteristics
  - The impact of virtualization (unused resource pooling)
- Performance
  - You're not helping
    - You're measuring the wrong stuff (single vs. general workload processing)
- Virtualization — need much more data
  - Restrictions by VMware
  - Cost of virtualization
    - Licenses, deployment, on-going management (tools/utilities)

CMG Value Computing

- Performance (the wrong stuff)
  - You're measuring the performance of single applications on various hardware implementations
    - Virtualization is moving the server world toward general purpose processing
  - You're not providing enough virtualization cost data
    - VMware restrictions
    - Comparative? Hyper-V, VMware, Xen, Open Source
    - Characteristics under varying simultaneous workloads

Are you focused on what IT buyers really want to know?????

CMG Value Computing

The Mainframe World

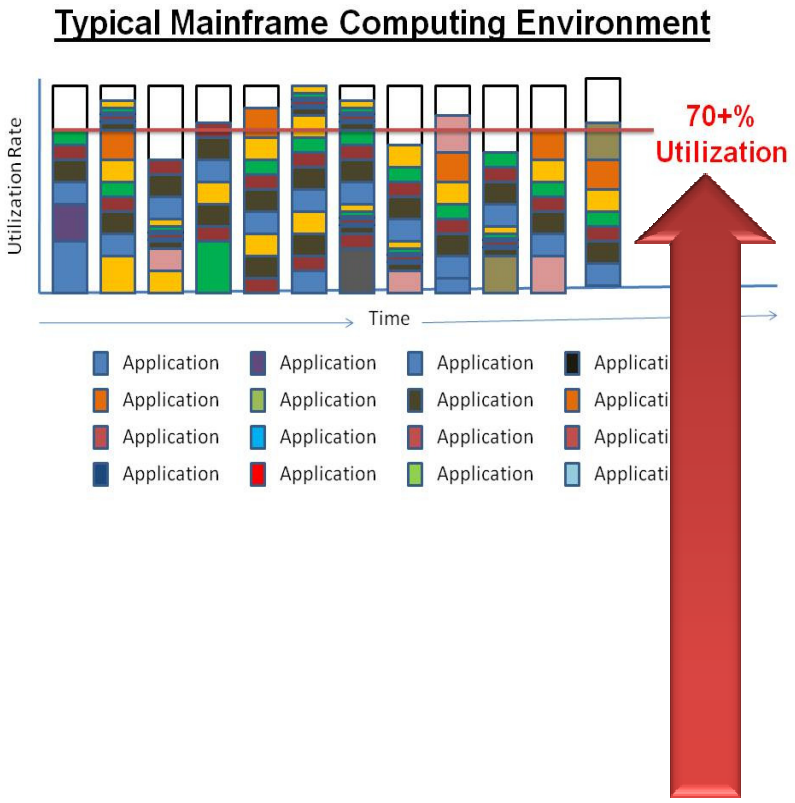

Excellent Virtualization!

The Distributed Computing World

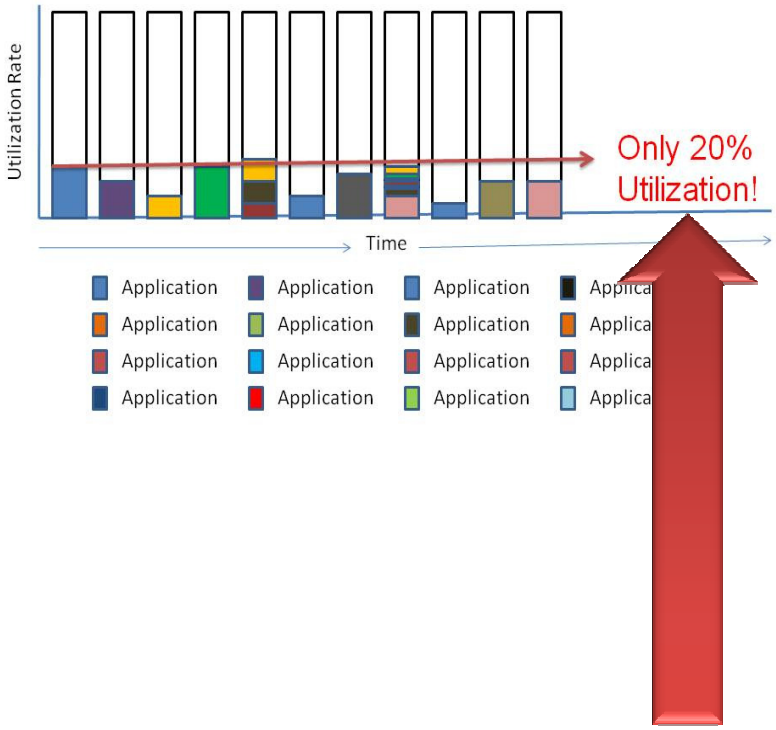

Typical Virtualization!

CMG Value Computing

IBM System z

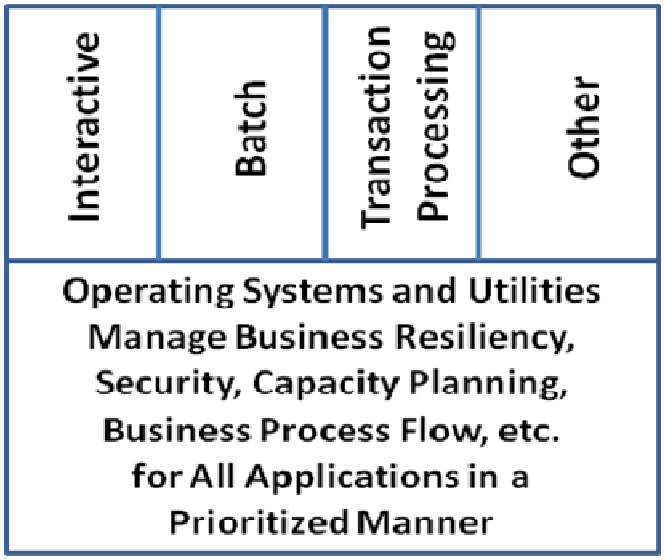

This design emphasizes **balanced performance** and consistent underlying service levels. Functions such as security and business resiliency (which carry performance overhead) are built into this architecture...

