# Best Practices for Virtualizing SAP BusinessObjects BI 4 on VMware 5

Technical Whitepaper

Version 0.9 DRAFT – SAP SAPPHIRE 2013 EDITION

***IMPORTANT NOTE:***

***This is a DRAFT document for SAP SAPPHIRE 2013 ATTENDEES and is NOT FINAL. Information in this version is subject to change until the final version is published.***

***The FINAL version of this document will be posted at www.sap.com/bivirtualization by June 2013.***

**Documentation in the SAP Service Marketplace**
You can find this documentation at the following address: http://service.sap.com/sizing

# Table of Contents


**1 INTRODUCTION ..... 6**

**1.1 Motivation ..... 6**
1.1.1 Background ..... 6
1.1.2 Hypothesis ..... 7
1.1.3 Scope ..... 7
1.1.4 Intended Audience ..... 7
1.1.5 Terminology ..... 8

**1.2 Products Used ..... 8**
1.2.1 SAP BusinessObjects BI 4 ..... 8
1.2.2 VMware ESXi 5.1 ..... 9
1.2.3 Supermicro Servers ..... 9

**1.3 Important Notes ..... 9**
1.3.1 Required Reading ..... 9
1.3.2 Lab Configurations Are Not Recommended For Production ..... 10
1.3.3 SAP's General Support Statement for Virtualization (SAP Note 1492000) ..... 10

**2 COIL TESTS ..... 11**

**2.1 Test Objectives ..... 11**

**2.2 Test Environment ..... 12**
2.2.1 SAP BusinessObjects BI 4 Environments ..... 12
2.2.2 Supporting Machines in Testing Landscape ..... 14

**2.3 Test Execution ..... 15**

**2.4 Analysis of Findings ..... 15**
2.4.1 Resource Contention Degrades Performance ..... 15
2.4.2 Improper System Configuration Can Significantly Impact Performance ..... 16
2.4.3 Synthetic Benchmarks Are Inappropriate For Virtualization ..... 16
2.4.4 Physical and Virtual Throughput Is Comparable ..... 17
2.4.5 Physical and Virtual Latency Is Comparable ..... 19
2.4.6 Physical and Virtual Degradation is Comparable ..... 20
2.4.7 BI-Specific Tuning Not Required In Virtual Environments ..... 21

**3 VIRTUALIZATION RECOMMENDATIONS ..... 21**

**3.1 Start With The Right Tools ..... 21**
3.1.1 Identify and Disregard Invalid Guidance ..... 21
3.1.2 Use the Latest Best Practice Guidance from VMware ..... 22

**3.2 Properly Configure All Hardware and Software Components ..... 23**
3.2.1 Intel Hyper-Threading and Turbo Boost Technologies ..... 23
3.2.2 Physical Hardware (BIOS level) ..... 24
3.2.3 Hypervisor ..... 25

3.2.4 Guest Operating System ..... 26
3.2.5 Application-Level Components (Database, Application Server, BI Platform) ..... 26

**3.3 Use Resource Reservations on all SAP BI Virtual Machines ..... 27**
3.3.1 Use Strict CPU Processing Reservations ..... 28
3.3.2 Use Strict Memory Reservations ..... 28
3.3.3 Addressing Concerns Using Reservations ..... 28

**3.4 Do Not Use Resource Shares, Limits, or Affinity ..... 29**

**3.5 Right Size Virtual Machines ..... 29**
3.5.1 Size VMs Large Enough To Meet Processing Requirements ..... 30
3.5.2 Size VMs Small as Possible ..... 30

**3.6 Always Deploy Virtualization Guest Tools ..... 31**

**4 ADDITIONAL TOPICS ..... 31**

**4.1 Architecture and Deployment Considerations ..... 31**
4.1.1 BI Sizing in a Virtual Environment ..... 31
4.1.2 Dealing with IT-Imposed Limitations ..... 32
4.1.3 Use Monitoring Metrics Wisely ..... 33

**4.2 Licensing Considerations When Virtualizing ..... 34**

**5 CONCLUSIONS ..... 35**

**5.1 Summary of Best Practice Recommendations ..... 36**

**6 REFERENCES ..... 38**

**About The Author**

Ashish C. Morzaria is a Director at SAP in the BI Solution Management organization. Ashish has worked for Business Objects and then SAP since 2004 driving a variety of initiatives such as platform strategy, virtualization, sizing and deployment, Cloud solutions, and next generation SAP HANA-based business intelligence.

**Acknowledgements**

This whitepaper is the culmination of many person-months of effort and while it is not possible to list everyone here, the author would like to thank the following people and partners for their invaluable contributions, without whom, this result would not be possible:

***SAP's Technology Innovation Platform (TIP) Vancouver P&R Team:***

At the core of any truly valid research project are experts who know how to break apart the problem, devise the right type of experiments to test theories, and most importantly, analyze the results to ensure the conclusions are valid. Derrick Wan, Jacques Buchholz, and Jonathan Chu-Chong tirelessly drove all lab testing and provided the oversight to ensure we reached valid conclusions. Additionally, without Corey Wilkie's belief from the beginning in the importance of this mission and his continued sponsorship, this whitepaper would still be a proposal.

***SAP's Palo Alto CO-Innovation Lab (COIL):***

Bringing together many parties to a single purpose is never an easy thing to do, but Kevin Liu and the SAP Co-Innovation Lab (COIL) boldly sponsored the project early on and ensured the project had the resources and partnerships we needed to succeed.

***VMware:***

Validating performance on another vendor's platform raises lots of questions and concerns, but VMware was a great partner and provided information, guidance, and support to the project even before it began.

***Supermicro:***

Generating confidence in test results requires enterprise grade hardware and a partner that goes the distance – Supermicro supplied all server hardware and the expertise on system configuration to ensure our team could focus on conducting the experiments instead of worrying about hardware issues.

***Other contributors and supporters:***

The author would also like to thank Ty Miller and Ingo Hilgefort who, while not directly involved in the project, were influential and believed in the importance of this project.

***IMPORTANT NOTE:***

***This is a DRAFT document for SAP SAPPHIRE 2013 ATTENDEES and is NOT FINAL. Information in this version is subject to change until the final version is published.***

***The FINAL version of this document will be posted at www.sap.com/bivirtualization by June 2013.***

# 1 INTRODUCTION

This whitepaper focuses on validating SAP's recommendations for deploying SAP BusinessObjects BI 4 systems in VMware environments. The content of this document is based on a series of tests performed in the SAP Co-Innovation Lab in Palo Alto, California in 2012. All tests were completed using the SAP BusinessObjects BI 4 suite and VMware ESXi 5.1 on Supermicro hardware. The guidelines contained in this whitepaper provide general recommendations and do not target any specific size of deployment, operating system, or solution architecture.

The topic of deployment is very complex and has many dimensions, so this paper will focus solely on virtualization best practices and assumes you have already understood how to size and architect high performing physical SAP BusinessObjects BI 4 deployments. All documents referenced in this whitepaper are noted with a *"[Ref. x]"* convention and are listed with source links in section 6 on page 38.

This document may undergo revisions and additions as feedback is collected, so you should always check http://www.sap.com/bivirtualization for the latest version of this document.

## 1.1 Motivation

### *1.1.1 Background*

Virtualization is a mature technology that has evolved to become the default deployment option for a large number of SAP customers. While SAP, and Business Objects before it, has supported virtualization for its business intelligence suite since 2006 (starting with Business Objects Enterprise XI R2), the question of performance degradation compared to physical systems has always been unanswered. If a BI system had a performance issue, virtualization was commonly cited as the primary cause with little effort or data to validate or dispute it.

Historically, it was believed that the performance delta between physical and virtual systems (termed "virtualization overhead" or "virtualization tax") was sizeable but acceptable given the other benefits of virtualization such as server consolidation and better resource utilization. Performance was also likely degraded because server hardware at the time was limited in both number of processors and amount of memory available, which forced system administrators who wanted high VM density to allocate less than the optimal amount of resources to each guest. There have been many attempts (both at SAP/Business Objects and externally) to determine the true performance delta, but while most succeeded in proving it did exist, there was no consensus on how big the difference really was.

Further examination of these studies shows that each was flawed in some way: in some, the software was not configured or licensed correctly, in others there was no true "apples-to-apples" comparison between physical and virtual instances, and finally many tests were simply either incorrectly executed or analyzed.

#### 1.1.2 Hypothesis

New advances in both virtualization software and extensions implemented in CPUs, chipsets, and interface cards suggest the “virtualization tax” is much smaller than previously thought. Additionally, server systems can be built with a much higher density of CPUs and memory than in the past, so it is also more feasible for administrators to allocate the proper amount of computing resources to ensure good performance.

This paper’s hypothesis can be summarized as:

> ***“Implementing SAP’s best practice recommendations for virtualizing SAP BusinessObjects BI 4 can result in a landscape that performs as well as, or with an acceptably small delta to, an identically configured physical deployment.”***

SAP has traditionally provided guidance in the form of presentations, webinars, and online forums, but customer feedback has been that more formal official guidance is required. This need is deeper than just a list of recommendations – BI teams also require validation of these recommendations in order to have a constructive dialog with the IT teams that maintain their BI system’s underlying virtualization infrastructure.

#### 1.1.3 Scope

This document makes specific recommendations for SAP BusinessObjects BI 4 systems based on the findings from our study and is not intended to be a replacement for general best practice guidelines from any of the referenced vendors, including SAP. The reader is assumed to have already read both the “SAP BusinessObjects BI 4 Installation Guide ” and “SAP BusinessObjects BI 4 Administration Guide” *[Ref. 4]* as well as VMware’s “Performance Best Practices for VMware vSphere 5.1” *[Ref. 7]* guide.

> ***Our SAP team performed all tests using VMware ESXi 5.1 and Microsoft Windows 2008 R2 on x86-based servers from Supermicro. While this whitepaper may refer to these vendors or their terminologies, the best practices for virtualizing SAP BusinessObjects BI 4 landscapes described here are transversal and applicable to other x86-based virtualization hypervisors and operating systems.***

#### 1.1.4 Intended Audience

This document is intended for both BI and IT infrastructure teams to understand how virtualization concepts and techniques can be used to implement virtualized SAP BusinessObjects BI 4 systems that can perform as well or at least similar to the same systems deployed physically.

You should already be familiar with the installation, deployment, and sizing aspects of the SAP BusinessObjects BI 4 suite as none of these topics will be covered here. This document is also not a substitute for working directly with SAP or an SAP partner that has real-world deployment or virtualization experience with the SAP BusinessObjects BI 4 suite.

#### *1.1.5 Terminology*

While this document refers to VMware specific terms and best practice documents, SAP's guidance on deploying SAP BusinessObjects BI 4 is hypervisor agnostic and can be applied to all supported hypervisors as listed in SAP Note 1492000 *[Ref. 5].*

The term "CPU" can be used in a variety of contexts, so this whitepaper will use the following definitions to differentiate and clarify:

- **CPU core:** This specifically refers to a single CPU core in a multi-core physical processor on the host system. For example, a dual socket, quad-core Xeon processor would have eight CPU cores.
- **CPU:** In most cases when referring to virtualization, this term is synonymous with a "CPU core" and not a physical processor. This is an important distinction as you may provision a "4 CPU" virtual machine, but this does not refer to 4 multicore processors, but typically refers to a virtual system with 4 cores. This whitepaper will use the "pCPU" and "vCPU" terms to avoid confusion wherever possible.
- **pCPU:** This is a short form for a "CPU core" as defined above.
- **vCPU:** This is a short form for a "virtual CPU core" that is provisioned in a virtual machine and has no mapping or relationship to a physical CPU or core in the host system. For example, an "8 vCPU VM" has eight virtual processors. The hypervisor is responsible for supplying the computational power for these vCPUs regardless of which or how many pCPUs are available on the host.

