IBM System x and Database Performance Analysis
June 2011

IBM System x3690 X5 with MAX5 and IBM solidDB v6.5:

Using the TATP Benchmark to Measure the Effect of Additional Memory Capacity on Database Performance

Dan Behman
Software Engineer Data Warehouse Performance
IBM Software Group

Joe Jakubowski
Senior Engineer System x Performance
IBM Systems and Technology Group

Ilkka Manner
Software Engineer for solidDB product family
IBM Software Group

Antoni Wolski
Chief Researcher for solidDB product family
IBM Software Group

Abstract

This paper presents the results of a study that examined how the additional memory capacity available to IBM System x® servers using IBM's eX5 server technology affects database performance with IBM solidDB v6.5. Measuring solidDB performance using a hybrid database on the x3690 X5 compared to an in-memory database using the x3690 X5 and the IBM MAX5 memory expansion system resulted in more than four times higher transaction throughput from the in-memory database than the hybrid database.

In the study, two server configurations were evaluated for the same size database. The first server configuration had insufficient physical memory to fully load the solidDB database into memory. Therefore, a hybrid database configuration with half of the database residing on disk was used. For the second server configuration, the IBM MAX5 memory expansion system was attached and the additional memory capacity was sufficient to completely load the solidDB database into memory. The IBM System x3690 X5 server's memory capacity can be doubled with the IBM MAX5 memory expansion system, which can extend the server's memory capacity to a maximum of 2TB when using 32GB DIMMs.

Optimizing server performance requires balancing the configuration of the primary server subsystems—CPU, memory and I/O—against the needs of the application and operating system. Application workloads such as Web, File, Database and Authentication/Security routinely operate against data sets that are many times larger than the physical memory installed on the server where the application is running. IBM internally published studies have demonstrated the performance benefit of increased memory capacity for these kinds of application workloads. The purpose of this study is to prove the benefit of additional memory capacity on solidDB database performance.

Introduction

The IBM MAX5 memory expansion system for System x is an IBM industry-first, memory-scaling technology that decouples memory from the processor, allowing expansion of memory capacity beyond industry-standard constraints. Memory-intensive workloads and application (e.g. Web, File, Database and Authentication/Security) and consolidation environments (e.g. virtualization and cloud computing) can take advantage of large amounts of memory to operate more efficiently.

To examine the effect of memory capacity on database performance, we used the Telecom Application Transaction Processing (TATP) Benchmark[1] to measure the throughput of an IBM solidDB database that was sized for two different scenarios:

- The solidDB database fit entirely in memory.
- The solidDB database did not fit entirely in memory, which required disk I/O to read portions of the database into memory as needed.

TATP measures the maximum Mean Qualified Throughput (MQTh) a server can sustain in terms of completed transactions per second (tps). TATP is designed for database performance analysis and system comparisons.

TATP Benchmark Workload[2]

The Telecommunication Application Transaction Processing (TATP) Benchmark is an open source workload designed specifically for high-throughput applications and is well suited for in-memory database performance analysis and system comparison.

The TATP benchmark simulates a typical Home Location Register (HLR) database used by a mobile carrier. The HLR is an application that mobile network operators use to store all relevant information about valid subscribers including the mobile phone number, the services to which they have subscribed, access privileges and the current location of the subscriber's handset. Every call to and from a mobile phone involves look-ups against the HLRs of both parties, making it a perfect example of a demanding, high-throughput environment where the workloads are pertinent to all applications requiring extreme speed: telecommunications, financial services, gaming, event processing and alerting, reservation systems, and so on.

The benchmark generates a flooding load on a database server. This means that the load is generated up to the maximum throughput point that the server can sustain. The load is composed of pre-defined transactions run against a specified target database.

The benchmark uses four tables and a set of seven transactions that may be combined in different mixes. The most typical mix is a combination of 80% database read transactions and 20% database modification transactions.

The implementation of the TATP load generator is database vendor agnostic. It has been successfully used with a variety of database systems.

The TATP distribution package consists of the TATP load generator source code, example test configuration files and SQL scripts that are useful to set up and use with the Test Input and Result Database (TIRDB). Some platform-specific packages include ready-to-run binaries as well.

