

# Benchmark Performance Results for Pervasive PSQL v11

---

A Pervasive PSQL White Paper  
September 2010

# Table of Contents

- Executive Summary..... 3
- Impact Of New Hardware Architecture On Applications..... 3
- The Design of Pervasive PSQL v11 ..... 4
- Configurations..... 6
  - Memory Caches ..... 7
- Benchmark Testing..... 7
  - Atomics Testing..... 8
- Conclusion..... 8

## Executive Summary

Advances in hardware technology previously meant advances in processing speed. Existing applications just ran faster with no extra effort to enhance them. Today, advances in computing technology mean increased parallelism and not increased clock speeds. Multiple cores provide the avenue for increased parallelism, but applications must be written with parallelized code to take advantage of multiple cores simultaneously. Pervasive PSQL v11 Server is specifically designed to increase scalability and performance on multi-core machines in multi-client environments.

This paper discusses the performance results from benchmark tests on Pervasive PSQL v11 Server. Test comparisons are shown between Pervasive PSQL v11 and the latest release in the 10.x family, Pervasive PSQL v10 SP3.

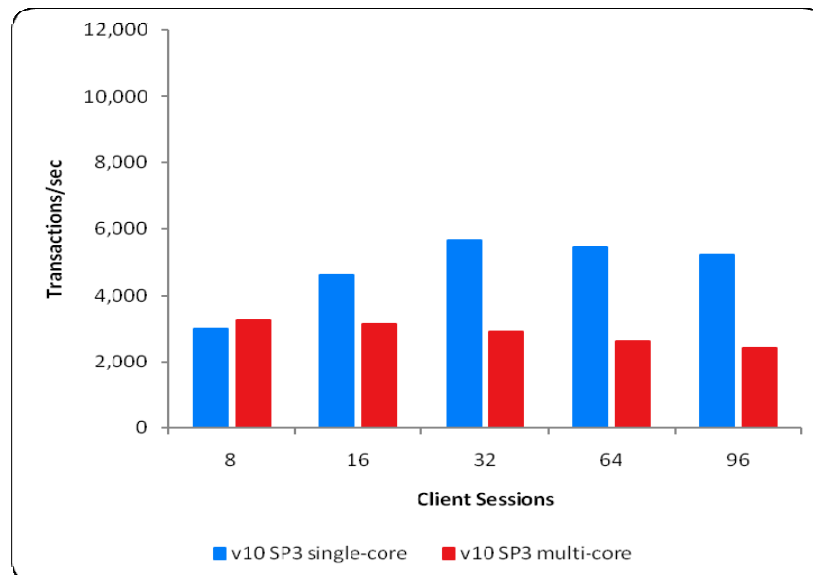
On multi-core machines with multiple clients accessing data, the performance of Pervasive PSQL v11 significantly exceeds that of Pervasive PSQL v10 SP3, sometimes over 300%.

## Impact Of New Hardware Architecture On Applications

Multi-core machines are now the norm. When you upgrade your single-core machines, your applications will have to function with a different hardware architecture. And for many applications, the multi-core architecture dramatically changes how the application performs.

As we explore this deeper, consider the following performance trend when Pervasive PSQL v10 Server moves from a single-core machine to a multi-core machine.

**Figure 1. Performance of Pervasive PSQL v10 Server on Single-core and Multi-core Machines**



As you can see, performance decreases with more client sessions. This behavior is not limited just to Pervasive PSQl v10 Server. *Any* complex multi-client application not specifically designed to take advantage of multi-core architecture—and many applications originated during the single-core era—will likely experience performance drops on multi-core machines.

Why? The technical reasons are complex and thoroughly explained in the white paper [\*The Multi-core Dilemma\*](#) by Dan Woods, CTO of CITO Research. But a brief summary here will clarify why it may seem counterintuitive that multithreaded applications would run slower on multi-core systems.

In a multithreaded application that shares data, the synchronization between threads consumes a lot of system resources on a multi-core machine. In fact, shared data is the bane of parallel computing. If multiple threads access the same data, access must be synchronized. And when you synchronize access to the data, that portion of code cannot be executed by more than one thread and therefore it is not concurrent. This section of the code then becomes the single door that everyone must line up to go through.

Caching does not improve this problem; it exacerbates it. If multiple cores or processors have caches that point to the same data and one core modifies the data, the cached data on the other core is no longer valid, and the caches must be synchronized. The accumulated overhead for all the synchronization of these operations is significant; ultimately it means that performance on multiple cores can be worse than on a single core machine that does not require such synchronization.

Where possible, each core should work on separate data; otherwise, there is overhead associated with synchronization and that overhead can slow down performance significantly.

