SOLUTION WHITEPAPER

bmcsoftware

ACTIVATE BUSINESS WITH THE POWER OF I.T.™

Implementing BMC® Performance Assurance in a VMware Environment

Table of Contents

Executive Summary

With BMC Performance Assurance Suite for Virtual Servers, you can manage the performance of VMware ESX servers. You can analyze performance information for the virtualized systems (Windows and Linux) and obtain a system-level view of the VMware ESX server for server sizing, modeling, performance management, and daily reporting.

BMC Performance Assurance Suite for Virtual Servers provides the ability to:

- collect VMware ESX server configuration and system-wide statistics
- analyze this data with the BMC Performance Analyzer
- report on the data
- model the VMware server with BMC Performance Predictor

You can also view VMware reports using BMC Performance Perceiver, a dynamic Web server that provides easy access, viewing, and export of BMC Performance Assurance data for non-console users.

The advanced modeling capabilities in BMC Performance Predictor enable continual performance management through modeling growth scenarios.

Illustrations and examples of analyzing and modeling the data collected from the virtual system environments described in this paper use the version 7.3.00 BMC Performance Assurance console on Microsoft Windows. Similar support is available on the UNIX console.

What data is collected from VMware?

The following sections describe the types of information collected.

System-level information: Why do you need to collect data from the VMkernel?

In general, the component closest to the actual hardware resources has the best view of system activity. The virtualization layer (the VMkernel) keeps track of the actual CPU usage and other resources used by the individual VMs and provides APIs to access these values. These metrics are the most reliable statistics for measuring system activity.

The end result of performance is individual workload response time. The response time is driven by two factors: service time and usage. An application running on an operating system inside the VM may still have acceptable response time, even if it reports high CPU usage. In this case, it is the physical CPU usage reported by the VMkernel that drives the response time, not the usage reported by the operating system. Therefore, it is important to collect data from the VMware service console, which is the only VM that has access to the VMkernel.

Application-level information: What data is available from inside the VM?

VMs do not “own” any of the physical hardware resources in the VMware environment. The virtualization layer (the VMware kernel) delegates all of the resource requests among all of the active VMs. However, the virtualization layer provided by VMware is completely transparent to the operating systems. To each VM, it appears that it is the only operating system that is using the physical hardware, and therefore the statistics collected from the VM reflect this solitary viewpoint.

For example, consider an implementation where there are five VMs, each with one processor in a two-CPU system. Through the virtualization layer, the VMs get time slices of physical processors. When all five VMs are active, only two of them have physical CPU service at any given time. The operating system instance running inside the VM does not have specialized timing support to account for wall-clock time, so it may not accurately report allocation of physical resources to individual processes. In this example, because each VM does not always have access to all of the cycles of a CPU, the operating systems will always report a higher CPU usage than that measured from the VMware service console.

How do I interpret data collected from inside the VM?

Performance data collected from within the VM's operating system may be used to display the relative resource usage of individual processes. It is important to note that because all VMs share the real, physical system resources, accurate measurements of the total resource usage are not available from inside the VM.

However, using the process data collected from inside the VM, you can perform basic reporting and analysis of specific workloads and users within the constraints of the virtual environment. For example, you could quickly identify performance problems causing contention among groups of users of a certain application in the VM, and could resolve the problem by moving the users to another VM.

What data is displayed?

Displaying VMware data in BMC Performance Analyzer

BMC Performance Analyzer has five reports that present information specific to VMware, including configuration and resource consumption:

- VMware Summary - Use this report to view summary information for the VMware service console and the VMs for the system.
- VMware Memory - Use this report to view memory statistics for the VMs.
- VMware Disk - Use this report to view disk statistics for the VMs.
- VMware Network - Use this report to view network statistics for the VMs.
- VMware Per Processor - Use this report to view CPU utilization information (per processor) for the system.

In these reports, each VM is identified by its display name used by VMware, and the guest operating systems are also displayed. Each operating system instance, including the VMware service console operating system, is shown as an independent computer.

