



Ignite IT Performance™

Tuna Helper Proven Process for SQL Tuning

Dean Richards
Senior DBA, Confio Software



Who Am I?

- Senior DBA for Confio Software
 - DeanRichards@confio.com
- Current – 20+ Years in Oracle, SQL Server
- Former – 15+ Years in Oracle Consulting
- Specialize in Performance Tuning
- Review Performance of 100's of Databases for Customers and Prospects
- Common Thread – Paralyzed by Tuning



Agenda

- Introduction
- Challenges
- Identify - Which SQL and Why
- Gather – Details about SQL
- Tune – Case Study
- Monitor – Make sure it stays tuned



- SQL Tuning is Hard
- This Presentation is an Introduction
 - 3-5 day detailed classes are typical
- Providing a Framework
 - Helps develop your own processes
 - There is no magic tool
 - Tools cannot reliably tune SQL statements
 - Tuning requires the involvement of you and other technical and functional members of team



Challenges

- Requires Expertise in Many Areas
 - Technical – Plan, Data Access, SQL Design
 - Business – What is the Purpose of SQL?
- Tuning Takes Time
 - Large Number of SQL Statements
 - Each Statement is Different
- Low Priority in Some Companies
 - Vendor Applications
 - Focus on Hardware or System Issues
- Never Ending



Identify – Which SQL

- Tracing a Session / Process
- User / Batch Job Complaints
- Highest I/O (LIO, PIO)
- SQL Performing Full Table Scans
- Known Poorly Performing SQL
- Highest Wait Times (Ignite, AWR, etc)




Identify – End-to-End

- **Business Aspects**
 - Who registered yesterday for SQL Tuning
 - Who uses this information?
 - Why does the business need to know this?
 - How often is the information needed?
- **Technical Information**
 - Review Tables, Indexes, Triggers, Views, etc
 - Understand Relationships
 - Know the Data (High Level)
- **End-to-End Process**
 - Understand Application Architecture
 - What Portion of the Total Time is Database



Identify – End-to-End Time





Wait Event Information

V\$SESSION

SID
USERNAME
SQL_ID
PROGRAM
MODULE
ACTION
PLAN_HASH_VALUE
ROW_WAIT_OBJ#

V\$SESSION_WAIT

SID
EVENT
P1, P1RAW, P2, P2RAW, P3, P3RAW
STATE (WAITING, WAITED...)

- Oracle 10g added this info to V\$SESSION

V\$SQL

SQL_ID
SQL_FULLTEXT

V\$SQLAREA

SQL_ID
EXECUTIONS
PARSE_CALLS
BUFFER_GETS
DISK_READS

V\$SQL_PLAN

SQL_ID
PLAN_HASH_VALUE

DBA_OBJECTS

OBJECT_ID
OBJECT_NAME
OBJECT_TYPE



Wait Event Information

V\$SESSION

SID
USERNAME
SQL_ID
PROGRAM
MODULE
ACTION
PLAN_HASH_VALUE
ROW_WAIT_OBJ#

V\$SESSION_WAIT

SID
EVENT
P1, P1RAW, P2, P2RAW, P3, P3RAW
STATE (WAITING, WAITED...)

- Oracle 10g added this info to V\$SESSION

V\$SQL

SQL_ID
SQL_FULLTEXT

V\$SQLAREA

SQL_ID
EXECUTIONS
PARSE_CALLS
BUFFER_GETS
DISK_READS

V\$SQL_PLAN

SQL_ID
PLAN_HASH_VALUE

DBA_OBJECTS

OBJECT_ID
OBJECT_NAME
OBJECT_TYPE



Wait Event Information

```
SELECT s.sql_id, sql.sql_text, sql.plan_hash_value,  
       DECODE(s.state, 'WAITING', s.event, 'CPU') waitevent,  
       s.p1, s.p2, s.p3  
FROM v$session s  
JOIN v$sql sql ON (  
    s.sql_id = sql.sql_id AND s.sql_address = sql.address  
)  
AND sql.sql_text LIKE 'SELECT%' -- substitute your own  
AND s.sid = 20                 -- if you know it  
AND <whatever else you know>
```



Wait Time Scenario

- Which scenario is worse?
- SQL Statement 1
 - Executed 100 times
 - Caused 100 minutes of wait time for end user
 - Waited 99% of time on “db file sequential read”
- SQL Statement 2
 - Executed 1 time
 - Caused 100 minutes of wait time for end user
 - Waited 99% on “enq: TX – row lock contention”



Identify – Simplification

- Break Down SQL Into Simplest Forms
 - Complex SQL becomes multiple SQL
 - Sub-Queries Should be Tuned Separately
 - UNION'ed SQL Tuned Separately
 - Get the definition of views
 - Are synonyms being used
- Use Execution Plan (later)
 - Helps isolate the portion of the query that is performing poorly



Identify – Summary

- Determine the SQL
- Understand End-to-End
- Measure Wait Time
- Simplify Statement
 - Based on Execution Plan



