ORACLE®

Why is my SQL slow

Bob Carlin
Consulting Member of Technical Staff
Oracle Real-World Performance

Why is My SQL Slow ?

Problem Query

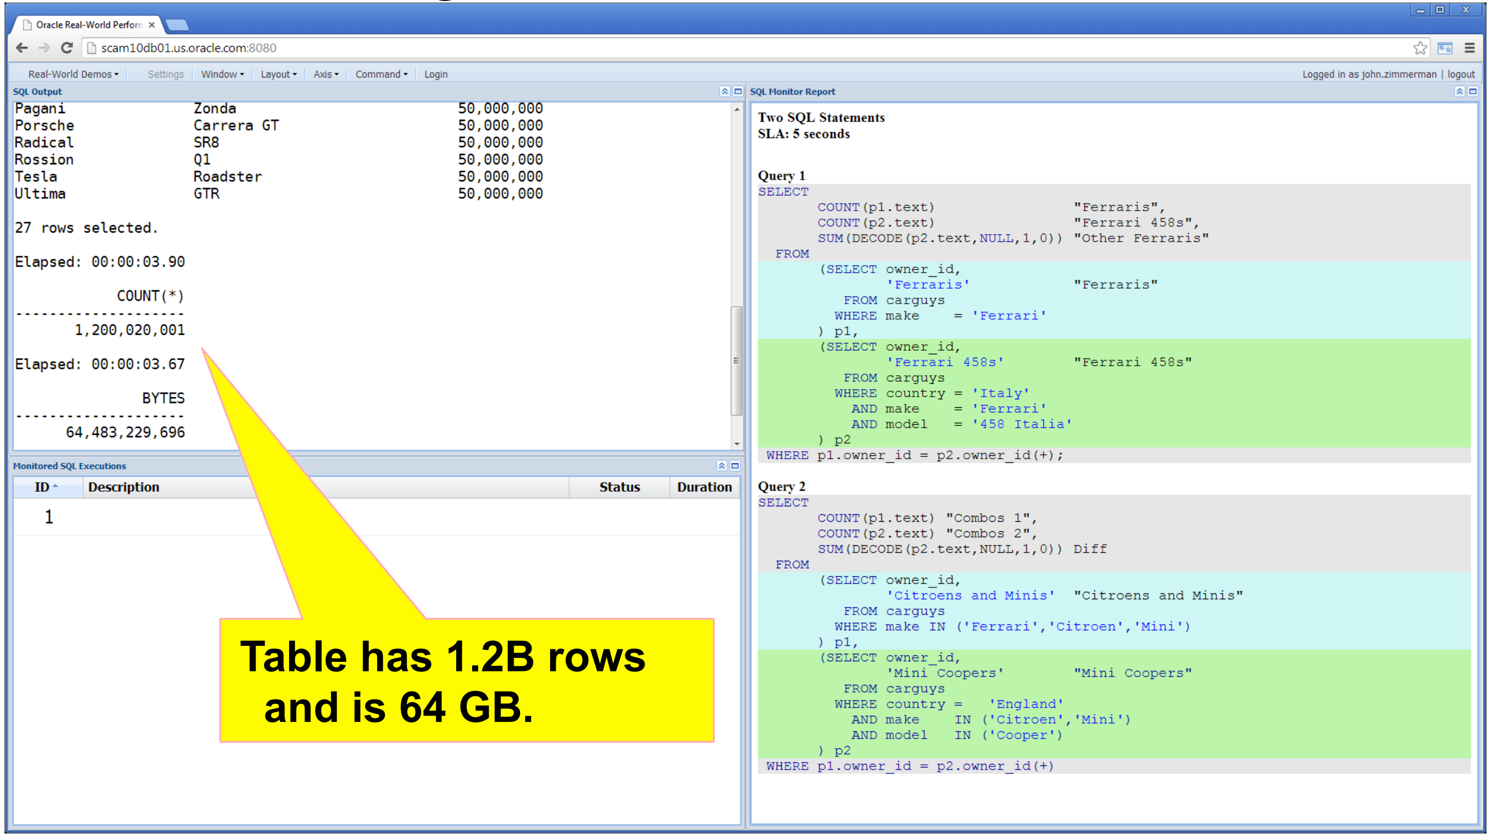

```
Pagani                Zonda                          50,000,000
Porsche               Carrera GT                     50,000,000
Radical               SR8                            50,000,000
Rossion               Q1                             50,000,000
Tesla                 Roadster                       50,000,000
Ultima                GTR                            50,000,000

27 rows selected.

Elapsed: 00:00:03.90

            COUNT(*)
--------------------
       1,200,020,001

Elapsed: 00:00:03.67

               BYTES
--------------------
      64,483,229,696
```

