April 2009

MANAGEMENT BRIEF

VALUE PROPOSITION FOR IBM POWER SYSTEMS

Comparing Costs of IBM PowerVM and x86 VMware for Enterprise Server Virtualization

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net

TABLE OF CONTENTS

List of Figures

EXECUTIVE SUMMARY

Technologies

In only a few years, server virtualization has moved to the forefront of the IT agenda in a growing number of organizations. A new field of opportunities has emerged to reduce numbers of physical servers, increase capacity utilization, improve flexibility of server provisioning, and realize other benefits.

Although industry debate has tended to focus on VMware and equivalents employed in the x86 server world, the spectrum of technologies is wider. In particular, UNIX server virtualization solutions are often more sophisticated and mature than their x86 counterparts, and are commonly employed to support applications and workloads for which VMware and equivalents would not be realistic candidates.

The IBM Power platform is the recognized industry leader in UNIX server virtualization. In contrast to VMware, which is a software overlay on x86 hardware, Power servers implement a virtualization architecture whose components are embedded in hardware, PowerVM firmware and operating system software. The capabilities of this architecture are significantly different and in many areas more advanced.

One major difference is that PowerVM enables use of "hard" partitions, which provide greater partition isolation than software-based techniques such as VMware Virtual Machines. Firmware-based logical partitions (LPARs) reduce the potential for performance bottlenecks and contribute to higher levels of availability and security than may be realized with software-based partitions.

Power server system and workload management capabilities are more granular and more closely integrated than is the case for VMware and equivalents. The Power platform also benefits from numerous industry-leading availability optimization features.

These and other distinctive capabilities have caused widespread adoption of Power servers to support transaction- as well as database-intensive systems whose performance and uptime requirements are significantly more demanding than the norm.

The importance of workload management should be highlighted. Partitioning creates the potential for high levels of capacity utilization. The extent to which this potential will be realized in practice, however, depends heavily on the mechanisms that allocate system resources between, and monitor and control workload execution processes across partitions.

If these mechanisms are ineffective, a high proportion of system capacity may be idle at any given time. Close integration of partitioning and workload management capabilities is also necessary to minimize risks that surges in workloads running in individual partitions will impact performance and availability.

For these and other reasons, organizations have found that VMware is less well optimized to handle the business-critical systems and complex multi-partition production environments for which Power servers are routinely employed.

Applications

Tens of thousands of organizations worldwide employ Power servers for production enterprise resource planning (ERP), customer relationship management (CRM), supply chain management (SCM) and other core business systems, as well as high-performance business intelligence and other high-end applications.

The VMware deployment picture is different. A recent survey by the International Technology Group found, for example, the distribution of VMware applications shown in figure 1.

Figure 1
Distribution of VMware Applications

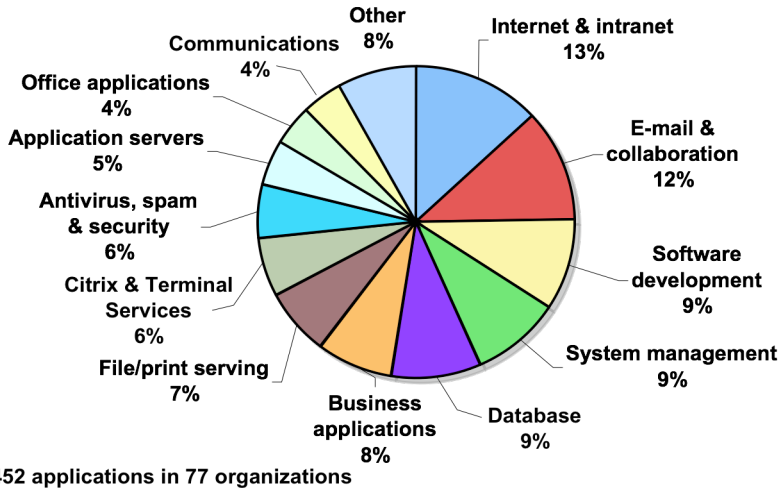

Base: 452 applications in 77 organizations
Source: International Technology Group

As these results suggest, VMware is most commonly employed for what are, by Power server standards, comparatively light-duty applications and workloads. Where databases are employed, these are typically small, non-critical or both. This is also the case for business applications.

This picture is confirmed by a 2008 customer survey conducted by VMware (the company). According to the survey, the most common production "ERP" system running on VMware was Great Plains, a Microsoft small business accounting package. (It was recently renamed Microsoft Dynamics GP).

