


The 5S Approach to Improving Database Performance

for DBA Managers



By Chuck Ezell
November 2013

ABSTRACT:

If your organization relies on data, optimizing the performance of your database can increase your earnings and savings. Many factors large and small can affect performance, so fine-tuning your database is essential. Performance Tuning expert and Senior Applications Tuner for Datavail, Chuck Ezell, sheds light on the right questions to get the answers that will help you move forward by using a defined approach, referred to as 5S.

This performance tuning white paper addresses each stage of this novel approach, as well as key performance issues: SQL, Space, Sessions, Statistics, and Scheduled Processes.

These five areas help organize the tuning approach and define the major concerns beyond the architecture, setup, and data model. It also addresses how performance tuning becomes less of a mystery if it can be measured, documented, affected, and improved.

Why Use a Defined Approach?

In an organization wholly reliant on data, ongoing database performance is critical. It is one of many tasks database administrators must juggle. If they are distracted and don't know the nuts and bolts of the database – beyond how it is operating – they can't really know how to troubleshoot and fine-tune its performance.

It's best to evaluate your system's health and diagnose database performance issues before moving forward. Often defining a root cause in an environment that is unstable leads to failed attempts at tuning performance problems.

It's too much to expect to get incredible performance from an overworked, underpowered database. Tuning shouldn't be pursued until hardware performance issues are resolved. Sure, you can limp along, but your database can only be truly optimized if its host environment is performing well.

You need an effective and comprehensive plan. Following a process to define, measure and test the performance you'll have documented evidence of improvement, validation for applying resources to the work and an understanding of why the problems were occurring.

Adjustments made in the tuning process will change the dynamic environment of the database. You will need to document your current system's performance status before beginning the performance tuning process. We call this a system health check. This health check can be your baseline to compare any improvements made during the tuning process. As you make small changes, measure the system performance each time, ensuring the change is bearing improvement. As a secondary benefit, if you need to stop and undo it is done easier if changes are incrementally made. Any and all changes made should be documented. We find the Six Sigma incremental process DMAIC is a significant benefit in providing a framework for managing this process.

Not to be confused with the 5S approach, the Six Sigma DMAIC process is defined as a 5-step

process for improvement that includes Define, Measure, Analyze, Improve and Control. Good performance tuning must involve these five elements if you are to render true, documented results at the end of this process. Remember, performance tuning is not voodoo; it should be measurable and explainable however complex the outcome or solution.

When Should You Conduct Performance Tuning?

There is much discussion about reactive and proactive tuning. If you feel your database performance is sub-par then it is highly likely you will find benefit in conducting performance tuning. In a reactive approach you're simply dealing with issues as they arise. Proactive would be incorporating tuning concepts during an implementation or a planning phase. However, the answer to the question of "when" is always both. Proactive and reactive. The ongoing refinement of database operations results in increasingly accurate and ever-faster response times to user queries. That information can be used to inform organizational decision-making, which may demand time-sensitive responses to address emerging issues, production timelines, or deadlines. A fast response to these can translate into earnings and savings.

What if you are uncertain about how to undertake or begin tuning? Problems can arise if you and your team don't know what to do. Knowledge about the right approach can make all the difference in how quickly you arrive at measurable results.

Think of the 5S approach as shining light into the dark corners of your environment. Performance issues vary widely. Some issues are obvious, but others may be subtle. Some nuanced link or dependency within the system, not immediately apparent to you, may be causing issues. Unless your DBA has experience with the complexities of performance tuning, developing effective solutions can be time-consuming and resource-intensive. Applications, databases and systems are precisely interconnected, so any changes influence other elements of the entire configuration.

We have found that, sometimes, only a few adjustments are needed to begin improving database performance. As few as a quarter of issues may account for 75% of a database's diminished performance, so small changes can address a significant portion of the problems you may be experiencing.

The 5S Approach

We have discovered, based on our ongoing work with database performance tuning, that most issues affecting database performance can be distilled into five key areas. From this, we have created our 5S Approach.

What are these five S's? SQL, space, sessions, statistics, and scheduled processes. While not exhaustive, these five aspects of your database directly affect its performance, from the SQL statement that's running in your database to varied aspects related to available disk or storage space. Sessions that are running in the database and scheduled processes can also influence the efficiency with which the database returns queries. You might not remember items that are scheduled to run, or that run so infrequently you forget they're scheduled until the task executes.

Structured Query Language is the actual code in the database that seeks the data you need. Among the keys to managing performance via SQL are knowing the best practices for using SQL, understanding when and what types of optimization can help boost performance, and when to make those changes within SQL. Database administrators typically feel comfortable with the code, but there can be variants in the particular flavor of SQL being used. You need to know what's running and how it's performing. You can, for example, fix persistent I/O issues simply by fixing some of your SQL statements.

Space is the place where the data lives. More than simply storage, the space within your database ecosystem must include sufficient capacity. However, in our approach we include the actual data integrity, indexing, data growth,

normalization and data management. Having adequate space is an important factor for performance tuning and managing that space properly with archiving, purging and delete procedures is crucial. Large data objects require proper indexing, and partitioning and wise approaches in the data model all play into performance tuning.