ID	Description	Status	Duration
1			

Two SQL Statements
SLA: 5 seconds

Query 1

```
SELECT
       COUNT(p1.text)                   "Ferraris",
       COUNT(p2.text)                   "Ferrari 458s",
       SUM(DECODE(p2.text,NULL,1,0))  "Other Ferraris"
  FROM
      (SELECT owner_id,
              'Ferraris'                "Ferraris"
         FROM carguys
        WHERE make     = 'Ferrari'
      ) p1,
      (SELECT owner_id,
              'Ferrari 458s'            "Ferrari 458s"
         FROM carguys
        WHERE country  = 'Italy'
          AND make     = 'Ferrari'
          AND model    = '458 Italia'
      ) p2
 WHERE p1.owner_id = p2.owner_id(+);
```

Query 2

```
SELECT
       COUNT(p1.text)  "Combos 1",
       COUNT(p2.text)  "Combos 2",
       SUM(DECODE(p2.text,NULL,1,0)) Diff
  FROM
      (SELECT owner_id,
              'Citroens and Minis'  "Citroens and Minis"
         FROM carguys
        WHERE make IN ('Ferrari','Citroen','Mini')
      ) p1,
      (SELECT owner_id,
              'Mini Coopers'           "Mini Coopers"
         FROM carguys
        WHERE country =   'England'
          AND make    IN ('Citroen','Mini')
          AND model   IN ('Cooper')
      ) p2
 WHERE p1.owner_id = p2.owner_id(+)
```

Table has 1.2B rows and is 64 GB.

Problem Query

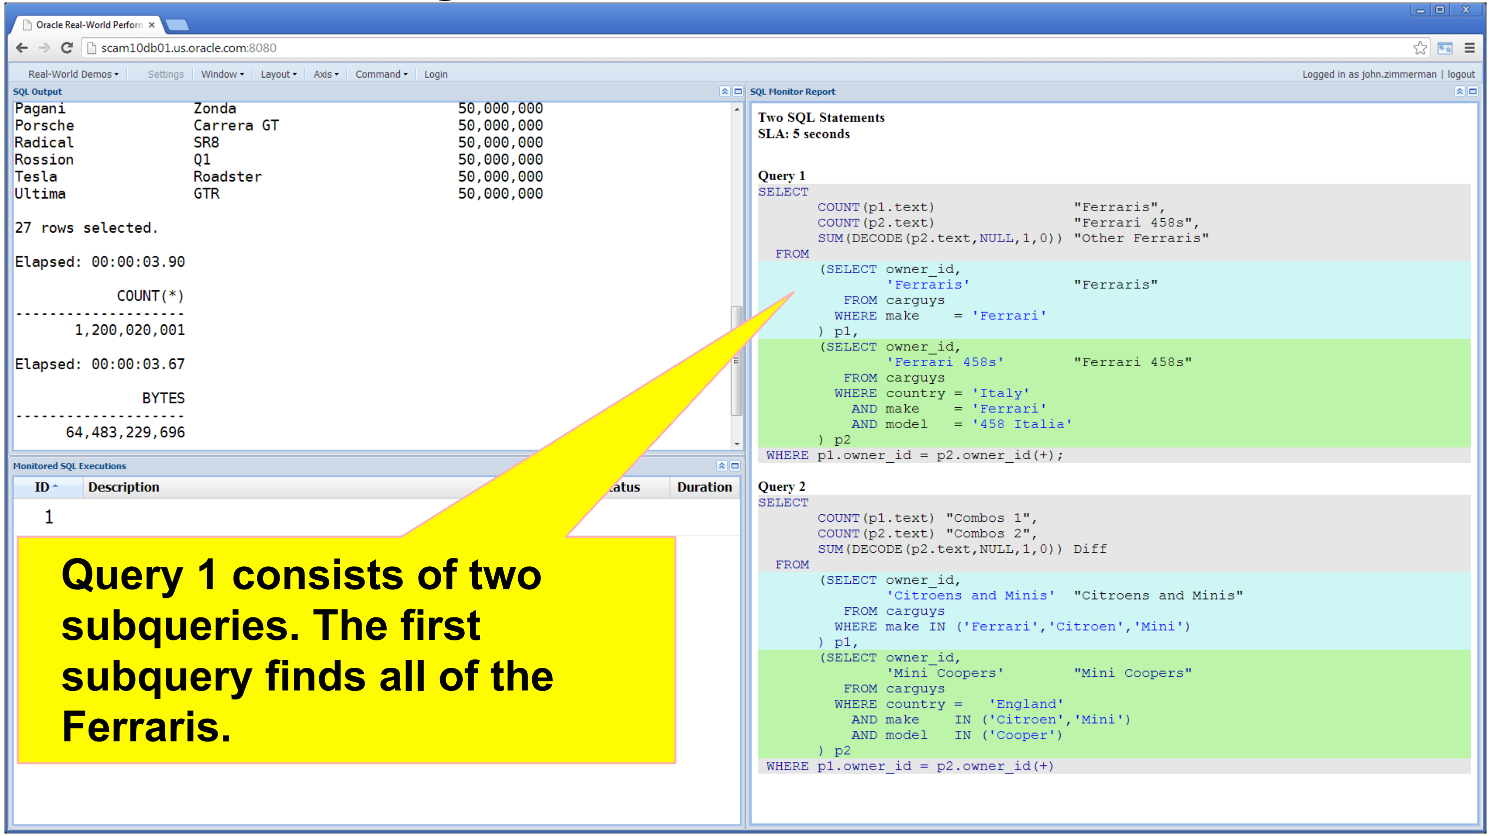

```
Pagani          Zonda                       50,000,000
Porsche         Carrera GT                  50,000,000
Radical         SR8                         50,000,000
Rossion         Q1                          50,000,000
Tesla           Roadster                    50,000,000
Ultima          GTR                         50,000,000

27 rows selected.

Elapsed: 00:00:03.90

            COUNT(*)
--------------------
       1,200,020,001

Elapsed: 00:00:03.67

               BYTES
--------------------
      64,483,229,696
```

