



IBM® eServer™ X3 Architecture™: Application Server Performance Gains

*By Chris Floyd, Dan Colglazier, and Phil Horwitz
IBM Systems and Technology Group*

Executive Summary

IBM® eServer™ X3 Architecture™ leverages technology exclusive to IBM: an integrated snoop filter that minimizes front-side bus (FSB) congestion and maximizes CPU performance, a key ingredient for higher performance on the new dual front-side bus architecture. The integrated snoop filtering speeds cache coherency and reduces FSB contention. X3 Architecture's performance benefits were observed consistently across the application scenarios examined in this white paper, even with *widely varying* workload characteristics. The measured performance gains showed an overall reduction in FSB traffic of *44% and 51%* with the use of the integrated snoop filter. This reduction means lower transaction latencies and increases in overall CPU throughput. Lowering FSB traffic becomes critical to overall performance with dual-core 64-bit Intel® Xeon™ Processors MP, since there are now twice as many processors initiating traffic on the FSB. The integrated snoop-filter significantly minimizes this traffic, reduces queuing and overall transaction latencies, and effectively provides greater head-room for dual-core systems.

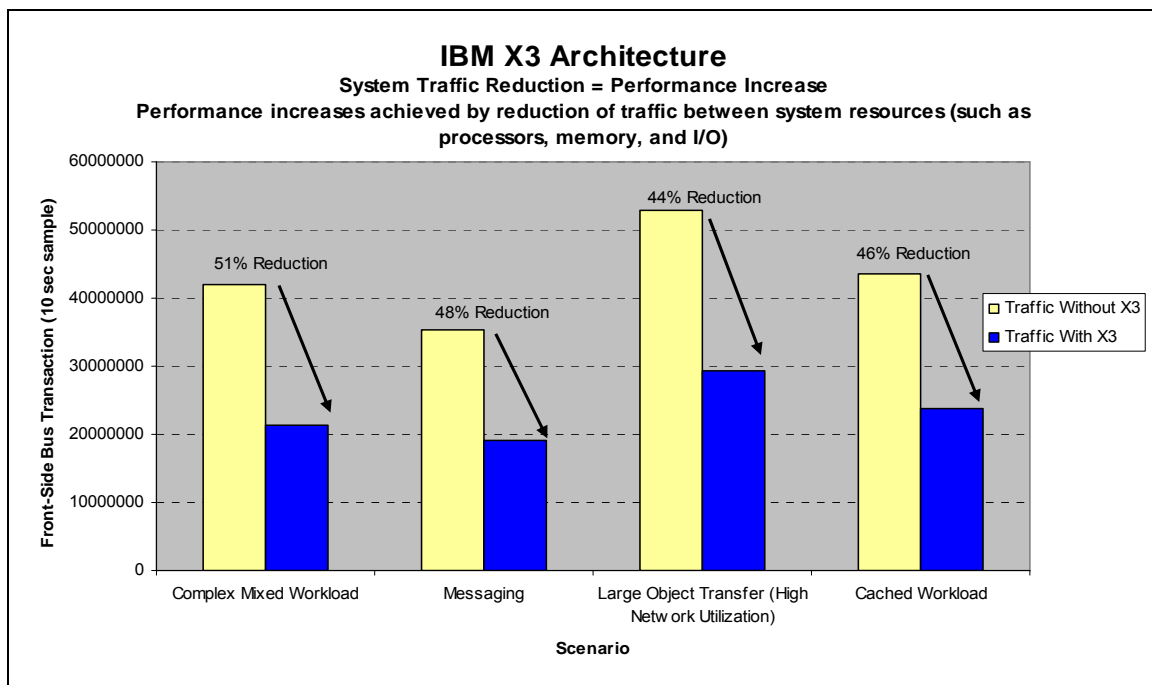


Figure 1: Snoop Filter Effectiveness

The X3 Architecture snoop filter uses intelligent caching to dramatically reduce system latencies to improve performance, as illustrated by the industry-leading benchmark results achieved by the x366 and x460.