[1] See http://tatpbenchmark.sourceforge.net.

[2] The TATP Benchmark description was copied from http://tatpbenchmark.sourceforge.net/.

Benchmark Bed Configuration

The benchmark bed used for this analysis consisted of an IBM System x3690 X5 server with the MAX5 memory expansion system as the system under test (SUT). The x3690 X5 server and the MAX5 memory expansion system were configured with 8GB memory DIMMs—a popular choice among today's clients because of the very competitive price per GB. The TATP benchmark drivers were loaded on the SUT, eliminating the need for external physical client driver machines. That corresponded to the current trend of vertical consolidation whereby both the applications and the database are run in a single system. Sufficient external disk storage was configured for database logging and checkpointing. The complete configuration is described below.

Server Hardware

- One IBM System x3690 X5
- Two Intel® Xeon® X7560 8-core processors
- Thirty-two 8GB Quad Rank x8 PC3-8500 CL7 ECC DDR3 1066MHz LP RDIMMs
- One QLogic QL2562 Fibre Channel host bus adapter (HBA)
- Two 146GB internal SAS disk drives (OS boot drives)

Storage Hardware

- One IBM System Storage® DS4800 controller
- Four RAID 0 LUNs each with sixteen 73GB 15Krpm 4Gbps Fibre Channel drives (64 total drives)
    - LUN 0 used for the database log
    - LUNs 1-3 used for the solidDB data tables

Memory Expansion

- One MAX5 memory expansion system
- Two Quick Path Interconnect (QPI) cables
- Thirty-two 8GB Quad Rank x8 PC3-8500 CL7 ECC DDR3 1066MHz LP RDIMMs

Software

- IBM solidDB v6.5 with Fixpack 5
- Red Hat Enterprise Linux® 6.0

The solidDB database was sized to fit in a memory capacity that was between the memory capacity supported by the x3690 X5 with and without the MAX5 memory system attached. This range of memory was 256GB without MAX5 and 512GB with MAX5. A database population of 230 million TATP benchmark subscribers was loaded. The resulting database size required roughly 400GB of memory to fit entirely in memory. The experiment where the database was divided between the disk subsystem and memory was possible because solidDB is equipped with two database engines: one maintains in-memory tables (M-tables) and the other maintains on-disk tables (D-tables). The following terminology was used to describe the two different database use cases:

> **In-memory database** – A TATP benchmark database built with 230 million subscribers that was exercised on a server *with* sufficient memory capacity so that the database was loaded entirely into M-tables. The database size was approximately 400GB.
>
> **Hybrid database** - A TATP benchmark database built with 230 million subscribers that was exercised on a server *without* sufficient memory capacity to load the database entirely into memory. Half of the database was loaded into M-tables and the other half into D-tables. The size of the page buffer pool of the disk-based database engine was set to 16GB. A properly sized Fibre Channel disk subsystem was used by the hybrid database for low-latency disk I/O access to database tables on disk. The database size was approximately 400GB.

The TATP benchmark was configured for 80% database table read transactions and 20% database table modification transactions (writes or updates to database tables). To achieve the highest performance levels, the transaction durability was set to "relaxed," which resulted in asynchronous log writing to disk. To reflect an actual production configuration, periodic database checkpoints were executed during the test runs.

Results

Results were obtained using the TATP benchmark to exercise the exact same solidDB database size on the x3690 X5 with and without the MAX5 memory system attached. The maximum Mean Qualified Throughput (MQTh) was measured and expressed in completed transactions per second (tps). The 90th percentile transaction response times were captured for each of the seven TATP benchmark transaction types.

Figure 1 plots the MQTh relative to number of TATP benchmark client drivers. A higher number of drivers represents a higher user load applied to the solidDB database. The peak MQTh was achieved at 32 client drivers. Figure 1 shows that the in-memory database achieved more than four times higher transaction throughput than the hybrid database when the user load was 32 client drivers.