There are some installations of SAP, PeopleSoft and other mainstream ERP systems on VMware. However, these are typically small and often involve development and test instances, or collaborative modules offered for these suites rather than transactional systems. Oracle, for example, does not support VMware for SAP production environments.

VMware and equivalents will clearly play major roles in enabling organizations to contain the effects of x86 server proliferation. But they do not possess the robustness and management capabilities of the Power virtualization environment, and they are not realistic alternatives for the core transaction- and database-intensive systems for which Power servers are normally employed.

Costs

More advanced does not necessarily mean more expensive. Even for the applications and workloads that are commonly deployed using VMware, Power servers may be cost-competitive.

This is illustrated by two sets of comparisons of five-year costs for x86 servers employing VMware, and Power servers employing IBM PowerVM that were prepared for this report.

Comparisons are based on composite profile installations in financial services, technology and government organizations with between 3,000 and 12,000 employees. Scenarios were developed for the use of VMware Infrastructure 3 on Dell Intel-based servers with Windows Server 2008, and PowerVM on IBM Power servers with IBM AIX 6.1, supporting the same applications and workloads.

Five-year IT costs as well as the costs of downtime – meaning business costs incurred by organizations due to outages – were then determined for each installation and scenario.

Results were as follows.

- ***IT costs*** for PowerVM scenarios ranged from 23 to 49 percent less, and averaged 34 percent less than those for VMware equivalents. Figure 2 illustrates these results.

Figure 2
VMware and PowerVM Scenarios: Five-year IT Costs

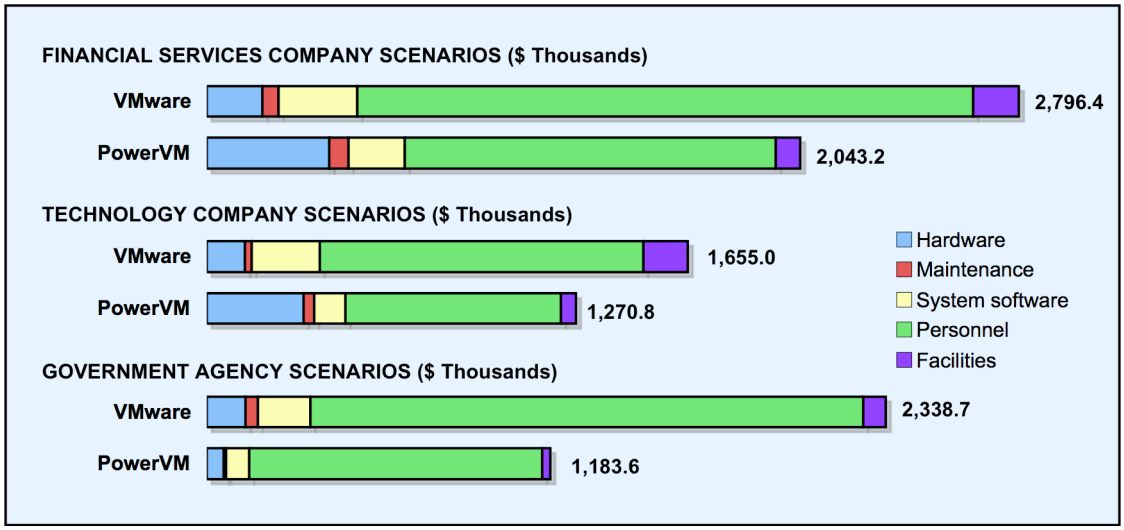

  Costs include acquisition of hardware and licenses for operating systems and virtualization software; hardware maintenance and software support; personnel costs for system administration and related functions, including VMware and PowerVM administration tasks; and facilities costs including data center occupancy, power and cooling, and energy consumption.

- ***Costs of downtime***, meaning business costs incurred by organizations due to outages, ranged from 59 to 77 percent less, and averaged 72 percent less than those for VMware equivalents.

  Costs are for lost user productivity resulting from service interruptions. Calculations were based on the frequency and duration of outages, and on numbers and average remuneration of users affected by these.

It should be emphasized that configurations and cost comparisons were based upon types of application typically deployed using VMware. It would not have been feasible to compare larger, more business-critical systems deployed upon Power servers because x86 VMware solutions could not have realistically handled these in most cases.

These results are presented and discussed in more detail Cost Calculations section of this report. Additional information on profile installations, scenarios, methodologies and assumptions employed, and cost breakdowns are also presented in this section.