HP Superdome & Large Power Systems

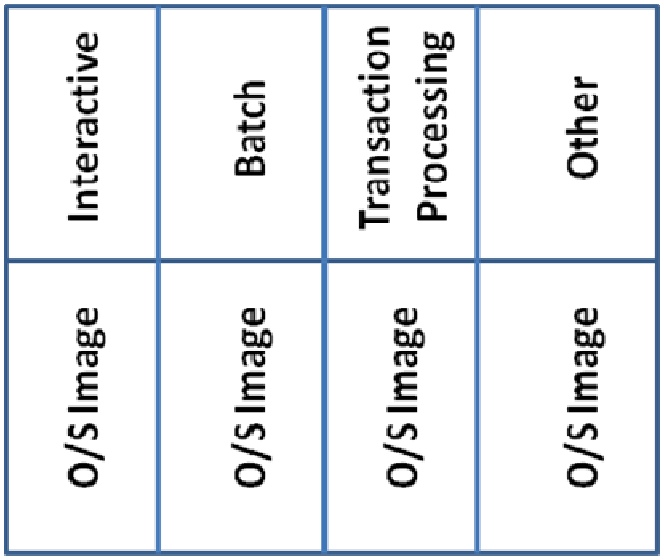

This design emphasizes **raw performance**. Functions such as security, business resiliency, business process management, etc. often require add-ons (that carry performance overhead but are not emphasized in raw performance benchmarks) ...

CMG Value Computing

Benchmarks	# Cores	GHz.	IBM System	POWER Result	Fastest Itanium	POWER Faster by	Fastest Itanium System
TPC-C 64-core	64	5	595	6,085,166	2,196,268	177%	Fujitsu Primequest
TPC-C 32-core	32	1.90	p5-595	1,601,784	1,245,516	28.6%	NEC Express5800
TPC-C 16-core	16	4.7	570	1,616,162	332,265	386.41%	HP rx8620
TPC-C 8-core	8	4.2	550	629,159	372,140	69%	HP rx6600
TPC-C 4-core	4	4.7	570	404,462	230,569	75.4%	HP rx6600
SAP SD 3-tier Overall	32	1.90	p5-595	168,300	100,000	68.3%	HP Superdome 64-core
SAP SD 2-tier 16-core	16	4.7	570	8,000	2880	177.78%	HP rx8620
SAP SD 2-tier 8-core	8	4.7	570	4,010	2150	86.51%	HP rx6600
SAP SD 2-tier 4-core	4	4.7	570	2,035	880	131.25%	HP rx4640
SAP SD 2-tier 2-core	2	2.10	p5-505	680	347	95.9%	Fujitsu Primergy RXI300
Oracle Apps Online 11.5.9	8	1.90	p5-570	15,004	DNP		
Oracle Apps. Std. Batch 11.5.9	8	1.90	p5-570	2,744,000	DNP		
SPECint_rate2000 4-core	4	2.10	p5-550	90.0	72.5	24.1%	HP rx4640-8
SPECfp_rate2000 4-core	4	2.10	p5-550	149	82.7	80.1%	SGI Altix 3000
SPECint_rate2000 8-core	8	2.20	p5-575	200	134	49.25%	HP rx7620-16
SPECfp_rate2000 8-core	8	2.20	p5-575	382	189	102.12%	Bull NovaScale
SPECint_rate2000 16-core	16	1.90	p5-575	314	266	18.05%	HP rx8620-32
SPECfp_rate2000 16-core	16	1.90	p5-575	571	373	53.08%	Bull NovaScale
SPECint_rate2000 32-core	32	1.65	p5-590	529	465	13.76%	NEC NX7700 i9510
SPECfp_rate2000 32-core	32	1.65	p5-590	870	766	13.6%	Fujitsu Primequest 480
SPECint_rate2000 64-core	64	2.30	p5-595	1,513	1108	36.5%	HP Superdome
SPECfp_rate2000 64-core	64	2.30	p5-595	2,406	1,257	91.4%	SGI Altix 3000
SPECint_rate2006 8-core	8	4.7	570	243	102	138.2%	HP rx6600
SPECfp2006	1	5	595	24.9	16.9	47.3%	HP rx6600
SPECsfs_R1.v3 SMP	8	2.20	p5-570	169,786	DNP		
SPECjbb2005 16-core	16	4.7	570	798,752	207,751	284%	Bull NovaScale
Lotus NotesBench R6Mail	16	1.65	i5-595	175,000	DNP		
Lotus NotesBench D7 R6iNotes	16	1.8	p5-560Q	55,000	DNP		
SPEC OMPM2001 (peak) 2-core	2	1.90	p5-520	8,174	2,637	209.97%	HP rx2600
SPEC OMPM2001 (peak) 4-core	4	4.2	520	20,443	6,886	196%	HP rx7620
SPEC OMPM2001 (peak) 8-core	8	4.2	550	40,773	12,762	219%	HP rx8620
SPEC OMPM2001 (peak) 16-core	16	4.7	570	94,350	25,789	265.8%	SGI Altix 3700
SPEC OMPM2001 (peak) Overall	64	5	595	242,116	63,037	284%	SGI Altix 3000
SPEC OMPL2001 base (64-core)	64	2.30	p5-595	1,005,583	507,602	98.1%	SGI Altix 4700
LINPACK HPC 2-core	2	1.90	p5-520	14.31	12.05	18.8%	HP rx1620
LINPACK HPC 4-core	4	4.7	570	61.56	21.71	183.56%	HP rx5670
LINPACK HPC 8-core	8	4.7	570	120.6	44.4	171.62%	HP rx7620
LINPACK HPC 16-core	16	4.7	570	239.4	88.8	169.59%	HP rx8620
LINPACK HPC 32-core	32	4.7	575	466.9	192.4	142.6%	HP x8640
LINPACK HPC 64-core	64	5	595	1032	342	201.7%	HP Superdome