Visit <http://www-03.ibm.com/servers/eserver/xseries/benchmarks> to view benchmark results.

Introduction

The ability of application servers to deliver high performance varies greatly, depending on the specifics of any given customer's workload. IBM raises the bar for scalable, high-performance application servers with the introduction of IBM eServer X3 Architecture—the ground-breaking third-generation of the IBM Enterprise X-Architecture XA-64e™ chipset. Incorporating this new chipset, the IBM eServer xSeries 366 and 460 have demonstrated unprecedented levels of performance and availability. X3 Architecture's performance advantage over the competition, given various typical application server functions, has been demonstrated by recently published record-setting benchmark results (<http://www-03.ibm.com/servers/eserver/xseries/benchmarks>).

Given a variety of typical application server functions, IBM eServer X3 Architecture offers great improvements, *in every instance*, in overall system performance through the use of features not available on any competitor's system. This paper will illustrate the benefits of IBM eServer X3 Architecture for typical application server functions, such as SSL encryption, messaging, Web service interactions (e.g., SOAP/XML serialization), HTTP services, and ODBC database interactions under various types of application server workloads. The paper presents the results of various test scenarios used to highlight the performance improvements that the reader can expect from systems based on X3 Architecture.

The term "application server" is applied broadly in the server world, but it typically pertains to a "middle-tier" server that receives requests from client systems, retrieves data from a separate database, and performs some data processing before returning a response. In this paper, the term "application server" refers to the commonly deployed, *managed application server environments*, such as Java™ application servers (e.g., IBM WebSphere®) and Microsoft® .NET application servers.¹ The performance characteristics described in this paper were obtained from statistics collected from a selection of Microsoft .NET 1.1 applications running on an IBM eServer xSeries 366.

This paper is organized in four sections:

- "Architecture and Concepts" presents an overview of the implementation of X3 Architecture in the x366 system.
- "Advantages of X3 Architecture" discusses the key role of the memory controller in achieving high performance.
- "Application Model and Test Scenarios" describes the application model used for all test scenarios and presents the results and analysis for each scenario.
- "Conclusion" summarizes the findings of the test scenarios.

¹ Applications that execute on the .NET managed environment include ASP.Net, C#, J#, and Visual Basic, and are often invoked via the IIS HTTP server as Web services or scripts.

Architecture and Concepts

IBM eServer xSeries 366 Architecture

The IBM eServer xSeries 366 was designed from the ground up to emphasize performance and manageability. It utilizes the IBM XA-64e chipset, which is a newly redesigned version of its predecessor chipset first introduced in the IBM eServer xSeries 440 in 2002.

Major features of the XA-64e chipset include:

- Support for latest Intel Xeon Processor MP with Extended Memory 64 Technology (EM64T)
- X3 Memory Controller - North Bridge and L4 Controller combined into a single high-performance controller
- Twin Front-Side buses
- High-performance DDR2 memory technology
- PCI-X support up to PCI-X-266MHz speeds