### 1.2 Products Used

#### *1.2.1 SAP BusinessObjects BI 4*

SAP BusinessObjects Business Intelligence suite (SAP BusinessObjects BI 4) is a flexible and scalable solution that provides the full spectrum of BI functionality including reporting and analysis, data exploration, dashboards, and self-service BI. It also gives IT departments a flexible means to share BI content throughout the entire organization, empowering business users to make effective, informed decisions via self-service access to information. Deployable on physical, virtual, and cloud environments, SAP BusinessObjects BI solutions are flexible enough to best fit the unique needs of each organization while providing a complete enterprise BI solution. Customers can decide better, perform better, and achieve better results throughout all areas of their business.

The solution is a multi-tier, server-based product that is comprised of a number of logical servers. Each report format has its own server modules that are controlled through a single management interface. These servers run as separate processes and they can all be installed on one machine or distributed across multiple machines, with multiple instances of individual servers able to run on each host. This allows for a very flexible and scalable architecture, with servers dedicated to specific processes and tasks, and of virtually any size, depending on how many concurrent users the environment needs to support.

A high level view of these logical servers is depicted below in green:

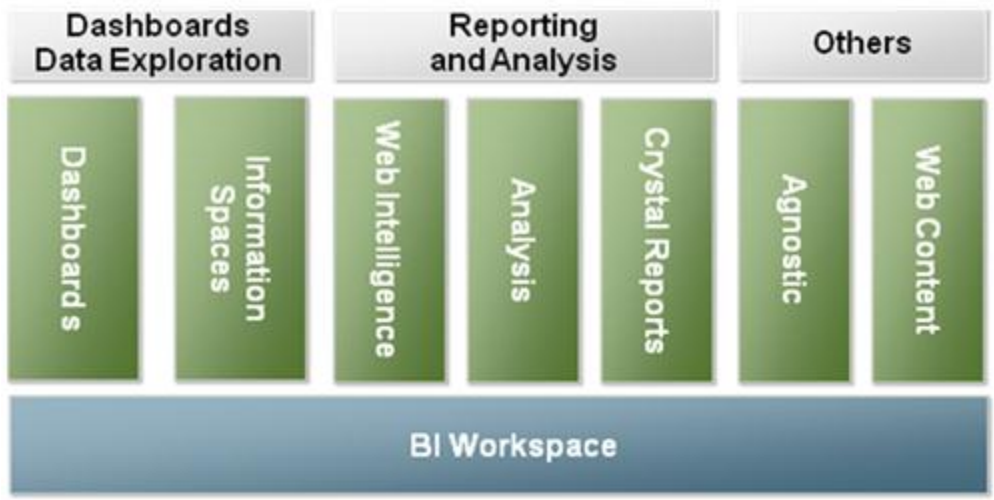


Storage of actual reports, dashboards, universes, and all other content is typically done on a NAS or SAN network storage infrastructure to take advantage of improved performance, reliability, and multipath I/O capabilities. Reporting databases and even the CMS database (CMS DB) are also typically on separate machines to increase the performance headroom of the BI system.

#### *1.2.2 VMware ESXi 5.1*

The VMware ESXi hypervisor is a core component of the vSphere product family that also includes the vCenter management suite of tools and added-value features such as Distributed Resource Scheduling (DRS) and High Availability (HA) among others. Customers can achieve tangible savings from this consolidation, and realize operational cost savings from reduced datacenter floor space, power, and cooling.

#### *1.2.3 Supermicro Servers*

Supermicro, a global leader in high-performance, high-efficiency server technology and innovation and is a premier provider of end-to-end green computing solutions for Enterprise IT, Data Center, Cloud Computing systems worldwide. The Supermicro TwinBlade is the most technologically advanced blade server system in the industry. Supporting up to 120 DP servers with 240 Intel® Xeon® 5600/5500 series processors per 42U rack, the TwinBlade achieves an unmatched 0.35U per DP node and is the Greenest, most energy-efficient blade server in the industry.

### 1.3 Important Notes

*This whitepaper focuses on validating certain virtualization recommendations and is not meant to be an exhaustive study of virtualization. It is instead meant to illustrate how applying these recommendations can result in a better performing system.*

#### *1.3.1 Required Reading*

Recommendations in this whitepaper are designed to be used in conjunction with, and are not intended to replace, guidance and best practices from SAP and third party vendors such as VMware and Microsoft. The following documents should be considered the bare minimum to ensure proper understanding of the SAP BI system and its components:

1. **SAP BI Platform Availability Matrix (PAM)** ***[Ref. 3]:*** The full list of stack components that are supported by SAP BusinessObjects BI 4 is always in the Platform Availability Matrix (PAM). This document is SAP's official statement of platform support for the BI suite.

2. **BI Platform Administrator Guide and BI Platform Installation Guide** ***[Ref.4]*****:** These standard documents contain architectural information and technical instruction on a wide number of topics which are intentionally not covered in this guide.

3. **SAP BusinessObjects BI 4 Sizing Guide** ***[Ref. 2]*****:** This document describes architectural concepts of the BI suite within the context of properly sizing a landscape and each of the BI services based on usage scenarios and user personas.

4. **SAP's BI 4 Virtualization microsite at www.sap.com/bivirtualization:** This microsite contains the latest version of this document as well as links to other resources related to virtualizing SAP BusinessObjects BI 4.

5. **Other documents at www.sap.com/bisizing:** This microsite contains links to other tools, documents, and presentations that provide additional explanation in adjacent topics.

Additionally, the following VMware-specific papers are required reading to ensure proper understanding of virtualization concepts and configuration:

1. **VMware's Best Practices for vSphere 5.1** ***[Ref. 7]*****:** This document is the definitive guide for how to configure, operate, and optimize VMware ESXi for general usage. The majority of guidance is applicable to SAP BI and should be considered additive to the guidance provided here.

2. **Evaluating Selected Java Best Practices for SAP BusinessObjects BI 4** ***[Ref. 9]*****:** This document is peripheral and focuses on the topic of whether to enable Java Large Page Support for the BI suite.

### *1.3.2 Lab Configurations Are Not Recommended For Production*

The CO-Innovation Lab (COIL) environment consisted of two identically configured hardware systems. This allowed a sharp focus on testing the effects of virtualization with minimal interference from environmental factors such as other virtual machines, congested network and I/O links, and CPU oversubscription. In a real-world deployment, the impact of these external factors can far outweigh the effects of virtualization overhead measured in a lab environment. It is therefore important to ensure your BI landscape is evaluated holistically instead of just focusing on virtualization overhead.

Furthermore, the lab configurations discussed in this document were designed specifically to test specific elements of the BI suite and should not be considered examples of good production architecture. Single node deployments do not provide the scalability and redundancy advantages of a properly architected multi-node architecture. Care has been taken to ensure the test results presented in this whitepaper and resulting recommendations are equally valid for multi-node deployments.

### *1.3.3 SAP's General Support Statement for Virtualization (SAP Note 1492000)*

In 2011, SAP adopted a single virtualization strategy across the majority of SAP products. This greatly simplifies configuration and support issues as different SAP products no longer have separate support policies. Full details are specified in SAP Note 1492000 *[Ref. 5]* and apply to all SAP products not explicitly excluded within this note.

SAP will troubleshoot and address performance issues of the SAP BI suite within a virtual environment, but the customer is responsible for ensuring the operating environment is configured correctly and operating optimally. During a support incident, SAP Support may not specifically check for virtualization-related issues and the root cause of an underperforming system may never be found.

***As a customer, you are fully responsible for:***

- Sizing and architecting the BI landscape correctly, including consulting with an SAP professional on BI architecture.
- Ensuring enough CPU, RAM, and I/O capacity for the host and guest are provisioned, configured/reserved, and available at all times.
- Optimizing the host and hypervisor environment for performance and stability as close as possible to bare-metal.

Many customers with physical deployments who experience performance problems typically have not allocated enough dedicated hardware resources or have designed a substandard BI system architecture – having either of these problems in a virtual deployment has far more performance implications and the effects will be far more noticeable.

> ***In the vast majority of cases, performance problems blamed on virtualization are usually a result of design decisions a customer has made and not virtualization or the BI suite.***

# 2 COIL TESTS

## 2.1 Test Objectives

Our objectives can be summarized as follows:

1. ***Validate:*** Our hypothesis is that best practice recommendations from SAP and VMware can be applied to create a virtualized SAP BusinessObjects BI 4 deployment with minimal virtualization overhead.

2. ***Disprove***: Previous test and publications supposedly proved an unusually large overhead when deploying SAP BI on VMware but there is evidence that most of these studies had critical flaws.

Our tests were devised to test our virtualization best practices and educate customers on how to apply these recommendations in their own environment. It is important to note that implementing best practices in production environments may not always result in increases in performance but can yield more important benefits by enabling the system to better survive peak loads, provide fast response times to end users, and demonstrate consistent performance across varying conditions.

The challenge in executing tests that do not focus solely on performance is that measuring the positive effects of specific configuration changes on overall system stability and consistency is very difficult to prove definitively. However, noticing significant deterioration or erratic behavior during the tests would be clear signs of a problem.

Keeping these goals in mind, our tests did not focus simply on raw performance but other criteria such as response time consistency and system behavior. For example, we closely monitor the resource usage consumption across the whole landscape and ensure the response time is consistent as we scale up the user loads.

> ***NOTE: The configurations described in this whitepaper are not typical or even recommended for production deployments. Each component here functions as a single point of failure, and was designed primarily to investigate the performance impacts of virtualization. By isolating the BI Platform from its CMS and reporting databases, we aimed to eliminate the potentially performance impacting load of these systems on our tests.***

## 2.2 Test Environment

The test environment was hosted at SAP's Palo Alto Co-Innovation Lab (COIL). All lab tests were performed by members of the SAP Technology Innovation Platform (TIP) BI P&R team.

The environment consisted of two physical systems, one representing a "bare metal" configuration and the other a VMware ESXi-hosted virtual configuration. Additional physical machines were also deployed for non-BI operations such as databases and load generators. All systems were completely dedicated to the lab tests to avoid any external factors.

### *2.2.1 SAP BusinessObjects BI 4 Environments*

Both the physical SAP BusinessObjects BI 4 system and the VMware host machine were configured identically and used default BIOS settings with the following settings explicitly enabled if they were not already system defaults:

- **Power Management Strategy:** High Performance
- **Hyper-Threading:** Disabled
- **Disk Controller Mode:** AHCI with DMA and write caching enabled

*Note: SAP recommends that production deployments always have Hyper-Threading enabled*. The feature was disabled during testing to eliminate potential performance impacts caused by differences in how Hyper-Threaded cores are utilized differently by physical machines and virtualization hosts.

#### *2.2.1.1 Physical SAP BusinessObjects BI 4 System*

The reference physical system (pss10147) was an SAP BusinessObjects BI 4.0 SP5 system equipped with 32 GB of memory running Microsoft Windows 2008 R2 on a 12 core (2 x 6-core) Xeon-based Supermicro server. To eliminate the performance impact of CMS DB and reporting database operations on the rest of the BI system, separate machines were used for all tests instead of relying on either the embedded database comes with SAP BusinessObjects BI 4 or an onboard reporting database.

