

SQL Server Performance Methodology with Oracle Applications

ISSN 1941-7802

TechNote Number: SQL-DISK-006
www.computationpress.com

© 2007 Computation Press, LLC. and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Computation Press disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Computation Press's research may discuss legal issues related to the information technology business, Computation Press does not provide legal advice or services and its research should not be construed or used as such. Computation Press shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The opinions expressed herein are subject to change without notice.

SQL Server Performance Methodology with Oracle Applications



By
Frank Earl McBath
frankmcb@hotmail.com

Team McBath: Looks Good, Smells Bad

(I.E. What Frank Believes)

- What did coach tell you? DIG! DIG! DIG! Can't? DROP AND GIVE ME 20!
We don't quit.
- Our customers buy answers, not problems.
- Queries drive almost all things bad.
- Run lean and get maximum output from minimum resources. We believe in scalability and HA.
- Let the computer tell you what the problems are.
- Keeps things simple.
- We're trained as nuclear arms negotiators: Trust but verify what customers tell us.
- We don't believe anything a vendors says.
- Never trust the sales guy.
- Automates as much as we can.
- We believe in McBath's: Oil Argument, Theory of Database Gravity and Theory of Two.
- We use command line and BCP, not SSIS or DTS.
- Do you pack the gear to be part of Team McBath?

Who I Am

- Global Database Technical Specialist for the Microsoft-Oracle Alliance
 - SQL Server worked with since 4.21.
 - Oracle since 5.0
- Been at Microsoft for 7 years in Consulting, Operations & Technical Sales
- Author/Co-Author/Contributor/Presenter
 - SQL Server Backup & Recovery, Prentice Hall
 - SQL Server 2000 High Availability, Microsoft Press
 - SQL Server 2000 Resource Kit, Microsoft Press
 - SQL Server 7 Resource Kit, Microsoft Press
 - SQL Server 2005 Administrators Companion, Microsoft Press
- Wrote most technical papers on
 - www.microsoft-oracle.com
- Database + Disk + Performance Blog
 - <http://databasediskperf.blogspot.com/>

Agenda

- Define Performance
- Oracle Apps & SQL Server
- Tools
- Methodology
- Case Study

Note...

- This deck will tend to focus on Query tools & tuning.
 - Will cover in workshop
- Disk Sub Systems cause most of the issues we see.
 - This is the next talk...

What is Performance?

- Not about how fast I can make one query run, but get consistent performance over time.
- Performance is an iterative process.
 - There will always be a “Top 10”
- Everyone wants sub-second... but no one defines an SLA with the business units...
- Finally, let the computer tell you what's wrong with itself... *and then fix itself...*

Why Are Oracle Apps Different? A Few Examples...

- Dynamic Query Generator
 - Often have to live with what you have...
- Usually more indexes than needed
- Server Side API Cursors
- Difficult to repro without understanding
- Batch Loading Tools

When I Go Onsite...

- First things I ask:
 - What is the SLA? (Everyone wants subsecond...)
 - What are the top 10... worst, run most often...
- This tells me a lot...
 - How realistic the people are...
 - How well you know your system...
 - Do they even know where the problem is?
 - Do they know what the tools are to monitor?
 - Do they know how to read the tools?
 - Tells me where I need to start working and how I need to talk with the customer...
- Look at database layout...
 - Look for cleanliness... Shows me how your shop runs... Operational hygiene.
 - SP_HELPDDB shows me a whole lot...

The Essential Tools

- Query Analyzer / Management Studio
- Database Tuning Advisor (DTA)
- SQL Server Profiler (profiler)
- Performance Monitor (perfmon)
- Dynamic Management Views (DMV) in SQL Server 2005
- Plan Guides in SQL Server 2005
- BACKUP to NUL:

Secondary...

- IOMeter
- SQLIOStress, SQLIOSim
- DISKPART
- HDTach
- OSTRESS

A Repeatable Methodology

- Capture the Data
 - Visually look for patterns: Number of Reads, Number of RESERVED, ... Is everything the same?
- Define what the Issue Actually Is
 - Look for “Top 10”
 - Quantify the results
 - Look at access patterns in the data
- Solve the Issue
 - Define plan of attack
 - Get a baseline
 - Iterate through the issues
 - Figure out issues: database? business rule?

