November 2009

# Analysis of IBM System x3850 M2 Performance and Scalability with VMware vSphere 4 and SAP Solutions

*Co-Authored by:*

*Param Singh*
*IBM Systems and Technology*

*Ken Barr*
*VMware, Inc.*

# Introduction

This paper provides an overview of the performance and scaling characteristics of the IBM System x® 3850 M2 running an implementation of an SAP® business solution in a virtualized environment using VMware® vSphere™ 4. Together these components comprise a highly scalable, flexible, and reliable virtualization infrastructure for mission-critical applications. This paper highlights how IBM X-Architecture® technology contributes to excellent performance scaling for a high server consolidation ratio.

These benefits are demonstrated using a production-class 64-bit software stack consisting of Microsoft® Windows® Server 2008 and Microsoft SQL Server 2008. SAP software includes SAP NetWeaver® 7.01 (ABAP) with a Unicode kernel and Enhancement Package 4 for SAP® ERP 6.0. Performance and scalability data are presented for both scale-up and scale-out scenarios. The scenarios highlight the relationship between CPU utilization and user response time, the importance of processor cache for the SAP application workload, and the ability of the vSphere scheduler to allocate resources equitably among multiple virtual machines running concurrently.

At the conclusion of the paper, tuning and optimization best practices are presented for implementing an SAP application workload on the IBM System x3850 M2 using IBM X-Architecture technology in a VMware vSphere 4 environment.

# Platform Used for the Performance Study

## *IBM System x3850 M2 and x3950 M2*

The 4-socket IBM System x3850 M2 and the scalable x3950 M2 are ideal as an x86 platform for virtualization. Designed to provide a rock-solid and flexible virtualization infrastructure, these servers scale from 2 to 16 sockets, 32 to 128 DIMM slots, and 7 to 28 PCI slots. With either the x3850 M2 or the x3950 M2, clients can start small and scale up processing power, memory capacity, and I/O bandwidth as their needs grow.

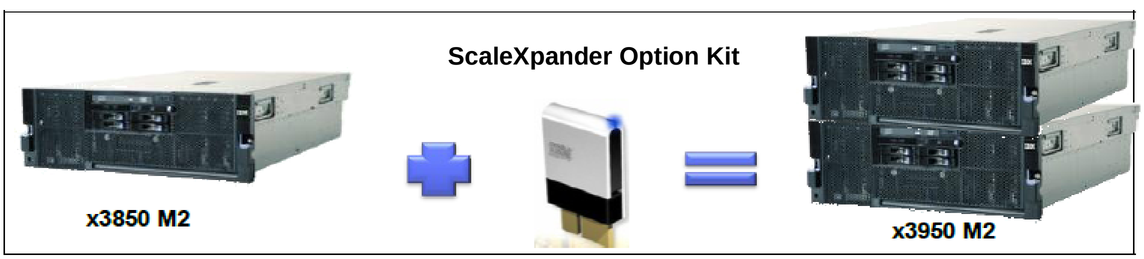


**Figure 1. IBM System x3850 M2 and 2-node x3950 M2**

The features that distinguish these servers from those sold by other vendors and that make them ideal for virtualization of mission-critical SAP business applications include:

- Fourth-generation IBM X-Architecture technology based on the eX4 chipset that provides:
  - Mainframe-like innovations for high reliability, availability, and serviceability (RAS)
  - Excellent scalability and performance with up to 16 processors (96 cores)
  - High power efficiency
- Easy scalability with the ScaleXpander Option kit, which enables the attachment of a second x3850 M2 chassis to double the number of processors, memory slots and PCI slots (See Figure 1). Adding two more chassis enables the x3850 M2 chassis to scale to 16 processors, 1TB of memory, and 20 PCI slots.

The x3850 M2 and x3950 M2 are designed to be an excellent platform for large-scale consolidation providing a high server-consolidation ratio as well as the headroom needed to meet peak-load demand and ensure excellent response time.

### *Key Highlights of the* eX4 Design Features for Server Consolidation in an SAP Application Environment

The eX4 chipset design offers numerous features per chassis to boost performance and reduce product and operating costs, including:

1. A built-in snoop filter to lower the amount of inter-processor communication needed in an SMP system. The snoop filter provides excellent scaling and balanced performance when multiple VMs are running concurrently.

2. Fast PC2-5300 DDR II ECC memory with Chipkill error correction, Memory ProteXion™ (redundant bit steering), and optional memory mirroring and hot-add/hot-swap memory protection provides speed and high availability.

3. Seven 64-bit high-speed (4Gbps) PCI-E x8 adapter slots offer investment protection by supporting high-performance adapters, such as 10Gb Ethernet, Fibre Channel and InfiniBand™ cards.