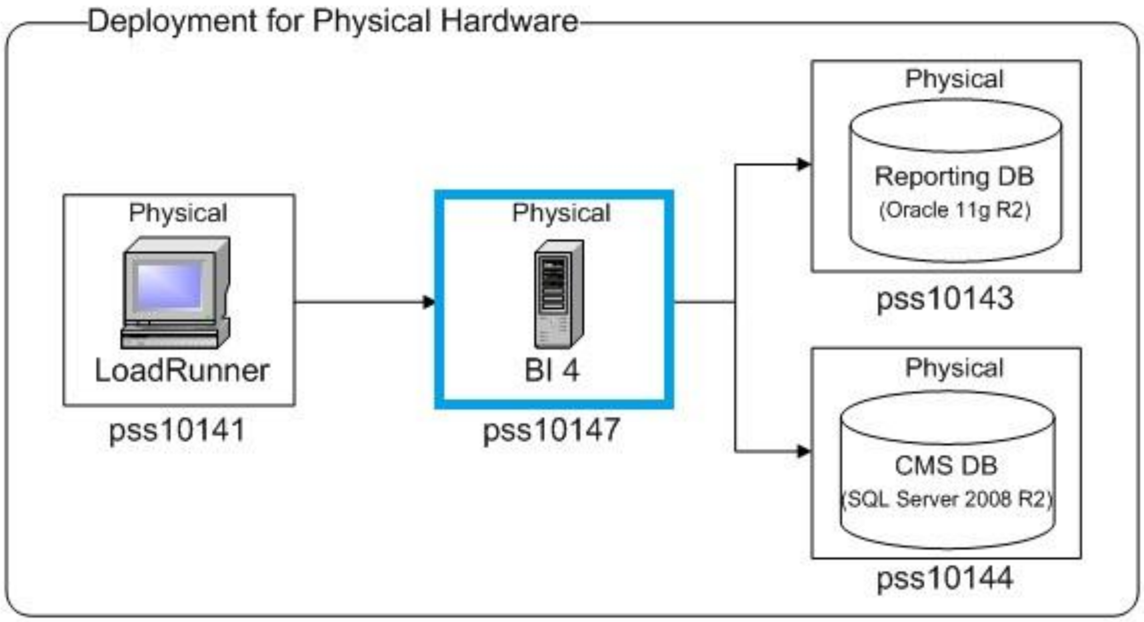


*2.2.1.2 Virtualized SAP BusinessObjects BI 4 System*

The virtualization host used VMware ESXi 5.1 using all installation defaults. This host was configured with 64 GB of RAM to enable a virtual machine with 32 GB of RAM to match specifications of the physical test machine.

As resource contention is a significant cause of performance degradation in virtual environments, our tests were conducted with only the single SAP BusinessObjects BI 4 virtual machine running on the host. While resource reservations are used to accomplish this in production environments, isolating the virtual machine physically also enabled us to eliminate resource contention at the network and storage levels as well.

The SAP BusinessObjects BI 4 virtual machine (pss10145) was built from scratch with an identical configuration to the physical BI 4 machine above. A “physical-to-virtual” (P2V) migration tool was not used to create the identical configuration to eliminate any potential performance impacts of using a different installation method for each instance.

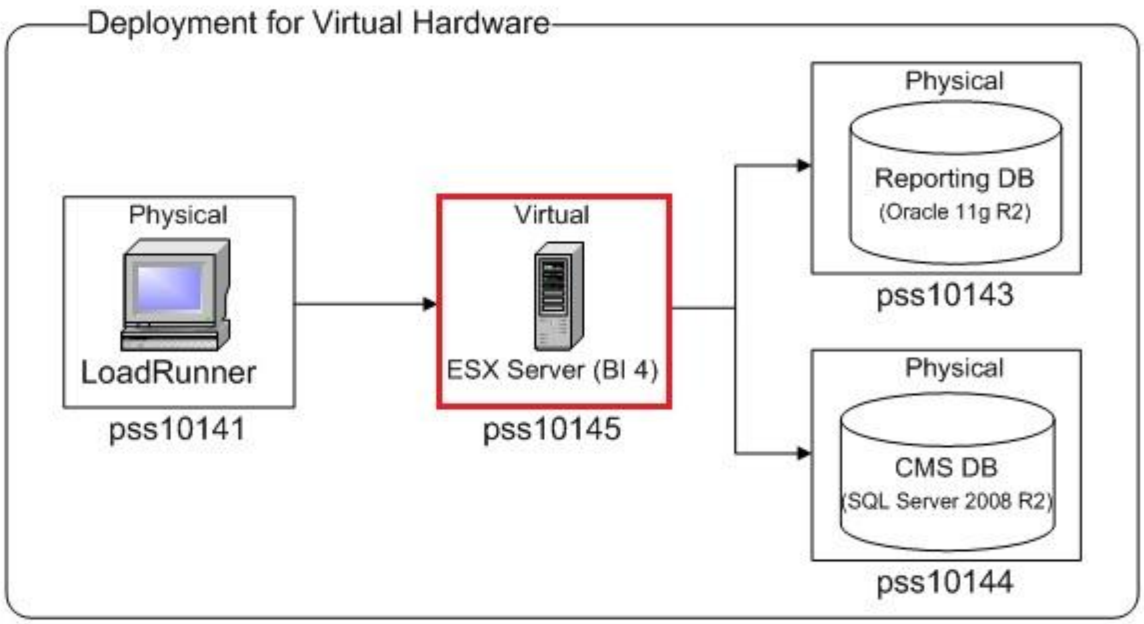


The virtual machine was configured as follows:

- Windows Server 2008 R2, with VMware VMtools installed
- 12 virtual vCPUs (2 sockets, 6 cores per socket)
- 32 GB of memory
- VMXNET3 and PVSCSI paravirtualized drivers

CPU processing and memory reservations were used, but since this virtual machine was the only VM on the host, there was no resource contention from other virtual machines on the host. Note that since the VM has the same number of vCPUs as pCPUs in the host, the hypervisor is forced to share processing with its guest. This means on a busy system, any significant resource usage by the hypervisor should theoretically impact the performance of the guest.

No virtualization specific tuning was performed at either the hypervisor or virtual machine levels other than installing VMware VMtools in the guest.

### *2.2.2 Supporting Machines in Testing Landscape*

All other systems in the testing landscape were physical machines with identical hardware specifications as the physical SAP BusinessObjects BI 4 system (12 core Xeon-based Supermicro TwinBlade servers). However to follow implementation best practices, Hyper-Threading was enabled on the supporting machines.

The servers listed below were used for both virtual and physical instance tests, ensuring the only variable in our tests was virtualization.

| System Role | System Name | OS | Software |
|---|---|---|---|
| **CMS Database** | pss10144 | Windows 2008 R2 | Microsoft SQL Server 2008 R2 |

| | | | |
|---|---|---|---|
| **Reporting Database** | pss10143 | Windows 2008 R2 | Oracle 11g R2 |
| **Load Generation** | pss10141 | Windows 2008 R2 | HP Loadrunner 11 |

## 2.3 Test Execution

SAP's technical teams developed a test BI architecture for both physical and virtual environments and then drove a simulated user load through each system to observe their behavior under different scenarios. It should be noted at the outset that the focus was not to test the effects of various virtualization optimization techniques of the hypervisor, but instead to determine how virtualization best practices could affect SAP BusinessObjects BI 4 performance**.** For all tests, we used out-of-box configurations of SAP BusinessObjects BI 4.0 SP 5 and VMware ESXi 5.1.

The test suite modeled common Crystal Reports 2011 (CR 2011) workflows which were generated with HP Loadrunner. Load testing started with a low number of users and we continued to add 10 users every 5 minutes until the system started experiencing problems. Measurements were made at the operating system, SAP BusinessObjects BI 4 application, and VMware ESXi hypervisor layers. To ensure consistency in the data, the entire experiment and all its tests were executed a minimum of three times. Statistically, the standard deviation across all tests was low which proves performance was consistent. All throughput and response time observations were taken at the $90^{th}$ percentile to normalize the results. This represents the worst case scenario users would experience in a similar real-world situation.

We continuously recorded the system (CPU, memory, disk, network), BI server, and VMware metrics during all tests. Having these metrics allows us to better understand how the resources are being consumed across the landscapes.

While it was not feasible to test each application within the SAP BusinessObjects BI 4 suite, the testing workflows using Crystal Reports 2011 were chosen to simulate extremely heavy CPU, RAM, network, and storage I/O usage. It is important to note that while these tests were using CR 2011 only, the majority of the BI Platform (including CMS, processing engines, and database) were exercised and the workflows chosen gives us confidence that results would be similar if other BI tools such Web Intelligence (using relational sources) or Dashboards had been tested instead.

## 2.4 Analysis of Findings

A number of conclusions can be derived from our tests, but it is important to remember that these tests were done on an isolated system under a specific load. Real-world environments are more complex and subject to external factors beyond the administrator's control. Therefore it is important to interpret these conclusions within the larger context of a multi-facetted approach to system architecture design and optimization.

### *2.4.1 Resource Contention Degrades Performance*

Many IT departments mandate host overprovisioning (where the total number of vCPUs provisioned to all VMs is greater than the number of pCPUs on the host) to increase hardware utilization. Some go as far as to consider a server underutilized if it has an average load factor of less than 80%. Previous testing by SAP has proven that even physical SAP BusinessObjects BI 4 systems operating at this load level can peak to 100% quickly when a large report is run or a large number of active concurrent users (ACUs) make requests to the system simultaneously.

It should be obvious that a physical machine with dedicated access to resources will outperform a similarly equipped virtual machine that must share with other guests, but it is impossible to quantify the consequences of resource contention because it depends on the aggregate demands of each VM at any given time.

Our initial tests on a heavily loaded virtualization host showed significant degradation that was roughly correlated to the resource scarcity in the system. *The results were conclusive enough that this hypothesis was proven and therefore removed from our experiment.*

The remainder of tests completely removed the possibility of resource contention or scarcity through CPU and RAM reservations and VM isolation to ensure the disk and I/O subsystems had no other guests to service.

> ***All sizing guidance from SAP assumes the BI system has full, dedicated, and explicit access to all CPU, RAM, disk, and network resources. If this is not the case, any sizing exercise will be rendered invalid and you should not expect acceptable performance under all conditions.***

#### *2.4.2 Improper System Configuration Can Significantly Impact Performance*

The initial configuration of a server is very important as any sub-optimal settings at the lower levels of the architecture will have a rippling effect across all layers above. While building the COIL test environment, we ran benchmarks at each step to understand the effects of each configuration change.

Validating configuration at each step was important to ensure we did repeat the common failings of previous tests that did not guarantee an apples-to-apples comparison or declared conclusions based on flawed data (as described on page 6 above).

After installing the operating system on both the physical and virtual systems, our OS-level synthetic benchmarks showed a significant (>10%) performance difference on both CPU and storage I/O metrics – surprisingly, with the virtual machine performing better. We had assumed each physical system was configured identically, but as they were repurposed from a previous COIL project, the BIOS settings were configured differently. The performance delta was eliminated once we examined the BIOS settings for each server and set them to correct, and in most cases, default values.

We also found similar problems at the operating system level when drivers or updates were applied to one system and not another. These very small details can be easy to miss so require vigilance at all levels from the BIOS up through the software stack and all the way through to the SAP BusinessObjects BI 4 system. Please refer to section 3.2 for more information on these details.

Clearly, BI administrators need to care about their underlying hardware and software configuration and cannot be solely interested in application level configuration.

