

# Top 10 SQL Server Questions Answered by SQL Stan

---

*written by  
Quest Software, Inc.*



White Paper

© Copyright Quest® Software, Inc. 2007. All rights reserved.

This guide contains proprietary information, which is protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software, Inc.

## **WARRANTY**

The information contained in this document is subject to change without notice. Quest Software makes no warranty of any kind with respect to this information. QUEST SOFTWARE SPECIFICALLY DISCLAIMS THE IMPLIED WARRANTY OF THE MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Quest Software shall not be liable for any direct, indirect, incidental, consequential, or other damage alleged in connection with the furnishing or use of this information.

## **TRADEMARKS**

All trademarks and registered trademarks used in this guide are property of their respective owners.

World Headquarters  
5 Polaris Way  
Aliso Viejo, CA 92656  
[www.quest.com](http://www.quest.com)  
e-mail: [info@quest.com](mailto:info@quest.com)

Please refer to our Web site for regional and international office information.

Updated—June 01, 2007

# CONTENTS

- INTRODUCTION ..... 1
- 1. WHAT ARE THE ADVANTAGES AND IMPLICATIONS OF  
MIGRATING FROM MICROSOFT ACCESS TO SQL SERVER? ..... 2
- 2. HOW CAN I ADD A COLUMN TO A TABLE WITHOUT TRUNCATING OR  
DROPPING THE TABLE? HOW WILL QUERY PERFORMANCE BE AFFECTED? . 3
- 3. HOW DO YOU AVOID DEADLOCKS IN A SQL SERVER DATABASE?..... 4
- 4. HOW DO YOU MOVE (OR DUPLICATE) A MICROSOFT SQL  
SERVER DATABASE FROM ONE PHYSICAL SERVER TO ANOTHER IN  
A DIFFERENT GEOGRAPHIC LOCATION? ..... 5
- 5. HOW CAN I CALL A DLL FROM SQL SERVER? ..... 5
- 6. HOW CAN I APPLY INDEXES TO A COLUMN THAT IS A NON-INTEGER  
DATA TYPE? ALSO, HOW CAN I USE A DIFFERENT ALGORITHM FOR  
AN INDEX? ..... 6
- 7. WHAT IS THE BEST WAY TO CREATE A UNIQUE TABLE THAT  
CONTAINS A COLUMN ID? ..... 6
- 8. CAN YOU EXPLAIN HOW PERMISSIONS WILL BEHAVE WHEN  
RESTORING A DATABASE TO ANOTHER SERVER? ..... 7
- 9. IS SQL MAINTENANCE THE SIMPLEST WAY TO MAINTAIN  
THE SIZE OF SQL SERVER TRANSACTION LOGS? ..... 7
- 10. HOW CAN I TRACE THE PERFORMANCE ISSUES FROM  
THE DATABASE LEVEL? ..... 8
- HOW TO CONTACT SQL STAN ..... 9
- ABOUT QUEST SOFTWARE, INC. .... 10
  - CONTACTING QUEST SOFTWARE..... 10
  - CONTACTING QUEST SUPPORT..... 10

# INTRODUCTION

SQL Stan, an online resource built on the thought leadership that Quest brings to the Microsoft SQL Server community, provides answers to your SQL Server related questions. SQL Stan is the embodiment of the collective SQL Server knowledge at Quest Software. This includes participation from revered SQL Server expert Kevin Kline as well as Quest SQL Server domain experts, product managers, solutions architects, systems consultants and development staff. Fueled by the needs of Quest's customers, partners and the whole SQL Server user community, SQL Stan was formed to leverage this existing knowledge base and thought leadership. This paper covers the top 10 questions submitted to SQL Stan and the answers based on SQL Stan's expert knowledge.