## **The Design of Pervasive PSQl v11**

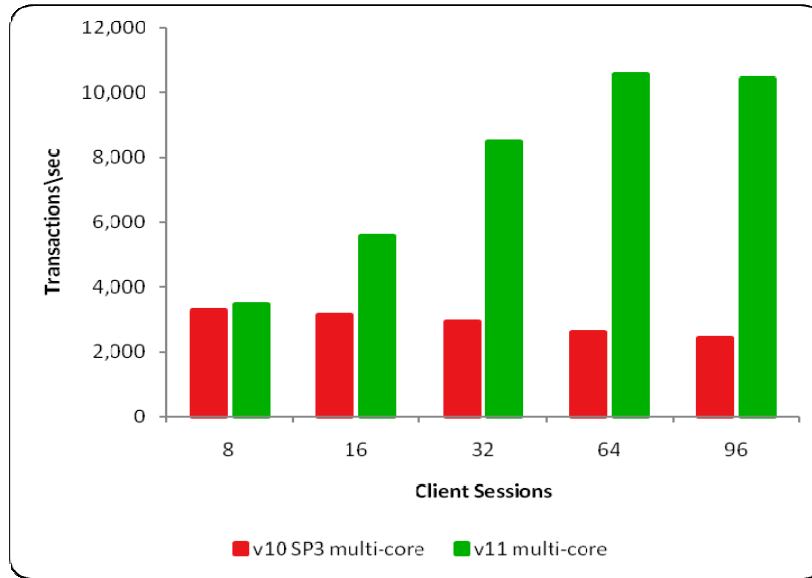
Pervasive PSQl v11 has been architected to provide parallel threads performing similar activities. The gains in increased parallel processing improve the throughput to the point that multiple processors are engaged. The result is that performance of the database engine increases in multi-core environments with multiple clients accessing a central server.

Pervasive PSQl v11 also provides enhancements to the low-level synchronizations mechanisms in the transactional interface. Multiple users can read the same cached file pages simultaneously and their operations can proceed on independent server CPUs. Non-user activity such as checkpoints and log management can also use additional server CPUs.

The scalability of Pervasive PSQL v11 has been enhanced through architecture designs made specifically for multi-core hardware. For example, multiple users accessing independent files can proceed on independent server CPUs. The database engine can also handle higher user loads with less overhead, resulting in steadier throughput.

How do those design changes improve performance? Let's compare the performance of Pervasive PSQL v11 on multi-core machines with that of Pervasive PSQL v10. Just as in Figure 1, the results are based on the transactional interface accessing 16 files fully cached in memory.

**Figure 2. Performance of Pervasive PSQL v11 Server and Pervasive PSQL v10 Server on Multi-core Machines**



As you can see, the architecture changes in Pervasive PSQL v11 provide a significant increase in performance. And the best news: your multi-client application can benefit from this increased performance without requiring you to recompile or rearchitect your code.

## Configurations

The performance tests discussed in this paper were conducted using the configurations described in the following table.

**Table 1. Configurations Used for Benchmark Tests**

Server Machine Processors	Dual Intel Xeon CPU E5420 @ 2.50 GHz, 4 Cores
Server Machine Total Cores	8*
Server Machine Total Memory	16 GB
Server Machine Operating System	Microsoft 2008 Enterprise Server (Service Pack 2), 64-bit
Server Software	Pervasive PSQL v11, Pervasive PSQL v10 SP3
Pervasive PSQL Server Settings	Installed with defaults
Client Machine Processors	Single Intel Core2 Quad CPU Q9400 @ 2.66 GHz, 4 Core (Two physical machines running against Pervasive PSQL Server with client sessions split evenly between the client machines.)
Client Machine Total Cores	4 (each machine)*
Client Machine Total Memory	4 GB (each machine)
Client Machine Operating System	Microsoft XP (Service Pack 3) 32-bit
Client Software	Pervasive PSQL v11 Client to Pervasive PSQL v11 Server Pervasive PSQL v10 SP3 Client to Pervasive PSQL v10 SP3 Server
Pervasive PSQL Client Settings	Installed with defaults
* Since commodity hardware today is typically 4 or 8 core machines, Pervasive Software chose such machines as configurations realistic to what our customers will experience. As the number of cores increases in commodity machines, Pervasive PSQL will	

continue to be optimized based on such changes.