#### *2.4.3 Synthetic Benchmarks Are Inappropriate For Virtualization*

Our test environment created an “apples-to-apples” comparison as much as possible – with virtualization being the only material difference, our goal was to measure what the raw performance difference of completely fresh operating system installs could be. For our hypothesis to be true, the performance difference at the BI application level would need to be less than this figure.

Synthetic benchmarks such as PCMark and HD Tune have long been a staple of testing in the physical world, but it is very important to remember that synthetic benchmarks (i.e. raw performance tests) are not indicative of real-world operation and grossly misrepresent the actual overhead a system can experience.

Often the workloads on a given host are not always CPU-bound – that is, guests may be waiting for user input, data transfer, or simply not using all available CPU cycles. In these cases, the overhead of operating the hypervisor may be absorbed by these spare CPU cycles and the system can provide throughput closer to native performance, although perhaps with slightly increased latency.

Therefore it stands to reason that the type of enterprise application being virtualized affects the possible performance delta between physical and virtual deployments. A heavily utilized enterprise application with CPU-bound workflows may resemble a synthetic benchmark and give a misleading representation while applications with higher I/O requirements may spend more time waiting for data and therefore show smaller deltas.

Our tests in the COIL bear this out as the synthetic benchmark deltas are quite different than that of real-world tests. To illustrate this, consider the following data:

- Using SAP's 2-tier OLTP Benchmark Certification tests (FT) *[Ref. 0]* where a standardized ERP suite is used, a virtualized system achieved within 6% of the performance of the identical test on physical hardware. This demonstrates that mission critical systems can be virtualized with little impact on the resources required if they are deployed and tuned correctly. This is shown in SAP SAPS Certification # 2011027 and 2011028.

- In our own COIL-based tests, synthetic benchmarks showed roughly a 10% difference between identical physical and virtual hosts across both CPU and storage I/O performance.

The SAP benchmark is an application level test that includes the network and disk I/O associated with the SAP Business Suite, so it is not surprising the synthetic benchmark has a higher delta. The Crystal Reports 2011 tests performed in our experiments were also application level tests representing realistic BI workflows on the system. It should be noted that the nature of business intelligence is inherently more network and storage I/O intensive than transactional systems such as the SAP Business Suite. To some extent, this also calls into question the applicability of the SAPS measurement for BI systems.

We conclude that synthetic benchmarking numbers are invalid and the true performance delta between physical and virtual systems must be conducted through application level benchmarks and heavily depend on which application and which workflow is tested. Our test environment attempted to eliminate any factors other than virtualization, so we could focus on a true "apples-to-apples" comparison. As described above, the performance delta at the BI application layer would need to be no greater than the 10% difference observed with a synthetic benchmark to properly prove our hypothesis.

#### *2.4.4 Physical and Virtual Throughput Is Comparable*

All simulated user tests were performed at various load levels and consistently showed a statistically insignificant throughput performance delta between physical and virtual SAP BusinessObjects BI 4 systems *when CPU utilization was less than 80%*.

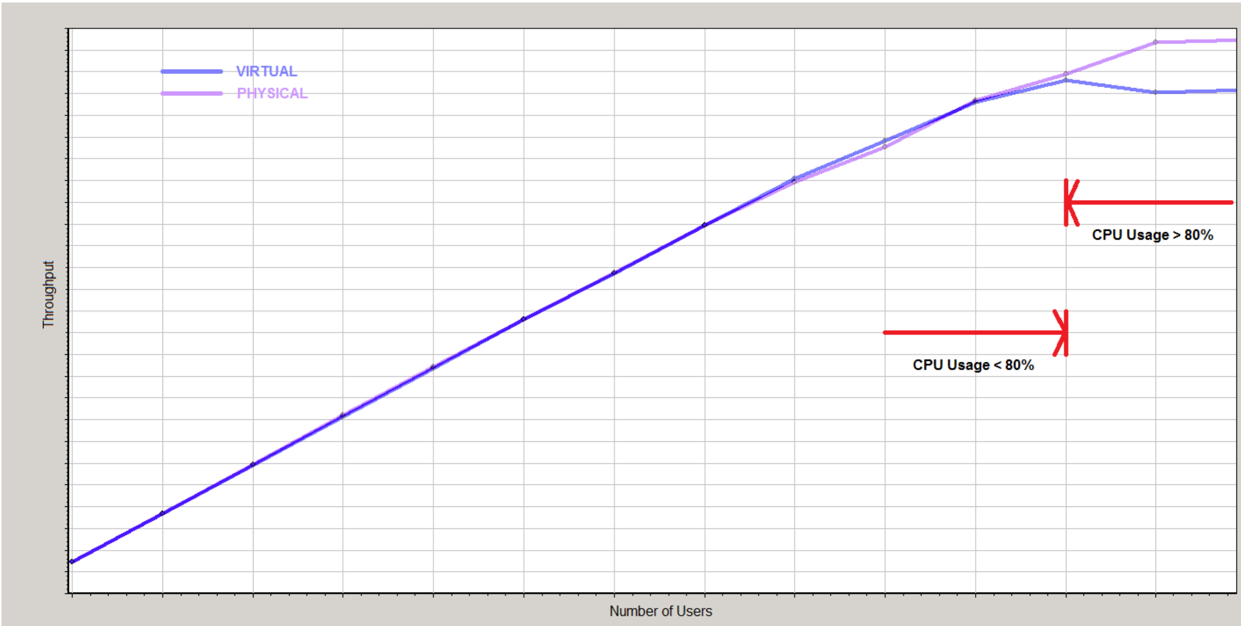


***Figure 1: Throughput: Physical vs. Virtual***

During the tests, we made the following observations and conclusions when the system was *under normal operating conditions* (CPU usage <80%) and SAP's recommendations were followed:

1. System throughput of the physical and virtual systems across increasing load levels was very close and indistinguishable from each other. When configured correctly and under normal conditions, a virtual SAP BusinessObjects BI 4 instance should perform as well as a physical system.

2. The throughput growth curve of the physical and virtual systems was similarly linear as more load was applied to each system. We conclude a properly configured virtual SAP BusinessObjects BI 4 should perform just as well as a physical system across a variety of load levels.

3. Resource usage of the physical and virtual systems was statistically similar across each test and at all load levels. We conclude that virtual SAP BusinessObjects BI 4 systems do not consume more resources than physical SAP BusinessObjects BI 4 systems.

After analyzing the test data from the hypervisor, OS, and BI layers, we concluded that while SAP BusinessObjects BI 4 is very processor intensive, the CPU likely still spends considerable time waiting for I/O. Therefore our virtualized BI system did not experience observable hypervisor overhead because the processing requirements of the host were handled in these spare CPU cycles. In the physical system, these CPU cycles were simply wasted as there was nothing else running at the time.

We conclude under normal load conditions where the virtualization host is properly configured and system resources are guaranteed, a virtualized SAP BusinessObjects BI 4 system can exhibit the similar performance to an identical physical deployment.

***It is important to reiterate here that our tests guaranteed the availability of all system resources for the virtual machine to ensure a comparable environment to our physical reference system. We already know that the story would be very different on a heavily loaded virtualization host where resources are overcommitted.***

### *2.4.5 Physical and Virtual Latency Is Comparable*

All simulated user tests were performed at various load levels and consistently showed a statistically insignificant latency performance delta between physical and virtual SAP BusinessObjects BI 4 systems *when CPU utilization was less than 80%*.

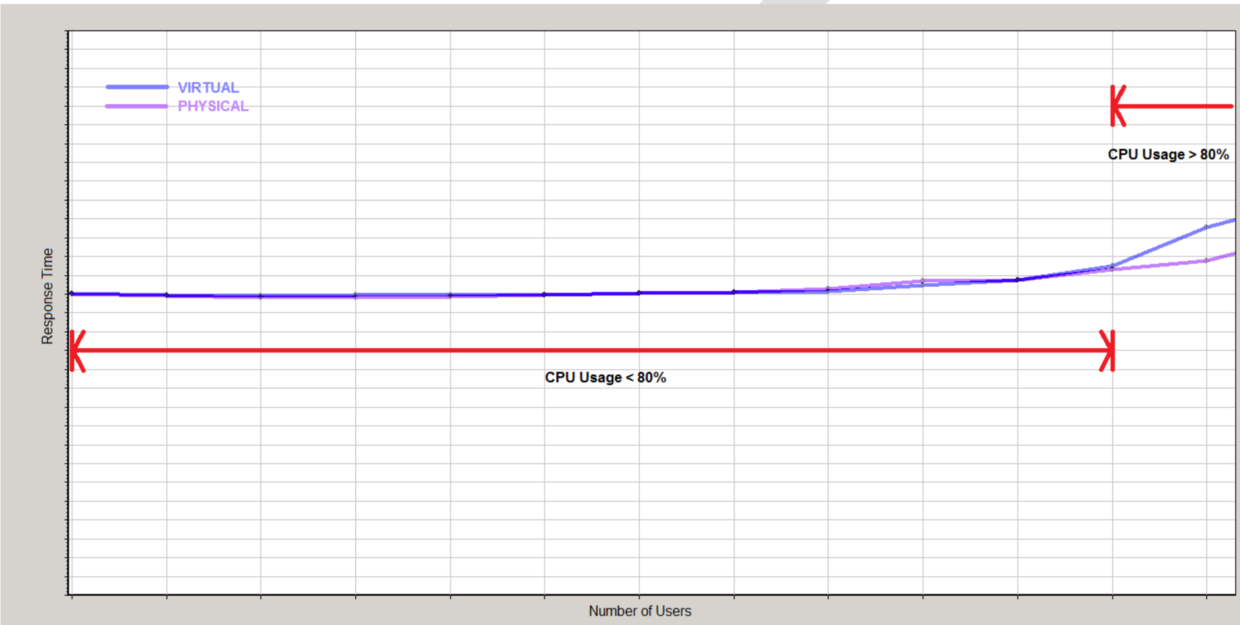


***Figure 2: Response Time: Physical vs. Virtual***

During the tests, we made the following observations and conclusions when the system was *under normal operating conditions* (CPU usage <80%) and SAP's recommendations were followed:

1. The response times at the 90th percentile between physical and virtual systems was within the timing margin of error and therefore exhibits equivalent latency performance profiles. When configured correctly and under normal conditions, a virtual SAP BusinessObjects BI 4 instance should perform as well as a physical system.

2. The latency growth curves of the physical and virtual systems were similar as more users were applied to the systems. We conclude a properly configured virtual SAP BusinessObjects BI 4 should perform just as well as a physical system across a variety of loads.

3. Resource usage of the physical and virtual systems was statistically similar across each test and at all load levels. We conclude that virtual SAP BusinessObjects BI 4 systems do not consume more resources than physical SAP BusinessObjects BI 4 systems.

We again analyzed the test data from the hypervisor, OS, and BI layers, and see that when CPU utilization <80%, virtualization does not impact response time because any latency introduced by virtualization is likely less than the other latencies in the system, such as waiting for I/O.

We conclude that under normal load conditions where the virtualization host is properly configured and system resources are guaranteed, a virtualized SAP BusinessObjects BI 4 system can exhibit the same latency performance as an identical physical deployment.

#### *2.4.6 Physical and Virtual Degradation is Comparable*

A typical IT principle is to consider scaling out a system once it reaches a specific load threshold (i.e. 80% CPU). While SAP BusinessObjects BI 4 is designed to exhibit a relatively linear performance curve under normal circumstances (i.e. CPU <80%), previous SAP tests in the CO-Innovation Lab (COIL) have proven that BI systems should be candidates for scale out when they reach an average CPU utilization of 65% to ensure enough headroom for large or intensive workloads without materially impacting users *[Ref. 10]*. *In fact, best practice recommendations for the majority of SAP software state that systems should be scaled out around the 65% mark.*

When system resources start to become scarcer, the system will start to suffocate, and in some cases thrash. The resulting drop in throughput and increase in response time can be dramatic and directly impacts user sessions.

