Microsoft SQL Server 7.0 and Macola Progression ERP- A Combination You'll Never Outgrow

White Paper

MACOLA®
SOFTWARE

Additional White Paper topics
are available. Please contact
Macola for a current list.

For more information on
Macola Software Products
and services please call
800-468-0834
or e-mail
response@macola.com
additional information is
available on the World Wide Web:
http://www.macola.com

Written by David A. Turbide
CFPIM, CMfgE, CIRM
Editor, *Midrange ERP* Magazine and *APS* Magazine

Table of Contents

INTRODUCTION

Choosing the right technology is a critical factor for manufacturers and distributors who prefer to focus on running their businesses rather than running information systems. Software developers consider it their responsibility to look ahead to technology advancements that will provide better value and benefit to their customers.

Macola has chosen the high-performance technology strategy of Microsoft's Windows NT/2000 and SQL Server to deliver its high-functionality ERP business solution. Macola's Progression SQL Server series leverages Microsoft's leading Windows database to bring business advantage and improved decision making to all levels of your organization.

MICROSOFT SQL SERVER 7.0

MORE THAN A NEW RELEASE

On November 16, 1998, Microsoft released version 7.0 of the SQL Server database. This is the latest and most significant upgrade to what is rapidly becoming the de facto industry standard for the small to medium-sized enterprise.

RDBMS - Relational Database Management System - quite simply, software that manages the information stored in the computer. An RDBMS follows strict rules and conventions to ensure data integrity and simplify access for both end users and programs.

Scalability - The ability to support a wide range of implementation sizes, from very small to very large

SQL Server, the heart of the Microsoft BackOffice suite, attains a new level with release 7.0, setting the standard for ease of use (automatic operations), reliability, scalability and performance. This is good news not only for end users but is welcomed by software developers and the application program vendor community. Not only will they benefit from the efficiencies and capabilities of SQL Server 7.0 in their own work, but they will also be able to provide their customers with a stable and effective, high-performance platform to underlie their applications.

Application software developers such as Macola bet their reputations and their futures on the performance of the entire system - not just the applications but the whole package. Customers may not know whether bad performance or system failure is due to poorly designed or written application code or the failings of the underlying technology. As a result, the software developer may be automatically deemed the culprit, which can be very damaging to the developer's reputation. Smart developers are choosing SQL Server 7.0 for its reliability, scalability, and performance. SQL Server 7.0 allows them to concentrate on the applications with full confidence that the underlying database system will deliver.

MYTH:

RDBMS = complexity, high cost and additional people to manage the database

This equation represents the conventional thinking with regard to relational database management systems (RDBMSs). Although they are vital to today's demanding application environments, many consider RDBMSs to be complex and expensive. In addition they often require expert database administrators and highly skilled systems managers just to perform routine tasks to keep the system running.

The truth is that an RDBMS is a required part of today's application environment; there is no other practical way to manage the huge mass of data and transaction volumes that are characteristic of today's business systems. RDBMSs work and have become essential in today's world. But high functionality doesn't necessarily have to mean high complexity and high cost.

SQL SERVER 7.0 CHANGES THE EQUATION...

Designed to be easy to use, reliable and scalable while delivering outstanding performance, SQL Server 7.0 is a new breed of database system.

- Ease of use - SQL Server 7.0 is transparent to the user because all standard operations are automatic. Routine tasks ordinarily requiring the attention of a Database Administrator (DBA) are either automatic or conveniently executed from a central location via menus and wizards. Using these automated processes, time required maintaining your Progression SQL Server data will be minimal.

- Reliability - Reliability is probably the primary requirement for a database system; to protect the valuable data on which your ERP system, and your company, rely so heavily. SQL Server 7.0 is greatly enhanced to provide data verification, simplified system management and automated operations to ensure that your data is protected and readily available when you need it.

- Scalability - Companies want only one database system for all their computing needs; therefore, the systems chosen must be able to work on stand-alone laptop computers, small-group networks, corporate-wide large networks and distributed environments covering multiple locations. SQL Server 7.0 can be deployed on Windows 95/98 laptops or server-based installations capable of handling the needs of a large corporation. Macola's "Software You'll Never Outgrow" now runs on the database you'll never outgrow.