4. Predictive Failure Analysis (PFA) designed to allow the x3850 M2/x3950 M2 to detect an impending failure of supported components (processors, memory, PCI-E slots, VRMs, power supplies, fans, and hard disk drives) *before* actual failure, and alert the administrator through IBM Director.

**eX4 Chipset “Snoop” Filtering**
One of the core features of the eX4 chipset that gives the platform its performance advantage is the “snoop” filter. The Xeon® cache coherence protocol relies on its participants to observe, or snoop, memory operations on a shared bus. When a processor attempts to operate on a data that may exist in another processor’s cache, the intent (read or write) of that processor and the data’s address are snooped by other processors in the SMP complex. These processors take action based on the snooped message such as invalidating local copies or providing the most recent data to the initiating processor. Cache coherence messages, cache-to-cache transfers, and write backs increase traffic on the front-side bus (FSB). The eX4 chipset contains 324MB of EDRAM (Embedded DRAM) within the Northbridge chip. All data is copied as it is written to the processor cache, allowing the Northbridge chip to respond directly to the processor snoop requests. This reduces the overall traffic across the FSB and allows the processor to service more requests.

**XceL4v Dynamic Server Cache**
Another performance feature of the eX4 chipset is the XceL4v L4 cache. When using a single node (chassis), the cache works with the snoop filter to help reduce FSB traffic. When more than one node is used, 256MB of virtual cache per node (taken from main memory) is used for interprocessor communications between chassis to keep data in synch. This reduces traffic between processors in multi-node chassis, allowing the processor to service more requests.

## *VMware vSphere 4*

VMware vSphere 4 can help customers manage a smarter SAP application-based landscape via template-based virtual machine cloning; distributed resource scheduling and power management (VMware DRS and DPM); and improved uptimes through VMotion, VMware HA, and Fault Tolerance [1]. Prior work demonstrates that these management and reliability features are achieved while maintaining excellent performance in a single virtual machine [2]. This paper focuses on the scalability features of vSphere 4 and includes results of SAP software running in a scale-out scenario with multiple virtual machines.

VMware vSphere 4 now supports larger virtual machines and physical hosts. The software supports up to 8 virtual CPUs and 255GB of memory per virtual machine. Up to 64 cores and 1TB memory are supported per host. Together with improvements made to Virtual SMP support and updates to storage and networking drivers, these sizes allow the SAP software to scale up as well as scale out with excellent performance.

## *Configuration of Tested System*

The performance of SAP solutions on vSphere was studied with a popular online transaction processing (OLTP) workload used to size SAP application deployments. The workload simulates sales and distribution scenarios with concurrent users creating orders, entering delivery information, posting a goods issue, and performing other typical sales and distribution tasks. There are 10 seconds of “think time” between each transaction.

SAP application deployments may be configured in a two-tier or three-tier configuration. All experiments in this paper used a two-tier configuration wherein the database and application server share the same virtual machine. An automated load generator served as the presentation tier. The load generator slowly ramps up the number of concurrent users until a preset maximum is reached. After ramp-up is complete, response times were measured for 10-15 minutes. The number of concurrent users that can be served while maintaining an average response time below 1 second was recorded. The experiments were conducted in good faith and run to completion without errors or warnings, but the results in this paper have not been certified by SAP.

The two-tier setup consisted of a database and application tier and a presentation tier as shown in Figure 2. This section refers to these tiers as the SAP Server and the Load Generator, respectively. The SAP Server was an x3850 M2 configured with four Intel® Xeon X7460 2.66GHz processors, 128GB memory, and VMware vSphere 4. Virtual machines were configured with Microsoft Windows Server 2008, Microsoft SQL Server 2008, and an SAP solution stack consisting of SAP ERP 6.0 EHP4 with Unicode and SAP kernel 7.01 (Patch Level 31). VMware Tools were installed in the virtual machine to provide optimized network and storage drivers (vmxnet3 and PVSCSI).

The following BIOS settings were used to optimize the x3850 M2’s performance:

| | |
|---|---|
| PowerExecutive Power Capping | [Disabled] |
| Processor Performance States | [Disabled] |
| Clustering Technology | [Logical Mode] |
| Processor Adjacent Sector Prefetch | [Disabled] |
| Processor Hardware Prefetcher | [Disabled] |
| Processor Execute Disable Bit | [Disabled] |
| Intel Virtualization Technology | [Enabled] |
| Processor IP Prefetcher | [Enabled] |
| Processor DCU Prefetcher | [Disabled] |
| C1E | [Disabled] |

The SAP server was connected to a storage-attached network (SAN) with a Fibre Channel network adapter. The SAN storage consisted of the IBM System Storage™ DS4800 with seven IBM System Storage EXP810 enclosures.

The Load Generator was implemented using an IBM System x3850 running SUSE Linux® Enterprise Server 10. Communication between the load generator and the x3850 M2 was over a 1Gb/second Ethernet link.