As we loaded more simulated users on the physical system to push CPU utilization past 80%, we observed significant system performance deterioration. Since our virtualized system performed just as well as the physical system when utilization was under 80%, we were curious if its degradation curve would be similar as well.

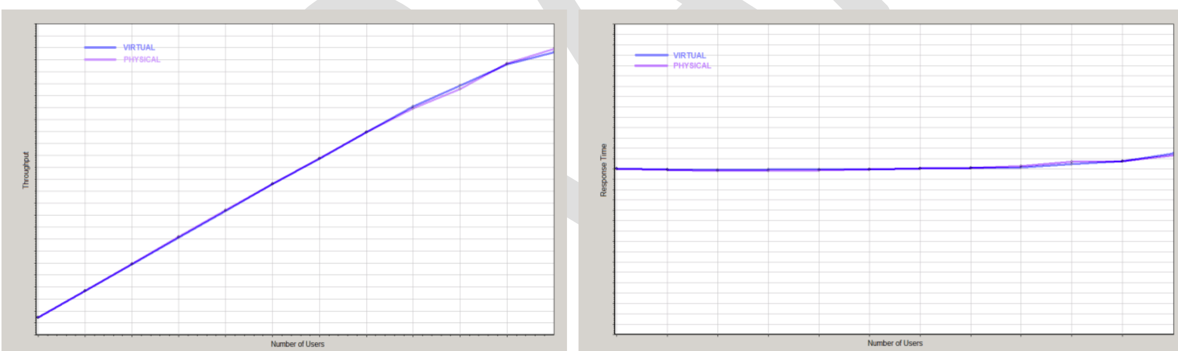


***Figure 3: Physical vs. Virtual Performance (CPU <80%)***

During the tests we made the following observations and conclusions:

1. Performance of both physical and virtual systems degrade when utilization is >80%. Note that as discussed in the previous section, the performance and growth curves of both systems are were consistent up to this level. This suggests scale-out guidance should be the same for physical and virtual systems.

2. SAP's recommendation to consider scale out of a BI node at 65% holds. Throughput measurements and response times increasingly degrade for both physical and virtual instances as CPU load continues to increase above 80%. This behavior is consistent with previous SAP tests using Web Intelligence which provides some reassurance that our Crystal Reports 2011 tests are somewhat representative.

3. Our tests started to see performance issues after CPU utilization went past 80%, and data from our load testing machines suggest that system performance becomes erratic and metrics calculated after this point have little value. It is interesting to note that the virtual system experienced performance degradation slightly before the physical system, but the degradation slope of both machines was the same.

These results suggest that when the system is heavily loaded the hypervisor's overhead starts to show and begins to impact observable performance. We therefore conclude that it is even more important to implement a proper scale-out strategy when deploying virtualized systems.

#### *2.4.7 BI-Specific Tuning Not Required In Virtual Environments*

The statistically insignificant delta between physical and virtual test systems suggests that additional tuning of the BI stack is not required to perform with near-physical performance within a VMware environment. The consistent performance and degradation profiles suggest that tuning is also not required to maintain performance across varying load levels.

This finding is consistent with that of a previous study which proved Java-specific tuning within a virtual SAP BusinessObjects BI 4 system is also not required, and can even be detrimental in some cases *[Ref. 9].* Specific guidance is discussed in section 3.1.2.1 below.

# 3 VIRTUALIZATION RECOMMENDATIONS

Modern hypervisors such as VMware ESXi 5 do an excellent job of enabling server consolidation without necessarily incurring a significant performance penalty, but the responsibility for properly configuring and optimizing the system rests solely with you. Realize that no amount of virtualization optimization can fix an inefficient database query or a poorly written report. You should focus on creating a good system architecture first, taking into account stack components, I/O bottlenecks, and processing tier distribution before you attempt any optimization at the virtualization level.

It also should be noted that the primary focus of the recommendations in this guide are not solely aimed at performance tuning or optimization. The material in this document is designed to help you create a scalable deployment that provides stable, consistent performance across highly variable and potentially severe conditions. There is no requirement to follow these recommendations, but they are a result of customer feedback, controlled lab testing, and SAP support experiences. Whatever negative consequences of implementing a "best practice" you experience will likely be infinitesimal compared to the consequences of not implementing it.

Your deployment scenarios may vary and how much you deviate from these suggestions will depend on your organization's comfort, education, and experience in deploying and operating mission critical business intelligence enterprise applications.

## 3.1 Start With The Right Tools

#### *3.1.1 Identify and Disregard Invalid Guidance*

The first challenge for many customers is to find the correct information for virtualizing the SAP BusinessObjects BI 4 suite. It is common to find conflicting information from different sources and even from SAP – therefore it is critical to ensure any guidance or best practice documents you follow are actually applicable and valid for your BI deployment. While this appears obvious, the reality is there is enough content available from outdated or untrusted sources that evaluating a document's validity is not always easy. When in doubt, always refer to www.sap.com/bivirtualization for the most current guidance.

1. **Ensure Guidance is for SAP BusinessObjects BI 4 (and not the SAP Business Suite)**

There is a significant amount of documented wisdom, studies, and guidance for virtualizing “SAP software” – however, the vast majority of these are referring to the “SAP Business Suite” or more accurately, applications based on the SAP NetWeaver technology platform. Such guidance is typically not applicable to BI systems.

2. **Ensure Guidance Is Validated on SAP BusinessObjects BI 4 or Later**

   Prior to this guide (based on SAP BusinessObjects BI 4 and later), a number of parties both at SAP and external to the company have issued guidance based on older software, improperly configured environments, or flawed analysis. This resulted in recommendations that can have severely negative impacts on system performance. Sizing guidance based on earlier versions (i.e. BOE 3.x or XI R2, etc.) is not valid for SAP BusinessObjects BI 4 and should be disregarded.

   Note that virtualization-specific recommendations within this document are equally applicable to SAP BusinessObjects BI 4 and BOE 3.x environment as long as you are using the most recent versions of the underlying operating system, software stack, and hypervisor.

3. **Ensure Virtualization Guidance is for VMware ESXi 5 or Later**

   VMware ESXi has undergone significant improvements over its product evolution and has even better performance, scalability, and configurability than previous versions. Therefore tests conducted using VMware ESX 3.5 or earlier should be completely discarded and those done with vSphere 4.x should be revisited with the latest version of ESXi available.

### *3.1.2 Use the Latest Best Practice Guidance from VMware*

The hypervisor orchestrates the operation of every virtual machine on the host so it is important to ensure it is configured properly. Choices made at the host and hypervisor levels have significant consequences that are impossible to correct elsewhere. On a lightly loaded host this may not be noticeable, but as the system experiences higher loads, performance can degrade exponentially with service affecting consequences.

VMware’s “Performance Best Practices for VMware vSphere™ 5.1” *[Ref.7]* is required reading for virtualization administrators, but it is also important to BI administrators to better understand the environment underlying their systems. Typical BI administrators do not administer the hypervisor environment directly, but should still be opinionated and ensure there is a clear understanding of what requirements need to be fulfilled for a proper virtualized BI deployment.

> ***All guidance in this whitepaper should be considered supplemental to VMware’s “*Performance Best Practices for VMware vSphere™ *5.1” [Ref. 7].***

#### *3.1.2.1 Special Note on Java Large Pages Optimizations*

Memory address translations use a Translation-Lookaside Buffer (TLB) which is a page translation cache that holds the most-recently used virtual-to-physical address translations. TLB is a scarce resource and a TLB miss can be costly as the CPU must read from the hierarchical page table which requires multiple memory accesses. Java-based applications can take advantage of a feature called “large pages” which allows them to establish large-page memory regions. By using a bigger page size, a single TLB entry can represent a larger memory range. Therefore, when using large pages, memory intensive applications could have better performance.

However, the use of large pages can also negatively impact system performance. If the working set of an application is scattered over a wide range of address spaces, the application is likely to experience thrashing of a relatively small number of TLB entries across a large number of pages and result in worse overall performance due to higher TLB miss rates. Additionally, applications that pin large amounts of memory may create a shortage of contiguous memory for other processes and negatively impact system performance due to memory fragmentation and eventual swapping.

Java large page support exists on physical machines as well because this feature is related to the Java Virtual Machine (JVM) and not machine virtualization. For physical deployments, SAP does not provide specific guidance on this feature and the majority of customers do not change it from the default setting of disabled. Customers are free to enable it as does not affect the functional operation of the suite, but depending on the workload, this feature can slightly increase or decrease performance.

To test the effect of enabling large page support in a virtual environment, SAP and VMware jointly performed tests in SAP's CO-Innovation Lab and produced a whitepaper entitled, "Evaluating Selected Java Best Practices For Sap BusinessObjects BI 4 On VSphere" [Ref. 9]

The result of these tests showed that no virtualization specific tuning of the Java stack is required for the SAP BusinessObjects BI 4 suite and performance was even slightly decreased during the tests by enabling this feature. Since the benefit of enabling large pages cannot be proven definitively, you should not enable Java large page support in virtual machines with the goal of optimizing performance.

## 3.2 Properly Configure All Hardware and Software Components

It should go without saying that evaluating the environment external to your BI system (networks, storage, DB servers, etc.) and optimizing them is necessary regardless of whether your landscape is physical or virtual. As proven in our COIL tests, ensuring proper configuration of each stack element is also important. Every deployment is different so it is not possible to have a definitive checklist, but the following are considerations that were researched and evaluated during our lab tests:

### *3.2.1 Intel Hyper-Threading and Turbo Boost Technologies*

Hyper-Threading technology allows a single physical processor core to behave like two logical processors. Unlike having two physical processors, Hyper-Threading this allows two independent threads to run simultaneously under certain circumstances. This boosts performance by keeping the processor's pipeline as full as possible. The actual performance gain can be slight to significant depending on how much parallelism is achieved. It is important to remember that these additional logical processors enable operating systems to parallelize threads and should be considered "real" processors or have the benefits of them.

Turbo Boost is the term for an Intel technology implemented in most current generation processors that allows the CPU to temporarily run above its base operating frequency to boost performance when requested by the operating system.

Both of these technologies increase CPU performance at the cost of energy and heat dissipation. Some server administrators have the philosophy that disabling Hyper-Threading in the BIOS will reduce the heat generated under normal operation and allow the CPU to operate at a higher clock frequency (Turbo Boost) for longer periods of time.

In previous generation CPUs, it was common practice to disable hyper-threading for production servers to enable the performance benefits of Intel's Turbo Boost technology. This feature enabled faster CPU clock rates while staying within the chip's thermal envelope, thereby boosting overall performance for more single threaded applications.

For both physical and virtual deployments of SAP BI, both Hyper-Threading and Turbo Boost should be enabled. The SAP BusinessObjects BI 4 suite has been architected to fully take advantage of Hyper-Threading and therefore will perform best when both features are enabled so the CPU can intelligently decide how to use these features. Disabling either feature can degrade performance.

Today's hypervisors will load physical processors before scheduling the logical ones backed by Hyper-Threading. Within the hypervisor you will see double the number of logical processors as you have physical, but you must remember that these extra logical processors will not delivery the full processing power of a physical one.