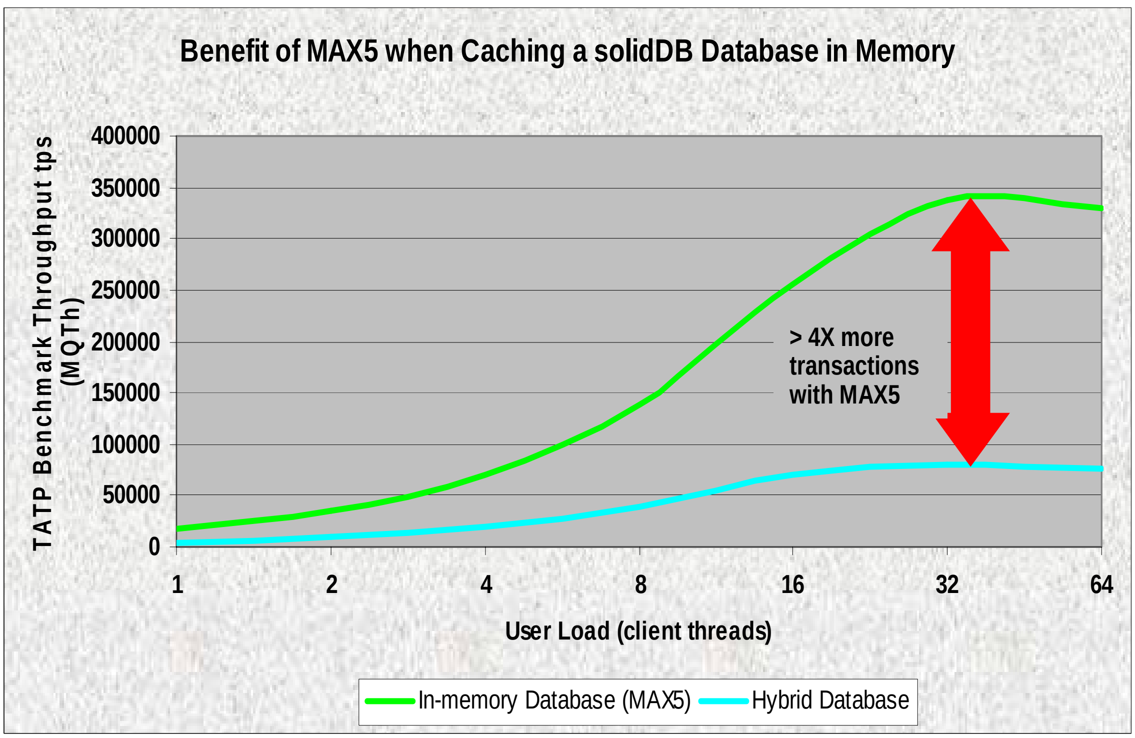

Figure 1 – Performance comparison between an in-memory and a hybrid database at varying user loads

Figure 2 provides more detail on the performance difference between the in-memory and hybrid databases at two different user loads. At a lighter load with 4 client drivers, the performance advantage of the in-memory database over the hybrid database is approximately 3.6 times. At a higher user load where the MQTh maximum was reached, the performance advantage of the in-memory database over the hybrid database is approximately 4.2 times. IBM solidDB achieves a consistent performance gain when run fully in memory when compared to a hybrid database that must use disk I/O to retrieve database table content that does not reside in memory.