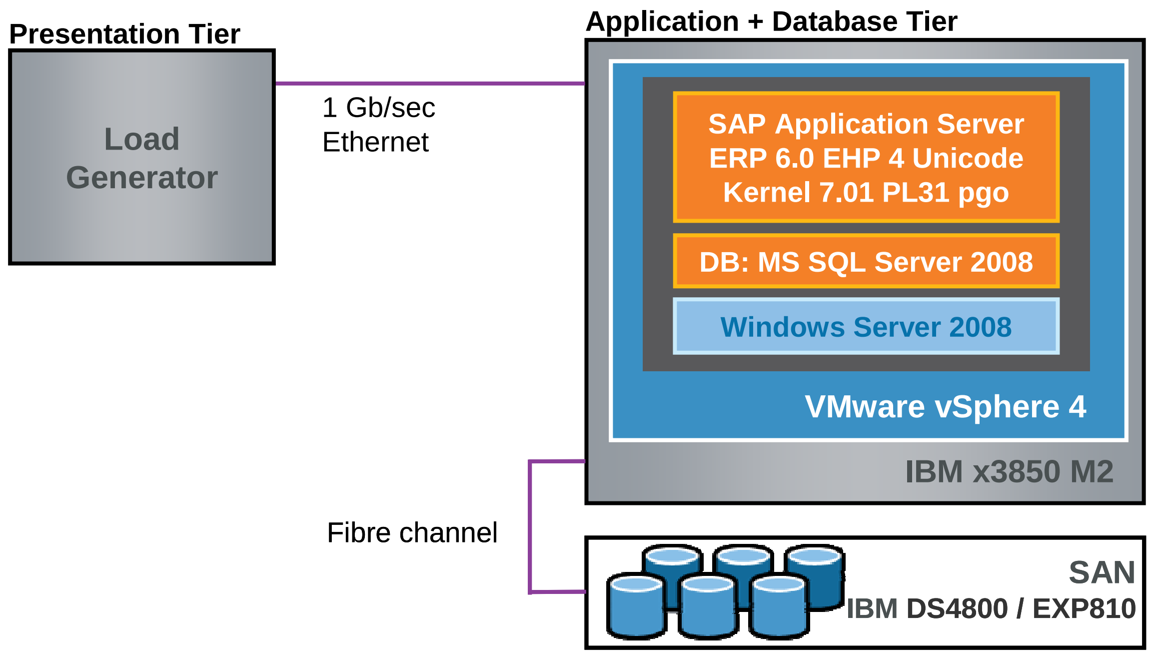


**Figure 2. SAP solution architecture used in this study**

## Results and Analysis

This section presents several test scenarios with analysis.

### *Scenario 1: User response time and CPU utilization*

The SAP application workload used in this paper is CPU-intensive and tuned to support as many users as possible within the 1-second response time guideline. As the number of simulated users is gradually increased, the CPU utilization begins to rise. When the CPU is 100% utilized, additional users cause response time to increase rapidly.

Figure 3 shows the characteristic behavior of the workload. An exponential curve is fitted to the data collected on a virtual machine with 2 vCPUs and 8GB of memory. Similar behavior is observed on native hardware.