Gather - Metrics

- Get baseline metrics
 - How long does it take now
 - What is acceptable (10 sec, 2 min, 1 hour)
- Collect Wait Time Metrics – How Long
 - Locking / Blocking
 - I/O problem, Latch contention
 - May be multiple issues
 - All have different resolutions
- Document everything in simple language



- **EXPLAIN PLAN**
 - Estimated execution plan - can be wrong for many reasons
- **V\$SQL_PLAN (Oracle 9i+)**
 - Real execution plan
 - Use DBMS_XPLAN for display
- **Tracing (all versions)**
 - Works when you know a problem will occur
ALTER SESSION SET tracefile_identifier = dean;
ALTER SESSION SET sql_trace = true;
- **Historical – AWR, Confio Ignite**



All Plans Not Equal

```
SELECT company, attribute
FROM data_out WHERE segment = :B1
```

- Wait Time – 100% on “db file scattered read”
- Plan from EXPLAIN PLAN

```
SELECT STATEMENT Optimizer=ALL_ROWS (Cost=1 Card=1 Bytes=117)
  TABLE ACCESS (BY INDEX ROWID) OF 'DATA_OUT' (TABLE) (Cost=1 Card=1 Bytes=117)
    INDEX (UNIQUE SCAN) OF 'IX1_DATA_OUT' (INDEX (UNIQUE)) (Cost=1 Card=1)
```

- Plan from V\$SQL_PLAN using DBMS_XPLAN

```
select * from table(dbms_xplan.display_cursor('az7r9s3wpqg7n',0));
```

```
-----
```

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT				370 (100)	
* 1	TABLE ACCESS FULL	DATA_OUT	1	117	370 (4)	00:00:05

```
-----
```

```
Predicate Information (identified by operation id):
```

```
1 - filter(TO_BINARY_DOUBLE("SEGMENT")=:B1)
```



- **V\$SQL_BIND_CAPTURE**
 - STATISTICS_LEVEL = TYPICAL or ALL
 - Collected at 15 minute intervals

```
SELECT name, position, datatype_string, value_string
FROM   v$sql_bind_capture
WHERE  sql_id = '15uughacxfh13';
```

```
NAME          POSITION  DATATYPE_STRING  VALUE_STRING
-----
:B1           1  BINARY_DOUBLE
```

- **Bind Values also provided by tracing**
 - Level 4 – bind values
 - Level 8 – wait information
 - Level 12 – bind values and wait information



- Use TuningStats.sql
 - <http://support.confio.com/kb/1534>
- Provides data on objects in execution plans.
 - Table sizes
 - Existing indexes
 - Cardinality of columns
 - Segment sizes
 - Histograms and Data Skew
 - Many things the CBO uses
- Run it for any table involved in query




- Who registered yesterday for SQL Tuning

```
SELECT s.fname, s.lname, r.signup_date
FROM student s, registration r, class c
WHERE s.student_id = r.student_id
AND r.class_id = c.class_id
AND UPPER(c.name) = 'SQL TUNING'
AND r.signup_date BETWEEN
      TRUNC(SYSDATE-1) AND TRUNC(SYSDATE)
AND r.cancelled = 'N'
```

- Execution Time – 12:38
- Wait Time – 95% on “db file scattered read”



Relationship





Execution Plan

Id	Operation	Name	Rows	Bytes	Cost
0	SELECT STATEMENT				95
1	NESTED LOOPS		1	167	95
2	NESTED LOOPS		1	138	94
3	NESTED LOOPS		7	357	87
4	VIEW	VW_SQ_1	201	7035	87
* 5	FILTER				
6	HASH GROUP BY		201	3417	87
* 7	FILTER				
* 8	TABLE ACCESS FULL	REGISTRATION	80000	1328K	76
* 9	INDEX UNIQUE SCAN	SYS_C0036920	1	16	0
* 10	TABLE ACCESS BY INDEX ROWID	CLASS	1	87	1
* 11	INDEX UNIQUE SCAN	SYS_C0036919	1		0
12	TABLE ACCESS BY INDEX ROWID	STUDENT	1	29	1
* 13	INDEX UNIQUE SCAN	SYS_C0036918	1		0

Predicate Information (identified by operation id):

```

5 - filter((MAX("SIGNUP_DATE")>=TRUNC(SYSDATE@!-1) AND
           MAX("SIGNUP_DATE")<=TRUNC(SYSDATE@!)))
7 - filter(TRUNC(SYSDATE@!-1)<=TRUNC(SYSDATE@!))
8 - filter("CANCELLED"='N')
9 - access("R1"."STUDENT_ID"="STUDENT_ID" AND "R1"."CLASS_ID"="CLASS_ID" AND
           "SIGNUP_DATE"="VW_COL_1")
   filter(("SIGNUP_DATE">=TRUNC(SYSDATE@!-1) AND "SIGNUP_DATE"<=TRUNC(SYSDATE@!)))
10 - filter(UPPER("C"."NAME")='SQL TUNING')
11 - access("CLASS_ID"="C"."CLASS_ID")
13 - access("S"."STUDENT_ID"="STUDENT_ID")

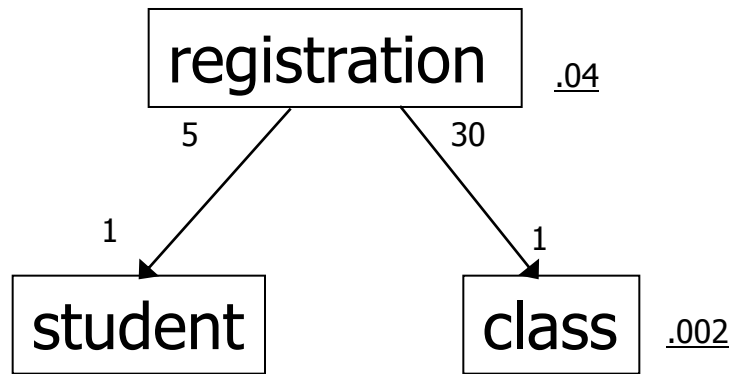
```



