A HURWITZ RESEARCH REPORT

Creating a Sense of Application Awareness in IT Virtualization Environments

Marcia Kaufman, COO
and Principal Analyst

HURWITZ
& ASSOCIATES
Insight to Action

Sponsored by IBM

Introduction

As IT organizations work to increase the flexibility, efficiency, and cost effectiveness of their data centers, many have focused on server virtualization. These companies are also leveraging server virtualization as an important starting point for their journey to adopt cloud computing. However, server virtualization alone is not sufficient to help companies realize the full compliment of cloud computing benefits. In essence, to achieve a well-managed data center requires a more application centric approach. One way to ensure that the end user is supported as part of the virtualization strategy is to make the resource allocation process application aware.

Why is this issue so important? The use of applications is not static. It is common that during certain business cycles, the demands on a particular application will increase and therefore require additional unanticipated computing resources. When this happens in a heavily virtualized environment, the virtual machines that are supporting the application may not be able to handle the changing workload.

Companies need a way to prioritize applications not only based on infrastructure requirements and policies, but with an understanding of the business context of the applications in question. In order for companies to ensure that the most business critical applications are given top priority, its service polices must reflect the relative business impact of each application. Server virtualization allows the addition of entirely new virtual machines to an application, but this takes time, and the spike could be gone before the new resources are ready. Furthermore, the action of adding a new virtual machine is only useful if the system knows when to use it. Server virtualization products can look at system metrics such as memory and CPU usage, but this gives only a very high level view of what's really going on, and completely ignores any difference in the priority of different workloads in periods of resource contention. As such, with server virtualization alone these additional applications may suffer poor performance even if they have a higher priority to the business.

… with server virtualization alone these additional applications may suffer poor performance even if they have a higher priority to the business.

Balancing server virtualization with application virtualization

The best way to balance the need for server consolidations with the need to support changing user expectations is to deploy server virtualization, using a hypervisor product such as IBM's PowerVM, IBM's z/VM, VMware's vSphere, or KVM in combination with application virtualization. It is becoming apparent that server virtualization and application virtualization are complementary approaches. When balanced, they will make it easier for developers and data center managers to support performance requirements and meet service levels. By combining application virtualization with server virtualization, a greater number and more complex applications can be managed by fewer servers. The resulting environment will be more stable and application requirements for high availability and redundancy can be met.

HURWITZ & ASSOCIATES
Insight to Action

This paper will explain the characteristics of application virtualization and the benefit of implementing application virtualization in connection with server virtualization. In addition, this paper will describe IBM's application virtualization solution, IBM WebSphere Virtual Enterprise, and how customers benefit from combining this solution with a hypervisor.

The case for application virtualization: a customer scenario

A mid-sized US bank was facing increased competition in its region. Several larger national banks had recently opened new branches and were offering promotions to encourage customers to switch to their bank. The bank had a long-standing reputation for excellent customer service, but some recent service outages and slow response times for their online services had led to many customer complaints. The senior executives were concerned and looked to the IT team to work out a solution. IT described how the bank's technology infrastructure had become increasingly complex as it was adapted to support a more diverse portfolio of business service offerings. In addition to keeping up with a steady flow of new business requirements, IT was also updating existing services on a more frequent basis. Initially, server virtualization had increased IT responsiveness and flexibility, but most recently IT experienced unexpected deterioration in service levels when they tested new application versions in production.

The bank's external facing applications were maintained on a cluster of servers. When the IT team needed to test and deploy a new application, they first needed to take the whole cluster of servers off-line. Only after they put in the new application would they be able to bring the cluster back into production. This approach made it almost impossible to maintain service quality. IT simply did not have the insight they needed into the load handled by each server in the cluster to understand the implications of their actions on the performance of highly critical customer applications.

IT found they could increase the cost saving benefits they were already receiving with server virtualization by implementing application virtualization as well. The automated features of the new system enabled routing decisions to be made based on information about the load experienced by each server in a cluster. This improved intelligence about the performance of each server led to major improvements in service levels. The bank was able to efficiently manage more applications without buying additional servers. Most importantly, the combination of server virtualization and application virtualization enabled IT to process transactions with the reliability needed to support and create additional value for the business.

IT simply did not have the insight they needed into the load handled by each server in the cluster to understand the implications of their actions on the performance of highly critical customer applications.

HURWITZ & ASSOCIATES
Insight to Action

Resource utilization optimization with an end-user focus

Application virtualization is a capability that allows IT to optimize resource utilization while at the same time optimizing the customer experience. This dual optimization can take place because application virtualization enables the automation of many of the underlying rules and processes. By codifying business and technical usage policies and automating workload management processes, organizations can ensure that applications will leverage virtual or physical resources in a predictable way. Instead of running each workload on a dedicated set of resources, IT can improve service levels by allocating a pooled set of resources to support a group of workloads. The key benefit of application virtualization is that resources are allocated based on the priorities and requirements specific to each application making your system application aware. As a result of this increased awareness of the requirements of each specific application, IT can increase server consolidation and control virtual machine sprawl by improving the management process. There are three key characteristics for application virtualization:

Application virtualization is a capability that allows IT to optimize resource utilization while at the same time optimizing the customer experience.