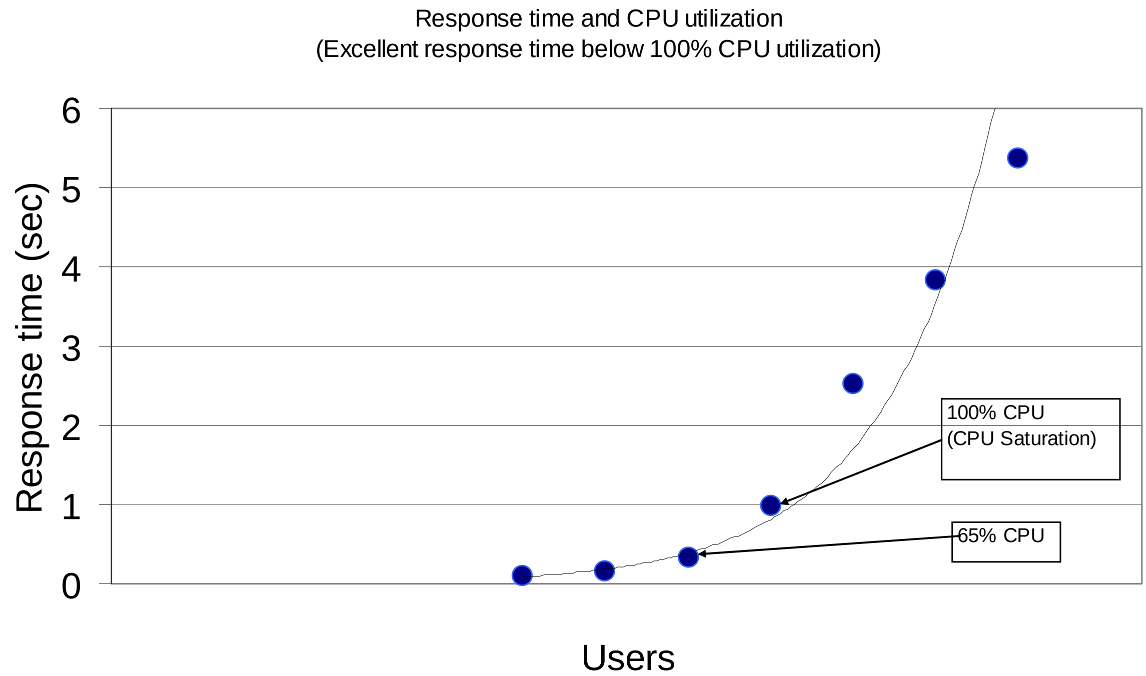


**Figure 3. User response time as CPU utilization increases for one VM configured with 2 vCPUs**

To support a greater number of users within the response time constraint, the virtual machine would require more CPU resources (e.g., additional vCPUs or faster processors). To balance good response time with high resource utilization while retaining the ability to handle peak loads, SAP and IBM recommend that the average CPU utilization be around 65%.

### *Scenario 2: SAP NetWeaver performance and processor cache*

In this scenario, the number of virtual machines running per processor is increased to measure the effect of processor cache on the performance of a virtual machine. The cache hierarchy of one Intel Xeon X7460 processor that populates an x3850 M2 server is shown in Figure 4. Each processor has six cores, and each core has 64KB of level one cache. Each pair of cores shares a 3MB level two cache. All cores share 16MB of L3 cache.

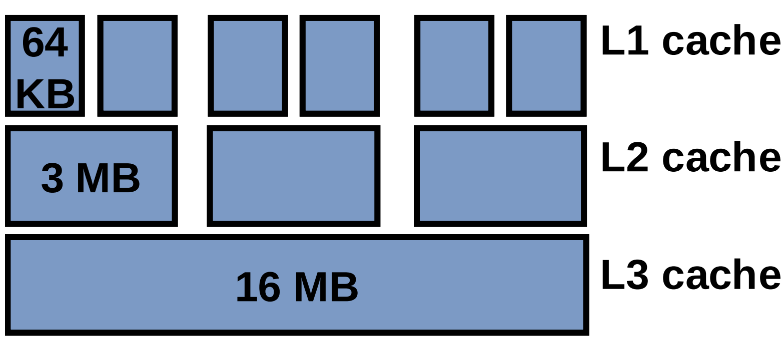


**Figure 4. Intel Xeon X7460 cache hierarchy**

Each six-core processor can support up to three 2-vCPU virtual machines without over-committing the physical processor. A single virtual machine running on the processor has the benefit of the entire L3 cache: up to 16MB of its most frequently used data can be kept in this cache to avoid more costly accesses to main memory. When a second virtual machine is scheduled on the same processor, the effective cache size per virtual machine is reduced to half. Similarly, when three virtual machines are scheduled on the processor, the effective L3 cache per virtual machine is one-third of processor cache.

Figure 5 shows the impact of this reduced effective cache space on the SAP application workload described earlier. When a single virtual machine runs on the processor, the response time is approximately 1 second. Halving the cache space per virtual machine by adding a second virtual machine increases the response time to 2 seconds, and adding a third virtual machine causes another doubling of response time as the effective L3 cache size per virtual machine is reduced to 5.3MB. Note that each virtual machine simulates the same number of users.