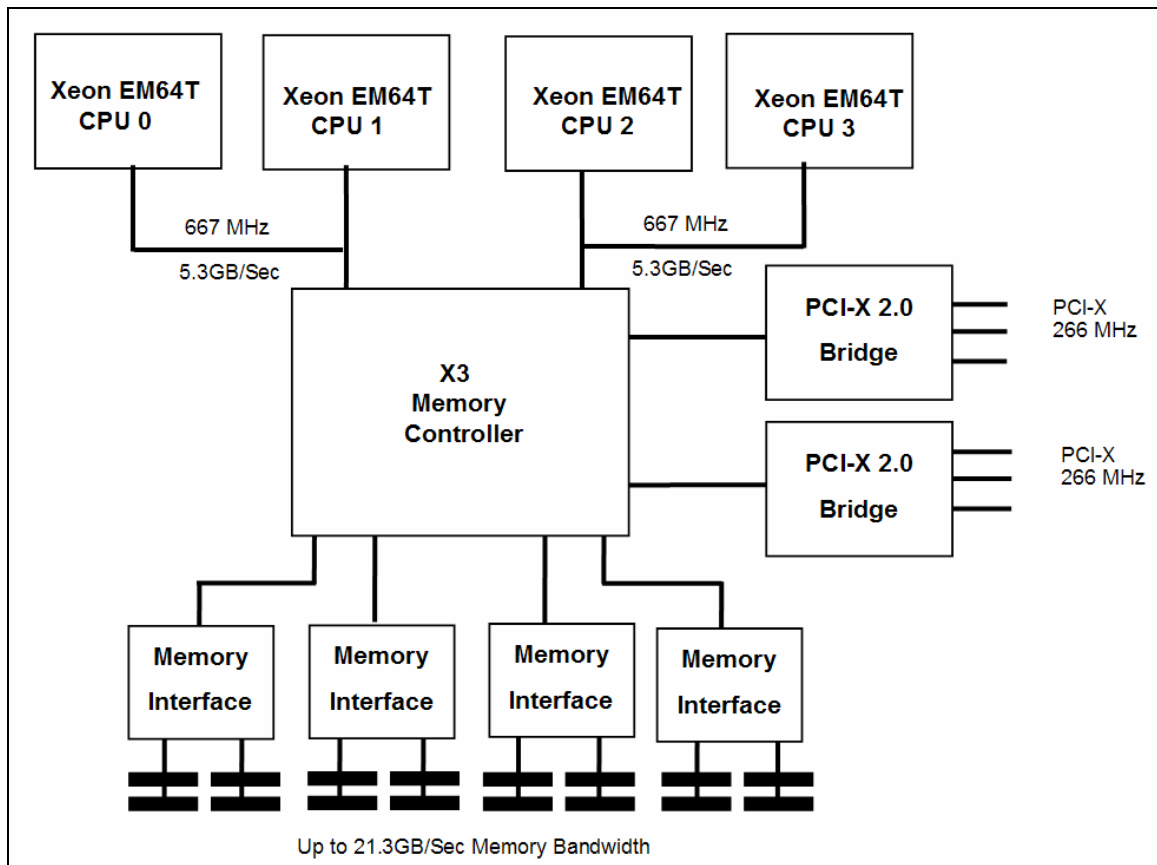


Figure 2: IBM eServer xSeries 366 Architecture

The x366 chassis accommodates up to four 64-bit Intel Xeon Processors MP and includes these major features:

- Enhanced microarchitecture
- 90 nm process technology
- HyperThreading Technology

- Support for up to four Intel Xeon Processors MP at 3.66GHz with 1MB L2 cache, or four Dual-Core Intel Xeon Processors 7040
- Front-Side Bus (FSB) enhancements that include a 667 MHz FSB on 133 MHz clock and speed that is 67% faster than current 4-way buses

To eliminate any disk performance bottlenecks for the tests described in this paper, the x366 included an IBM ServeRAID-6M adapter attached to an IBM EXP400 External SCSI Storage Enclosure. The 14 SCSI drives were partitioned into physically separate arrays to isolate functions such as the OS, HTTP log, messaging (MSMQ) log, and a storage for a collection of 100,000 JPEG images used in some of the tests.

The x366 is designed to be upgraded with the Dual-Core Intel Xeon Processor 7040, which is currently available. This enables investment protection for older versions of the x366.

For more information about the IBM eServer xSeries 366, visit IBM's Web site:

<http://www-1.ibm.com/servers/eserver/xseries/>

The advantages of IBM eServer X3

At the heart of IBM server X3 architecture is the memory controller, which eliminates significant traffic on the system front-side bus (FSB). This task is performed via a "snoop filter" that stores the memory addresses for data contained in every processor's L2 cache. To compare, for the competitor's architecture, all reads that are initiated by any processor (due to a miss in the processors cache) must check to see if that data is in any other processors' caches before it can obtain and use it from main memory. These cache "snoops" contribute to a significant amount of traffic on the FSB, since for each processor cache miss, the memory controller has to initiate transactions to and from all of the other processors. On a system under load, this level of traffic can have drastic effects on FSB transaction latencies due to the number of snoops performed, and the amount of queuing of these snoops that occurs on the memory controller.