COST CALCULATIONS

Basis of Calculations

The comparisons of five-year costs presented in this report are based on composite installations. Scenarios were constructed using data on applications, server configurations, use of virtualization tools, utilization levels, and other variables supplied by 23 organizations in the same industries and approximate size ranges, with generally similar business profiles.

Organizations employed VMware, PowerVM or both. VMware users had deployed VMware solutions on a range of x86 servers from Dell, HP, IBM and other vendors, primarily running Windows Server 2003. PowerVM users employed various Power and System p models with AIX Versions 5 and 6.

Where companies employed older versions of platforms, configurations were updated to current-generation models using vendor or industry comparative performance data. Comparable configurations were then developed for the other vendor's platforms based on industry comparative performance data.

Server configurations were equipped with Windows 2008 or AIX 6.1 operating systems, along with VMware Infrastructure 3 Standard or Enterprise Edition, or PowerVM Standard or Enterprise Edition. IBM BladeCenter systems were also equipped with BladeCenter Open Fabric Manager (BOFM).

Installations and scenarios for these comparisons are summarized in figure 3.

In VMware scenarios, VMware Virtual Machines (VMs) were employed to host between 142 and 303 operating system instances.

In comparison to VMware, which employs a single software-based partitioning method, PowerVM enables use of two complementary partitioning technologies – firmware-based LPARs and software-based Workload Partitions (WPARs), which allow users to create multiple partitions within a single operating system instance.

In PowerVM scenarios, combinations of these were employed to handle the same applications and workloads. WPARs were in many cases employed as alternatives to VMs.

For reasons discussed earlier, the applications upon which comparisons were based were not database-intensive or business-critical in nature. They included a variety of departmental, messaging and collaboration, Internet and intranet infrastructure, and other comparatively light-duty applications, as well as development, test and other non-production instances for these and other systems.

Five-year IT costs were then calculated for each scenario. Database software costs were not included in comparisons. In the organizations upon which profile installations were based, VMware was not widely employed for database serving. Where it was employed in this role, it was typically with Microsoft SQL Server, which is not supported for AIX.

IT Cost Disparities

IT cost disparities were primarily due to higher personnel costs for VMware scenarios, reflecting both the use of fewer physical servers in PowerVM scenarios, and the higher productivity of AIX and PowerVM administrators compared to their Windows and VMware counterparts.

Although PowerVM hardware costs were for the financial services and technology companies, the effects of more expensive Power servers were to a considerable extent offset by the higher capacity utilization levels realized by these. This was due to the use of multiple forms of partitioning to handle varying application requirements, as well as to more effective system and workload management capabilities.

Figure 3
Profile Installations and Scenarios

COMPANY	FINANCIAL SERVICES	TECHNOLOGY	GOVERNMENT
Business Profile	Regional operations of diversified global bank 5,000+ employees 15 subsidiary & local offices	Data communications equipment manufacturer $1 billion revenues 3,000 employees	National government agency 12,000 employees 12 regional offices
Focus of Comparisons	**Business applications** Departmental applications, document management, financial reporting, project management, payments **Infrastructure** Access control, antivirus, e-mail, firewalls, gateways, Internet & intranet servers, system management **Other applications** Development, test & staging, sandboxes, user acceptance & training instances	**Business applications** Application servers, departmental applications, collaboration & office applications, various **Infrastructure** End user computing, Internet & intranet infrastructure, network management, system management **Other applications** Development, test, training & other instances	**Business applications** Custom applications, departmental applications, statistical analysis & reporting, workforce management, various **Infrastructure** Antivirus, Apache, firewalls, FTP servers, security applications, Sendmail, SMTP gateways, system management, various **Other applications** Development, test, training & other instances
VMWARE SCENARIOS			
Servers & software	PowerEdge M1000e 1 x M600 2/8 x 3.0 GHz 1 x M600 2/8 x 2.83 GHz 2 x M600 2/8 x 2.5 GHz 6 x M600 2/8 x 2.0 GHz 1 x R900 4/24 x 2.67 GHz 2 x R900 4/24 x 2.4 GHz 1 x R900 4/16 x 3.0 GHz 3 x R900 4/16 x 1.6 GHz 4 x 2900 III 2/8 x 2.0 GHz 2 x 2900 III 2/8 x 3.0 GHz 1 x 2900 III 2/8 x 2.5 GHz 5 x 1950 III 2/8 x 2.0 GHz Windows Server, VMware Totals: 29 servers 273 VMs	2 x R900 4/24 x 2.4 GHz 1 x R900 4/16 x 2.13 GHz 4 x R900 4/16 x 1.6 GHz 1 x 2900 III 2/8 x 3.0 GHz 2 x 2900 III 2/8 x 2.66 GHz 3 x 2900 III 2/8 x 2.33 GHz 11 x 2900 III 2/8 x 2.0 GHz Windows Server, VMware Totals: 24 servers 142 VMs	2 x PowerEdge M1000e 2 x M600 2/8 x 3.16 GHz 3 x M600 2/8 x 2.66 GHz 2 x M600 2/8 x 2.5 GHz 7 x M600 2/8 x 2.33 GHz 11 x M600 2/8 x 2.0 GHz Windows Server, Linux, VMware Totals: 27 servers 303 VMs
Personnel	3.9 FTEs	2.05 FTEs	3.5 FTEs
POWERVM SCENARIOS			
Servers & software	BladeCenter H 7 x JS23 2/4 x 4.2 GHz 3 x JS12 1/2 x 3.8 GHz 5 x 550 4/8 x 4.2 GHz 3 x 520 2/4 x 4.7 GHz AIX, PowerVM BOFM (BladeCenter) Totals: 18 servers 128 LPARs 176 WPARs	4 x 550 4/8 x 4.2 GHz 5 x 520 2/4 x 4.7 GHz AIX, PowerVM Totals: 9 servers 42 LPARs 115 WPARs	BladeCenter H 2 x JS43 4/8 x 4.2 GHz 7 x JS23 4/8 x 4.2 GHz AIX, Linux, PowerVM, BOFM Totals: 9 servers 20 LPARs 303 WPARs
Personnel	2.15 FTEs	1.25 FTEs	1.7 FTEs