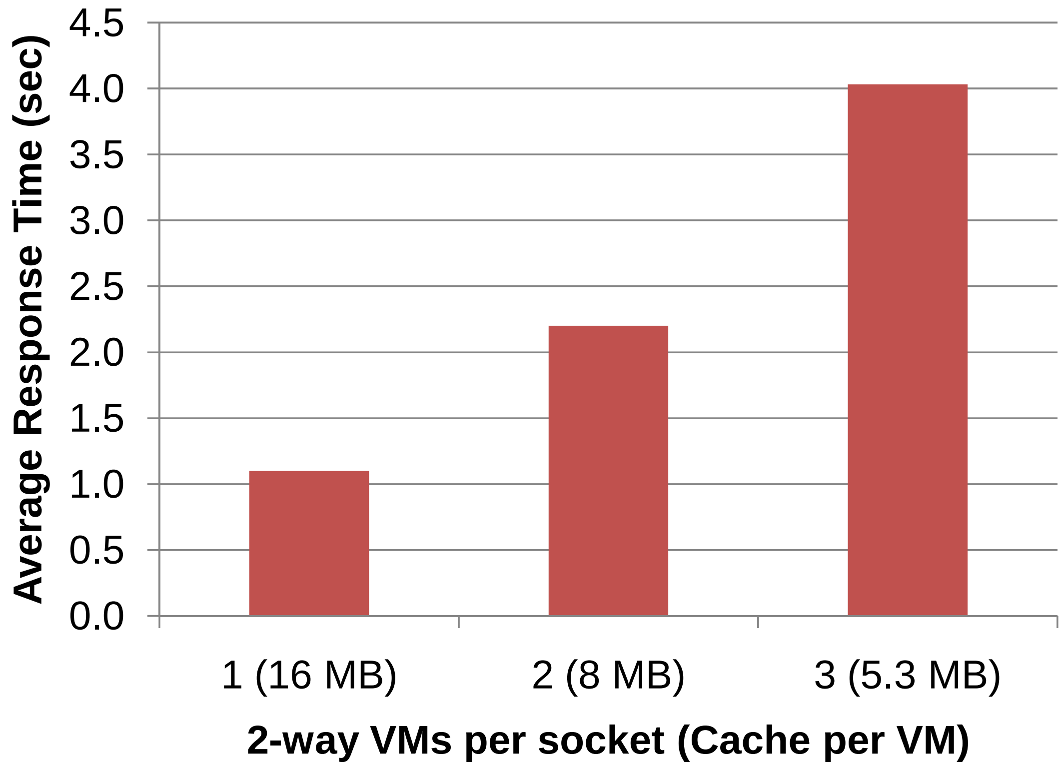


**Figure 5. Increase in user response time as VMs are added to the x3850 M2's six-core processor**

This result, which demonstrates the cache sensitivity of this workload, motivated the cache-aware scheduler in vSphere. vSphere is designed to recognize processor cache hierarchy and allocate and assign virtual CPUs to physical cores such that available cache is maximized [3].

### *Scenario 3: Scale up – Performance scaling from 2 to 8 vCPUs*

A scale-up scenario in a virtualized environment refers to increasing the size of single virtual machine. In this particular case, the size of the virtual machine was increased by increasing the number of virtual CPUs and the amount of virtual memory. A larger virtual machine can support additional SAP work processes. In this scenario, one SAP application instance with several work processes is deployed for every two vCPUs in the virtual machine.

The results in Figure 6 show excellent scaling behavior as the virtual machine grows. The number of users in each case is compared to a baseline that uses 2 vCPUs, 8GB of memory, and a single SAP application instance. Labels on the graph show the number of users with respect to ideal scale-up of the 2-vCPU virtual machine. The experiment shows that smaller implementations are somewhat more efficient (in terms of users per vCPU) than larger implementations. This scaling behavior is similar to that seen without virtualization on an x3850 M2.