# 1. WHAT ARE THE ADVANTAGES AND IMPLICATIONS OF MIGRATING FROM MICROSOFT ACCESS TO SQL SERVER?

SQL Server is a high-performance relational database engine that supports hundreds or thousands of users accessing data simultaneously. SQL Server includes advanced security mechanisms to secure, back up, and recover your data. Microsoft Access includes a client-“friendly” but limited database engine, as well as the ability to create applications using its front-end development environment. When you move from Access to SQL Server, you are moving only the Access backend (the database portion of Access). SQL Server does not include features to create applications; for that, you can use numerous development tools and languages, such as Access, Visual Studio, and Java.

Microsoft published an informative article on migrating from Access to SQL Server 2000. Although the paper refers to SQL Server 2000, the concepts will apply to a migration to SQL Server 2005. You can find the article at

<http://www.microsoft.com/technet/prodtechnol/sql/2000/deploy/accessmigration.mspx>.

Another paper from Microsoft, titled “When to Migrate from Microsoft Access to Microsoft SQL Server,” may provide additional information:

[http://download.microsoft.com/download/5/d/0/5d026b60-e4be-42fc-a250-2d75c49172bc/when\\_to\\_Migrate\\_from\\_Access.doc](http://download.microsoft.com/download/5/d/0/5d026b60-e4be-42fc-a250-2d75c49172bc/when_to_Migrate_from_Access.doc).

These articles should give you a good start on your migration to SQL Server.

You might also consider investigating the SQL Server Migration Assistant, a tool created by Microsoft to assist with migrating from any one of several database platforms (such as Oracle and Access) to SQL Server.

## **2. HOW CAN I ADD A COLUMN TO A TABLE WITHOUT TRUNCATING OR DROPPING THE TABLE? HOW WILL QUERY PERFORMANCE BE AFFECTED?**

The first step in modifying your table is to use the ALTER TABLE statement to add a column to the database object. There may be many rules about the different kinds of new columns you can define, but we'll stay fairly general for this case. The new column you want to add either must allow NULL or can be NOT NULL with an assigned DEFAULT value. When the table is actually altered, the new column is populated with either the NULL value (if NULL is allowed and no default value is assigned) or the default value.

Regarding the question about performance, the table you modify will be locked until the ALTER TABLE statement completes. This operation might take a while if you're adding data to the new column of a very large table and consequently adding data to the data pages, which means that SQL Server has to write new data to disk. Thus, it is essential to be aware of traffic on the database. After the ALTER TABLE statement completes, any new columns can immediately be used in indexes and otherwise behave like regular columns in terms of performance.

### 3. HOW DO YOU AVOID DEADLOCKS IN A SQL SERVER DATABASE?

This is an interesting technical question because deadlocks are, to some degree, unavoidable in any database application with more than one user. On the other hand, if your application experiences an inordinate number of deadlocks, you can take a couple courses of action.

Your first course of action, short of redesigning the system, is to evaluate your indexing strategy. If clustered indexes exist in your environment; ensure that they store their data on 8K pages far away from one another. You might even need to confirm that Fill Factors (and Pad Indexes) are set to a non-default value, perhaps 75 to 80 percent, to inject greater space between your data pages. As a side effort, make sure your transactions take out the minimal amount of locks required. Also, try to limit the use of wildcards (such as `SELECT * FROM...`) and especially limit the amount of time the `INSERT`, `UPDATE` and `DELETE` transactions remain open and uncommitted.

Second, you can set one transaction, or set of transactions, as the all-time loser (or winner in SQL Server 2005) by using the command `SET DEADLOCK_PRIORITY`. Read more about this type of transaction at <http://msdn2.microsoft.com/en-us/library/ms186736.aspx>.

Finally, you can alter the default locking behavior for a given connection by using the `SET TRANSACTION ISOLATION LEVEL` command (read about that at <http://msdn2.microsoft.com/en-us/library/ms173763.aspx>) or on a query-by-query basis by using query hints such as `NOLOCK` to modify the behavior of an individual query. I don't recommend that you do this without carefully considering the implications of altering the native behavior of your SQL Server database.