Displaying VM CPU usage in Visualizer

The VMware Hierarchy report (See “Figure 1: Visualizer VMware Hierarchy report” on Page 5) available in Visualizer, the BMC Capacity Database (CDB), shows the breakdown of activity on the VMware service console. The Hierarchy Graph displays the VMware server at the top of the tree, with the VMs as branches.

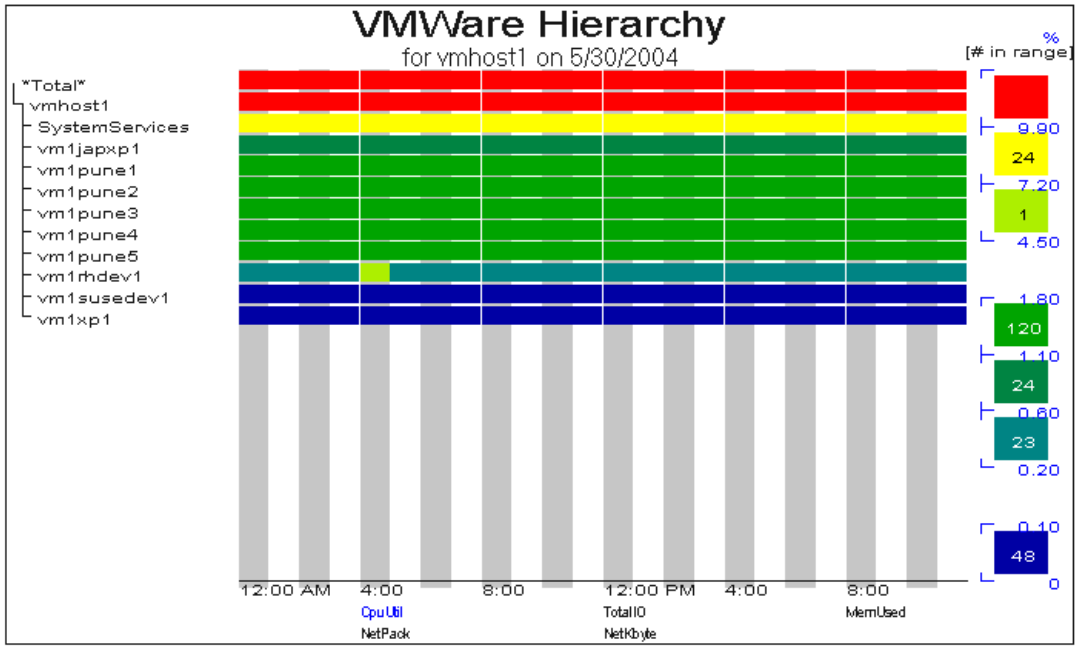

Figure 1: Visualizer VMware Hierarchy report

Modeling a VMware server

Using BMC Performance Predictor, model a VMware system to identify which virtual machines on the VMware server negatively affect overall performance. You can use this information to experiment with different configurations of VMware servers and virtual machines to see how to achieve the best mix. You can also use the Scenario Planner capability to perform long-term capacity planning.

Once you have established a valid baseline model, you can change the VMware host configuration or the virtual machine configuration and investigate how the changes will affect system performance. The following section explores a sample “What if...?” scenario involving VMware systems.

Example scenario

In this example, a company is looking to reduce the number of dedicated redundant servers by managing the failover of their 24 physical servers to a single ESX Server. They have decided they want the ability to have up to four servers simultaneously failover to one ESX Server, but do not know how big the VMware ESX Server and individual virtual machines need to be.

The following example demonstrates how BMC Performance Assurance can be used to accurately size an ESX Server to accommodate the workloads from four physical machines.

1. Collect performance data from the machines of interest. In this example performance data was collected from all of the 24 physical servers. To understand the performance of the systems and capacity, you need to continually collect performance data. In this example data was collected over a period of time which included the periods of peak demand. This peak demand period relates to the four most highly utilized servers in the server farm.

2. Define workloads that represent the total activity on each of the four machines.

3. After you've collected performance data over a meaningful period and created workloads for each machine in the study. Visualizer is the component that provides scheduled and ad-hoc graphical reports on application and system performance.