#### *3.2.1.1 Hyper-Threading Example:*

To illustrate how Hyper-Threading allows a single core to be viewed as two logical processors, you can use the "msinfo32" application on Windows 2008 R2 and look at the "System Summary" for details. On our Supermicro servers, this is what is displayed:

| | | |
|---|---|---|
| Processor | Intel(R) Xeon(R) CPU | X5675 @ 3.07GHz, 3066 Mhz, 6 Core(s), 12 Logical Processor(s) |
| Processor | Intel(R) Xeon(R) CPU | X5675 @ 3.07GHz, 3066 Mhz, 6 Core(s), 12 Logical Processor(s) |

*Figure 1: System Summary from msinfo32*

In this case, the operating system will "see" 24 logical processors even though only 12 physical cores are present. Each core is responsible for scheduling processes onto its "real" and "logical" threads.

You can also see how this appears at the operating system level by opening the Windows Task Manager and looking at the "Performance" tab:

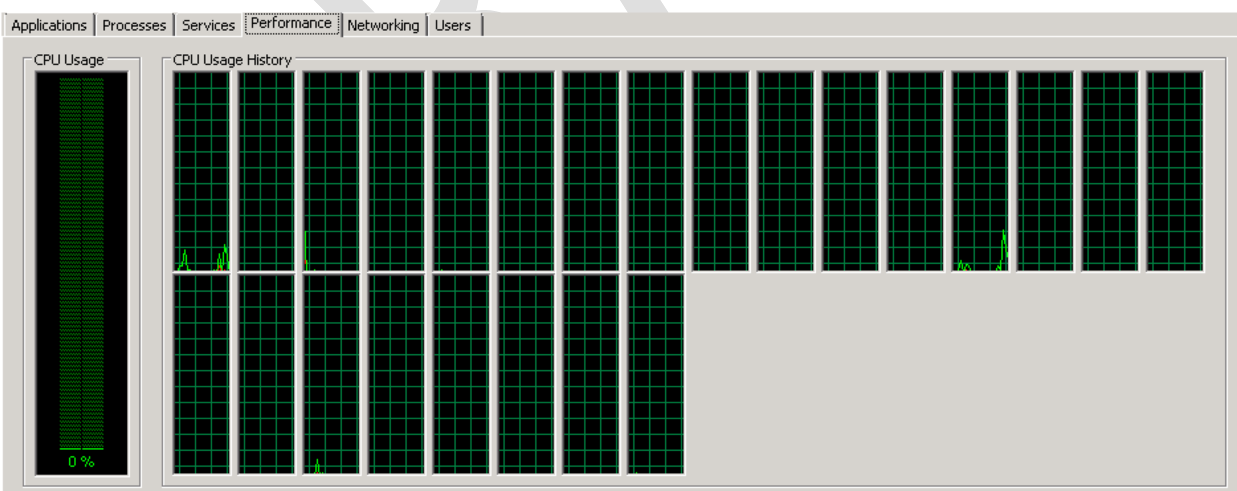


*Figure 2: 24 Logical Processors on a 12-Core Hyper-Threaded Server*

It is extremely important to remember that each of the additional 12 logical processors will not provide the performance of a physical core.

### ***3.2.2 Physical Hardware (BIOS level)***

| BIOS Feature | Preferred Setting | Description |
|---|---|---|
| **Power Management Strategy** | • Balanced<br>OR<br>• Performance | • Refer to your hardware manufacturer's manual on the definition of each and compare it to the power management policies in VMware's best practices document.<br>• Note power management strategy definitions and defaults have changed between version 4.x and 5.x of VMware's hypervisor. |
| **Intel Turbo Boost** | • Enabled | • This feature dynamically clocks the CPU faster when possible while staying within the CPU's thermal envelope.<br>• It is a misconception that this should be disabled in favor of enabling Intel Hyper-Threading. When both features are enabled, the system will dynamically choose either or both options based on which will provide the best performance while staying within the CPU's thermal envelope. |
| **Intel Hyper-Threading** | • Enabled | • This feature allows a single CPU core to behave like two logical processors, thereby enabling the core to execute two independent threads at once.<br>• It should not be disabled in an attempt to reduce the CPU's thermal envelope to allow faster CPU clocking with Intel Turbo Boost. |
| **Other BIOS Settings (AHCI, ACPI)** | • Specifically set or use BIOS defaults as appropriate | • Ensure other BIOS settings are set to intelligent and reasonable values. For example, local storage controllers should use AHCI, power management be configured to use ACPI, etc.<br>• Features should not be manually changed from BIOS defaults without a sufficient reason. |

### *3.2.3 Hypervisor*

| Best Practice | Description |
|---|---|
| **Use hypervisor-level monitoring and telemetry tools** | • VMware has a number of utilities, including “esxtop” that provides hypervisor-level metrics that show not only resource utilization, but wait times and other critical information.<br>• Third party tools (free and paid) are available as well, but should be used in conjunction with VMware's own tools such as “esxtop”. |

| | |
|---|---|
| **Use advanced networking features of the hypervisor** | • Implementing NIC Teaming, Distributed vSwitches, etc. can significantly boost performance while reducing load on the host.<br>• See VMware's best practice document [Ref. 7] for more information. |
| **Use paravirtualized devices in all VMs** | • Using paravirtualized network cards such as VMXNET3 reduce host CPU overhead while avoiding artificial bottlenecks and overhead imposed by the need to emulate physical network cards<br>• Using paravirtualized storage controllers reduces the host's processing overhead required for disk read and writes. |
| **Use VM BIOS defaults** | • Changing the virtual machine's BIOS settings should be unnecessary as choosing the right machine type and OS during VM creation will set these values correctly. |
| **For best performance, do not use snapshots** | • Snapshots create additional overhead for the file system and this overhead continues to exist even when all snapshots for a virtual machine have been deleted. |

### *3.2.4 Guest Operating System*

| Best Practice | Description |
|---|---|
| **Use VMware's VMtools in every VM on each host** | • Installing guest tools in each VM on the host enables the hypervisor to get a full picture of resource usage and enables it to make better resource allocation decisions.<br>• Do not disable any VMware features even if they appear to be impacting resource usage (i.e. the VMware Balloon Driver allocating RAM) as these are mechanisms the hypervisor uses to manage resource usage. Instead, make the changes at the hypervisor level (i.e. VMs with memory reservations do not need to use the balloon driver as they do not need to return memory back to the host). |
| **Use the latest version of your chosen OS** | • More recent versions of most operating systems have been architected to work better in a virtual machine and may handle memory allocations and resource usage better when in virtual environments. |
| **Be careful with virtualization-specific optimizations** | • SAP recommends following only VMware's specific guidance in its best practices document [Ref. 7].<br>• Disregard non-SAP BusinessObjects BI specific virtualization guidance. |

### *3.2.5 Application-Level Components (Database, Application Server, BI Platform)*

| Best Practice | Description |
|---|---|

| | |
|---|---|
| **Focus on creating a good architecture instead of optimizing a poor one** | • Benefits from tuning individual parameters at any application layer will be inconsequential compared to the negative performance impact of a poor architecture or deployment. |
| **Follow vendor best practices for virtualization** | • For each stack component (database, application server, etc.), follow the vendor's own recommendations for virtualization.<br><br>• Do not follow "general" virtualization optimization techniques unless they are derived from your vendor's own guidance or VMware's original best practices guide. |
| **Do not implement Java-specific virtualization optimizations unless specifically for SAP BusinessObjects BI** | • General tuning advice for Java applications within VMware is not applicable for the SAP BusinessObjects BI 4 suite.<br><br>• A recent SAP and VMware whitepaper shows evidence that virtualization-specific tuning of Java is unnecessary and in some conditions degrades performance [Ref. 9] |

### 3.3 Use Resource Reservations on all SAP BI Virtual Machines

Hypervisors are very intelligent and generally do a good job of executing each virtual machine when resources are freely available, but the very nature of virtualization guarantees resource contention when the system becomes loaded.

When multiple virtual machines ask for the same resource and there isn't enough to go around, requests will start to be queued and the host will service each request in the queue as resources become available. In the case of CPU power, the requests from each VM are interleaved with each other so that all virtual machines can run, but each with reduced performance. However memory cannot be shared through interleaving, so when there is contention, there is always a winner and a loser. The consequence for the loser can be quite severe – it may be forced to page out to disk itself or wait for the hypervisor to do the same.

Therefore it is very important that when the system experiences resource contention, the hypervisor knows the resource requirements and priority of each virtual machine so that it can make intelligent scheduling and resourcing decisions. The hypervisor will first ensure guests with resource reservations are properly serviced and will then execute the remaining VMs with whatever resources are left.

In VMware ESXi, a ***resource reservation*** specifies the guaranteed minimum resource allocation for a virtual machine. This means no other virtual machine can use those resources ***if*** the reserving guest requires them. Many IT departments are under the misconception that reserved resources are blocked off and are wasted if not used. This is incorrect - if at any time the resources are not immediately being used, they may be used by other virtual machines on the host as long as these resources are immediately relinquished on demand by the reserving guest. However virtual machines that must free their resources under these conditions have to suffer the consequences such as reduced performance and memory-to-disk page swapping.

Some important things to remember regarding reservations:

- A ***resource reservation*** represents a fixed amount of resource and does not change when you alter the environment – such as adding or removing other virtual machines on the host. This is different than a ***resource share*** which specifies a virtual machine's resource rights as a fraction of total capacity.

- Specifying a resource reservation that is equal to the configured amount for the virtual machine effectively ensures the hypervisor will try to give the VM all the resources it has been provisioned for.

- Remember that resource reservations do not prevent other virtual machines from using those resources if the reserving guest is not using them.

If your BI system has any risk of sharing the resources it needs, or worse, having to release them due to other guests owning reservations, the BI system cannot be expected to perform well. BI administrators need to ensure there is enough visibility into the virtualization host's operation to confirm that resource reservations are strictly implemented as described in this whitepaper.

> ***Resource sharing can greatly impact performance and therefore, all sizing estimations and architectural decisions should be considered invalid unless the amount of processing power and memory is fully reserved. In general, customers with systems that do not implement reservations should not expect performance similar to physical systems, nor should they expect technical support from SAP for performance issues until such reservations are in place.***

For more detail on resource management features in ESXi, please refer to VMware's "vSphere Resource Management Guide for ESXi 5.1:" guide [Ref. 8]

### *3.3.1 Use Strict CPU Processing Reservations*

Setting a CPU reservation for each of your SAP BusinessObjects BI 4 virtual machines is crucial to ensure proper performance of the entire system. Without such a guarantee of processing power, the system can lose CPU cycles to competing virtual machines on the same host and can make performance highly variable. Under these conditions it is unreasonable to expect the system to perform as well as a physical system, or even as consistent as one.

More importantly, when resource availability is not guaranteed, all sizing estimations and the resulting architecture decisions you have made to this point should be considered invalid. All sizing guidance from SAP assumes that each of your BI nodes have the full processing power of the CPU cores assigned to them.

### *3.3.2 Use Strict Memory Reservations*

Setting a memory reservation for each of your SAP BusinessObjects BI 4 virtual machines equal to the total amount of memory configured allows your system to avoid unnecessary memory paging when it becomes heavily loaded. Note that business intelligence systems are memory intensive and can easily allocate large amounts of RAM without notice. When this RAM is not available, the system will need to page memory to disk at the guest or even at the host level. Performance can be further degraded in the case where the page file used is on a SAN or other network-based storage.