Application prioritization. Application virtualization requires IT to set service level policies related to the required response times of the mission critical applications. These policies should be defined based on business context and customer requirements. In application virtualization environments, many aspects of application performance levels are monitored automatically. For example, response times can be monitored at a fine-grained level in accordance with the knowledge of which application a given request is for or which action is required for a specific application. As a result of this increased understanding of the application environment, automatic adjustment of resources can be made to keep customer service levels in balance.

Automation of resource optimization. Each application in an application virtualization environment gets the resources it needs based on resource availability and the application's priority to the business. This customer-focused approach to resource allocation dramatically reduces the chance that users or customers will experience performance degradation in mission critical applications and put business revenue at risk. The goal is to optimize the delivery of services to customers while using fewer servers more efficiently and cost effectively.

Setting policies and rules for high availability. Application virtualization environments are desined to run applications across a pool of resources to ensure that each application gets additional resources when needed based on the priority of the application. This means that any resource can run any workload. If there are resource limitations, the application with the lowest business priority at that point in time is allocated with the fewest number of resources. To facilitate changing requirements, workloads can be automatically moved to a different server if a server fails. The benefit of these capabilities is that upgrades can be managed with out outages and corrective actions can be taken when problems are detected.

HURWITZ & ASSOCIATES
Insight to Action

Improving the Service Level Agreement

The benefits of application virtualization are similar to server virtualization in that they both help to improve service level to customers. These two technologies each have a role to play in improving the utilization of IT resources, serving to isolate applications and data when needed, and increasing data center flexibility. However, the two approaches are complementary and deliver a far better result when used in combination. While server virtualization helps to improve overall data center economics, application virtualization helps to improve resource allocation based on business requirements and priorities for specific workloads.

While server virtualization helps to improve overall data center economics, application virtualization helps to improve resource allocation based on business requirements and priorities for specific workloads.

There are many very good reasons why companies have implemented server virtualization, using a hypervisor product like VMware's vSphere or IBM's Power VM. For example, many companies with a focus on cloud computing want to begin by eliminating a long standing source of discontent between IT and its end-users – lack of an efficient and fast provisioning process for IT resources. Server virtualization helps companies move forward on this path. User demand for IT resources is supported by rapidly spinning up a new virtual machine rather than using the time and funds required to provision a new physical machine for each request. In addition to faster provisioning, users also benefit from increased operational flexibility based on the mobility and isolation capabilities of virtual machines. For example, if an unanticipated surge in demand for IT resources leads to slow response times for a bank's online banking customers, the operations team can relieve the stress on the physical machine by moving a virtual machine from one physical machine to another. If this can be done on demand with an early warning system in place, then the end user will receive continuous service and may not even be aware of the problem. In addition, applications and processes are independent from each other in a virtual machine, which is important in hosted environments and for server maintenance. While these benefits are significant, companies with successful server virtualization implementations need to recognize that there are times when this approach falls short from an application perspective.

While companies that have deployed server virtualization are able to leverage the inherent management capabilities of the hypervisor, they are not able to monitor and control application specific performance. With server virtualization, the provisioning and management of hypervisors and virtual machines is done without any insight into the requirements and priorities of each unique application. As a result, there is no guarantee that each application will have the resources required to meet user demand on an ongoing basis. By pooling infrastructure resources, application virtualization ensures that server resources can be adjusted dynamically so that each application gets the resources needed at the right time. While server virtualization enables data center operations to easily move virtual machines from one host physical machine to another if needed for maintenance, mobility is not always so smooth in the case of unplanned server outages. In this situation, the benefit of application virtualization is that any one of the pooled virtual machines can process the application.

HURWITZ & ASSOCIATES
Insight to Action

One of the key benefits of server virtualization – rapid provisioning – is much more effective when companies use application virtualization to automate the control of the provisioning process. With server virtualization alone, the ease with which virtual machines can be instantiated can lead to unchecked growth in these machines. Each new virtual machine image requires memory and disk resources. When the number of virtual machines grows out of control, companies end up spending more time and money on disk, storage, and memory resources than was anticipated. If customers do not have access to the applications they require on a timely basis, then those dramatic cost savings realized from server virtualization will not look so valuable any more. Service level agreements cannot be met. Companies use policies set by the business to establish when resources can and cannot be provisioned.

One of the key benefits of server virtualization – rapid provisioning – is much more effective when companies use application virtualization to automate the control of the provisioning process.