ID	Description	atus	Duration
1			

Two SQL Statements
SLA: 5 seconds

Query 1

```
SELECT
       COUNT(p1.text)                   "Ferraris",
       COUNT(p2.text)                   "Ferrari 458s",
       SUM(DECODE(p2.text,NULL,1,0))    "Other Ferraris"
  FROM
      (SELECT owner_id,
              'Ferraris'                "Ferraris"
         FROM carguys
        WHERE make    = 'Ferrari'
      ) p1,
      (SELECT owner_id,
              'Ferrari 458s'            "Ferrari 458s"
         FROM carguys
        WHERE country = 'Italy'
          AND make    = 'Ferrari'
          AND model   = '458 Italia'
      ) p2
 WHERE p1.owner_id = p2.owner_id(+);
```

Query 2

```
SELECT
       COUNT(p1.text)  "Combos 1",
       COUNT(p2.text)  "Combos 2",
       SUM(DECODE(p2.text,NULL,1,0))  Diff
  FROM
      (SELECT owner_id,
              'Citroens and Minis'  "Citroens and Minis"
         FROM carguys
        WHERE make IN ('Ferrari','Citroen','Mini')
      ) p1,
      (SELECT owner_id,
              'Mini Coopers'        "Mini Coopers"
         FROM carguys
        WHERE country =   'England'
          AND make    IN ('Citroen','Mini')
          AND model   IN ('Cooper')
      ) p2
 WHERE p1.owner_id = p2.owner_id(+)
```

Query 1 consists of two subqueries. The first subquery finds all of the Ferraris.