Quest Magazine

- Page 58:
 - How To Find Queries That Are Killing Your System
 - Step by Step...
 - Cross posted on
 - www.computationpress.com

Methodology: The Game Plan

- The Process



Adobe Acrobat 7.0
Document

- The Script



C:\tmp\
methodology\queries.sql

Easy, Hard & Very Difficult

- Easy:
 - LRQ
 - Example: leading wildcards, third party app hitting database directly (ex. BizTalk)
- Hard:
 - Query running 10,000 times a day in 1.5 seconds
 - Example: A workflow query that is run after every screen.
- Very Difficult:
 - Query running 100,000 times a day in subsecond
 - Example: SAP table and packing data closer

How the Apps are Same

- All the apps are Server Side API Cursors. Some are JDBC (JDE). Some are ODBC (PSFT & SEBL). All look the same in SQL Profiler.
- All prepare, execute, fetch then reuse the plan. This allows them to mimic the behavior of stored procedures.
- All can be nebulous to filter, read, repro, and debug due to the fetching mechanisms.
 - For example, you can't filter on a statement that has not completed. Hence, you may miss the executing SELECT statement and only see the "fetch" as the offender.

How the Apps are Different

- JD Edwards shows up as relatively simple joins and uses "sp_execute" to fetch the result set back.
- Siebel has complicated joins from 25 to 80 tables and uses "sp_cursorfetch" to bring back the data.
- PeopleSoft does simple joins, uses sp_execute, but does not "parametrize" queries. Hence, tool doesn't group by "common" WHERE clauses well.
- Siebel tends to use 1:1 users, whereas PeopleSoft & JD Edwards use one proxy user.

Key Concepts in EA Perf Work

- What Bad Queries Do (Next Slide)
- Query Repro
 - How to Read A Plan
 - API Cursors: Fast Forward, Dynamic, Key Set, ...
- Disk Architecture
 - One big file, all on RAID 5, etc...
- Memory Architecture: 32 bit vs. 64 bit
 - Not going to solve problems, maybe mask them.
- Look for 80/20, Look for Patterns

Bad Queries

- Poor Queries artificially create problems:
 - They hammer the disk,
 - Read more data than they need causing excess memory to be consumed,
 - Consume more CPU than needed.

Query Repro

- Everyone does it wrong...
 - They spool the query, cut & paste, then hard code the variables...
- What you need to do...
 - Capture the RPC:STARTING, cut & paste this into Query Analyzer
 - If you can't... cut and paste the header from any RPC:STARTING onto the top of an RPC:COMPLETE
- Use STATISTICS IO and STATISTICS PROFILE

Impact of Architecture

- Everything sits on top of disk. If disk is slow, everything else slows down.
- Bad disk aggravated by poor queries.
 - Goal: 10ms response times for data and log.
 - I don't care how you do it... RAMSAN⁽¹⁾... magic wand... whatever... just give me 10ms or better.
- McBath's Oil Argument

<http://www.superssd.com/products/ramsan-400/>

Disk Architecture

- Technical Note 9: Best Practices for Microsoft SQL Server Disk Layout with Siebel Applications

This note seeks to cover some of the common issues that we see in database/disk configuration with Oracle Apps and types of configurations that you can use when planning the layout of your database. It is not meant to substitute for careful analysis or working with Oracle Technical Services, your hardware vendor or Microsoft.