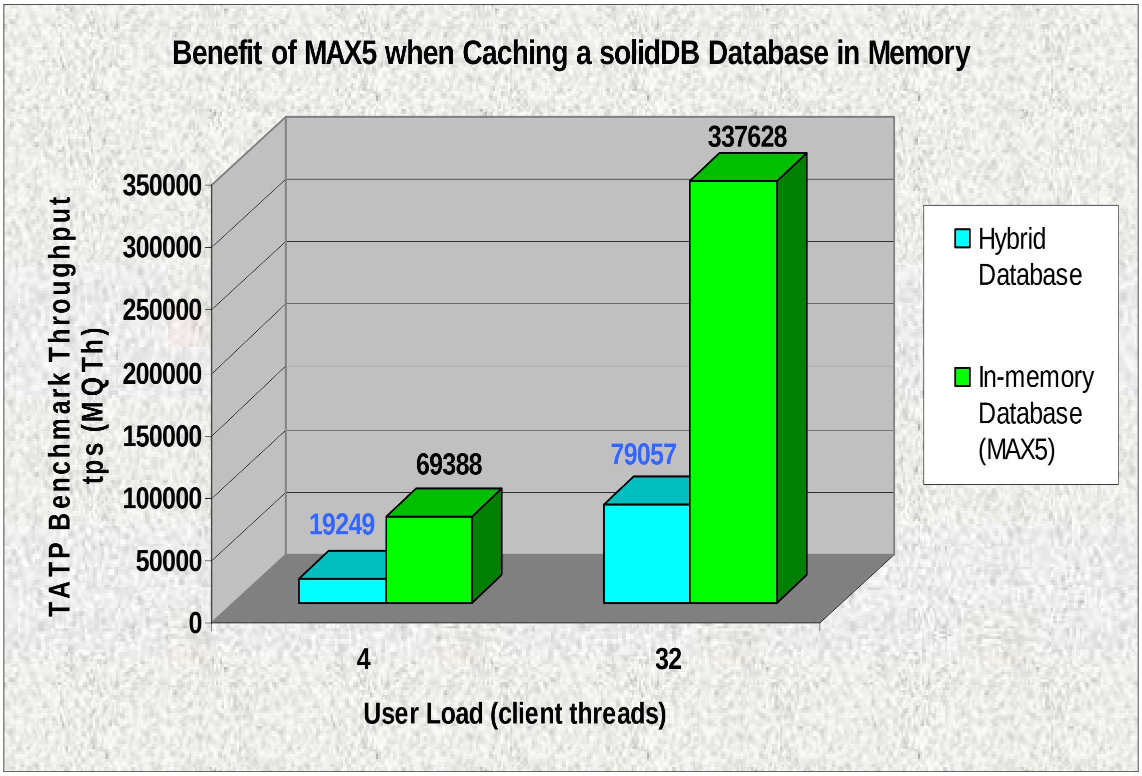

Figure 2 – Performance comparison between an in-memory and a hybrid database at minimum and maximum user loads

Figure 3 shows a significantly lower price per transaction under light load with 4 client drivers and at high load with 32 client drivers when MAX5 was used. The in-memory database with MAX5 achieved $1.29 per transaction compared to $5.26 per transaction for the hybrid database with 32 client drivers.

Parts	**Hybrid Database**	**In-memory Database with MAX5**
IBM System x3690 X5 with one Intel Xeon X7560 processor, Redundant 675W power supply and kit, 16-DIMM Internal Memory Expansion, 2 x 146GB 2.5" SAS HDDs	$12,476	$12,476
Second Intel Xeon X7560 processor	$4,999	$4,999
MAX5 memory expansion system	N/A	$7,495
MAX5 cable kit	N/A	$1,299
32 x 8GB Memory	$10,848	$10,848
32 x 8GB Memory for MAX5	N/A	$10,848
QLogic QL2562 HBA (P/N 42D0510)	$1,999	$1,999
IBM System Storage DS5300 expansion unit (open bay) [1]	$80,000	$80,000
4 x EXP5000 model D1A 16-Pak 146.8 GB/15K DDM [2]	$81,160	$81,160
IBM solidDB 6.5 with FixPack 5	$224,000	$224,000
Total Price	$415,482	$435,124
$ /Transaction (4 TATP client drivers)	$21.58	$6.27
$ /Transaction (32 TATP client drivers)	$5.26	$1.29

Figure 3 – Price per transaction comparison between an in-memory and a hybrid database
[1] – DS4800 controller is no longer available. Pricing is based on the currently available model.
[2] – 73GB Fibre Channel HDDs are no longer available. Pricing is based on the currently available 146GB Fibre Channel HDDs.

In addition to the higher transaction throughput measured on the solidDB in-memory database, the TATP benchmark transaction response times were also significantly lower when compared to those measured with the hybrid database. Figure 4 shows lower transaction response times for all transaction types measured on the in-memory database relative to the hybrid database.