- Performance - A software company such as Macola must be very careful when selecting the technology that will underlie its applications. It is only the overall performance and utility of the entire package that the user sees and judges. To preserve its own reputation and to satisfy its customers fully, Macola must choose a platform that will provide the kind of performance that makes the system nearly invisible -- making information available without noticeable wait. SQL Server 7.0 has passed a number of benchmark tests during its testing and roll-out, as well as real-life field tests using Macola Progression applications, to prove that it has "the right stuff" to provide the performance Macola users demand.

Macola is one of the first application software developers to commit to SQL Server 7.0. The company participated in the Microsoft SQL Server 7.0 product launch. As a leading enterprise software developer for mid-size manufacturers and distributors, Macola Software maintains an uncompromising commitment to technology decisions that make sense for customers. Based on a thorough and continuing analysis of industry developments and emerging technologies, Macola made an early commitment to SQL Server 7.0.

SQL Server 7.0 is comprised of two components:

- Storage engine
- Query processor

Each component supports the design objectives mentioned above (usability, reliability, scalability and performance) through the incorporation of advanced features and capabilities.

The storage engine is responsible for the management of the data. This portion of the relational database system manages the physical placement of data, its organization, integrity and control. The storage engine determines how and where data is placed within the computer's file (storage) system which, in turn, determines the speed and efficiency of storage and access.

But data is not valuable until it is accessed and used. That's where the query processor comes in. The query processor is the buffer between the user and database. It translates "query" commands (SQL statements) into database requests and optimizes those requests for the most efficient execution and best performance.

STORAGE ENGINE FEATURES AND DESIGN

The SQL Server 7.0 storage engine offers comprehensive data management capabilities. Data within SQL Server 7.0 is organized as "pages," each containing space for 8 kilobytes of information. Each page contains a header entry that identifies the type of page, the amount of free space remaining and the owner of the page. A page can contain rows of data, index entries, log records (records data changes for use in recovery operations), text/images, allocation maps (to keep track of multiple-page "extents"), page free space (information about free space available on pages), or index allocation maps (information about extents). For data row pages, there is also an offset entry that tracks where each row starts on the page.

Here are other key storage engine features:

SQL Server 7.0 stores and manages many kinds of information, including images and text. This flexibility will allow Macola to expand object management capabilities within the Progression Series software application and enhance its ability to work with multimedia objects and applications.

An extent is made up of eight contiguous pages. For large tables, extents represent a basic unit of space used for storage. For smaller tables, however, this could cause great inefficiency. SQL Server 7.0 provides mixed extents, which can be owned by more than one object (e.g., contain pages from multiple owners), for tables that are smaller than eight pages. Space is added to a table one page at a time. When the table size exceeds eight pages, it will automatically switch to allocating uniform (e.g., owned by one object) extents.

The database structure uses space efficiently so users won't be burdened with extra costs for inefficiently used resources.

Space is allocated a page at a time, or 8KB per page, but Windows NT/2000 Server reads and writes in 512-byte segments. Because of this difference, it is possible that a power failure or other disruption could cause an incomplete update. To protect against such an eventuality, SQL Server 7.0 tracks each segment as it is written, using a mask of bits. If updating is incomplete, the mask will clearly show what portion of the page is incomplete. This is called "torn page protection."

Many safeguards are built into SQL Server 7.0 to protect information and ensure against loss due to power failure or other disaster. Macola SQL Server 7.0 users can remain confident that their valuable company data is secure and complete.

Files and file groups improve database performance by allowing a database to be distributed across multiple disks and multiple disk controllers, which allows parallel I/O operation. While most databases work best with a single data file and single log file, spreading a multiple-file database across multiple disks and/or controllers improves both performance and data integrity.

A very scalable database, SQL Server 7.0 is designed to perform in all sizes. Built-in design features such as this make the most of the resources available. A user's company grows, and the system expands, SQL Server 7.0 can grow gracefully. Distributed storage improves performance by spreading access across multiple read-write heads (parallel access) and improves data integrity by limiting the impact of a single disk failure on them.