Problem Query

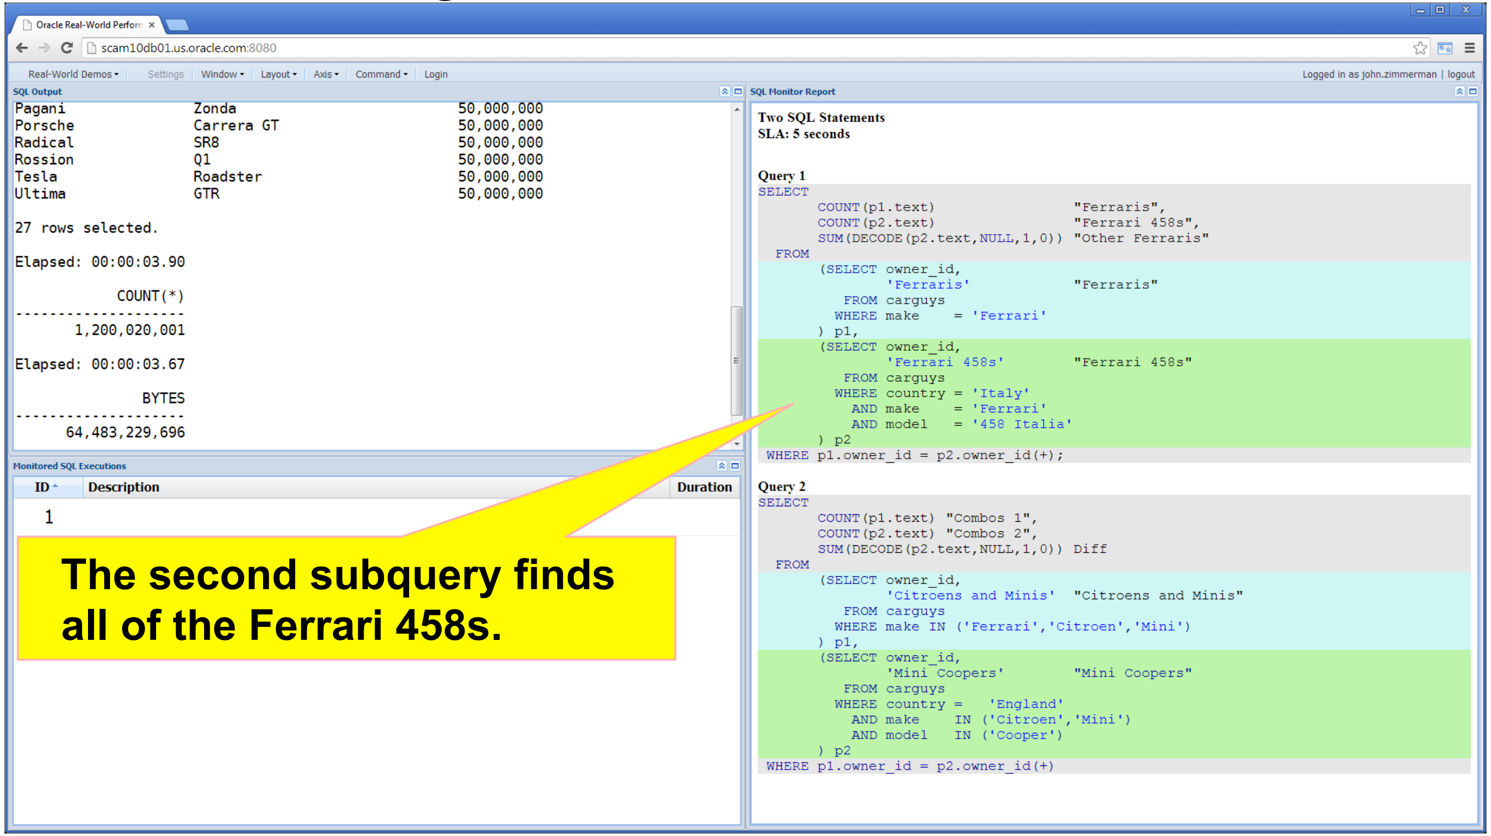

```
Pagani                Zonda                          50,000,000
Porsche               Carrera GT                     50,000,000
Radical               SR8                            50,000,000
Rossion               Q1                             50,000,000
Tesla                 Roadster                       50,000,000
Ultima                GTR                            50,000,000

27 rows selected.

Elapsed: 00:00:03.90

            COUNT(*)
--------------------
       1,200,020,001

Elapsed: 00:00:03.67

               BYTES
--------------------
      64,483,229,696
```

ID	Description	Duration
1		

The second subquery finds all of the Ferrari 458s.

