



SYBASE for IBM RS/6000 SMP Servers

By Josh Bersin

This article discusses Sybase's SQL Server 10 architecture for the IBM RISC System/6000 PowerPC-based SMP servers. It describes the scalable high-performance SYBASE architecture for AIX and how it can be used in the development of large-scale, mission-critical client-server applications.

SYBASE SQL Server 10 is the latest generation of Sybase's client-server Relational Database Management System (RDBMS) servers. SQL Server has been available since 1987, and through a joint development and marketing relationship with Microsoft®, it has been widely adopted in the industry on over 150,000 RISC and PC server systems. Symmetric Multiprocessor (SMP) support became available with Release 4.8, which shipped in 1990.

Sybase's view of the client-server market is still unique in the industry. The key to high performance client-server applications is the ability to deliver multithreaded high performance in the server with advanced capabilities for transaction integrity. Stored procedures, triggers, and database Remote Procedure Calls (RPCs) have made this possible.

Stored procedures and triggers are server-based applications that encompass shared business logic similar to CICS™ transactions. They are compiled and resident in shared memory for maximum performance. They implement SQL and other functions, and protect access to the data from across the network.

In the SYBASE architecture, all database users are viewed as clients. Clients can execute SQL or RPCs to access data directly or to invoke stored procedures. Clients can run on PCs, UNIX workstations, Macintoshes®, or even IBM mainframes.

This powerful architecture, coupled with Sybase's unique multithreaded RDBMS server, has

created an entire development environment for new applications. The RDBMS environment offers much more than just relational access—it offers scheduling, multithreading, memory management, online backup, built-in networking, and a variety of security functions. On a system such as the RISC System/6000 (RS/6000) PowerPC SMP, SQL Server provides a total client-server programming and deployment environment.

SMP Hardware Architectures

SMP systems implement multiple processors sharing a single memory. To deliver a single systems image, they have *cache coherence* features incorporated—ensuring that one processor does not write over the memory being used by another. The key challenge in the design of these systems is scalability. The hardware and operating system must be designed to minimize contention and locking as additional processors are added. Most UNIX SMP systems today scale well up to four processors, but then tail off significantly after that, as shown in Figure 1.

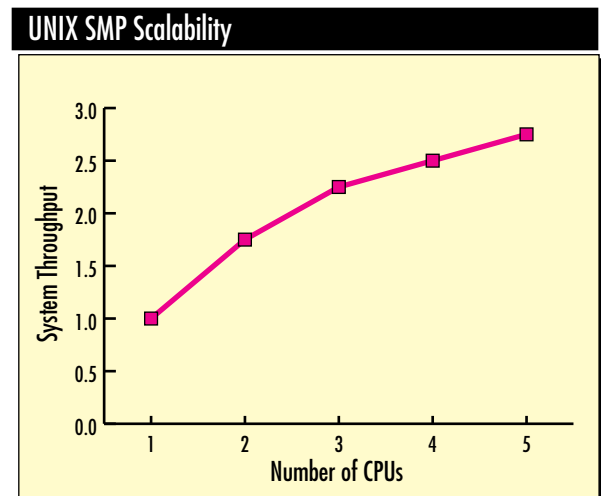


Figure 1. Typical UNIX SMP scalability—operating system throughput



Josh Bersin

Most measurements of scalability measure operating system-level tasks. The key features in client-server applications are scalability of transactions-per-second and number of users, while maintaining response time at a subsecond level. Sybase, realizing the demands of SMP architectures, specifically designed an implementation of SQL Server for optimum scalability in the Online Transaction Processing (OLTP) environment.

SYBASE Architecture

The key innovation in SQL Server is the single-process multithreaded architecture (see Figure 2). SQL Server uses a single database process per CPU with an optimized thread model within the process. The SQL Server itself handles network connections, scheduling, and memory management within the database process.

Scalability

This single-process multithreaded architecture results in an RDBMS server that scales to thousands of concurrent users on RS/6000 systems. As workload is added, SQL Server achieves uniform subsecond response time with almost no degradation.