## 4. HOW DO YOU MOVE (OR DUPLICATE) A MICROSOFT SQL SERVER DATABASE FROM ONE PHYSICAL SERVER TO ANOTHER IN A DIFFERENT GEOGRAPHIC LOCATION?

There are actually two answers to this question. First, if you don't need the data from the database, you can use the Generate SQL Server Scripts Wizard to create a Transact-SQL script of all the objects within the database. After generating the script, you simply create the database on the target location (making sure it's a large enough size to avoid lots of auto-growth) and run the script in the new database. I suggest as a best practice running the script twice because the script will be generated by the wizard in alphabetic order, causing some objects to be created before objects they depend upon.

Second, if you need to duplicate the data as well as the database, the easiest option is to use the **sp\_detach** stored procedure to put the database into a quiescent state. Then use a file copy to move the database to the new location. Use the **sp\_attach\_db** stored procedure to activate the database on the new server.

## 5. HOW CAN I CALL A DLL FROM SQL SERVER?

There are two ways to call a dynamic-link library (DLL) from SQL Server. Each typically relies on the location of the called DLL. If the DLL is located within SQL Server, use the extended stored procedures. If you need to call a DLL located outside of SQL Server, use **xp\_cmdshell**. However, only members of the **sysadmin** role may use **xp\_cmdshell**. If security is a concern; this may not be the best option. Also, be careful of a poorly written DLL, which can cause SQL Server to crash. That's one of the main reasons Microsoft introduced common language runtime (CLR) programming: it protects you from crashes.

## **6. HOW CAN I APPLY INDEXES TO A COLUMN THAT IS A NON-INTEGER DATA TYPE? ALSO, HOW CAN I USE A DIFFERENT ALGORITHM FOR AN INDEX?**

SQL Server provides a lot of flexibility on the datatypes that can be indexed. You can index just about any type of column in SQL Server except for IMAGE and TEXT columns. As for the second question, because SQL Server always uses b-tree indexing, you cannot apply other indexing algorithms.

## **7. WHAT IS THE BEST WAY TO CREATE A UNIQUE TABLE THAT CONTAINS A COLUMN ID?**

For this type of situation, you can use either the uniqueidentifier type or the integer type with an identity property. Let's take a look at the pros and cons of each of these choices.

Uniqueidentifier creates a value that is guaranteed to be unique on any instance of SQL Server. This is highly advantageous when you have a large collection of SQL servers that share duty, such as a replicated group of servers or servers organized into a server farm. On the downside, uniqueidentifier values are not easily readable — and they're huge! One more suggestion: Because uniqueidentifier values are not necessarily guaranteed to be sequential, use the NEWID function to assign a uniqueidentifier value.

The alternative is to use the integer type with the identity property, commonly defined like this: `int[identity(1,1)]`. While this avoids the size issues associated with the uniqueidentifier option, you really don't have the same reliability of uniqueness when you use an integer type with an identity property. It might be unique for the table it resides in, but it's very likely to be duplicated elsewhere. Plus, identity properties require additional administrative care.

## **8. CAN YOU EXPLAIN HOW PERMISSIONS WILL BEHAVE WHEN RESTORING A DATABASE TO ANOTHER SERVER?**

Database permissions are controlled on a server-by-server basis. Once a backup is restored to a separate server, the permissions that exist on the target database will prevail, rather than the server-level permissions from the source database. That's why restoring a database to a new server can sometimes be problematic; users and their privileges are stored in the master database and probably haven't been restored along with the user database. For example, let's say Admin A has full SA privileges on Server1 and Admin B has full SA privileges on Server2. If you restore a database from Server1 to Server2, Admin A will not be able to log on to Server2 because this administrator has no privileges there. Also, Admin B with full SA privileges on Server 2 will be able to administer the database restored on this server. From a security standpoint, this type of situation is definitely not appealing.