Note: IBM slide from Aug. 2008

POWER vs. Fastest Itanium

Comparing the best available Itanium results vs. POWER

64-core (32/64/128) IBM Power 595 TPC-C result of 6,085,166 tpmC, $2.81/tpmC, avail. 12/10/08
64-core (32/64/128) Fujitsu Primequest TPC-C result of 2,196,268 tpmC, $4.70/tpmC, avail. 04/30/08
32-core IBM p5-595 TPC-C result of 1,601,784 tpmC, $5.05/tpmC, avail. 04/20/05
32-core (16/32/64) NEC Express5800 TPC-C result of 1,245,516 tpmC, $4.57/tpmC, avail. 04/30/08
16-core (8/16/32) IBM Power 570 TPC-C result of 1,616,162 tpmC, $3.54/tpmC, avail. 11/21/07
16-core HP rx8620 TPC-C result of 332,265 tpmC, $4.48/tpmC, avail. 07/15/05
8-core (4/8/16) IBM Power 550 TPC-C result of 629,159 tpmC, $2.49/tpmC, avail. 04/20/08
8-core (4/8/16) HP rx6600 result of 372,140 tpmC, $1.81/tpmC, avail. 06/11/07
4-core (2/4/8) IBM Power 570 TPC-C result of 404,462 tpmC, $3.50/tpmC, avail. 11/26/07
4-core (2/4/8) HP rx6600 TPC-C result of 230,569 tpmC, $2.63/tpmC, avail. 12/01/06

Sources:
http://www.spec.org
http://www.tpc.org
http://www.sap.com/benchmark/
http://performance.netlib.org/performance/html/PDSreports.html
All results are as of 08/01/08

TPC-C results with processor chip/core/thread.
SPEComp results: IBM cores = 2x chip, threads = 4x chip.
SAP certification numbers can be found in SAP section of charts.
Linpack results are SMP only.

CMG Value Computing

- What *SHOULD* be measured?
  - What happens when an important e-business application experiences a sudden spike in activity when the system is already 100% busy?
  - What happens when an LPAR running the mission critical work in a Sysplex fails?
  - With six different workloads running across five LPARs sharing 16 physical CPUs, what happens when a CPU fails?

These kind of measurements get to the truth of value computing!

CMG Value Computing

- What should be measured? (cont'd)
  - Platforms have different "balance points"
    - Reliability/availability characteristics;
    - Performance characteristics;
    - Scalability/capacity;
    - Memory management and memory capacity;
    - Virtualization;
    - Power management and heat dissipation

CMG Value Computing

- What *SHOULD* be measured? (cont'd)
  - The Mettle Test
    - Captures real operational performance that highlights behavior that is not generally exposed by standard industry performance benchmarks
    - It was designed to illustrate the ability of the z/OS Workload Manager (WLM) and WebSphere to distinguish between:
      - High priority work and low priority work running in the same system, and
      - Manage workload priorities across systems in a Parallel Sysplex using the Intelligent Resource Director (IRD)
  - The self healing capabilities exposed by the Mettle Test include recovery from a processor failure, system failure, and application failure
  - View a demo:
    - http://www-01.ibm.com/software/webservers/appserv/zos_os390/mettle.html

CMG Value Computing

- The TPC
  - TPC-App
    - "TPC-App showcases the performance capabilities of application server systems"
  - TCP-C
    - "TPC-C simulates a complete computing environment where a population of users executes transactions against a database"
  - TPC-E
    - "The TPC-E benchmark uses a database to model a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research"
  - TPC-H
    - "The TPC Benchmark™H (TPC-H) is a decision support benchmark. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications"

These are single workload benchmarks. With virtualization, we're talking about a server world that processes general, varied workload…

CMG Value Computing

- The TPC-C
  - **Sun SPARC Enterprise T5440 Server Cluster**
    - 18,083,745 USD
  - **IBM Power 595 Server Model 9119-FHA**
    - 17,111,788 USD
  - **Bull Escala PL6460R**
    - 17,127,928 USD
  - **HP Integrity Superdome-Itanium2/1.6GHz/24MB iL3**
    - 11,978,134 USD

The setup and testing of this environment is incredibly expensive. And, is this activity really measuring what's important to IT users?

CMG Value Computing

- New topic: reliability/availability data
  - Wouldn't it be great if we could get availability ratings for the systems that we buy?

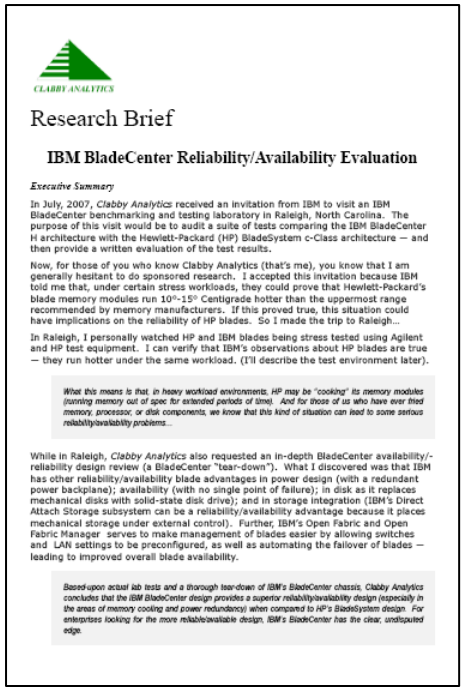

CLABBY ANALYTICS

Research Brief

IBM BladeCenter Reliability/Availability Evaluation

Executive Summary

In July, 2007, *Clabby Analytics* received an invitation from IBM to visit an IBM BladeCenter benchmarking and testing laboratory in Raleigh, North Carolina. The purpose of this visit would be to audit a suite of tests comparing the IBM BladeCenter H architecture with the Hewlett-Packard (HP) BladeSystem c-Class architecture — and then provide a written evaluation of the test results.