Isolation in virtual environments is another important characteristic that is more impactful when the technologies of application and server virtualization are combined. With server virtualization alone, you have a more limited form of isolation or protection of your applications from errors that may occur with other applications. The applications in one virtual machine are protected against problems that may take place in other virtual machines. But, what do you do when there are multiple applications running in the same virtual machine and one of these applications is compromised? Or what if one application erroneously demands so many resources from the virtual machine that other applications cannot meet SLAs? The workload management capabilities of application virtualization will help you to ensure that these problems do not occur. You can make sure that applications are sufficiently isolated from problems both in the virtual machine and across all virtual machines in the environment.

IBM WebSphere Virtual Enterprise

IBM WebSphere Virtual Enterprise (WVE) is IBM's solution for application virtualization. It provides policy-based workload management capabilities that are intended to provide a greater level of control and consistency to the way organizations manage and improve performance and response times of applications. An important aspect of its functionality is the automatic management of the creation and deletion of virtual machines. The broad set of capabilities includes application server virtualization, resource management, and advanced operational facilities, such as performance visualization, health monitoring, and application editions. WVE can be implemented as a stand-alone product or as a component of the IBM WebSphere Extended Deployment package of application infrastructure products. It is intended to enhance and extend the benefits of server virtualization and includes support for virtualization technologies such as VMware ESX Server, IBM PowerVM and IBM z/VM virtualization technologies, and IBM WebSphere Portal.

There are four key foundational services within WVE including:

Policy based workload management. This capability is the essence of application virtualization. The organization sets policies and service level

HURWITZ & ASSOCIATES
Insight to Action

agreements to differentiate between the relative priorities of each application so their demands for resources can be balanced against how critical each one is to the business. WVE uses these policies to monitor application performance and automatically and intelligently alter the allocation of resources as required to meet SLAs.

Application health management. A range of application server health conditions such as number of time-out requests, response times, memory usage, and run times are defined and then monitored to anticipate and correct problems. The organization can create very fine-grained health policies so that server issues can be automatically corrected or even pre-empted before users are impacted. An alert system is used to keep administrators aware of status and to make manual adjustments if needed.

Operational management. This service is designed to provide IT operations with the visual alerts and in-depth information on IT resources that provide a higher level of operational control over the running of applications. In real-time. How stable is the resource-what actions need to be taken by the operator?

Application edition management. The capability is intended to help organizations to overcome a major source of downtime – the inaccessibility of the right version of an application. Therefore, this service enables multiple versions of an application to be in production at the same time. Therefore, an organization can perform an upgrade with out interruption.

Setting workload priorities

The ability to accurately set workload priorities so that business requirements and service levels are met on a consistent basis is central to achieving the benefits of a virtualization strategy. Operations managers can do some workload optimization with a manual approach but there are many limitations. Simply put, an operations manager cannot always anticipate changes in priorities and the impact of decreased response time. IT operations needs an automated way of reviewing response times across multiple applications to keep ahead of errors before they cause a noticeable decrease in response times for customer. How do you know which application should get top priority at any given moment when using a manual approach? An automated approach that gives access to historical performance will provide more realistic policies.

WVE includes an On Demand Router (ODR) designed to automate workload prioritization. The ODR maintains a list of workload requests and keeps track of all the resource requirements for each workload. IT operations managers can set policies to establish a hierarchy of priorities for specific workloads. Customers will benefit by getting more information about what to expect in terms of application response times. This doesn't mean that every customer will receive a fast response for his workload request. It simply means that the right workloads get the resources they need at the right time. For example, if IT resources

The ability to accurately set workload priorities so that business requirements and service levels are met on a consistent basis is central to achieving the benefits of a virtualization strategy.

HURWITZ & ASSOCIATES
Insight to Action

are very tight and a user is requesting access to a lower priority application, IT operations will be notified to delay taking action until resources become available.

Conclusion

To deliver a superior customer experience, IT organizations need to focus their attention on providing the right level of manageability for those applications that important customers rely on. This is easier said than done since critical workloads often compete for the same limited IT resources. In server-virtualized environments, these applications compete for resources without any indication of the relative importance of each workload. Incorporating application virtualization into your environment helps to create a more balanced and application aware approach. By combining server and application virtualization companies can more efficiently and cost effectively manage resources while respecting the business rules and business priorities.

By combining server and application virtualization companies can more efficiently and cost effectively manage resources while respecting the business rules and business priorities.

HURWITZ & ASSOCIATES
Insight to Action

About Hurwitz & Associates

Hurwitz & Associates is a consulting, market research and analyst firm that focuses on how technology solutions solve real world business problems. The firm's research concentrates on disruptive technologies, such as Cloud Computing, Service Oriented Architecture and Web 2.0, Service Management, Information Management, and Social and Collaborative Computing. We help our customers understand how these technologies are reshaping the market and how they can apply them to meet business objectives. The team provides direct customer research, competitive analysis, actionable strategic advice, and thought leadership. Additional information on Hurwitz & Associates can be found at www. hurwitz.com.

175 Highland Avenue, 3rd Floor • Needham, MA 02494 • Tel: 617-597-1724
www.hurwitz.com