SQL Server 7.0 will automatically shrink and grow files as needed to accommodate storage needs and free as much space as possible. Periodically, in the background, the system checks for empty space and moves rows from pages at the end of files to pages at the beginning of files. Pages freed at the end of the file are returned to the file system. Similarly, files can be allowed to grow as needed. SQL Server 7.0 will add space as much as necessary unless a limit is placed on the file or file growth is disabled. By automatically expanding the size of the database, there will be no delay in processing in Progression SQL Server. The automatic reduction of the database will aid in recovering the space available after purging occurs in Progression SQL Server.

No routine administration is required to keep the database in prime condition. SQL Server 7.0 automatically takes care of the details such as growing and shrinking files.

Locking in SQL Server 7.0 is dynamic and automatic, tuned to provide the utmost concurrency consistent with data integrity requirements. Row-level locking is used for inserting, updating and deleting data where possible. Page locking is used for data scans. The storage engine cooperates with the query processor at run-time to determine the lowest-cost locking strategy based on the characteristics of the database and the query. As a result, there is no need for a database administrator to manage lock-escalation thresholds, nor does the application developer need to be concerned about database locking administration.

Providing ready access to information is the primary function of a database system, second only to keeping the data safe and secure. Locking ensures data integrity, but it can get in the way of access needs. Intelligent, flexible locking balances these sometimes conflicting needs.

SQL Server 7.0 maintains log files, which include

- Serial records of all modifications that have occurred in the database
- Which transaction performed each modification
- The ability to recover the database if there is a system failure
- The ability for SQL Server to roll transactions forward to recreate transactions, or backward to reverse an incomplete transaction

When disaster strikes or mistakes are made- it's nice to know that complete records exist, and that they can be applied to the database to correct or recreate your information.

To improve performance, SQL Server 7.0 manages memory usage dynamically. This allocates and frees memory, which minimizes swapping (either within SQL Server or by Windows) and thereby optimizes the use of system resources. In addition, the system monitors the use of (memory) buffers, cache, as well as the use of I/O operations to optimize performance. Also, the system uses read-ahead logic based on the general pattern of reads to be performed.

The database is not the only thing running on your system. System resource allocation and usage must help all applications and utilities perform at their peak for users to get the response they expect and deserve. SQL Server 7.0 allocates memory dynamically for optimal performance.

All in all, SQL Server 7.0 is packed with improvements and automation to simplify operation while maximizing performance and optimizing the use of resources. These tools and features combine to deliver a new level of ease of use, reliability, scalability and performance.

Storage Engine Ease of Use

A company buys an application system such as Macola Progression Series ERP to solve business problems and support business operations. Unfortunately, business applications often come with considerable baggage in the way of technical and administrative support for the underlying technology. SQL Server 7.0 greatly reduces or eliminates these burdens through a number of features, including

- Automation of routine tasks.
- Transparent server configuration.
- Streamlined and simplified options.
- Wizards to help the user through tasks such as creating a database, scheduling backups, and importing or exporting data.
- Memory and lock resources that are adjusted dynamically in SQL Server 7.0.
- File sizes that grow or shrink automatically within configurable limits, minimizing the need for DBA intervention. This means there is no need to preallocate space and manage data structures in SQL Server 7.0.
- Autotuning features that guarantee consistent performance under variable load conditions.
- The database that monitors its own performance, availability, and security status and issues notification (and initiates recovery) when it encounters policy-based thresholds and security levels.

ODBC - Open Data Base Connectivity, a de facto standard for communicating with any compliant database

OLE - Object Linking and Embedding, a Windows standard protocol allowing dynamic, interactive connection between objects (programs)

APIs - Application Programming Interfaces, predefined and published interactive "hooks" into and out of an application program

Replication - the process of duplicating and synchronizing information in a number of parallel databases

- Automatic job scheduling and execution that offers great flexibility, including the ability to use Visual Basic Scripting, Java Scripting, Windows NT/2000 Commands, and custom ODBC and OLE DB Programs.
- Drag-and-drop and single commands can be used to implement changes across groups of servers. Management is simplified through the use of a repository that maintains schema, profiles and data transformation metadata for all servers in the enterprise.

- Security administration is integrated with Windows NT/2000 security for authentication, support for multiple groups, a grant/revoke/deny model and the dynamic use of groups.