Now, for those of you who know Clabby Analytics (that's me), you know that I am generally hesitant to do sponsored research. I accepted this invitation because IBM told me that, under certain stress workloads, they could prove that Hewlett-Packard's blade memory modules run 10°-15° Centigrade hotter than the uppermost range recommended by memory manufacturers. If this proved true, this situation could have implications on the reliability of HP blades. So I made the trip to Raleigh...

In Raleigh, I personally watched HP and IBM blades being stress tested using Agilent and HP test equipment. I can verify that IBM's observations about HP blades are true — they run hotter under the same workload. (I'll describe the test environment later).

What this means is that, in heavy workload environments, HP may be "cooking" its memory modules (running memory out of spec for extended periods of time). And for those of us who have ever tried memory, processor, or disk components, we know that this kind of situation can lead to some serious reliability/availability problems...

While in Raleigh, *Clabby Analytics* also requested an in-depth BladeCenter availability/-reliability design review (a BladeCenter "tear-down"). What I discovered was that IBM has other reliability/availability blade advantages in power design (with a redundant power backplane); availability (with no single point of failure); in disk as it replaces mechanical disks with solid-state disk drive); and in storage integration (IBM's Direct Attach Storage subsystem can be a reliability/availability advantage because it places mechanical storage under external control). Further, IBM's Open Fabric and Open Fabric Manager serves to make management of blades easier by allowing switches and LAN settings to be preconfigured, as well as automating the failover of blades — leading to improved overall blade availability.

Based-upon actual lab tests and a thorough tear-down of IBM's BladeCenter chassis, Clabby Analytics concludes that the IBM BladeCenter design provides a superior reliability/availability design (especially in the areas of memory cooling and power redundancy) when compared to HP's BladeSystem design. For enterprises looking for the more reliable/available design, IBM's BladeCenter has the clear, undisputed edge.

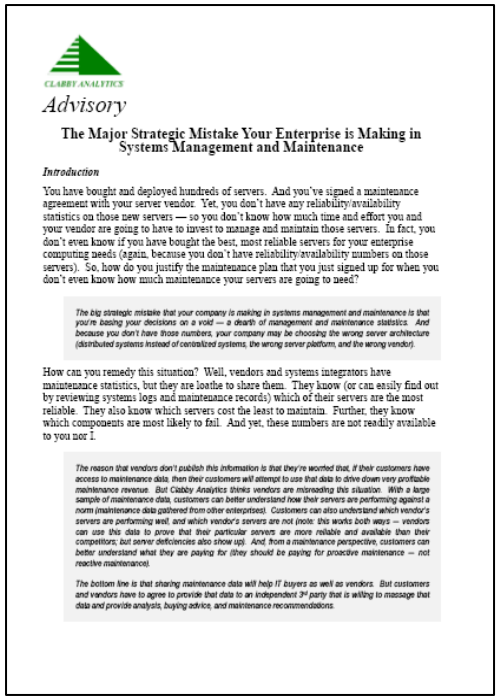

CLABBY ANALYTICS

Advisory

The Major Strategic Mistake Your Enterprise is Making in Systems Management and Maintenance

Introduction

You have bought and deployed hundreds of servers. And you've signed a maintenance agreement with your server vendor. Yet, you don't have any reliability/availability statistics on those new servers — so you don't know how much time and effort you and your vendor are going to have to invest to manage and maintain those servers. In fact, you don't even know if you have bought the best, most reliable servers for your enterprise computing needs (again, because you don't have reliability/availability numbers on those servers). So, how do you justify the maintenance plan that you just signed up for when you don't even know how much maintenance your servers are going to need?

The big strategic mistake that your company is making in systems management and maintenance is that you're basing your decisions on a void — a dearth of management and maintenance statistics. And because you don't have those numbers, your company may be choosing the wrong server architecture (distributed systems instead of centralized systems, the wrong server platform, and the wrong vendor).

How can you remedy this situation? Well, vendors and systems integrators have maintenance statistics, but they are loathe to share them. They know (or can easily find out by reviewing systems logs and maintenance records) which of their servers are the most reliable. They also know which servers cost the least to maintain. Further, they know which components are most likely to fail. And yet, these numbers are not readily available to you nor I.

The reason that vendors don't publish this information is that they're worried that, if their customers have access to maintenance data, then their customers will attempt to use that data to drive down very profitable maintenance revenue. But Clabby Analytics thinks vendors are misreading this situation. With a large sample of maintenance data, customers can better understand how their servers are performing against a norm (maintenance data gathered from other enterprises). Customers can also understand which vendor's servers are performing well, and which vendor's servers are not (note: this works both ways — vendors can use this data to prove that their particular servers are more reliable and available than their competitors; but server deficiencies also show up). And, from a maintenance perspective, customers can better understand what they are paying for (they should be paying for proactive maintenance — not reactive maintenance).

The bottom line is that sharing maintenance data will help IT buyers as well as vendors. But customers and vendors have to agree to provide that data to an independent 3rd party that is willing to massage that data and provide analysis, buying advice, and maintenance recommendations.

CMG Value Computing

- The need for reliability/availability data
  - What does failure cost? Go to:
    - http://www.bin95.com/equipment_failure_cost.htm
    - Additional:
      - Capital
      - Equipment
      - Labor
      - Services
      - Materials
      - Product
    - Loss of revenue opportunities

CMG Value Computing

- The need for reliability/ availability data
  - What if you knew which vendor makes the most reliable systems
    - What if one vendor's blades are 2x more likely to fail than IBM blades?
    - http://www.bin95.com/equipment_failure_cost.htm