4. Use Visualizer to identify the time interval of peak demand for each server in the study, in order to estimate the size of the target VMware ESX server. You can draw performance graphs which enable you to understand the performance of the physical system(s). "Figure 2: Visualizer Node CPU Utilization" on Page 6 shows the total demand over a 24-hour period for the four servers in this scenario. The intervals between 10:00 - 12:00 can clearly be identified as the periods of peak demand.

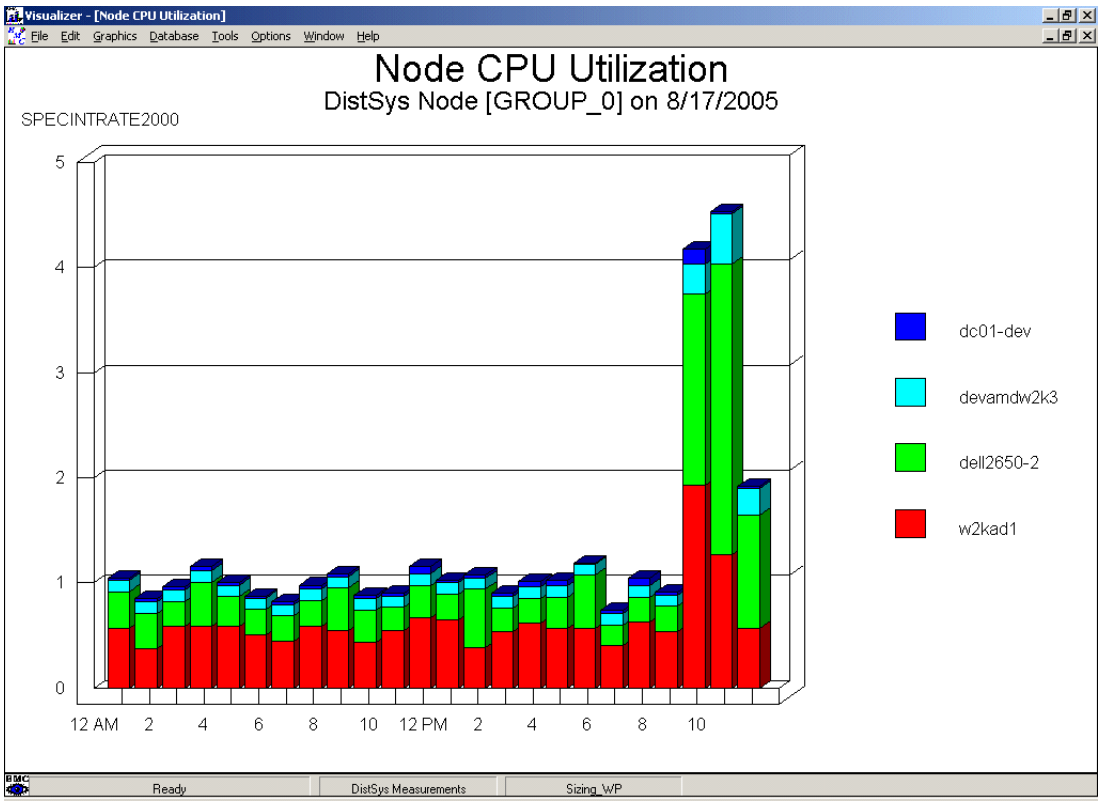

Figure 2: Visualizer Node CPU Utilization

5. Model the migration of the four physical system workloads to a VMware ESX server, as virtual machines. The workloads from each of the four physical machines can be imported into a new model, where you create the new 'target' ESX Server and four virtual machines.

6. Review the performance of the VMware Host system and all virtual machines, using the BMC Performance Predictor – VMware Summary report (See "Figure 3: BMC Performance Predictor – VMware Summary report" on Page 7).