Looking at the same situation using X3 Architecture, a processor cache miss will invoke a lookup into the memory controller's snoop filter while at the same time requesting the data from the main system memory. If the memory address is found in the snoop filter, the memory controller will request the data from the processor that contains the address. However, *more often than not*, the memory address *is not found* in the snoop filter, meaning no additional traffic is placed on the FSB, and the memory controller will use the data returned from the main system memory. This dramatically reduces the amount of traffic on the FSB, freeing the processors from excess snoop traffic, and greatly reduces transaction queue depth (and hence, the latencies) on a memory access.

The following application server examples illustrate several typical application server scenarios, and demonstrate the performance advantages and effectiveness of X3 Architecture. The x366 system was configured with four single-core 3.66GHz Intel Xeon Processors MP with a 1MB L2 cache. Hyperthreading and hardware prefetch were enabled for all of the tests.

Application Model

For all of the scenarios, all requests to the application server were invoked as Web service requests (HTTP/SOAP requests). The system tested was running Microsoft Internet Information Services (IIS) as the HTTP server, and used Microsoft .NET 1.1 Framework for the business logic. The business model used for the tests is that of an online retailer hosting B2B ordering, product catalog information, and customer information services. The application program interfaces (via ODBC and stored procedures) to a separate system running a Microsoft SQL Server 2000 database. For distributed transactions, the Microsoft Distributed Transaction Coordinator (DTC) is used on the database server system. The majority of distributed transactions were between the DBMS and the message queue service, Microsoft Message Queue (MSMQ).

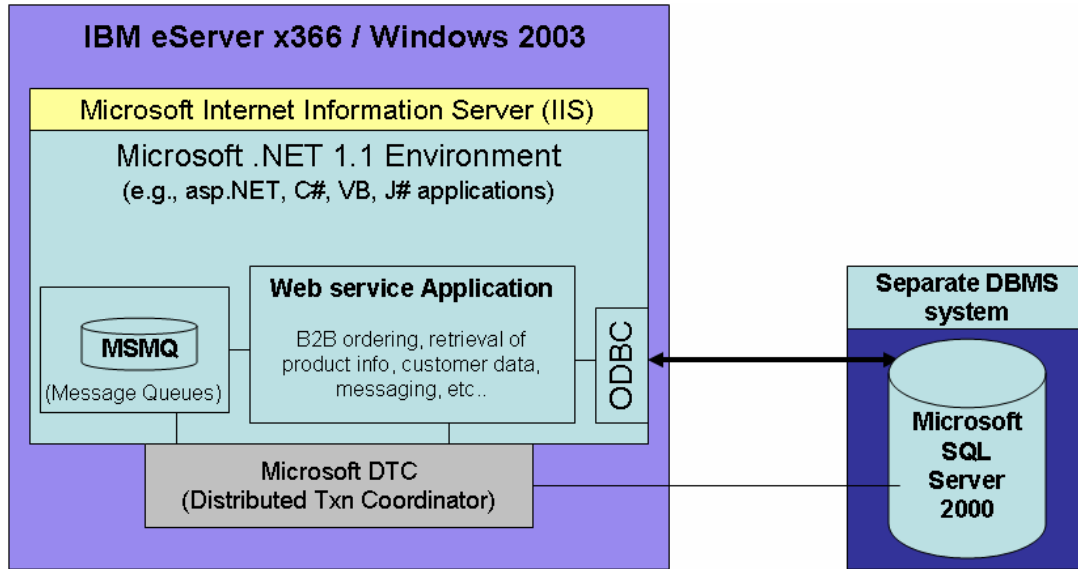


Figure 3: Application Environment Model

Scenario 1: Mixed Workload, Secure

For this scenario, a wide variety of application server functions was performed and invoked via secure Web services (SSL over HTTP over SOAP). Additionally two message queue applications simultaneously processed messages placed on a shipping warehouse queue and a stock management queue, as new queue messages arrived with each new order. Each shipping warehouse message invoked a distributed (2-phase commit) transaction between the message queue and the DBMS. The majority of the requests were for new orders, followed by product catalog queries, followed by customer information queries. The intent of this test was to show how X3 Architecture performs in an environment with complex and heavily computational business logic.