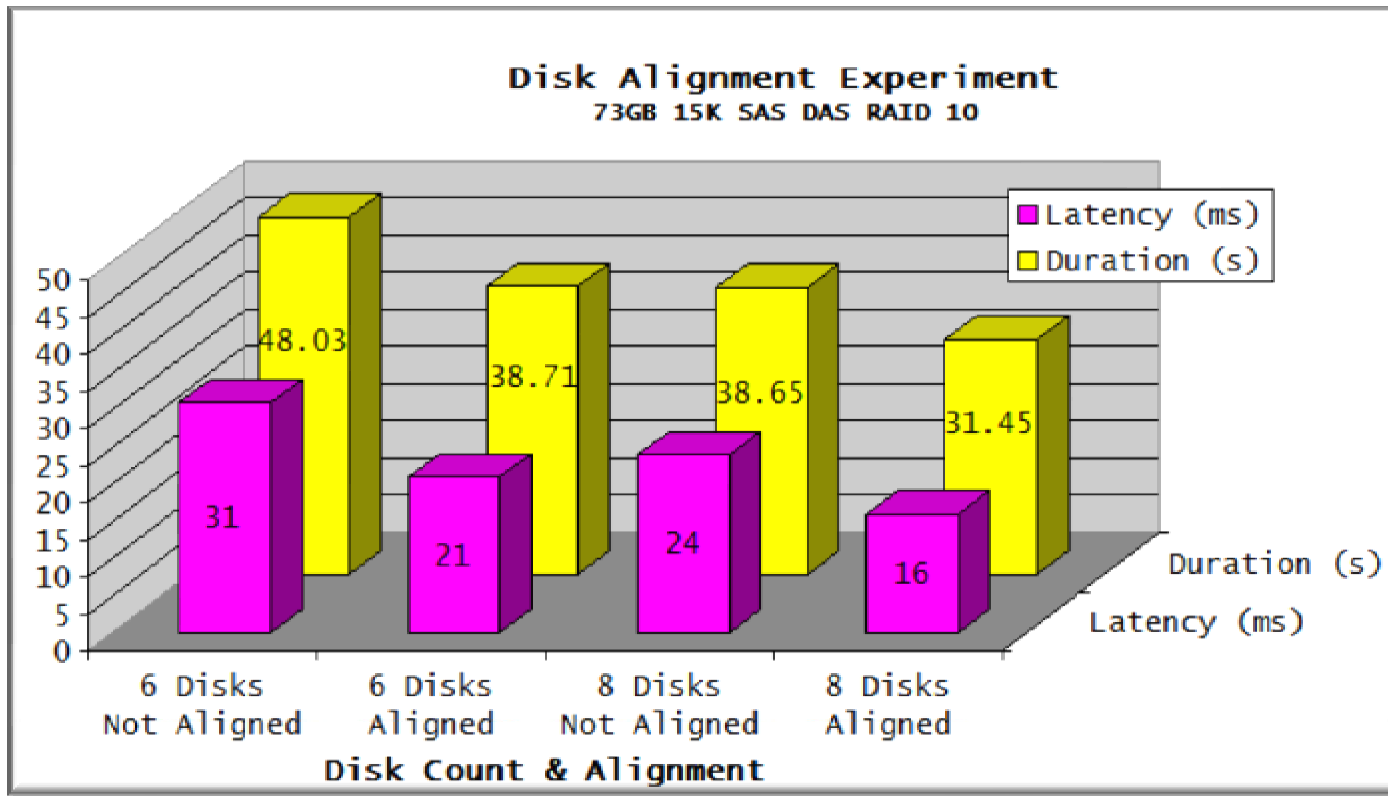
[Read More](#)

- <http://www.microsoft-oracle.com/Siebel/Pages/TechnicalNotes.aspx>

DISKPART: Free Money

- Windows utility to configure disk partitions and align with the correct offset.
- It is NOT a default in Windows Server 2003!
 - Is default in Vista and greater. 1M Offset.
 - See:
 - HKLM\SYSTEM\CurrentControlSet\Services\VDS\Alignment
- 18% to 20% increase in query response time without changing any syntax or business logic in the following examples.
- Full discussion outside the scope of this discussion. See:
 - Video: Database Disk Performance and IO for Enterprise Applications
 - <http://www.microsoft-oracle.com/Siebel/Pages/Workshops.aspx>

DISKPART: Performance Increase



Source: Jimmy May, Microsoft Consulting Services

Oracle Apps & Database Memory

- Technical Note 21: Memory Addressing for Siebel Servers

The matrix and explanation of memory settings in this TechNote can serve as a quick guide to making sure your server's memory parameters are set properly for a Oracle Apps Database Server running Microsoft SQL Server (32 bit).