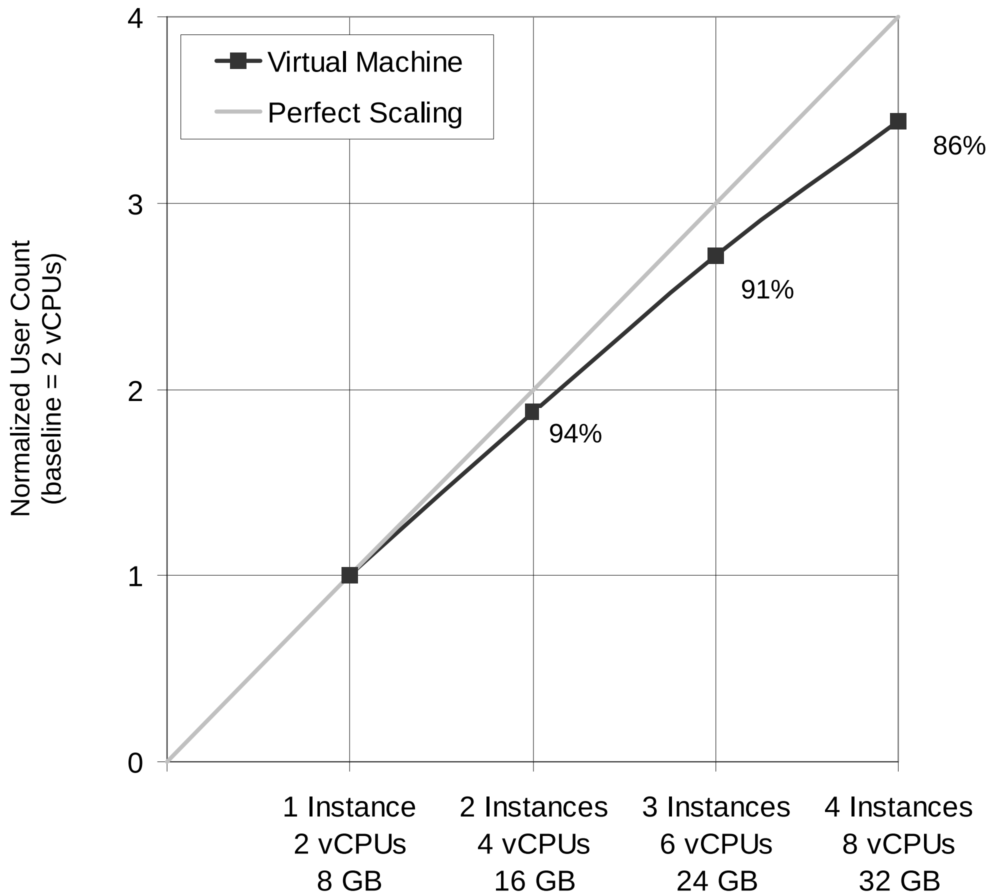


**Figure 6. Performance scales up as vCPUs are added to a VM**

The scale-up performance is influenced by many factors. Software factors include whether the software in the guest can take advantage of multiple processors and how efficiently the guest operating system can schedule multiple processes. Hardware factors such as processor type, number of cores per processor, processor cache, and memory controller affect scale-up performance as well. Given the observed performance at each data point, clients running SAP applications using VMware vSphere 4 on the x3850 M2 have the flexibility to configure virtual machines with as many vCPUs as supported by vSphere based on sizing, configuration requirements, and licensing costs.

### *Scenario 4: Scale out – Performance comparison for multiple concurrent virtual machines*

In this scenario, multiple virtual machines were run concurrently to utilize the 24 physical cores available on an x3850 M2. Out of the many possible combinations of virtual machines of various sizes, three examples are shown in Figure 7. The figure shows performance scaling with twelve 2-vCPU virtual machines, four 6-vCPU virtual machines, and three 8-vCPU virtual machines. All cases fully utilize the 24 available physical cores on the x3850 M2. The workload performed best in the example with 12 virtual machines due to the slight fall-off in scaling with larger virtual (and native) machines. However, the performance of all three examples is within 6%, allowing flexibility to configure virtual machines based on sizing and configuration requirements as well as licensing costs.

With 12 concurrent virtual machines, the aggregate number of users was eight times that achieved in a single 2-vCPU virtual machine. This sublinear scaling is largely due to the observation in Scenario 2: a reduction in effective cache size results in less cache per virtual machine when multiple virtual machines are scheduled on the same physical resources. Figure 8 shows the resulting effect as up to twelve 2-vCPU virtual machines are deployed. The curve is extrapolated from datapoints measured with 1, 2, 8, and 12 VMs.

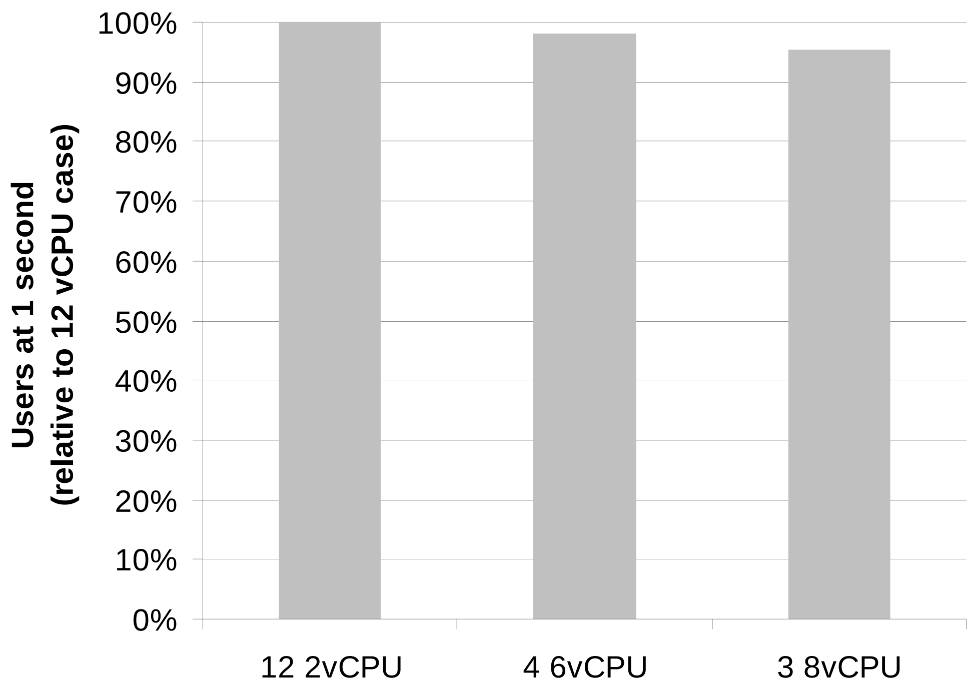


**Figure 7. Performance comparison between virtual machines with different numbers of vCPUs in a scale-out configuration**