## **9. IS SQL MAINTENANCE THE SIMPLEST WAY TO MAINTAIN THE SIZE OF SQL SERVER TRANSACTION LOGS?**

SQL Maintenance Plans are an easy way to keep your databases under a routine of preventative maintenance. However, I don't recommend that enterprise DBAs use them. It is far better to have separate, individual preventative maintenance tasks scheduled for each important activity than to rely solely on Maintenance Plans. These tasks include DBCC, CHECKALLOC, DBCC CHECKDB, index de-fragmentation; fill factor restorations, backups of databases and transaction logs.

If you are going to rely on Maintenance Plans, you may want to include these other maintenance tasks so that you're being as comprehensive as possible; erring on the side of caution makes sense, especially when the availability of a production database is at stake. I've lost count of the number of times novice SQL Server users have been burned by not covering all of the bases. Doing things reliably may require a few more steps, but you'll almost never be sorry.

## 10. HOW CAN I TRACE THE PERFORMANCE ISSUES FROM THE DATABASE LEVEL?

Without knowing the specific symptoms of this particular performance issue, I can provide only a general answer. If you have some sort of resource bottleneck, I would try to find all causes of that bottleneck. For example, if the symptom was slow throughput due to excessive CPU usage (that is, you have a CPU bottleneck), I would suggest starting by identifying the top five CPU consumers on the server. If the symptom was slow throughput due to saturated I/O (that is, you have an I/O bottleneck), I would start with identifying the top five I/O consumers on the server. Is blocking on locks (lock contention) a symptom? I would start with analyzing what resources (tables and indexes) are being contended for and identifying the top five SQL statements involved in that contention.

Although the question doesn't state whether the issue stems from SQL Server 2000 or SQL Server 2005, the methods of identifying the causes of different bottlenecks differ slightly between the platforms. But these general rules of thumb should cover your approach. On either platform, you can limit the scope of your investigation to a particular database.

## HOW TO CONTACT SQL STAN

SQL Stan is Quest's online resource for answering any questions related to Microsoft SQL Server. The questions that have been shared through this document are the top 10 that have been submitted to SQL Stan—which now contains hundreds of questions and answers. SQL Stan is available through a dedicated Web page provided by Quest Software and located at [www.Quest.com/SQLStan](http://www.Quest.com/SQLStan). We look forward to your interaction with SQL Stan, a channel for the SQL Server community to communicate directly with the thought leadership Quest Software brings to this community.

## ABOUT QUEST SOFTWARE, INC.

Quest Software, Inc. delivers innovative products that help organizations get more performance and productivity from their applications, databases and Windows infrastructure. Through a deep expertise in IT operations and a continued focus on what works best, Quest helps more than 50,000 customers worldwide meet higher expectations for enterprise IT. Quest Software can be found in offices around the globe and at [www.quest.com](http://www.quest.com).

## Contacting Quest Software

Phone:	949.754.8000 (United States and Canada)
Email:	<a href="mailto:info@quest.com">info@quest.com</a>
Mail:	Quest Software, Inc. World Headquarters 5 Polaris Way Aliso Viejo, CA 92656 USA
Web site	<a href="http://www.quest.com">www.quest.com</a>

Please refer to our Web site for regional and international office information.

## Contacting Quest Support

Quest Support is available to customers who have a trial version of a Quest product or who have purchased a commercial version and have a valid maintenance contract. Quest Support provides around the clock coverage with SupportLink, our web self-service. Visit SupportLink at <http://support.quest.com>

From SupportLink, you can do the following:

- Quickly find thousands of solutions (Knowledgebase articles/documents).
- Download patches and upgrades.
- Seek help from a Support engineer.
- Log and update your case, and check its status.

View the ***Global Support Guide*** for a detailed explanation of support programs, online services, contact information, and policy and procedures. The guide is available at: [http://support.quest.com/pdfs/Global Support Guide.pdf](http://support.quest.com/pdfs/Global%20Support%20Guide.pdf)