## Memory Caches

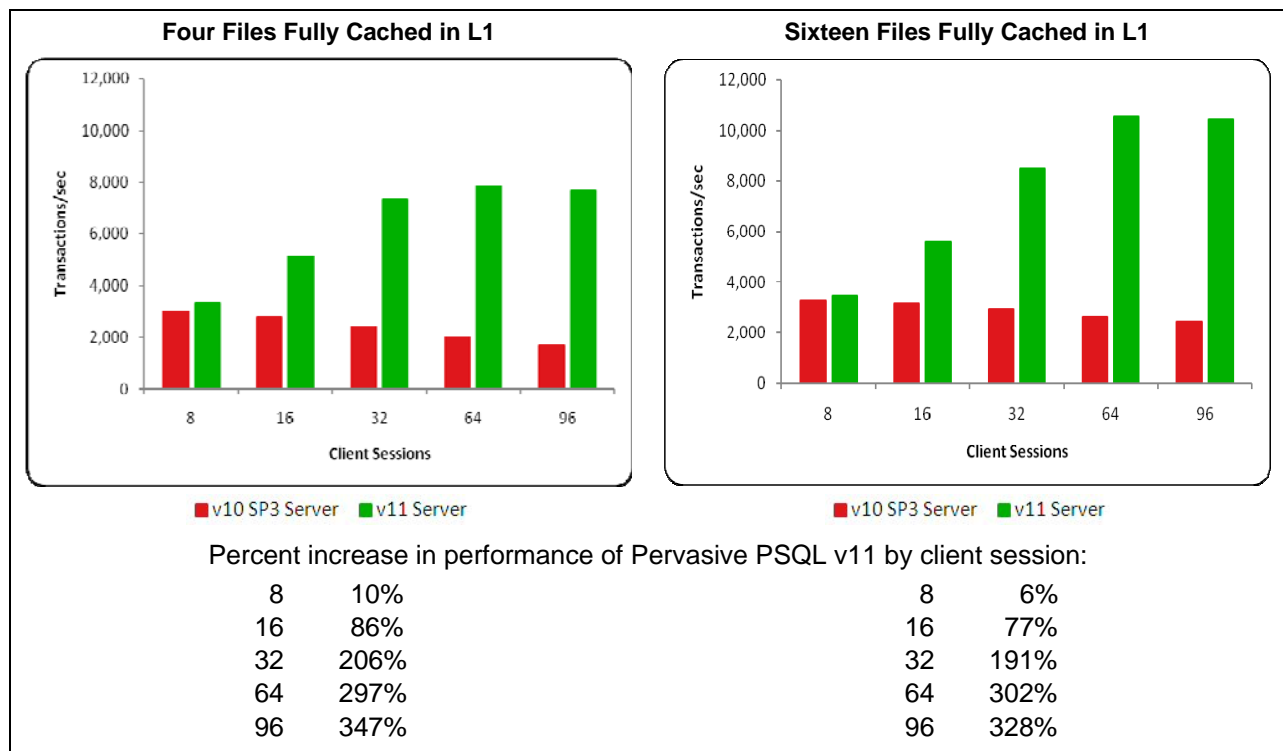
Pervasive PSQL contains two primary caches, Level1 and Level2, termed L1 and L2 respectively. The L1 cache is a fixed size based on the Cache Allocation Size setting. It does not expand or contract based on database operations. The benchmark tests show results for data fully cached in L1.

## Benchmark Testing


To demonstrate the effect of different client session loads, Pervasive Software ran a test harness with a standard installation of Pervasive PSQL on the server and on the client (see "Configuration" section above). The test harness simulates a processing load similar to the industry standard TPC-B benchmark. A single transaction performs a READ and UPDATE on a record in each of three files, then performs an INSERT of a record to a fourth file.

For details about TCP-B from the Transaction Processing Performance Council, see <http://www.tpc.org/tpcb/default.asp>.

**Figure 3. Performance Results for Transactional Interface (Btrieve)**



**Figure 4. Performance Results for Relational Interface (ODBC)**



## Atomics Testing

In addition to the TCP-B benchmark tests, atomics testing of READ and UPDATE actions were performed. The atomics results on multi-core machines are similar to the TCP-B results. They show that Pervasive PSQL v11 outperforms Pervasive PSQL v10 SP3. If you would like to receive the results of the atomics tests, e-mail [database@pervasive.com](mailto:database@pervasive.com) or telephone 1 800 287 4383.

## Conclusion

The numbers speak for themselves. Pervasive PSQL v11 Server demonstrates increased performance over Pervasive PSQL v10 Server on multi-core machines. The majority of increases are over 100%, with some exceeding 300%.

The more data files involved, the better the performance as the increased parallel processing in Pervasive PSQL v11 engages multiple processors. Multiple users can access independent files as well as read the same cached file pages simultaneously while such operations proceed on independent server CPUs.

In contrast, the performance of Pervasive PSQL v10 Server decreases on multi-core machines as client sessions increase. Any complex multiuser application that originated during the single-core era will likely experience performance drops on multi-core machines. The numerous

design considerations required for optimal performance on multi-core hardware simply could not have been anticipated at the time.

Multi-core support is a primary feature of Pervasive PSQL v11 because the new multi-core hardware dramatically changes how most applications perform. That support is of primary importance to you as you transition your multi-client applications into multi-core environments.