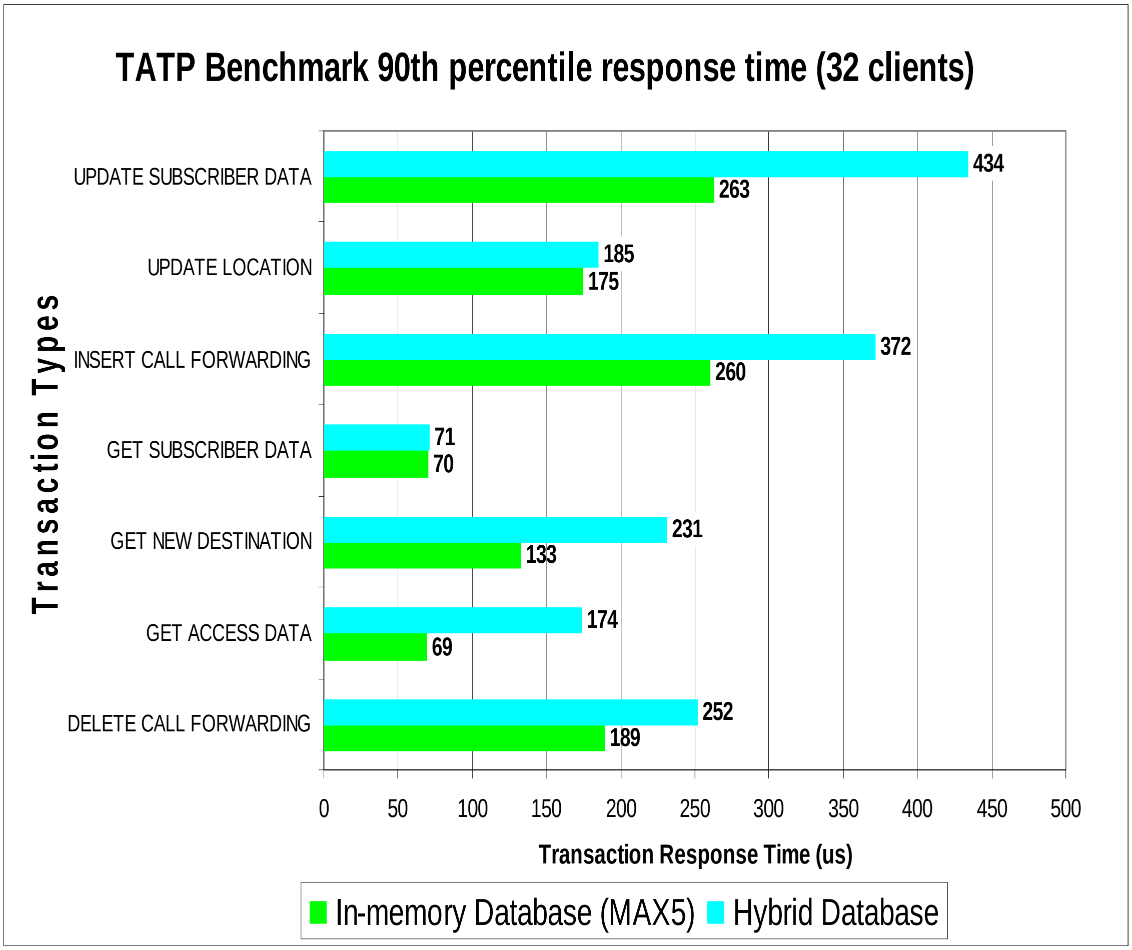

Figure 4 – TATP benchmark transaction response time comparison between an in-memory database and a hybrid database

Conclusion

Extending server memory capacity on eX5 technology-based IBM System x servers with the IBM MAX5 memory expansion system increases the performance of many application workloads, including Web, File, Database and Authentication/Security. The same benefit can be achieved in other memory-intensive consolidation environments such as virtualization and cloud computing. This paper focused on the increased performance that can be achieved for database environments when sufficient memory capacity is provided to fully load the entire database in memory. IBM’s solidDB v6.5 database achieved up to four times higher transaction throughput at measurably lower transaction response times when the database was loaded entirely in memory compared to a hybrid database that had to retrieve some database table content from disk. The overall price performance of the hardware/software configuration was significantly advantaged for the in-memory database compared to the hybrid database. At maximum TATP benchmark throughput, the price/performance of the in-memory database configuration was less than one fourth that of the hybrid database configuration.

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

Produced in the USA
June 2011

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

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.

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.

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.