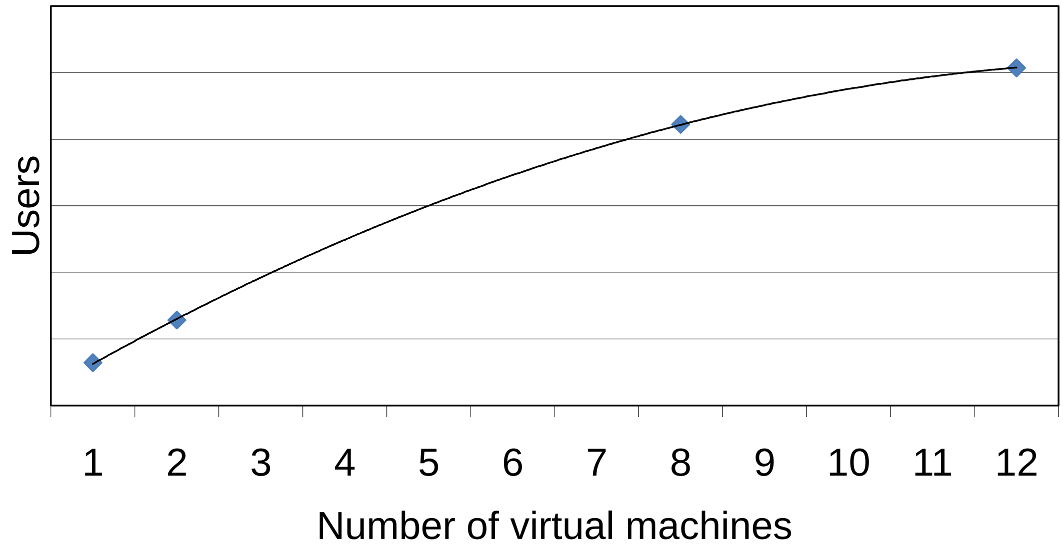


**Figure 8. Aggregate supported users as virtual machines are added to the x3850 M2 to utilize all 24 physical cores**

These measurements demonstrate that the x3850 M2 with a high core count and the ability to support large memory and I/O bandwidth delivers similar performance for various configurations of virtual machines running with VMware vSphere 4.

### *Scenario 5: Resource distribution fairness*

The purpose of this scenario was to demonstrate how resources are fairly distributed when multiple virtual machines are running concurrently to utilize all available 24 physical cores on the x3850 M2. All the virtual machines are configured with two vCPUs and 8GB memory and simulate the same number of users.

As shown in Figure 9, the variation in response time is small across all the virtual machines even in this situation of 100% processor utilization. Only one of the twelve virtual machines differs by more than the 0.8 - 1.0 second target range that we have set for this workload. This demonstrates the fair distribution of resources by vSphere.

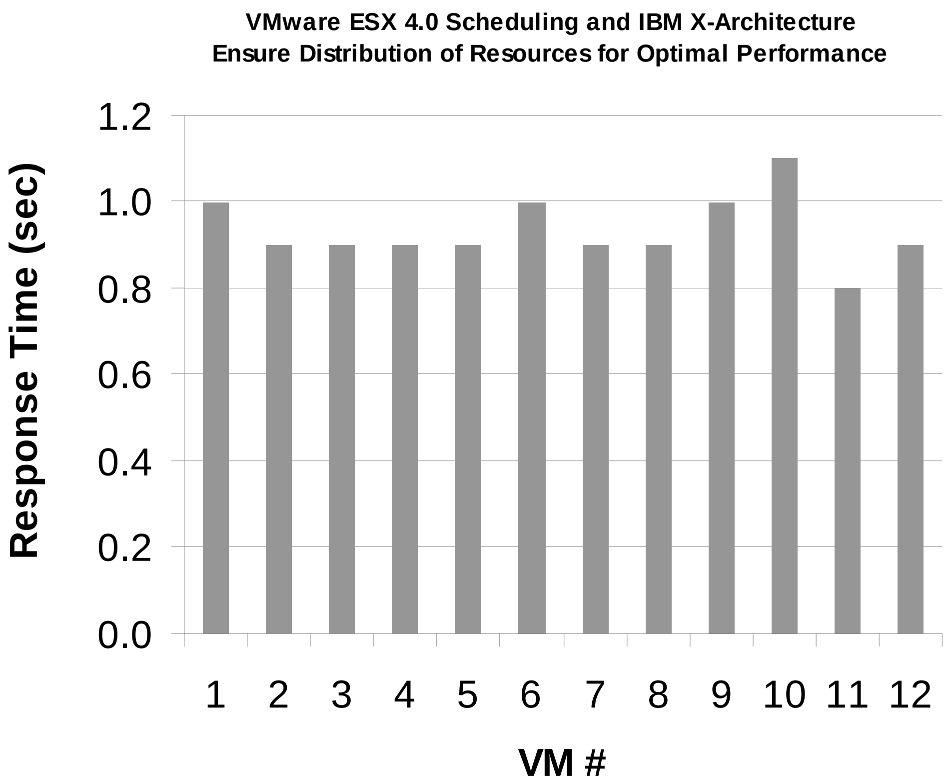


**Figure 9. Response time with 12 virtual machines running concurrently**

# Conclusion and Best Practices

The IBM System x3850 M2 and VMware vSphere 4 provide a highly compatible platform for a scalable, flexible, and reliable virtualization infrastructure that can support mission-critical applications and can achieve a high server consolidation ratio. A set of best practices can be derived from the experience of running the scenarios detailed in this paper.