Two SQL Statements
SLA: 5 seconds

Query 1

```
SELECT
       COUNT(p1.text)                  "Ferraris",
       COUNT(p2.text)                  "Ferrari 458s",
       SUM(DECODE(p2.text,NULL,1,0))   "Other Ferraris"
  FROM
        (SELECT owner_id,
                'Ferraris'             "Ferraris"
           FROM carguys
          WHERE make     = 'Ferrari'
        ) p1,
        (SELECT owner_id,
                'Ferrari 458s'         "Ferrari 458s"
           FROM carguys
          WHERE country  = 'Italy'
            AND make     = 'Ferrari'
            AND model    = '458 Italia'
        ) p2
 WHERE p1.owner_id = p2.owner_id(+);
```

Query 2

```
SELECT
       COUNT(p1.text)  "Combos 1",
       COUNT(p2.text)  "Combos 2",
       SUM(DECODE(p2.text,NULL,1,0)) Diff
  FROM
        (SELECT owner_id,
                'Citroens and Minis'   "Citroens and Minis"
           FROM carguys
          WHERE make IN ('Ferrari','Citroen','Mini')
        ) p1,
        (SELECT owner_id,
                'Mini Coopers'         "Mini Coopers"
           FROM carguys
          WHERE country =    'England'
            AND make    IN ('Citroen','Mini')
            AND model   IN ('Cooper')
        ) p2
 WHERE p1.owner_id = p2.owner_id(+)
```

Problem Query

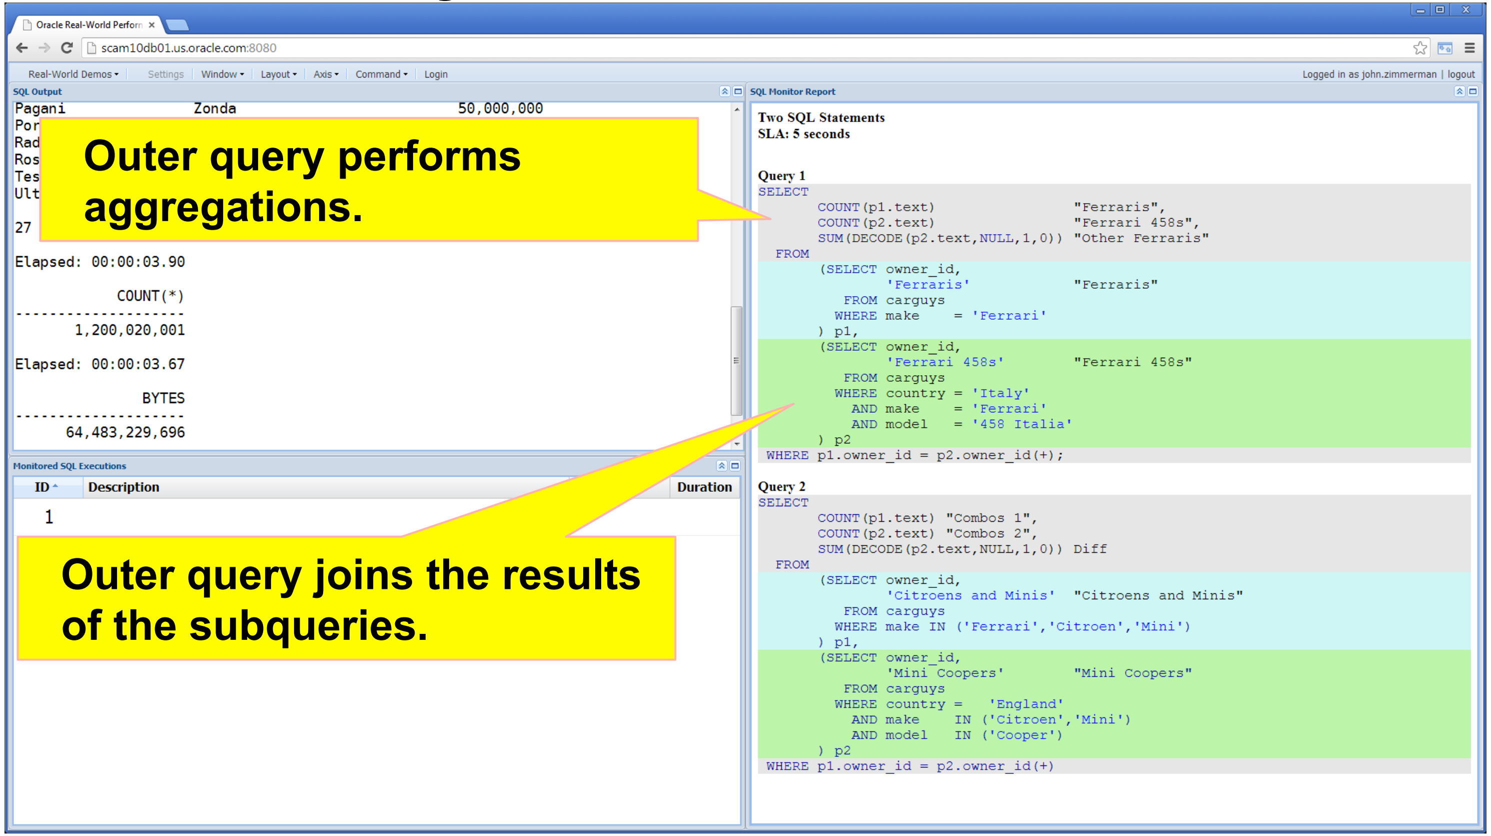

```
Pagani                 Zonda                             50,000,000
Por
Rad
Ros
Tes
Ult

27

Elapsed: 00:00:03.90

             COUNT(*)
--------------------
        1,200,020,001

Elapsed: 00:00:03.67

                BYTES
--------------------
       64,483,229,696
```