Power server workload management capabilities play an important role in ensuring availability. Highly granular and responsive mechanisms ensure that surges in workloads running in individual partitions do not impact the performance of others, or crash the entire server. Concerns that this might occur have deterred many organizations from employing VMware for production applications that are sensitive to response time, availability or both.

In a Power server environment, partitioning is tightly integrated with system and workload management functions built into PowerVM firmware and AIX 6.1. LPARs and WPARs are widely employed even for large-scale, business-critical production systems.

Users may also dedicate I/O resources to specific LPARs. This approach, which is normally employed for applications that are particularly sensitive to I/O throughput, availability or both, minimizes risks that insufficient I/O resources will be available to the system running in the LPAR in the event of an unexpected workload surge.

Costs of Downtime

Costs of downtime are for the same applications and configurations as IT cost calculations. Results are summarized in figure 4.

Figure 4
Five-Year Costs of Downtime

ORGANIZATION	FINANCIAL SERVICES	MANUFACTURING	GOVERNMENT
VMWARE SCENARIOS			
Availability levels	99.73% – 99.93%%	99.54% – 99.89%	99.63% – 99.9%
Five-year costs ($000)	**706.9**	**142.7**	**423.3**
POWERVM SCENARIOS			
Availability levels	99.95% – 99.99%	99.6% – 99.98%%	99.7% – 99.98%
Five-year costs ($000)	**163.2**	**58.3**	**129.1**

In this presentation, availability percentages reflect hours of system-level outages relative to annual hours of operation of companies, or business areas within companies that are supported by specific applications.

Costs are for lost user productivity. This was measured in terms of (1) idle time, reduced productivity or both for periods when applications were not available to users during working hours and (2) reduced productivity following an outage; e.g., a one-hour outage causing a 60 percent productivity reduction might be followed by a two-hour period in which a 20 percent productivity loss occurs.

Productivity loss calculations were developed for each user community supported by applications shown in figure 3. Productivity losses were quantified based on the number of individuals using each application who would be affected by outages, their average remuneration per hour, and percentage values for reductions in their productivity during and subsequent to outages.

Average remuneration per hour was calculated based on industry median salaries, benefits and other compensation for the relevant occupational groups in financial services and manufacturing companies, and government agencies in the same general size range as profile organizations. Average numbers of hours worked per year varied by industry and occupational group.

Although VMware supports clustered failover, live partition movement and other availability optimization functions, these are generally less effective than PowerVM equivalents, and require more extensive administrator intervention.

The number of patches that must be applied to VMware and Windows is also greater than for PowerVM and AIX, which resulted in longer and more frequent planned outages for VMware scenarios. Use of WPARs, which simplify operating system maintenance and patching tasks (only one instance is required for multiple applications), also contributed to lower PowerVM costs of downtime.