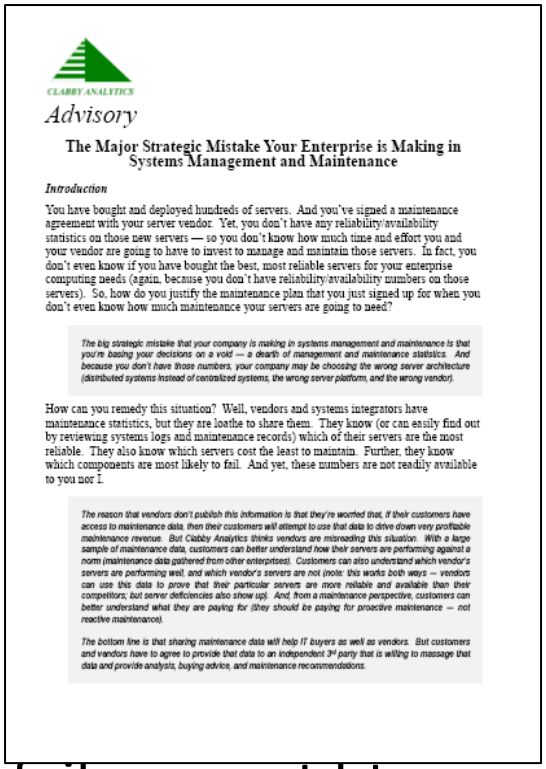

CLABBY ANALYTICS

Advisory

The Major Strategic Mistake Your Enterprise is Making in Systems Management and Maintenance

Introduction

You have bought and deployed hundreds of servers. And you've signed a maintenance agreement with your server vendor. Yet, you don't have any reliability/availability statistics on those new servers — so you don't know how much time and effort you and your vendor are going to have to invest to manage and maintain those servers. In fact, you don't even know if you have bought the best, most reliable servers for your enterprise computing needs (again, because you don't have reliability/availability numbers on those servers). So, how do you justify the maintenance plan that you just signed up for when you don't even know how much maintenance your servers are going to need?

The big strategic mistake that your company is making in systems management and maintenance is that you're basing your decisions on a void — a dearth of management and maintenance statistics. And because you don't have those numbers, your company may be choosing the wrong server architecture (distributed systems instead of centralized systems, the wrong server platform, and the wrong vendor).

How can you remedy this situation? Well, vendors and systems integrators have maintenance statistics, but they are loathe to share them. They know (or can easily find out by reviewing systems logs and maintenance records) which of their servers are the most reliable. They also know which servers cost the least to maintain. Further, they know which components are most likely to fail. And yet, these numbers are not readily available to you nor I.

The reason that vendors don't publish this information is that they're worried that, if their customers have access to maintenance data, then their customers will attempt to use that data to drive down very profitable maintenance revenue. But Clabby Analytics thinks vendors are misreading this situation. With a large sample of maintenance data, customers can better understand how their servers are performing against a norm (maintenance data gathered from other enterprises). Customers can also understand which vendor's servers are performing well, and which vendor's servers are not (note: this works both ways — vendors can use this data to prove that their particular servers are more reliable and available than their competitors; but server deficiencies also show up). And, from a maintenance perspective, customers can better understand what they are paying for (they should be paying for proactive maintenance — not reactive maintenance).

The bottom line is that sharing maintenance data will help IT buyers as well as vendors. But customers and vendors have to agree to provide that data to an independent 3rd party that is willing to massage that data and provide analysis, buying advice, and maintenance recommendations.

Decisions are being made in a void — a dearth of management and maintenance statistics. And because stats aren't easy to obtain enterprises may be choosing the wrong server architecture (distributed systems instead of centralized systems, the wrong server platform, and the wrong vendor).

CMG Value Computing

- How can this situation be remedied?
  - Vendors and systems integrators have maintenance statistics, but they are loathe to share them
    - Why????
    - They know (or can easily find out by reviewing systems logs and maintenance records) which of their servers are the most reliable
    - They also know which servers cost the least to maintain
    - Further, they know which components are most likely to fail. And yet, these numbers are not readily available…

CMG Value Computing

- Why is reliability data important?
  - To make more informed decisions about systems maintenance costs
    - With maintenance statistics in hand, IT buyers can better decide what kind of contract they want with their vendor (break fix, proactive, reactive, time-and-materials, etc.)
      - If your systems prove to be highly reliable, IT buyers may be able to change the focus of their maintenance contracts — moving away from a reactive maintenance contract to a more proactive contract (such as the best-in-class proactive maintenance found on mainframe servers)
  - To make more informed decisions on systems architecture choices
    - We believe those numbers will show huge availability benefits by adopting centralized systems as opposed to distributed systems
  - Make better decisions on the servers that it chooses
    - Maintenance statistics will show your IT organization which servers are the most/least reliable.
      - This data is extremely important to users of distributed systems — you don't want to buy a large number of unreliable servers if you can avoid it.
  - To improve the way it manages systems, storage, telecommunications devices, and PCs (we think you'll move from reactive to proactive maintenance
  - To understand how your numbers compare to a norm
    - Is your shop excellent or not?

CMG Value Computing

- What needs to be tracked?
  1. The *maintenance rate*
     - *The maintenance rate indicates how many months your server environment runs before a maintenance action is required. When this rate is compared to a norm (based-upon statistics gathered from other customers), this rate can also be used to measure the health of your environment*
  2. *The availability rate*
     - *The availability rate measures how many months a specific hardware product can operate before experiencing a hardware disruption that stops it from functioning. This rate can be used to determine how reliable your servers are; how available your servers are; and how your servers compare to other vendor's servers.*

CMG Value Computing

- What we are finding…
  - Centralized computing models are far more highly available than distributed computing models (better RAS designs, multiple layers of redundancy, etc.). And if that is the case, then we will have solid proof that enterprises should migrate away from distributed computing deployments whenever possible; and,
  - Distributed servers are not being managed as efficiently as mainframes and other scale-up designs (we think that we will find that distributed servers are not being managed "proactively" — but rather "reactively"). And we think that the numbers will show that managing servers proactively is a better management practice than waiting for failures. And this, too, should drive more IT organizations to move away from distributed server deployments whenever possible.
- In addition to these findings, we also expect to find that:
  - Enterprises that are paying for reactive maintenance are paying too much. Maintenance contracts that are based on reactive maintenance are substantially overpriced. The maintenance that you pay for should be proactive maintenance because proactive maintenance not only ensures that your systems are kept up-and-running, it also includes engineering changes that prolong the life of your investment. (Imagine refreshing your servers every eight years instead of every four)!