ID	Description	Duration
1		

Outer query performs aggregations.

Outer query joins the results of the subqueries.

Two SQL Statements
SLA: 5 seconds

Query 1

```
SELECT
       COUNT(p1.text)                    "Ferraris",
       COUNT(p2.text)                    "Ferrari 458s",
       SUM(DECODE(p2.text,NULL,1,0))     "Other Ferraris"
  FROM
       (SELECT owner_id,
               'Ferraris'                "Ferraris"
          FROM carguys
         WHERE make    = 'Ferrari'
       ) p1,
       (SELECT owner_id,
               'Ferrari 458s'            "Ferrari 458s"
          FROM carguys
         WHERE country = 'Italy'
           AND make    = 'Ferrari'
           AND model   = '458 Italia'
       ) p2
 WHERE p1.owner_id = p2.owner_id(+);
```

Query 2

```
SELECT
       COUNT(p1.text)  "Combos 1",
       COUNT(p2.text)  "Combos 2",
       SUM(DECODE(p2.text,NULL,1,0))  Diff
  FROM
       (SELECT owner_id,
               'Citroens and Minis'   "Citroens and Minis"
          FROM carguys
         WHERE make IN ('Ferrari','Citroen','Mini')
       ) p1,
       (SELECT owner_id,
               'Mini Coopers'          "Mini Coopers"
          FROM carguys
         WHERE country =   'England'
           AND make    IN ('Citroen','Mini')
           AND model   IN ('Cooper')
       ) p2
 WHERE p1.owner_id = p2.owner_id(+)
```