Additional Detail

For all installations and scenarios, calculations include hardware and software license acquisition, as well as hardware maintenance, software update and support subscriptions, and personnel and facilities costs.

All hardware maintenance and software support costs are for 24x7 coverage. Costs for VMware scenarios include Microsoft Software Assurance and VMware Platinum Support.

Detailed breakdowns of IT costs are presented in figure 5.

Figure 5
Five-year IT Cost Breakdowns

ORGANIZATION	FINANCIAL SERVICES	TECHNOLOGY	GOVERNMENT
VMWARE SCENARIOS			
Hardware	192.4	130.5	132.8
Maintenance	54.5	25.3	43.7
Systems software	271.8	233.4	180.9
Personnel	2,120.5	1,114.6	1,903.0
Facilities	157.2	151.2	78.3
TOTAL ($000)	**2,796.4**	**1,655.0**	**2,338.7**
POWERVM SCENARIOS			
Hardware	420.5	333.5	58.0
Maintenance	67.7	36.7	8.9
Systems software	193.7	107.4	78.4
Personnel	1,276.4	742.1	1,009.3
Facilities	84.9	51.1	29.0
TOTAL ($000)	**2,043.2**	**1,270.8**	**1,183.6**

Costs for server hardware acquisition and maintenance, and for software licenses and update and support subscriptions were calculated based on vendor list prices. These were discounted to reflect prevailing "street" prices for installed bases and new deployments comparable to those in profile installations.

Personnel costs were calculated based on annual salaries of $72,640 for Windows server and VMware system administrators, and $79,316 for PowerVM and AIX system administrators. Salaries were increased by 49.7 percent to allow for benefits, bonuses, training and other personnel-related overhead, and multiplied for a five-year period.

Facilities costs include data center occupancy, power and cooling equipment and energy costs over a five-year period. Occupancy cost calculations were based on EIA 42U rack mount units and service clearances for these, plus allowance for inactive areas. A conservative assumption for annual cost per square foot for existing facilities was employed (i.e., costs do not include new facilities construction).

Costs for power and cooling equipment were based on configurations of such equipment appropriate for the servers employed in each installation and scenario. Costs were calculated for acquisition and maintenance over a five-year period, using discounted list prices for equipment from leading vendors.

Energy costs were calculated using vendor electricity consumption values for servers, as well as power and cooling equipment. Specific utilization levels and hours of operation for each profile installation were applied, and a conservative assumption for average price per kilowatt/hour was employed to determine five-year costs.

All cost values are for the United States.

ABOUT THE INTERNATIONAL TECHNOLOGY GROUP

ITG sharpens your awareness of what's happening and your competitive edge
. . . this could affect your future growth and profit prospects

The International Technology Group (ITG), established in 1983, is an independent research and management consulting firm specializing in information technology (IT) investment strategy, cost/ benefit metrics, infrastructure studies, deployment tactics, business alignment and financial analysis.

ITG was an early innovator and pioneer in developing total cost of ownership (TCO) and return on investment (ROI) processes and methodologies. In 2004, the firm received a Decade of Education Award from the Information Technology Financial Management Association (ITFMA), the leading professional association dedicated to education and advancement of financial management practices in end-user IT organizations.

The firm has undertaken more than 100 major consulting projects, released approximately 160 management reports and white papers, and delivered nearly 1,800 briefs and presentations to individual clients, user groups, industry conferences and seminars throughout the world.

Client services are designed to provide factual data and reliable documentation to assist in the decision-making process. Information provided establishes the basis for developing tactical and strategic plans. Important developments are analyzed and practical guidance is offered on the most effective ways to respond to changes that may impact or shape complex IT deployment agendas.

A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources. Customized client programs involve various combinations of the following deliverables:

Status Reports	In-depth studies of important issues
Management Briefs	Detailed analysis of significant developments
Management Briefings	Periodic interactive meetings with management
Executive Presentations	Scheduled strategic presentations for decision-makers
Email Communications	Timely replies to informational requests
Telephone Consultation	Immediate response to informational needs

Clients include a cross section of IT end users in the private and public sectors representing multinational corporations, industrial companies, financial institutions, service organizations, educational institutions, federal and state government agencies as well as IT system suppliers, software vendors and service firms. Federal government clients have included agencies within the Department of Defense (e.g. DISA), Department of Transportation (e.g. FAA) and Department of Treasury (e.g. U.S. Mint).

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net