- Full compliance with ANSI/ISO SQL-92 Entry Level standards.

- Fully automated upgrade utility for conversion from version 6.X to SQL Server 7.0. This allows users to take advantage of new features with minimal impact on operations.

- Graphical interface for building and managing schema and other database objects.

Storage Engine Reliability

Security and control are foremost concerns with any database system. SQL Server 7.0 has been significantly enhanced, not only to provide reliability and tight management of data integrity, but also to do so with a minimum amount of attention from users, system managers and administrators. Many functions are automatic and logic-driven to ensure full access, yet tight control.

- Full row-level locking for both data rows and index entries ensures data integrity with minimal impact on other users. Dynamic locking automatically chooses the optimal level of lock (row, page, multiple page, table) for all database operations. This provides improved concurrency with no tuning. A particular level of locking can be forced if desired.

- Simplified data structures and algorithms improve concurrency, scalability and control. Run-time checks of critical data structures make the database much more robust, minimizing the need for consistency checks.

- Database Consistency Checker (DBCC) statements check the logical and physical consistency of a database, memory usage and performance statistics. Running DBCC is a valuable tool and a good standard practice for peace of mind, but it is not considered mandatory for normal database operation. DBCC commands are fast, running at near-disk speed, and many commands have been consolidated to simplify the database verification process. DBCC can automatically repair some of the errors and inconsistencies it finds during verification.

- SQL Server 7.0 includes a broad spectrum of innovative replication technologies, including a simplified user interface, wizards, improved monitoring, scripting and troubleshooting.

- Standard published APIs support bidirectional replication with other databases including Oracle, DB2, Sybase and Informix.

- Changes to a subscriber's data can be immediately propagated to the publisher via two-phase commit and then to other subscribers using transactional or snapshot replication. Replication functions are built in, not added on.

- Multisite replication update allows updates of multiple copies of the same data at different locations.

- Replication to hundreds of servers and thousands of users is supported through a streamlined architecture that reduces contention on replication tables.

Storage Engine Scalability

Macola's "Software You'll Never Outgrow" must be supported by a database you'll never outgrow. SQL Server 7.0 provides the kind of scalability only dreamed of before. From a small, mobile, single-user laptop configuration to terabytes of data and millions of transactions, SQL Server 7.0 is truly the database for all sizes of applications and companies.

Simplified mapping of database objects to files eases management and enables tuning flexibility. More efficient memory management helps SQL Server 7.0 work better on small machines and improves performance in larger applications.

Among the greatly increased capacities of SQL Server 7.0 are the following upper limits:

- Memory addressing capability greater than 4 gigabytes
- Maximum database size: 1,048,516 terabytes
- Maximum number of files per database: 32,767
- File size (data): 32 terabytes
- Rows per table limited only by available storage

Megabyte - One thousand bytes

Gigabyte - One thousand megabytes

Terabyte - One thousand gigabytes

SQL Server 7.0 is scalable. It can grow large enough for massive data warehousing tasks, yet it can be compact enough to run on a laptop. It is truly the database you'll never outgrow.

Storage Engine Performance

Performance is always a concern; applications only work as well as their underlying databases. SQL Server 7.0 has been tested extensively, and benchmark tests place it among the industry leaders. Many features have been added to enhance system performance:

- Dynamic memory allocation and usage improves performance while minimizing contention with other resource managers.
- Data import/export uses OLE DB and works in conjunction with the query processor to plan and optimize queries. The result: dramatic improvements in speed.
- Bigger I/O sizes support higher throughput rates. Smart I/O technology enhances performance with big I/O sizes. More efficient read-aheads, physical row-order scans and parallel I/O also improve performance.
- New indexing strategies, including the use of multiple indexes in a single table or multiple tables, multicovered and join indexes, parallel index creation on the same table, and automatic statistics maintenance all contribute to improved performance.
- New hash and merge join strategies improve performance for certain types of data retrieval, in addition to nested loop joins. The query processor recognizes certain join types (such as star query joins) and optimizes for these join types.

- Performance improves with intra-query parallel execution of a single query across multiple processors. Steps in a single query are executed in parallel, delivering optimum response time.

- Sort speed is improved.