Predict - VMware Summary

Model Name: VMtest
Analysis Interval: Sep-20-2005.18.15 -- Sep-20-2005.19.15.00
Created On: Wed Sep 21 12:55:36 2005

Number of rows to display: 15 Update

Currently sorted by: VMware Host / VM System

VMware Host	VM System	CPU Model	Vendor	Processor	System SVC	Util	% Utilization out of	Total Disk I/O	Total Memory	Overall Response Time Wait	Other VM Response Time Wait
						%		Kbyte/sec	MB	%	%
VHdevesx25	[Total rows: 1]	1750@3066	DELL	1.00	1.22	2.26	100.00	4.52	4,031.46		
	vmDevESX25Host2			1.00		1.04		4.52	512.00	1.04	0.00
VHesx2	[Total rows: 10]	6650@2000	DELL	4.00	32.80	77.24	400.00	236.53	15,871.50		
	esx2XPQA1			1.00		3.15		40.67	256.00	0.04	0.04
	vm-2003Ent			1.00		9.08		69.28	384.00	0.02	0.02
	joydeng			1.00		4.52		19.12	256.00	0.04	0.04
	esx2chris			1.00		2.92		15.05	1,024.00	0.04	0.04
	vmexception			1.00		5.45		45.44	1,024.00	0.04	0.04
	vmVZbldsvr			1.00		6.89		14.03	512.00	0.03	0.03
	esx2Win2kQA1			1.00		2.19		2.57	384.00	0.04	0.04
	vmxp-class			1.00		0.00		0.00	512.00		
	esx2Win2003QA1			1.00		3.06		11.02	384.00	0.04	0.04
	vm-ipredict1			1.00		7.17		19.35	512.00	0.04	0.04
VHesx3	[Total rows: 4]	6650@2000	DELL	4.00	3.85	11.43	400.00	29.98	15,871.50		
	esx3Win2003Dev1 (Chris Thornley)			1.00		2.14		11.48	1,000.00	0.00	0.00

Figure 3: BMC Performance Predictor – VMware Summary report

When reviewing the VMware Summary Report, pay special attention to the values in Overall Response Time Wait column and the Other VM Response Time Wait column. Overall Response Time Wait is the percentage of CPU wait time due to all virtual machines (including itself) on the ESX server. Other VM Response Time Wait indicates the percentage of CPU wait time due to activity from other virtual machines.

Comparing overall wait time and wait time due to other VMs helps you identify if you have balanced the VMs on a the ESX Server properly. The CPU shares and utilization both impact the wait time.

Once you have reviewed the report, you can also consider the following course of action:

- If performance is unacceptable, use the virtual machine properties page to reallocate memory, the number of virtual processors, and CPU shares until performance is acceptable.
- Use the Virtual System Allocation dialog to reallocate resources to all of the virtual machines.
- Use the VMware Host System properties page to change the configuration of the ESX Server (for example, to upgrade or downgrade CPU, memory, and disk).

For more information

BMC Performance Assurance Suite for Virtual Servers is the product for comprehensive performance management of virtualized server environments, specifically VMware ESX servers. It allows the analysis and reporting of performance information for the virtualized servers (Windows and Linux) and a system-level view of the VMware ESX server for server sizing, management, and daily reporting.

Refer to the following website for additional information:

www.bmc.com/products/proddocview/0,2832,19052_19429_20611260_112215,00.html

Helping you maintain advantage

BMC Software Education Services offers a strategic investment for your business, maximizing the value for your employees and Business Service Management initiatives. Education ensures successful product implementation, promoting mastery of all product capabilities and highest productivity with your BMC Software solutions. To explore our education offerings, visit our web page at http://www.bmc.com/bmceducation, or contact BMC Software Education Services by telephone or e-mail:

- North America
  Telephone: 800 574 4262
  E-mail: education@bmc.com

- Asia Pacific
  Telephone: +61 3 9657 4404
  E-mail: ap_education@bmc.com

- Europe, Middle East, and Africa (EMEA)
  Telephone: 00800 26233822
  E-mail: emea_education@bmc.com

About BMC Software

BMC Software helps IT organizations drive business value through better technology management. Our industry-leading Business Service Management solutions ensure that everything IT does is prioritized according to business impact, so IT can proactively address business requirements to lower costs, drive revenue, and mitigate risk. BMC solutions share BMC® Atrium™ technologies to enable IT to manage across the complexity of diverse systems and processes — from mainframe to distributed, databases to applications, service to security. Founded in 1980, BMC Software has offices worldwide and fiscal 2005 revenues of more than $1.46 billion. BMC Software. Activate your business with the power of IT. For more information, visit www.bmc.com

61851