[Read More](#)

- <http://www.microsoft-oracle.com/Siebel/Pages/TechnicalNotes.aspx>

Impact of Hints

- What are Hints?
- Where are hints bad?
 - EIM
- When are they good?
 - Non uniform data
 - Bad Plans
- How to Add them:
 - SQL Server 2000
 - Not clean
 - Have to invalidate the plan
 - SQL Server 2005
 - Plan Guides
 - Force recompiles, Add join hints, etc...
- How get rid of hints:
 - Use the 8602 database trace flag.
 - Technote 33: Hint Removal
 - <http://www.microsoft-oracle.com/Siebel/Pages/TechnicalNotes.aspx>

Hints: Example (I)

```
/*
UPDATE   BT
SET      BT.END_DT = IT.AS_END_DT,
         BT.NAME = IT.AS_NAME,
         BT.START_DT = IT.AS_START_DT,
         BT.X_CHANGE_CODE = IT.X_CHANGE_CODE,
         BT.X_CHANGE_DATE = IT.X_CHANGE_DATE,
         BT.X_CHANGE_TYPE = IT.X_CHANGE_TYPE,
         BT.X_POLICY_TYPE = IT.X_POLICY_TYPE,
         BT.X_PREMIUM = IT.X_PREMIUM,
         BT.X_PRINTED_FLG = IT.X_PRINTED_FLG,
         BT.X_PRODUCT_DESC = IT.X_PRODUCT_DESC,
         BT.X_PRODUCT_TYPE = IT.X_PRODUCT_TYPE,
         BT.X_RATE_PLAN_CD = IT.X_RATE_PLAN_CD,
         BT.X_SOURCE_SYSTEM = IT.X_SOURCE_SYSTEM,
         BT.LAST_UPD = @P1,
         BT.LAST_UPD_BY = @P2,
         BT.MODIFICATION_NUM = BT.MODIFICATION_NUM + 1
FROM      dbo.S_ASSET BT (INDEX = S_ASSET_P1),
         dbo.EIM_ASSET IT (INDEX = EIM_ASSET_M1)
WHERE (BT.ROW_ID = IT.T_ASSET__RID AND
       IT.IF_ROW_BATCH_NUM = 10410001 AND
       IT.IF_ROW_STAT_NUM = 0 AND
       IT.T_ASSET__EXS = 'Y' AND
       IT.T_ASSET__UNQ = 'Y' AND
       IT.T_ASSET__DUP = 'N' AND
       IT.T_ASSET__STA = 0)
*/
/*
WITH HINTS:
Table 'S_ASSET'. Scan count 1273, logical reads 4038, physical reads 0, read-ahead reads 0.
Table 'EIM_ASSET'. Scan count 1, logical reads 5875, physical reads 0, read-ahead reads 0.
WITHOUT HINTS:
Table 'S_ASSET'. Scan count 1273, logical reads 4038, physical reads 0, read-ahead reads 0.
Table 'EIM_ASSET'. Scan count 1, logical reads 1774, physical reads 0, read-ahead reads 0.
*/
```

Hints: Example (II)

WITH HINT:

Table 'S_CONTACT'. Scan count 1142, logical reads 8008, physical reads 0, read-ahead reads 0.

Table 'EIM_CONTACT'. Scan count 1, logical reads **3162**, physical reads 0, read-ahead reads 0.

WITHOUT HINT:

Table 'S_CONTACT'. Scan count 1142, logical reads 8008, physical reads 0, read-ahead reads 0.

Table 'EIM_CONTACT'. Scan count 1, logical reads **231**, physical reads 0, read-ahead reads 0.

WITH HINT:

Table 'S_APPLD_CVRG'. Scan count 1, logical reads **394774**, physical reads 0, read-ahead reads **280810**.

Table 'EIM_ASSET5_FN'. Scan count 1, logical reads 366, physical reads 0, read-ahead reads 0.

WITHOUT HINT:

Table 'S_APPLD_CVRG'. Scan count 1268, logical reads **10203**, physical reads 697, read-ahead reads **0**.

Table 'EIM_ASSET5_FN'. Scan count 1, logical reads 366, physical reads 0, read-ahead reads 0.

In Workshop: Case Studies

- Let's walk through a few and use our tools and new methodology.
 - Bad Data Loading
 - Leading Wild Cards
 - Non-indexed columns
 - Intrusive third party applications
 - Poor Disk Architecture

Questions?