We need help proving these theories…

CMG Value Computing

- New topic: Cloud Computing
  - The current IT distributed computing model cannot be sustained
    - The current IT distributed computing model calls for constantly adding more people as capacity expands
    - People are expensive (salary, benefits, sick time, etc.)
    - Labor costs are now nearing 50% of data center cost of operations — and as more people are added, this percentage is rising!
    - People make errors
      - Lost productivity; lost revenue; and lost opportunity
  - **Extreme virtualization/provisioning with automated management is needed (a.k.a.: cloud architecture)**

CMG Value Computing

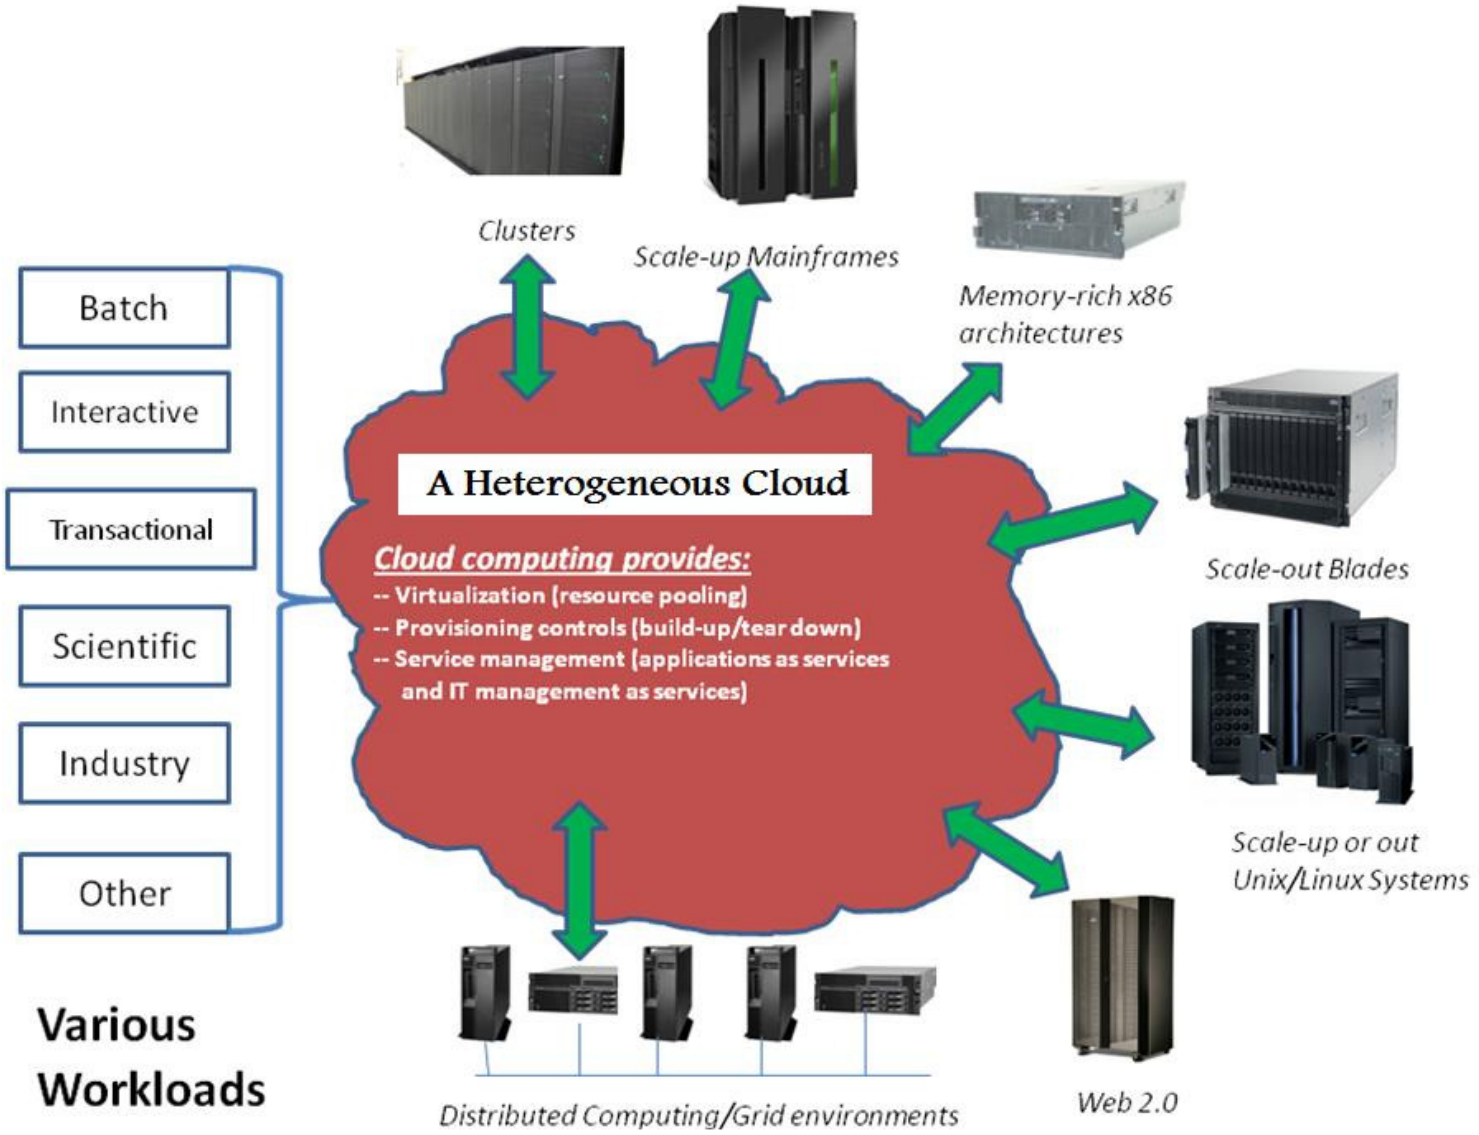

CMG Value Computing

- The clouds of the future
  - Homogeneous vs. heterogeneous
  - Public, private, hybrid
- What you need to understand
  - Server market is consolidating
    - Three microprocessor architectures will dominate
      - X86-based multi-cores (post-Nehalem); POWER; z
  - The concepts of virtualization, provisioning, workload management