- Parallel backup and restore utilities scale at device speeds. This minimizes impact on operational systems and provides very high server transaction processing rates during full online backup.

QUERY PROCESSOR FEATURES AND DESIGN

Microsoft SQL Server 7.0's query processor accepts SQL syntax (statements or requests), determines how to execute that syntax and executes the request. The user or program interacts with the query processor (using SQL statements), and the query processor interacts with the storage engine. This provides a buffer between the user (or programmer) and the details of the database and storage engine. The user specifies the result; the query processor determines how to obtain that result.

Database queries can be complicated, especially if the user is not trained in the specifics of query language and syntax. SQL Server 7.0 provides a buffer between the user and the database, interpreting and optimizing the request to gather the information in the most efficient and effective way.

The query processor's operation consists of two components: query optimization and query execution. Optimization is where the logic within the query processor determines how best to go about the task:

- Which indexes to use (if any)

- How best to join files together

- How to handle conditions and clauses in the query request

- What procedure and algorithms to be used to execute the request most efficiently. This is determined from statistics about file sizes and organization, past experience, etc.

Query execution follows the optimizer's instructions and creates the actual response. Think of the optimizer as the brain, and execution as the muscle behind the query processor.

The SQL Server 7.0 query optimizer is what's known as a cost-based optimizer. Rather than simply taking the query request directly as it is stated and processing it as written, the optimizer develops a plan for execution that considers the size and organization of the files. To make the most of this approach, the optimizer first takes a quick look at the request and determines if a simple approach will deliver quick results. If it will (e.g., the projected execution time is below a threshold) no further optimization is attempted. For more complicated requests, optimization is warranted.

Optimization includes a long list of techniques and capabilities designed to make query processing more efficient. While most query processors are able to store predefined execution plans (for reuse), SQL Server 7.0 can also handle dynamic SQL with automatic parameters. Basic information about the query is stored, and constants are turned into parameters. This provides the efficiencies of stored procedures with the flexibility of ad-hoc queries — even for applications that can't take advantage of stored procedures.

In addition to the traditional nested-loop iteration join, which is fine for row-to-row navigation, SQL Server 7.0 also employs hash joins (better for large, nonindexed tables) and merge joins for when join fields are indexed and the columns represented in the index cover the query. In the new (and exclusive) hash team approach, intermediate

results (from inner join, merge join and aggregation) are passed to the next phase in sorted order, thereby avoiding the need to execute a sort operation. This approach, called interesting ordering, is available in all major query processors today. However, only SQL Server 7.0 applies this concept to hash joins, delivering faster processing of complex queries.

These are but a few of the advanced techniques and tools used by SQL Server 7.0 to optimize the query for most effective execution. SQL Server 7.0 also uses "group by" clauses to process the requests in the most efficient order, partitioning using filegroups for lowest practical granularity, special techniques for handling queries against star schemas and more.

The query execution engine uses a number of techniques to minimize I/O, the most costly part of the query. Large page and extend sizes (8K and 64K respectively) allow it to read in a substantial amount of data per disk access, scanning read-aheads and prefetch hints. These tell the storage engine about information that will be needed soon, which helps improve performance. Sort algorithms have also been dramatically improved in release 7.0.

Query Processor Ease of Use

The built-in advanced optimization, automatic statistics and wide range of tools all work together to make SQL Server 7.0 simultaneously one of the easiest database systems to use as well as the leading edge of technology and efficiency. Its features and capabilities include

- Fully automated query optimization.
- No database administration required for standard operations.
- Fast and simple programming model for developers.
- New query processor features that add "industrial strength" for data warehousing applications.
- Improved tools for diagnosing and fixing problems including Profiler, Query Analyzer and Index Tuning Wizard. These tools pinpoint any problems and often suggest remedies.

Query Processor Reliability and Scalability

A database is a utility — something the user should not have to worry about. Likewise, the application software developer relies heavily on the functionality, reliability and performance of the underlying database system — including the query processor — to make the applications run well and look good. SQL Server 7.0 sets the bar for leading-edge performance in a package that makes life easier for developers, users and system administrators.