The x366 was driven to ~95% CPU utilization using 50 "business emulation" threads executing from a separate driver system. The x366 was able to process ~175 Web service requests per second.

As expected, the complexity and variety of the business logic for this example showed a fairly high processor L2 cache-miss ratio of 5.93%. Although 5.93% may not seem like much, small increases in L2 cache-miss rates result in a relatively large impact to the FSB, causing a large number of processor read requests to be handled by the X3 Architecture memory controller / snoop filter. **For this test scenario, the snoop filter eliminated 82% of these cross-FSB snoops.** By examining the X3 Architecture chipset trace-analysis, we know how much traffic occurred on the FSB, and we know how many snoops the snoop filter eliminated. By comparison of the two, **the snoop filter decreased overall FSB traffic by 51%.** This reduction in traffic results in greatly reduced transaction latencies and increases in overall system throughput.

Scenario 2: Messaging

For this scenario, the business logic includes only the processing of messages in a message queue. No HTTP server, SSL, or Web services are active in this test case. As each message is processed, the application invokes an ODBC stored procedure on a separate DBMS system, and invokes a Web service request to obtain a shipping label (e.g., FedEx or UPS label) from a separate system. The purpose of this example is to illustrate the characteristics of X3 Architecture for a system running “standalone” applications such a message queue process or other isolated applications that are not exposed as Web services.

During this test, the x366 processed ~210 messages per second, with the CPUs at ~97% utilization. **Again, the snoop filter was effective in this scenario, removing 79% of snoops from the front side bus.** Given the measured traffic compared to the number of FSB snoops that would have occurred without the filter, the **FSB traffic was reduced by ~48%, reducing queuing, and overall front-side-bus transaction latency.**

Scenario 3: Product Catalog, Data Retrieved from DBMS

For this scenario, the requests are of only one type, requesting the product catalog details for a given set of catalog items from a separate DBMS system. The details for each item include a JPEG image (average size ~20k). The purpose of this test scenario is to show the performance characteristics of an application workload retrieving a large amount of data from the DBMS and returning large amounts of data to the client systems.

The x366 was able to return ~280 Web service responses per second with the CPU utilization at 82%. The network traffic was relatively high in this case, utilizing 65% of the 1 gigabit network adapter's capacity. When examining the chipset performance for this test case, we find that the snoop filter is again effective, this time **reducing cross-bus snoops by 73%, which corresponded to a 44% decrease in FSB traffic.** **These reductions of cross bus snoops reduce front-side bus congestion and allow for more ‘real work’ to be completed by the processor.**

Scenario 4: Product Catalog, an “in-memory” Scenario

For this scenario, the requests are of only one type, requesting the product catalog details for a single item. The details for the item also include a JPEG image (as in Scenario 3). The entire working set size for the product catalog (including the images themselves) is approximately 2GB. To further decrease the complexity of this scenario, the Web service responses were cached in memory prior to the test. This means that for each Web service request, IIS simply checks its HTTP cache for the request, is guaranteed to find it in the IIS cache, and simply returns it to the requestor. This is a fully “in-memory” test case with a minimal amount of business logic, no ODBC requests, no SSL, no messaging, and so on. The purpose of this test scenario is to illustrate that even with very simple application logic, and a relatively high L2 processor cache-hit ratio, the X3 Architecture chipset still shows substantial and measurable performance advantages over the competitor's architecture.