CMG Value Computing

- The clouds of the future (cont'd)
  - What the move to clouds means to you
    - Need to produce cloud-related data
      - What does it cost to operate a cloud vs. trad approaches

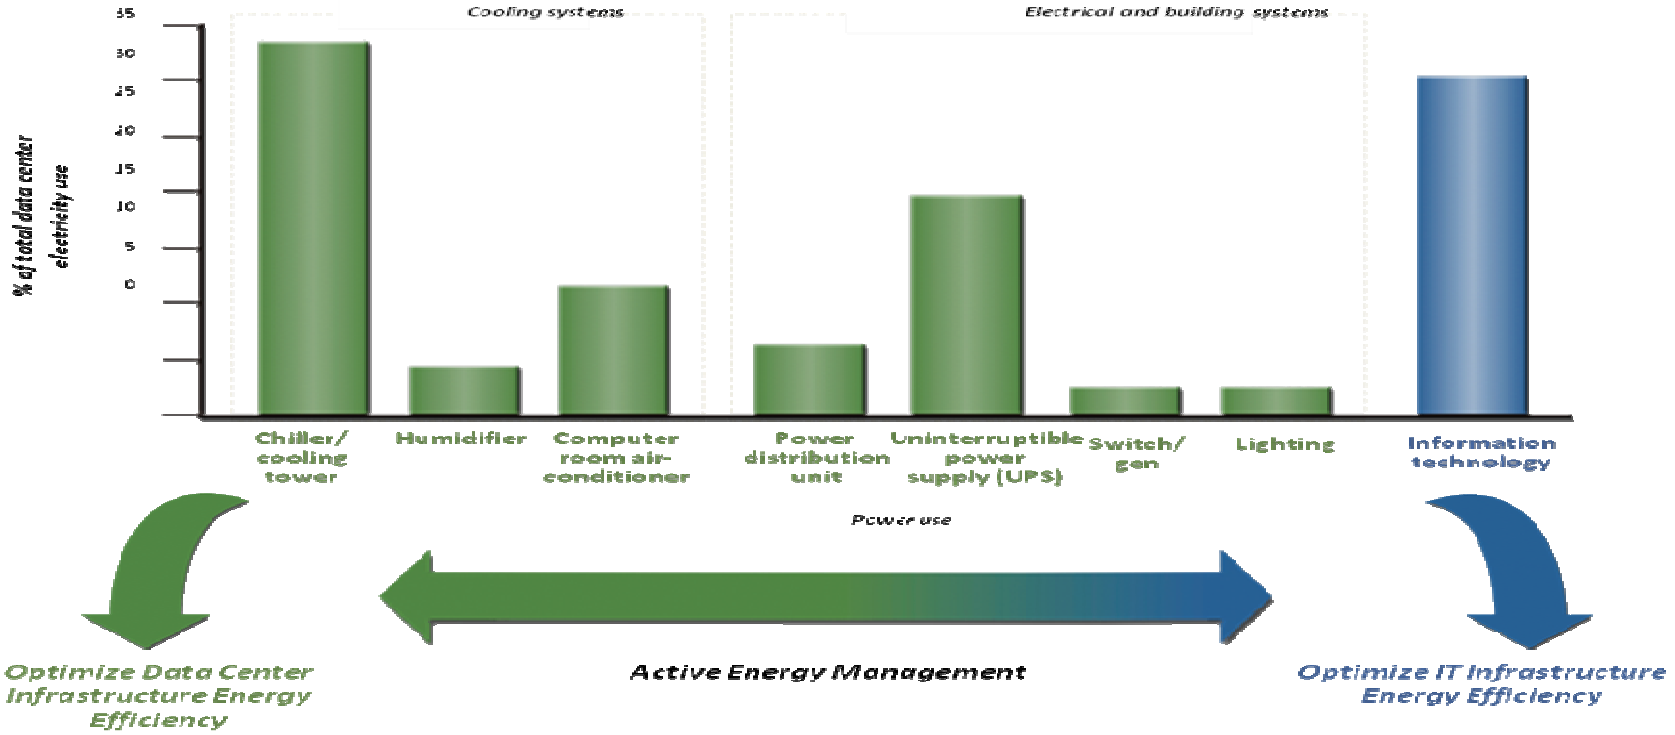

CMG Value Computing

- The clouds of the future (cont'd)
  - We need all sorts of new data

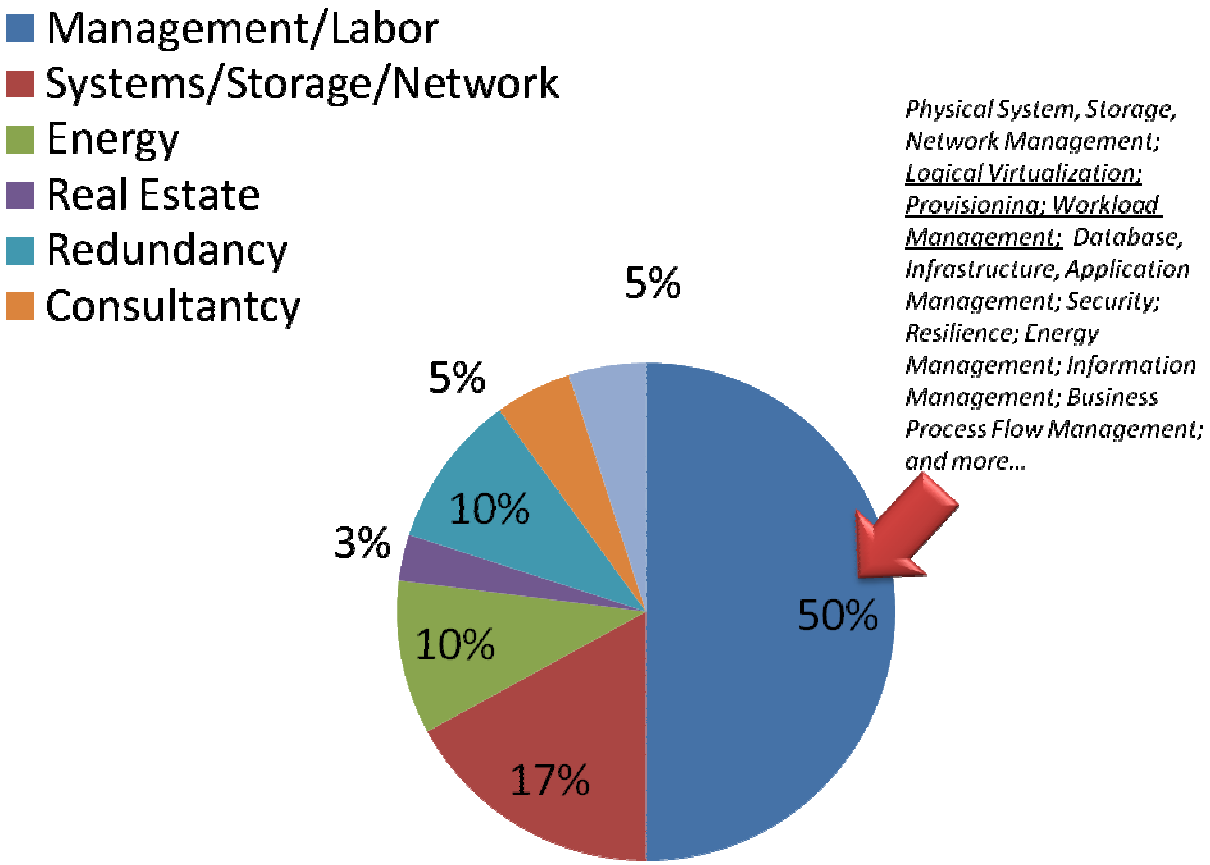

CMG Value Computing

- The clouds of the future (cont'd)
  - And you need to become familiar with the term "service management"
    - Could be an info gathering tool

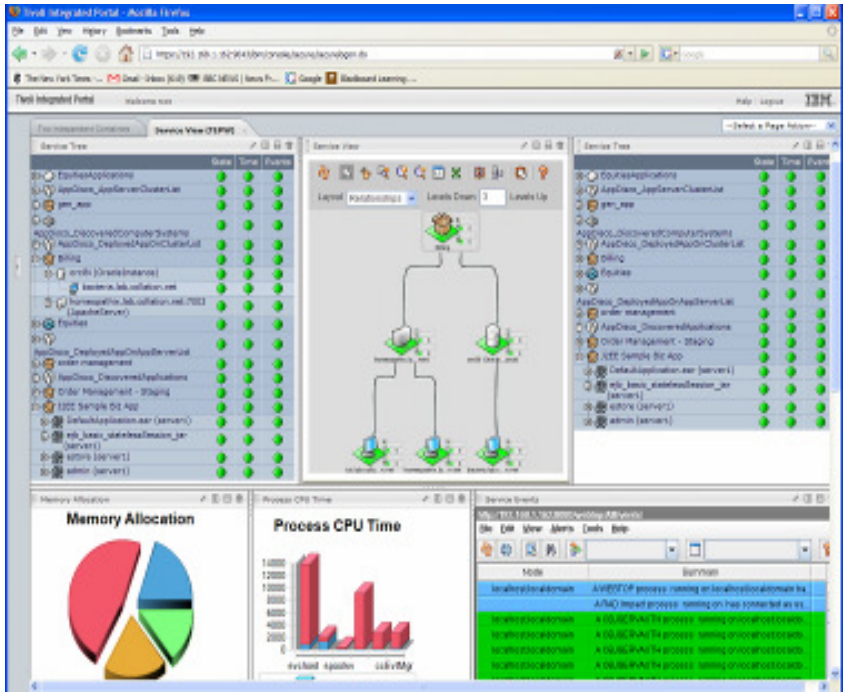

CMG Value Computing

- Summary Observations
  - New emphasis on reduced acquisition costs AND reduced operational costs
  - Acquisition
    - Focus on getting more out of existing investments
      - Maximize performance through virtualization
      - Reduce redundancy through virtualization
  - Operational (a.k.a — Total Cost of Ownership)
    - Reducing costs related to human-oriented, management, energy, security, resiliency, etc.

CMG Value Computing

- Summary Observations (cont'd)
  - To better analyze IT environments, we need:
    - Acquisition
      - Better reliability/availability numbers
      - Better virtualization comparisons (especially cost)
    - TCO
      - Better management cost analysis
        - » For instance, how much can service management save an organization in people-related management costs
      - Better cloud analytics
        - » Cost to move to a cloud (virtualization, provisioning, etc.)
        - » Cost models (when to deploy which type of cloud)
        - » More…

CMG Value Computing

- Summary Observations (cont'd)
  - Most of all, we have got to improve our measurement systems for general workload processing
    - The cloud is all about automated, general workload processing on available, underlying, self-provisioned, virtualized information systems
  - We need an industry standard "Mettle Test"
    - Virtualization is making it possible for dedicated application servers to become general workload processors

CMG Value Computing

- Questions and Answers

For Additional Information on Tivoli Service Management for System z

http://www-01.ibm.com/software/tivoli/solutions/zsmc

http://www-01.ibm.com/software/os/systemz/itsm

Please join the Service Management for System z Community:
https://www-950.ibm.com/communities/service/html/communityview?communityUuid=9051592f-7640-466e-8524-6ae7dcc20c79