Alternative database architectures implement multiple processes, producing much higher system overhead and greater resource utilization.

The net result is a performance characteristic that provides predictable, scalable throughput, and subsecond response time as user load is added, shown in Figure 3.

High-Performance OLTP on PowerPC SMP

Based on benchmarks being performed at the time this article was written, SYBASE is already achieving very high performance and throughput on the new IBM SMP systems (Model J30 4-Way SMP with TPC-C transactions), as shown in Figure 4.

Support for Many Online Users

SYBASE's architecture, coupled with compiled and memory-resident stored procedures, allows very large numbers of users to be supported with predictable subsecond response time. A key feature making this possible is Sybase's internal threading architecture. The SQL Server's internal threads allow individual users to run within the RDBMS process, using less than 60 KB of memory per user.

The SQL Server can support hundreds to thousands of users in an online environment on the RS/6000 SMP products. Based on expected performance of the PowerPC family, we have esti-

Multithreaded Architecture Versus Traditional RDBMS Design




Figure 2. SYBASE multithreaded architecture versus traditional RDBMS design on uniprocessors and SMP

SYBASE Performance Characteristics




Figure 3. SYBASE performance characteristics on RS/6000 PowerPC SMP

TPC-C Workload




Figure 4. RS/6000 with SYBASE System 10 running TPC-C workload




Figure 5. Estimated online user capacity RS/6000 SMP servers running the TPC-C workload (1.5 sec response time)

SYBASE System 10 Backup Server




Figure 6. SYBASE System 10 Backup Server architecture

estimated the online performance of SQL Server on the PowerPC to scale to thousands of users, as shown in Figure 5.

Support for Very Large Databases (VLDBs)

One major problem facing customers as their applications grow is the support of large databases (tens to hundreds of gigabytes online). The critical issues to support such databases include high-speed online backup, support for mixed workloads (OLTP and queries), and complex queries that access large amounts of data.

Over the last ten years, Sybase has developed features and capabilities to support VLDBs. Most

important is the SYBASE Backup Server, which is shipped with every copy of SQL Server. The SYBASE Backup Server is a separate UNIX process that can run on any SMP processor. It provides online parallel database backup to a variety of tape devices. Backup Server runs while the SQL Server is operating, backing up over 90% of the transactions that execute during the backup process. Backup Server has been benchmarked on SMP systems backing up online databases at over 40 GB per hour. Figure 6 shows the Backup Server architecture.

Support for Complex Queries and Mixed Workloads

As databases grow larger, it becomes increasingly important to support mixed workload: high-volume OLTP and growing loads of queries and batch reports. This environment requires the RDBMS to interact efficiently with the operating system so the SMP capabilities of the system can be efficiently utilized to balance the workload. SQL Server's Virtual Server Architecture (VSA) efficiently accommodates this requirement.

Sybase Interactive Query Accelerator

A different strategy is often needed for complex queries. When a database reaches tens of gigabytes in size, random queries can often monopolize an entire CPU and I/O channel for extended lengths of time. This naturally contends with system needs for online applications.

Sybase recently announced the availability of the SYBASE Interactive Query Accelerator (Expressway Technologies), which provides a fully indexed copy of the SQL Server database for direct query access at very high speeds. The Query Accelerator allows queries to access indexes against the database on any field, regardless of query, to minimize both CPU and disk access for queries. It exploits the large memory and SMP capacity of the PowerPC server to maximize throughput—and provides speed-up from 400 to more than 1,000 times normal for SQL queries against the database.

SQL Server VSA for SMP

In the RS/6000 SMP, the SQL Server VSA enables the customer to load multiple servers (up to the number of processors) that function as a single server. Each process can send threads to any CPU, allowing the system to maximize use of the computing resources. Since the customer can allocate as many database engines as needed, the

system can be partitioned between RDBMS processing and other system use.