### *3.3.3 Addressing Concerns Using Reservations*

Resource reservations are typically unpopular because a common misconception is that the requirement for the resource to be available at start prevents any other VM from using the resource. The truth is once the reserving virtual machine starts up and it is clear that these resources are not immediately being used, the hypervisor is free to allocate some of those resources to other virtual machines until the reserving virtual machine requires them again. However the total amount of memory reserved across all virtual machines on a host cannot exceed the physical memory of the host as there is always possibility that all guests will request their allocations at the same time.

> ***SAP does not explicitly require resource reservations as a condition of support due to the flexibility needed by many customers, however these customers also take on the risk of poorly sized and poorly performing systems requiring remediation before receiving support for performance issues.***

Resource reservations do have a cost in the sense they can reduce VM density on a host as well as restrict the movement of virtual machines to other hosts using live migration techniques such as VMware VMotion. However it is important to emphasize to your infrastructure team that physical systems implicitly guarantee CPU and memory resources because they are not sharable and there should be a similar expectation for virtual systems as well.

### 3.4 Do Not Use Resource Shares, Limits, or Affinity

Virtualization introduces a number of other features to manage resource usage by guests so it is important to understand the terminology behind these terms and why these features should not be used. All features that restrict resource availability or hinder the hypervisor in making positive decisions in favor of your SAP BusinessObjects BI guest virtual machines should be avoided.

- ***Resource shares*** specify the proportion of resource each virtual machine should receive relative to other VMs. This is expressed as a number of "shares" as a fraction of the total number of shares distributed. For example, a 4-core host with 4 guests, each with 4 vCPUs and 1,000 shares may get as little as 25% (1,000/4,000) of processor cycles if all VMs become CPU-intensive at once. If a fifth VM with 1,000 shares is added, each VM's share is now 20% even though the administrator did not explicitly reduce the resource allocation.

- ***Resource limits*** specify the total amount of resource a virtual machine can consume, regardless of the amount configured. Limits are expressed as concrete resources. For example, a VM configured with 8 GB of RAM and a memory limit of 6 GB means that while the guest will "see" 8 GB of RAM, but only 6 GB of RAM can actually be used. Once this 6 GB is exhausted, the system will not be able to allocate more without swapping memory to disk. This feature is meant to operate software that checks for large amounts of memory but does not use it. This should never be used for SAP BI systems.

- ***CPU affinity*** is a feature that can "pin" or fix a virtual machine's vCPUs to specific pCPUs on the host. Note this feature is for affinity and not exclusivity – therefore it does not prevent other virtual machines from using these "pinned" cores if it would improve their performance. However the hypervisor cannot move the processing for a pinned virtual machine to other cores, even if those cores are less busy and could improve its performance. This feature should never be used for SAP BI systems.

### 3.5 Right Size Virtual Machines

Virtual deployments are very flexible because you can size each VM according to its needs rather than being limited to the physical configuration of a server. This also gives you the ability to more intelligently decide on the size of each BI node rather than having them all sized equally. Sizing each node for its requirements reduces the type of waste unavoidable with fixed size physical machines.

However virtualization also enables administrators to provision as many virtual CPUs as they would like - which leads typical environments to having far more virtual CPUs (vCPUs) provisioned than actual physical CPUs (pCPUs) in the machine. The obvious consequence is that multiple virtual machines may expect the full power of a full physical core but be forced to time-share with each other.

SAP recommends a simple philosophy: Provision the right amount of resource required and get an iron-clad guarantee from the hypervisor that what is asked for will be delivered. When creating virtual machines, you must realize that for some parameters, these are “requests” that are not actual promises by the hypervisor. In times of severe host stress, the virtual machine may not get the resources it thinks it has, or may have to pay a higher price to get it such as having to suffer increased latency or reduced throughput.

### *3.5.1 Size VMs Large Enough To Meet Processing Requirements*

The SAP BusinessObjects BI 4 Sizing Guide specifies the *minimum* requirement for BI nodes in production landscapes to be 8 CPU cores and 16 GB of RAM. While it is possible to run SAP BusinessObjects BI 4 on fewer cores and less RAM, this should only be done in development and test environments as these landscapes do not have the same expectation for capacity and performance as production deployments.

Most modern multi-socket servers have more than 8 cores, but it can sometimes be a challenge to convince an IT department that VMs of this size should be allocated. SAP’s internal testing shows that the resource overhead of three separate 4-core machines versus one 12-core machine is large enough to justify deploying a single 12-core or larger virtual machine even if it reduces VM density.

Contrary to previous generations of hypervisors, products like VMware ESXi 5 do not incur significant additional overhead when allocating large amounts of memory for a virtual machine. Therefore there is no technical reason why a VM cannot be configured with 32 GB, 64 GB, or even more if required and configured properly. However there may be other reasons why large amounts of RAM in a single VM may not be a good idea - such as when it is more appropriate to split the number of BI service instances across more nodes rather than stacking them all on a single machine.

> ***Always remember the consequences of not giving the BI system enough resources (i.e. RAM) are exponentially worse than in a physical world. Weigh the cost of additional RAM or additional CPU cores against the impact of poor system performance for hundreds of users over the lifetime of your BI system.***

It may be tempting to reduce the memory allocation on a virtual machine that reports “active memory usage” or “average memory usage” significantly lower than the amount allocated – don’t do it. BI systems dynamically allocate and de-allocate large amounts of memory based on workload and will trigger swapping to disk if there isn’t enough memory. Tracing performance problems back to this mistake can be difficult as the dynamic nature of BI means the problem can exist but be difficult to reproduce.

### *3.5.2 Size VMs Small as Possible*

It is a virtualization best practice to allocate the minimum resource required for a virtual node to reduce resource contention and scheduling conflicts that can cause “ready waits”, or situations where the system may have to halt a VM until all allocated resources are physically available. VMware ESXi 5.x allows up to 32 virtual CPUs in a single VM, but each additional vCPU imposes a small resource requirement that translates into real CPU consumption on the host. If a VM is configured for more vCPUs than required, the host overhead will increase unnecessarily.

SAP BusinessObjects BI 4 is engineered to leverage all CPU and memory available and has proven to scale to well over 10,000 active concurrent users (representing an organization of over 100,000 users) *[Ref. 10]* SAP does not recommend virtual machines with a large number of vCPUs – while there are exceptions, there are few situations where allocating more than 16 virtual CPUs in a single VM will result in better performance and stability compared to a few, smaller VMs. For example, it is likely that splitting a single 24-vCPU virtual machine that has multiple Web Intelligence instances spread throughout multiple VMs will not only reduce resource and scheduling conflicts, but also provide a better ability to monitor resource usage and succeed in your scale out strategy.

### 3.6 Always Deploy Virtualization Guest Tools

The magic of virtualization is that each hardware element an operating system expects in a physical system is properly emulated when running in a virtual machine. Since any hardware function can be called at any time, the hypervisor must present a complete system to the guest - right down to floppy drives, physical card slots, and even BIOS interrupts. For example, when VMware's virtual environment emulates a standard network card in VMware such as the Intel E1000, it also emulates the PCIe slot it is plugged into and communication with this card is done through system interrupts just as it is done in a physical system. This emulation is so perfect that a virtualized guest can typically use the same device driver as it would for a physical E1000 network card.

Guest virtualization tools, such as VMtools in VMware, allow the hypervisor and guest be aware of each other and co-operate together. Using this interface, the virtual machine reports OS-level telemetry metrics and a set of interfaces that allow the host and guest to communicate with each other without requiring hardware emulation. This allows the hypervisor to make better judgments when its guests request resources. Additionally, many virtualization specific features such as memory overcommit, live migration, and cut-and-paste are not possible at all without these guest tools.

In the example described above, a virtual machine with VMtools could use the VMXNET3 virtual network card instead of the Intel E1000 card to improve performance. VMXNET3 is a paravirtualized driver that allows an operating system to use a standard device driver interface to co-operate with the hypervisor and participate on a network without requiring the hardware emulation and related overhead when emulating a physical network card.

For optimal performance, SAP recommends you always install the latest version of VMware VMtools in each guest virtual machine on every host. As these tools are occasionally updated, it is equally important to ensure you upgrade these tools within each VM when you upgrade your VMware hypervisor to a newer version.

# 4 ADDITIONAL TOPICS

This section includes additional topics relevant to SAP BusinessObjects BI 4 customers but is not specifically related to optimally deploying their systems in a virtual environment. These topics are not covered exhaustively here but are included for completeness.

### 4.1 Architecture and Deployment Considerations

In general, your virtual deployment should be fairly similar to what you would architect in a physical landscape. You should be able to more efficiently distribute a physical host's processing power by segmenting its resources into multiple, smaller virtual machines. If you intelligently plan your scale out strategy, you will be able to individually monitor each BI node and determine when you need to scale out to an additional virtual machine and which services you should deploy on it. Conversely, in a larger more monolithic architecture, it will be difficult to isolate performance bottlenecks and you will likely be forced to scale out with another large system that has more services running on it than you may need.

#### *4.1.1 BI Sizing in a Virtual Environment*

Sizing an SAP BusinessObjects BI 4 landscape is an involved exercise and ensuring system performance in a virtual environment can add another dimension to an already difficult problem. As it has been discussed in previous sections, resource sharing and contention can drastically vary the performance of a virtual machine without warning or recourse. In such systems, it is impossible to determine if a performance issue is caused by the BI system itself or an external factor such as another virtual machine executing a large workload on the same host.

In SAP sizing terminology, this means the effective "SAPS Rating" or measure of a virtual machine's computing capacity can be highly variable and therefore cannot be accurately calculated. For example, a VM that could deliver between 50%-100% of it's expected performance is similar to saying, "a BI landscape requires somewhere between 4 and 8 nodes" – the variability is so large, the measurement is not very helpful. Without guaranteeing system resources in the hypervisor, all SAPS calculations, sizing guidance, and performance expectations are invalid.

The primary defense to counter these effects is to employ CPU processing and memory reservations for each virtual machine running SAP BusinessObjects BI 4. While these reservations do not guarantee performance or even performance consistency, they are the best and most basic requirements that clearly signal to the hypervisor that these resources belong to your BI system and should not be shared if at all possible.

In practical terms, how well you configure your systems and optimize your environment will significantly affect its real-world performance. Our COIL tests show that no virtualization overhead calculation is necessary in an optimal system *under optimal conditions*. However it is impossible to optimize each system perfectly and it is not always possible due to existing IT policies, virtualization inexperience, or external factors. You may consider factoring in an overhead measure of 10-15% to account for this. In general, customers who are claim their virtualization is greater than 10% typically have an oversubscribed (overloaded) system or have not properly configured their environment.

***When strict CPU and memory reservations are used and resource contention is kept to a minimum, all sizing guidance and information contained in the SAP BusinessObjects BI 4 Sizing Guide [Ref. 2] is fully applicable to virtual systems as well.***

***Official guidance for sizing and deploying the SAP BusinessObjects BI 4 suite is always at www.sap.com/bisizing. Always ensure you are obtaining sizing tools and documents from this site at the time you need them as they may have been updated since you last looked there.***

No changes to the Sizing Estimator *[Ref. 2]* process are needed for a virtualized system configuration—follow the same process when sizing for either a virtual or a physical environment. While there can be a certain amount of overhead for virtualization, it varies greatly based on the exact system configuration, workload, I/O requirements, and even which CPU is used. As discussed in the previous section, you may want to factor in a 10-15% buffer to the sizing to account for this.

Note that the benchmarking data used to create the BI 4 Sizing Estimator is based on a defined data and report set, with testing done at different load levels. Therefore it is important to remember that *the numbers it returns are bare minimums and only starting points for you* – your real-world scenario requires a proper sizing strategy and methodology. Do not simply enter numbers into the Sizing Estimator and consider your sizing exercise done!