The same SQL Server 7.0 database engine can handle laptops on the low end, to the central corporate database or data warehouse on the high end. This extraordinary scalability (ability to serve users and applications of a wide range of sizes and the ability to grow gracefully) is a defining characteristic of SQL Server 7.0.

Query Processor Performance

Without question, SQL Server 7.0 outperforms its competitors. Many features and functions, as outlined above, have been redesigned specifically to deliver better response and more effective use of all system resources. Benchmark tests prove that SQL Server 7.0 has the "right stuff" for today's demanding applications.

MACOLA SOFTWARE AND SQL SERVER 7.0

Macola's Progression SQL Server series is optimized to take full advantage of the Microsoft SQL Server environment. For example

- The use of stored procedures greatly increases the performance of Progression SQL Server.
- Progression SQL Server has been designed with a multiconnection lock architecture to greatly increase scalability within the SQL engine.
- Macola's use of an SQL-based search engine with SQL-based visual definitions and editor means more efficient searches.
- Progression SQL Server's higher usage of complex SQL queries to enhance specific areas of the application takes advantage of SQL Server 7 multi-index query capability.
- Progression SQL Server uses database derived indexes to increase performance.
- Data access and reporting activities in Progression SQL Server enjoy a multifold performance increase over non-SQL Server databases.

Application software developers live in the world of computers, operating systems, databases and networking technology. Their customers are at least one step removed from that world. Therefore, it is the developer's responsibility to make technology choices that are in the best interest of their customers. The truth of the matter is that the customer only sees the complete system and will judge its performance in its entirety. All of the components must work together to provide users with a system that supports their needs by being responsive and providing easy access to information. In other words, developers know the result is a combination of their own code leveraging the capabilities of the system and the database. However, the user sees only one result.

A CUSTOMER SPEAKS

Linden Industries, a small (50 employees) manufacturer of engineered-to-order industrial equipment, has been a Macola Software customer for about four years. While the software functionality was adequate, Linden found the Macola Progression Series system before Progression SQL Server a bit too rigid. In addition, access to data was not as convenient or flexible as desired.

After researching the market, Linden decided SQL Server 7.0 was a requirement for any replacement ERP solution under consideration. After learning Macola Software was committed to early availability on SQL Server 7.0, Linden Industries decided to implement the newest release of Progression Series (7.5) to leverage existing investments in SQL Server and Macola functionality and use.

Linden implemented Macola Progression and SQL Server as a controlled-release installation in the fall of 1998. Ted Gentsch, Linden's MIS Manager, describes the new Progression release as "almost a complete reversal" of the previous version. "Macola has leveraged Microsoft technology," he adds. "The searches, flexibility and capabilities are exactly what we were looking for."

Gentsch goes on to describe how Progression Series 7.5 and SQL Server 7.0 allow Linden to pull information from Macola data files quickly, easily and in a format that delivers the information to run the business. He believes Linden can now really take advantage of Macola Progression applications and is extending the use of the system to more areas of the business. The company is even eliminating a number of Access-based custom applications that are no longer needed since the Macola data is now so much more valuable thanks to SQL Server 7.0.

Linden is just one manufacturer discovering how ERP systems should work — providing flexible access to vital company information in support of better business management.

SUMMARY

Technology is changing at mind-boggling speed. Each new development presents new opportunities and new challenges. It's important to stay on the leading edge of technology — to provide the best capabilities to users — but it's equally important to choose carefully so that users are not locked into technologies made obsolete by the next advancement. Making the right decision is a very difficult task.

Macola Software, along with its peers and competitors, faces these technology issues every day. When Macola makes a commitment to a technology tool, it has to be for the long term, providing solid capabilities that will withstand the test of time. Macola Software considers its commitment to Microsoft SQL Server 7.0 to be the right technology decision. In turn, Microsoft has invested heavily in advanced database technology and has clearly proven its commitment to deliver industry-leading functionality and ease of use. In addition, Microsoft has worked very closely with leading application developers such as Macola to ensure that the latest release of SQL Server provides the reliability, scalability and performance Macola and its customers demand.

333 E. Center St.
P.O. Box 1824
Marion, OH 43301-1824

Sales: 800.468.0834
tel: 740.382.5999
fax: 740.382.0239
http://www.macola.com

MAS1184-990115