Gather – Summary

- Execution Plan
 - V\$SQL_PLAN
 - Do not use EXPLAIN PLAN
 - DBMS_XPLAN
- Bind Values
 - V\$SQL_BIND_CAPTURE
 - Tracing
- Table and Index Statistics
- ERD

- SQL Tuning – Dan Tow
 - Great book that teaches SQL Diagramming
 - <http://www.singingsql.com>



```
select count(1) from registration where cancelled = 'N'  
and signup_date between trunc(sysdate-1) and trunc(sysdate)
```

```
3562 / 80000 = .0445
```

```
select count(1) from class where UPPER(name) = 'SQL TUNING'
```

```
2 / 1000 = .002
```




New Plan

create index cl_uname on class (upper(name));

- Index on registration was (student_id, class_id)

create index reg_alt on registration (class_id);


Id	Operation	Name	Rows	Bytes	Cost
0	SELECT STATEMENT				10
* 1	FILTER				
2	NESTED LOOPS		1	132	7
3	NESTED LOOPS		1	103	6
* 4	TABLE ACCESS BY INDEX ROWID	CLASS	1	87	5
* 5	INDEX RANGE SCAN	CL_UNAME	4		1
* 6	INDEX RANGE SCAN	REG_ALT	1	16	1
7	SORT AGGREGATE		1	17	
* 8	TABLE ACCESS BY INDEX ROWID	REGISTRATION	1	17	3
* 9	INDEX RANGE SCAN	REG_ALT	1		2
10	TABLE ACCESS BY INDEX ROWID	STUDENT	1	29	1
* 11	INDEX UNIQUE SCAN	SYS_C0036918	1		0



■ Who cancelled classes within the week

```
SELECT s.lname, c.name, r.signup_date cancel_date
FROM   registration r, student s, class c
where  r.signup_date between sysdate and sysdate-7
AND    r.cancelled = 'Y'
AND    r.student_id = s.student_id
AND    r.class_id = c.class_id
```

- 30% of rows are dated within last week
- No index on CANCELLED column = FTS
- Will an index on CANCELLED column help?
 - Why or why not?



```
select count(1) from registration where cancelled = 'Y'  
and signup_date between trunc(sysdate-1) and trunc(sysdate)
```

622 / 80000 = .0077

```
select count(1) from registration where cancelled = 'Y'
```

638 / 80000 = .0079

```
select count(1) from registration  
where signup_date between trunc(sysdate-1) and trunc(sysdate)
```

11598 / 80000 = .1449



Query 2 Column Stats

```
create index reg_can on registration(cancelled);
```

```
select cancelled, count(1)  
from registration group by cancelled;
```

```
C    COUNT(1)  
-  -  
Y          638  
N       79345
```

- Oracle will not use an index on this column
 - Unless it has more information
 - CBO assumes an even data distribution
- Histograms give more information to Oracle
 - Based on skewed data, CBO realizes an index would be beneficial
 - Works best with literal values
 - Bind Variables – Oracle peeks first time only



Query 2 - Histogram

```
dbms_stats.gather_table_stats(  
  ownname    => 'STDMGMT',  
  tabname    => 'REGISTRATION',  
  method_opt=>'FOR COLUMNS cancelled SIZE AUTO')
```

```
-----  
| Id | Operation | Name | Rows | Bytes | Cost |  
-----  
| 0 | SELECT STATEMENT | | | | 7 |  
|* 1 | FILTER | | | | |  
|* 2 | TABLE ACCESS BY INDEX ROWID | REGISTRATION | 1 | 17 | 7 |  
|* 3 | INDEX RANGE SCAN | REG_CAN | 754 | | 2 |  
-----
```

- Monitor the improvement
 - Be able to prove that tuning made a difference
 - Take new metrics measurements
 - Compare them to initial readings
 - Brag about the improvements – no one else will
- Monitor for next tuning opportunity
 - Tuning is iterative
 - There are always room for improvements
 - Make sure you tune things that make a difference
- Shameless Product Pitch - Ignite



- Developer of Wait-Based Performance Tools
- Igniter Suite
 - Ignite for SQL Server, Oracle, DB2, Sybase
- Provides Help With
 - Identify
 - Gather
 - Monitor
- Based in Colorado, worldwide customers
- Free trial at www.confio.com