#### *4.1.2 Dealing with IT-Imposed Limitations*

Even with the configuration flexibility that virtualization provides, many IT organizations have standardized on predefined or "T-shirt" size VM specifications such as "small", "medium", and "large" instances and will not allow arbitrary virtual machine sizes. If your organization has such limitations, you will need to do further planning to ensure your deployment spans enough VMs to provide enough resources for all services. These restrictions not only force you to think carefully about the resources you need, but also force you to think about how you would scale out rather than scale up your system.

Some IT departments may also impose specific limits on the number of virtual CPUs or RAM that can be allocated to a single virtual machine. If this limit is less than the minimum requirements to deploy SAP BusinessObjects BI 4, you will need to justify an exception. For example, if your IT department has a restriction for virtual machines to have no more than 4 vCPUs and 8 GB of RAM, you may be able to deploy test systems, but you will simply not be able to deploy a production SAP BusinessObjects BI 4 system that hopes to perform acceptably.

There are a number of strategies to deal with this: you may need to fund the acquisition of dedicated hardware for only your virtual machines or obtain approval from your CIO for the exception on shared hardware.

> ***It is quite common for customers that cannot obtain virtual machines large enough to properly support their SAP BusinessObjects BI 4 system to complain about performance and incorrectly blame virtualization as the cause. A physical system with less than the minimum requirements will also underperform – however this is usually not noticed because modern servers typically ship with specifications well above the minimum requirements for productive SAP BusinessObjects BI 4 systems.***

### *4.1.3 Use Monitoring Metrics Wisely*

Monitoring metrics in a virtual environment can be misleading and you should take great care in interpreting the values you receive. It is important you consider metrics from the hypervisor layer as well as within the VM – but remember that monitoring a virtual machine from *within itself* can yield different results from monitoring at the hypervisor level.

Guest monitoring tools that are not hypervisor aware rely on measuring the effective usage of the resource the virtual machine has been given. For example, the guest may report 100% CPU utilization, but if the host is only giving it 50% of the CPU cycles the VM expects, you will not be able to conclude whether the BI node is overloaded and should scale out or if there is a performance issue with the host.

Hypervisor metrics should also be evaluated with care. At the macro level, the VMware administrator is focusing on keeping the resource usage within certain limits to ensure there are enough CPU cycles or memory for all virtual machines during *average usage*. While this can be appropriate for the vast majority of applications such as email servers, database servers, and even high throughput ERP systems, activity on BI systems is I/O intensive and driven dynamically by user activity. A number of users returning from lunch and running CRM reports at the same time can quickly overwhelm an under-provisioned BI system. Therefore, BI virtual machines should be monitored for their *maximum,* not average CPU usage (and memory).

> ***BI users rarely appreciate "average" performance. A BI system for the Finance department that is idle most of the day but is heavily used during the afternoon may report "average" CPU usage most of the day, but can provide unacceptably long response times when users need the system the most. Poor performance of BI systems is commonly blamed on virtualization – the real cause is often a misunderstanding that provisioning for "average" resource usage instead of peak usage is acceptable.***

### 4.2 Licensing Considerations When Virtualizing

When planning your virtual deployment, you must also consider its licensing consequences as well. This will largely depend on your SAP license agreement and the virtualization rights specified within it. As these rights are occasionally updated, it is important to check your company's own license agreement and then contact your SAP account representative to clarify.

From a virtualization perspective, the least restrictive license type is the "Concurrent Session Based License" (CSBL). In this model, the system is licensed by the number of concurrent users connecting to the system regardless of how many machines or cores are deployed. This provides the most freedom to deploy additional machines in a scale out strategy or to split larger monolithic BI nodes into a more reasonable architecture. The "Named User" license also provides this type of freedom, albeit with the restriction that specific users need to be licensed.

Customers with a "CPU-based" license have an even greater need to understand their virtualization rights properly. While rights have changed over time, the general concept is you have a number of CPU licenses and can only deploy on calculated (formula-based) number of cores on the host to remain license compliant. While CPU and memory reservations are important for all license types, customers with CPU licenses can drastically increase their software licensing costs without even knowing it if resource reservations are not strictly enforced.

Consider this over-simplified example: A customer has a license allowing deployment on 16 cores and deploys two 8-core virtual machines on a completely empty host and does not use resource reservations. As more virtual machines are put on the system and load increases, the BI system eventually starts to perform unacceptably – specifically, queries and reports are taking much longer than expected. The BI administrator, using OS-level and BI-level monitoring in the guest, sees high CPU usage and determines that both nodes are overloaded and a third node is required. Naturally, this would require licensing for 8 more cores, presumably on another, less loaded host.

In this case, the BI administrator likely does not know the original host is oversubscribed, and without CPU and memory reservations, the host is only able to provide 50% of the power each VM is expecting. Adding a third VM would indeed add processing power (by moving from two nodes to three), but eight additional cores would need to be licensed for the additional BI node.

Configuring the host to give the promised computing resource in the first place would be a far more cost effective alternative. Even in the case where the BI team would have to buy dedicated hardware to guarantee the BI system gets full power, the hardware and related costs would likely be much cheaper than the combined cost of additional software licenses and the TCO of an additional BI node.

> ***Customers who have CPU-based licenses but do not enforce resource reservations are effectively increasing their cost-per-CPU as the throughput of each licensed CPU will be less than expected. To achieve the required throughput by adding more underpowered nodes will be more expensive than properly provisioning the existing virtual machines.***

# 5 CONCLUSIONS

***The hypothesis that SAP BusinessObjects BI 4 can be virtualized with little or no performance impact is proven to be true provided the recommendations in this whitepaper are implemented.***

We also conclude that customers who do not follow these recommendations risk suffering significant performance penalties and significantly raising their TCO through inefficient hardware use and increased licensing costs.

Results from multiple test runs of our COIL study also strongly correlate with feedback from customers and other internal tests conducted within SAP, suggesting a sufficient level of confidence the recommendations in this study are representative of best practices.

Comparing the methodology and results of our COIL experiments with earlier internal and external studies, we also conclude these previous studies are invalid due to inherent flaws in their testing methodologies and analysis.

It is also clear that since straying from these recommendations can easily create poorly performing systems, BI teams have an important responsibility to work with their IT teams to ensure SAP's recommendations are properly understood and implemented. Deviating from these recommendations does not impact the support a customer can expect from SAP, but customers should expect instruction to implement the recommendations in this paper before further diagnosis can be done on their systems for performance related issues.

Following the recommendations in this paper will not guarantee a high performance system as there are many architecture, sizing, and configuration choices that also have significant impacts on performance, but there is sufficient evidence that these recommendations are key success factors to achieving such a system.

## 5.1 Summary of Best Practice Recommendations

Below is a high level summary of best practice recommendations as a result of our COIL tests and analysis. Page numbers to the relevant sections are included for reference.

| Recommendation | Key Points |
|---|---|
| Design and ensure a good system architecture before performing optimizations at any level (page 31). | • Virtualization gives freedom to make a better architecture than typically possible in a physical deployment.<br>• Optimization cannot compensate for poor system design – poor designs will suffer the same bottlenecks on physical systems as well. |
| Follow guidance specifically for SAP BusinessObjects BI 4 and disregard previous studies (page 21). | • Disregard virtualization guidance for “SAP Software” unless it is specifically for SAP BusinessObjects BI 4.<br>• Disregard previous SAP and non-SAP virtualization studies for BI as they may be on previous versions or be flawed/incorrect. |
| Use the latest version of VMware and follow VMware’s best practices guide (page 22). | • Guidance in this whitepaper is supplemental to VMware’s best practices.<br>• Use the latest version of ESXi to take advantage of performance improvements at all layers of the stack (hardware, OS, application, etc.) |
| Review configuration of each level of your hardware & software stack (page 23). | • Ensure configuration at all levels have sensible values.<br>• Use defaults for each stack element unless following guidance by the specific vendor for that element.<br>• Always use the latest versions of all components whenever possible to take advantage of performance improvements. |
| Use strict CPU (processing power) reservations for each BI VM on every host (page 28). | • VMs without CPU reservations may share CPU time with other processor intensive guests and negatively impact performance.<br>• All standard BI sizing and performance expectations are invalid if CPU reservations are not strictly enforced. |
| Use strict memory reservations for each BI VM on every host (page 28). | • VMs without memory reservations may be forced to page to disk if the host experiences memory pressure.<br>• All standard BI sizing and performance expectations are invalid if memory reservations are not strictly enforced. |
| Do not use shares, limits, affinity, or other artificial mechanisms to divide VM resources on the host (page 29). | • Any mechanism that divides or shares resources enables the host to not provide all promises resources.<br>• All standard BI sizing and performance expectations are invalid if resource sharing features are employed. |
| Size VMs large enough for BI workloads (page 30). | • Size your VMs the same way you would on physical nodes. If your IT department does not allow for large enough VMs, seek exceptions or expect the system to struggle. |

| | |
|---|---|
| | • Ensure your VMs meet at least the hardware minimum requirements as stated in the BI 4 Sizing Guide *[Ref. 2].* |
| Use more, smaller VMs rather than a few very large (>16 CPU) nodes (page 30). | • Split very large VMs into smaller ones to better monitor and manage resource usage.<br>• Individual VMs still have inherent bottlenecks regardless of how powerful or well provisioned the host is. |
| Always deploy VMtools in every guest (page 31). | • VMtools provides the mechanism for the host and guest to work efficiently together.<br>• Without VMtools, the system may still run, but all standard BI sizing and performance expectations should be considered invalid. |
| Use strict resource reservations and other recommendations from this guide to ensure your BI sizing is valid (page 31). | • Minimize the “virtualization tax” to avoid invalidating your BI sizing.<br>• Conservatively factor in some overhead to provide a performance buffer. |
| Use hypervisor-level monitoring tools (page Use Monitoring Metrics Wisely33). | • Guest-level monitoring tools are not hypervisor aware and cannot report the true system statistics.<br>• Work with your IT department to get access to hypervisor-level tools for monitoring or work with them to do a performance analysis with them using tools like “esxtop”. |
| Ensure your landscape reflects your SAP software license and virtualization rights (page 34). | • Maintaining license compliance can affect how you deploy in a virtual environment.<br>• Ensure you review your specific virtualization rights – if you have CPU-based licensing, ensure you are calculating the host’s cores correctly and getting full value from each one through reservations. |

# 6 REFERENCES


1. SAP BI 4 Virtualization Resources
   http://www.sap.com/bivirtualization

2. SAP BI 4 Sizing Resources
   http://www.sap.com/bisizing

3. SAP BusinessObjects BI 4 Platform Availability Matrix (PAM)
   http://www.sap.com/pam

4. SAP BusinessObjects BI 4 Documentation Microsite
   http://help.sap.com/bobip40

5. SAP Note: 1492000: General Support Statement for Virtual Environments
   https://service.sap.com/sap/support/notes/1492000

6. SAP SD Standard Application Benchmark Results, Two-Tier and Three-Tier Configurations
   http://www.sap.com/solutions/benchmark/index.epx

7. Performance Best Practices for VMware vSphere™ 5.1
   http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf

8. vSphere Resource Management Guide for ESXi 5.1:
   http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-resource-management-guide.pdf

9. Evaluating Selected Java Best Practices For Sap BusinessObjects BI 4 On VSphere
   http://scn.sap.com/docs/DOC-29008

10. Large-Scale SAP BusinessObjects BI 4 Platform Deployment Over SAP Sybase ASE and SAP Sybase IQ Databases
    http://scn.sap.com/docs/DOC-31865