This architecture has many advantages:

- ◆ Maximizes use of the system by dispatching low-overhead threads to any available CPU.
- ◆ Allows customers to tune the system. If only two CPUs are needed for SQL Server, customers only create two engines. As more are needed, the system can be reconfigured to add capacity.
- ◆ Allows SQL Server to efficiently manage mixed workloads of OLTP and queries by dispatching tasks to separate CPUs and minimizing contention.

Compared to other RDBMS architectures, only a minimum number of AIX processes are used, allowing excellent performance and scalability.

Managing SYBASE System 10 on the IBM SMP

In the SYBASE SQL Server VSA architecture (see Figure 7), the system can be tailored for any particular workload. The client systems and administrator see the VSA as a single SQL Server accessing a single database. When the SQL Server is brought up, the administrator determines the number of engines to be loaded. Each engine is an SQL Server process, sharing single thread and run queues.

There should be no more engines than CPUs, and the number of engines should be increased until optimum performance is reached. The engines dispatch client tasks to their appropriate processors using a single queue and shared memory, presenting a single virtual server. As additional engines are added, system performance will improve—increasing demand on I/O and other system resources. As the number of engines is increased, the system should have enough balanced disk to handle the increased workload.

In this symmetric architecture, each engine can run on any CPU. If a CPU is busy, available tasks will be dispatched to any available CPU. If the number of SQL Server engines is less than the total number of CPUs (usually recommended), there will always be available CPUs for additional applications on the system.

SQL Server Task Management in SMP

The SYBASE SQL Server documentation describes the actual task management in the SMP. Here is a simple description of the process:

1. A client application issues a login request. In response, SQL Server creates a user task to handle work from the client.

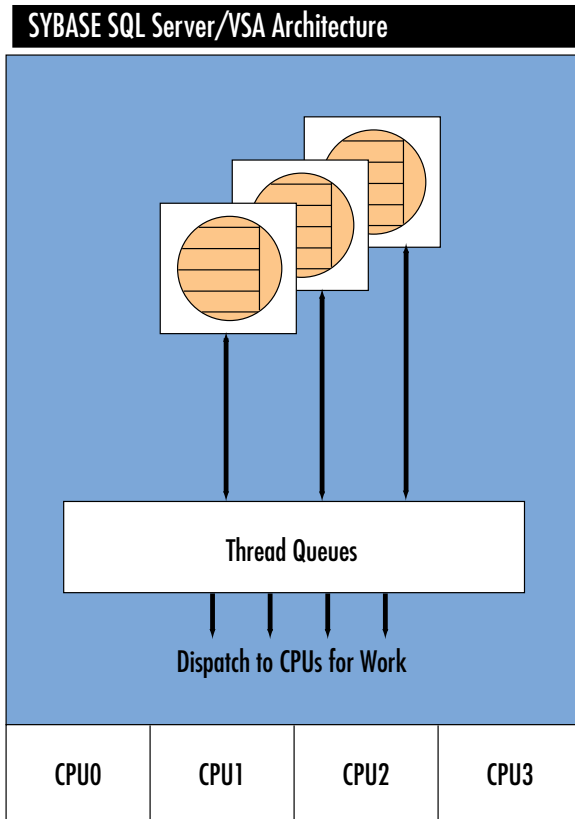


Figure 7. SYBASE SQL Server VSA architecture

2. The client presents SQL Server with work to do—a series of Transact-SQL commands.

3. SQL Server adds the client's user task to the runnable task queue. The server engines compete for the user task at the head of the task queue.

4. The engine executes each step until the task completes or blocks, while waiting for I/O or locking. When the task blocks, it yields the server engine to run other user tasks. Once the block is resolved, the user task is again added to the runnable task queue.

5. After the task blocks for the last time, it continues executing until it finishes. Then the user task yields the server engine and moves to the sleeping task queue until the client presents the server with more work.



Josh Bersin, Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608. Internet: joshua.bersin@sybase.com. Mr. Bersin manages Sybase's product and marketing strategy for the IBM product line. He has a BS from Cornell University, an MS from Stanford University, and an MBA from the University of California at Berkeley.