Impressively, the x366 was able to return more than 1,500 Web service responses per second to the client systems in this case. As expected the network traffic was again quite high, utilizing ~50% of the gigabit network adapter's capacity. In this case, with a relatively low L2 processor cache-miss rate of 4%, **the snoop filter was able to eliminate 70% of the cross-bus snoop traffic, which effectively reduced the total FSB traffic by 46.81%.**

Conclusion

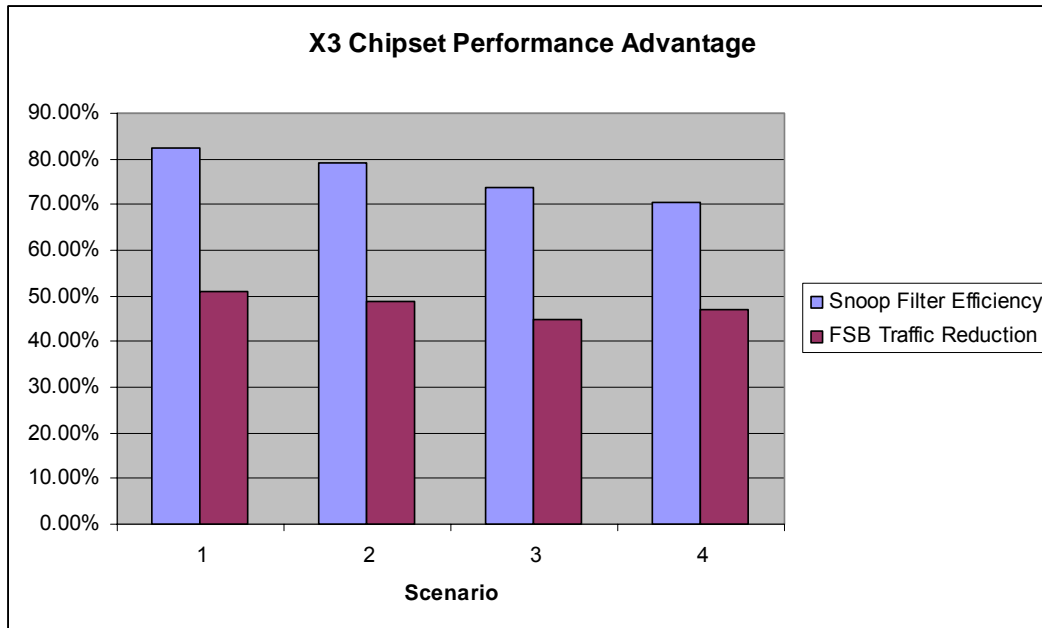


Figure 4: Snoop Filter Efficiency / FSB Traffic Reduction

For each application workload scenario examined, we find that X3 Architecture is effective at eliminating between 70% to 82% of processor front-side bus "snoops," resulting in a 44% to 51% decrease in total front-side bus traffic. The *total* FSB utilization for the observed workloads varied not only from workload to workload, but also varied with time within each workload substantially. For the four workload scenarios, the FSB was fully saturated 4.4% to 7.8% of the time. Without the snoop filter, the FSB would have been saturated (meaning no more "work" getting done) much more frequently, with a certain performance degradation impact.

With the introduction of the Dual-Core Intel Xeon Processor, the impact on performance resulting from the increased traffic on the front-side buses is dramatic. Since each processor core shares a front-side bus with the other core, and the potential number of processors effectively doubles, the cross-bus snoops would drastically decrease overall performance without the X3 Architecture snoop filter. However, as shown in the preceding examples, with IBM eServer X3 Architecture and the snoop filter in place, this otherwise severe increase in FSB traffic is greatly diminished, lowering latencies and queue depths, and boosting overall system performance. This illustrates one of the many reasons why IBM eServer xSeries servers, such as the x366 and x460, based on X3 Architecture, are excellent choices for application tier systems.



© IBM Corporation 2006

IBM Systems and Technology Group
Department 23UA
Research Triangle Park NC 27709

Produced in the USA.

01-06

All rights reserved.

IBM, the IBM logo, the eServer logo, eServer, xSeries, ServeRAID, X-Architecture, X3 Architecture, XpandOnDemand, Active Memory, Xcel4, Copper Diagnostics are trademarks or registered trademarks of IBM Corporation in the United States and/or other countries.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation.

Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

Performance is based on measurements using industry standard or IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve performance levels equivalent to those stated here.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.