Do you know what **sessions** are running at any given time? You have many participants vying for your data that can compete for system resources. Their sometimes unmerciful database requests can lock up objects or, perhaps, a developer has connected to production and is running a "test" query and it is consuming all the CPU. Whatever the root cause, different databases need different strategies to deal with the problems these competing sessions may pose, such as contention, locking, and concurrency. Honing performance requires understanding the specific needs and demands of your various sessions.

Statistics help queries move effectively through a database's data tables. These can enable a SQL engine to make quick estimations of which index to use, whether the query needs to scan the whole table to find the data or just select the first few rows of a table. The more accurate your statistics are, the better your SQL execution plans will perform. Ideally, rather than scanning an entire table at runtime, simply the table or index statistics can be read. For example, environments that have a high degree of data manipulation (updating, inserting & deleting) could also have poor statistics. This might lead the database optimizing engine to execute a full table scan on a table of 10 million rows because the outdated statistics show only 10 rows.

Scheduling is another important issue that can infringe on performance. Much like directing traffic through a bustling section of a city, the database administrator needs to be aware of the flow of traffic, including periods when traffic is heavy or light. The database administration manager needs to be aware of when users are competing for resources and know when peak use occurs. Processes that were set to run at a

specific time may no longer be needed. Could there be a maintenance program or update that is slowing traffic? Was that backup process scheduled for 12 a.m. or 12 p.m., right in the middle of the sales workday? Scheduling processes is an important part of enhancing your database performance.

One thing to realize is database problems are not always related to the database you're interested in optimizing. Database professionals tend to automatically point their fingers at the database vendor, for instance, when the problems originate elsewhere in the system. Maybe your database needs shipping information from a third party shipping company. If it's waiting to confirm a large group of addresses, this could cause a big performance hit as your database waits for that information to be returned. Those sorts of external dependencies may cause delays, as can applications, such as one that may be running a package that invalidates objects.

These are but a few examples of potential issues. Your needs may correspond to these examples, or look nothing like them. We know that every organization is different. The 5S approach helps evaluate your system to determine the changes and improvements needed to optimize performance. Following a measured process while documenting your changes, provides you tangible evidence that can be presented and understood.

Conclusion

What should you expect from performance tuning? First, your expectations need to be reasonable. As database administration managers know, performance tuning is an iterative, ongoing process that takes place over time. There is no fixed, one-size-fits-most solution that resolves performance issues in one magic pass. There are many possible factors slowing your system. A tiny change may inadvertently affect how another portion of the system works.

There are several benefits to following the 5S approach. First, you get to root cause much faster by eliminating the distractions of all the

possibilities that are extraneous. You have greater degree of insight after arriving to a solution; no tuning voodoo required. Time and money spent looking is now spent finding root cause with better use of resources (people & hardware). The 5S approach works both in reactive & proactive situations. Outcome provides measured response and thus root cause reporting is much easier.

But, it is important to understand that performance tuning cannot be a panacea for poor data or programming in your organization's relational database management system.

The approach Datavail takes is unique. We get to know your organization, your operating environment, and its processes to better understand the challenges and problems you are facing. Our experts ask the right questions, discovering the factors at the heart of the issues you face. Then, we present a solution tailored specifically to your operation.

To learn more about our database performance tuning services, call Datavail toll-free at (866) 828-7843. Prefer to chat online? We have experts available 24x7x365 to answer your questions on our chat line at Datavail.com.

BIOGRAPHY

Chuck Ezell Senior Applications Tuner for Datavail



Chuck supports some of the world's largest retailers with database performance tuning for their internal applications. Chuck helps Datavail locate and eliminate bottlenecks in database administration, while uncovering opportunities to improve performance. He excels at optimizing and customizing EBS systems. Chuck works with Java, SQL, .NET, and other languages on systems by Oracle, HP, IBM, among others. He uses AppDynamics, Splunk and Visual VM, along with other tools, for database tuning operations.

ABOUT DATAVAIL

Datavail Corporation is one of the largest providers of remote database administration (DBA) services in North America, offering database design and architecture, administration and 24x7 support. The company specializes in Oracle, Oracle E-Business Suite, Microsoft SQL Server, MySQL, MongoDB and DB2, and provides flexible onsite/offsite, onshore/offshore service delivery options to meet each customer's unique business needs.

CONTACT US

General Inquiries: 1-877-634-9222

Fax Number: 303-469-2399

Email: info@datavail.com

Corporate Headquarters:

Datavail Corporation
11800 Ridge Parkway
Suite 125
Broomfield, CO 80021

Database Operations Control Center:

Datavail Infotech Pvt. Ltd
3rd Floor, Unit No. B-3
Ashar IT Park, Road No. 16Z
Wagale Estate
Thane (West), Thane 400604
Direct Telephone Number: 022-61517000

Bangalore Office

Majestic Terrace - Plot 62B
Phase 1 - Opposite Post Office/Police Station
Electronic City
Bangalore - 560100, Karnataka

Seattle Office

1408 4th Avenue, Suite 200
Seattle, WA 98101

New York Office

112 W 34th Street, Suite 1403
New York, NY 10120



datAvail

DATABASE SERVICES

datavail.com | 877.634.9222