1. Scale up or scale out an SAP application deployment based on vendor sizing recommendations, license cost, and ease of configuration rather than performance concerns.
   a) The IBM System x3850 M2 running VMware vSphere 4 provides excellent scale-up performance because of optimizations to VMware Virtual SMP performance, a large 16MB L3 cache per processor, 6 cores per processor, and a large snoop filter built into the eX4 chipset.
   b) The IBM System x3850 M2 running VMware vSphere 4 provides excellent scale-out performance due to the fair CPU scheduler of vSphere, the high core count of the x3850 M2's processors, and performance enhancements built into the eX4 chipset.
2. Treat each virtual machine as an ordinary machine for installation, configuration, and tuning purposes. Follow SAP best practices.
3. Follow VMware application-independent best practices [4,5].
4. Refer to SAP Note 9942 to determine the best balance between number of SAP application instances and number of work processes. The balance depends on the hardware used and the behavior of the workload. With the test bed described in this paper, the best performance was achieved using one SAP application instance per two vCPUs. Each instance had 4 dialog processes and 1 update process.
5. No CPU affinity setting at the Guest or vSphere level was used to produce these results. Performance improvement due to such a setting is small and rarely outweighs the limitations imposed, so it is not recommended by VMware or IBM.
6. Do not saturate the physical CPUs. Response time increases rapidly after CPU saturation. Allocate more vCPUs to a virtual machine or add more virtual machines to ensure optimal peak load response time. To balance good response time with resource utilization while retaining the ability to handle peak loads, SAP and IBM recommend that the average CPU utilization be around 65%.
7. Avoid memory paging:
   a) The application working set should fit in the virtual machine's memory; balance the virtual machine's memory size with SAP application and database profile parameters.
   b) The active working set of all virtual machines, virtual machine memory overhead, and vSphere system memory must not exceed host physical memory.
   c) Use memory reservation if necessary.
   d) Add memory. The x3950 M2 and VMware vSphere 4 support up to 1TB memory.

## References

1. ***SAP Solutions on VMware vSphere: High Availability***
   http://www.vmware.com/resources/techresources/10031
2. ***Virtualized SAP Performance with VMware vSphere 4***
   http://www.vmware.com/resources/techresources/10026
3. ***VMware vSphere 4: The CPU Scheduler in VMware ESX***
   http://www.vmware.com/resources/techresources/10059
4. ***Performance Best Practices for VMware vSphere 4***
   http://www.vmware.com/resources/techresources/10041
5. ***Understanding Memory Resource Management in VMware ESX Server***
   http://www.vmware.com/resources/techresources/10062
6. ***IBM System x - Enterprise servers X-Architecture platform***
   http://www-03.ibm.com/systems/x/hardware/enterprise/xarchitecture.html

# Additional Resources

### *SAP Notes (for Windows platform)*

674851: Virtualization on Windows
1104578: Virtualization on Windows: Enhanced Monitoring
1056052: Windows: VMware ESX Server 3.x or vSphere configuration guidelines
1260719: Detailed virtualization data using saposcol

### *SAP Notes (for Linux platform)*

1122388 – Linux: VMware ESX Server 3.x or vSphere configuration guidelines
1122387 – Linux: Supported Virtualization technologies with SAP
171356 – Virtualization on Linux: Essential information

### *SAP, VMware (and other) benchmarks*

http://www.sap.com/solutions/benchmark/sd2tier.epx

### *Introductory document "Virtualizing SAP applications on Windows" by SAP*

https://www.sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/70f63258-bff1-2a10-9db6-cda6ef202bfc

### *IBM Virtualization Solution*

Infrastructure solutions – Virtualization on System x and BladeCenter
http://www-304.ibm.com/jct03004c/systems/x/solutions/infrastructure/virtualization/

SAP Applications with VMware ESX Server on IBM System x3850/x3950 M2
ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/xsw03023usen/XSW03023USEN.PDF

### *VMware and IBM*

http://www.ibm.com/virtualization/vmware
http://www.vmware.com/partners/alliances/oem/ibm.html

### *SAP and IBM*

http://www.ibm.com/solutions/sap

### *VMware and SAP*

White Papers, Success Stories, Webinars, Links, etc.
http://www.vmware.com/sap

SAP VMware Blog
http://communities.vmware.com/blogs/SAPsolutions

IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
November 2009