Problem Query

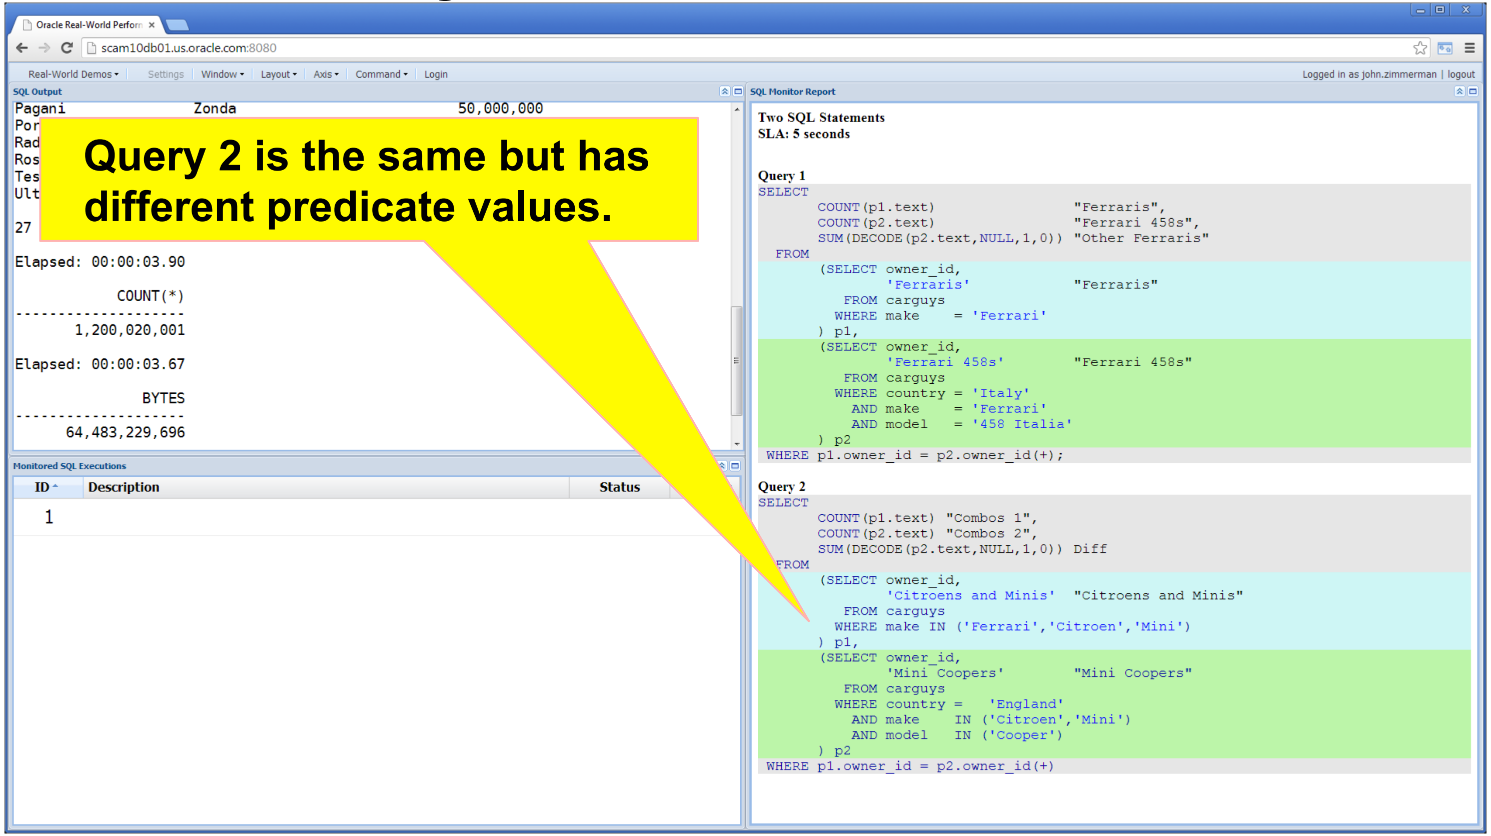

```
Pagani                   Zonda                                50,000,000
Por
Rad
Ros
Tes
Ult

27

Elapsed: 00:00:03.90

             COUNT(*)
--------------------
        1,200,020,001

Elapsed: 00:00:03.67

                BYTES
--------------------
       64,483,229,696
```

ID	Description	Status
1		

Query 2 is the same but has different predicate values.

Two SQL Statements
SLA: 5 seconds

Query 1

```
SELECT
       COUNT(p1.text)                  "Ferraris",
       COUNT(p2.text)                  "Ferrari 458s",
       SUM(DECODE(p2.text,NULL,1,0))  "Other Ferraris"
  FROM
      (SELECT owner_id,
              'Ferraris'               "Ferraris"
         FROM carguys
        WHERE make     = 'Ferrari'
      ) p1,
      (SELECT owner_id,
              'Ferrari 458s'           "Ferrari 458s"
         FROM carguys
        WHERE country  = 'Italy'
          AND make     = 'Ferrari'
          AND model    = '458 Italia'
      ) p2
 WHERE p1.owner_id = p2.owner_id(+);
```

Query 2

```
SELECT
       COUNT(p1.text)  "Combos 1",
       COUNT(p2.text)  "Combos 2",
       SUM(DECODE(p2.text,NULL,1,0))  Diff
  FROM
      (SELECT owner_id,
              'Citroens and Minis'  "Citroens and Minis"
         FROM carguys
        WHERE make IN ('Ferrari','Citroen','Mini')
      ) p1,
      (SELECT owner_id,
              'Mini Coopers'           "Mini Coopers"
         FROM carguys
        WHERE country =    'England'
          AND make     IN ('Citroen','Mini')
          AND model    IN ('Cooper')
      ) p2
 WHERE p1.owner_id = p2.owner_id(+)
```

Default Statistics

Default Statistics

REAL-WORLD PERFORMANCE

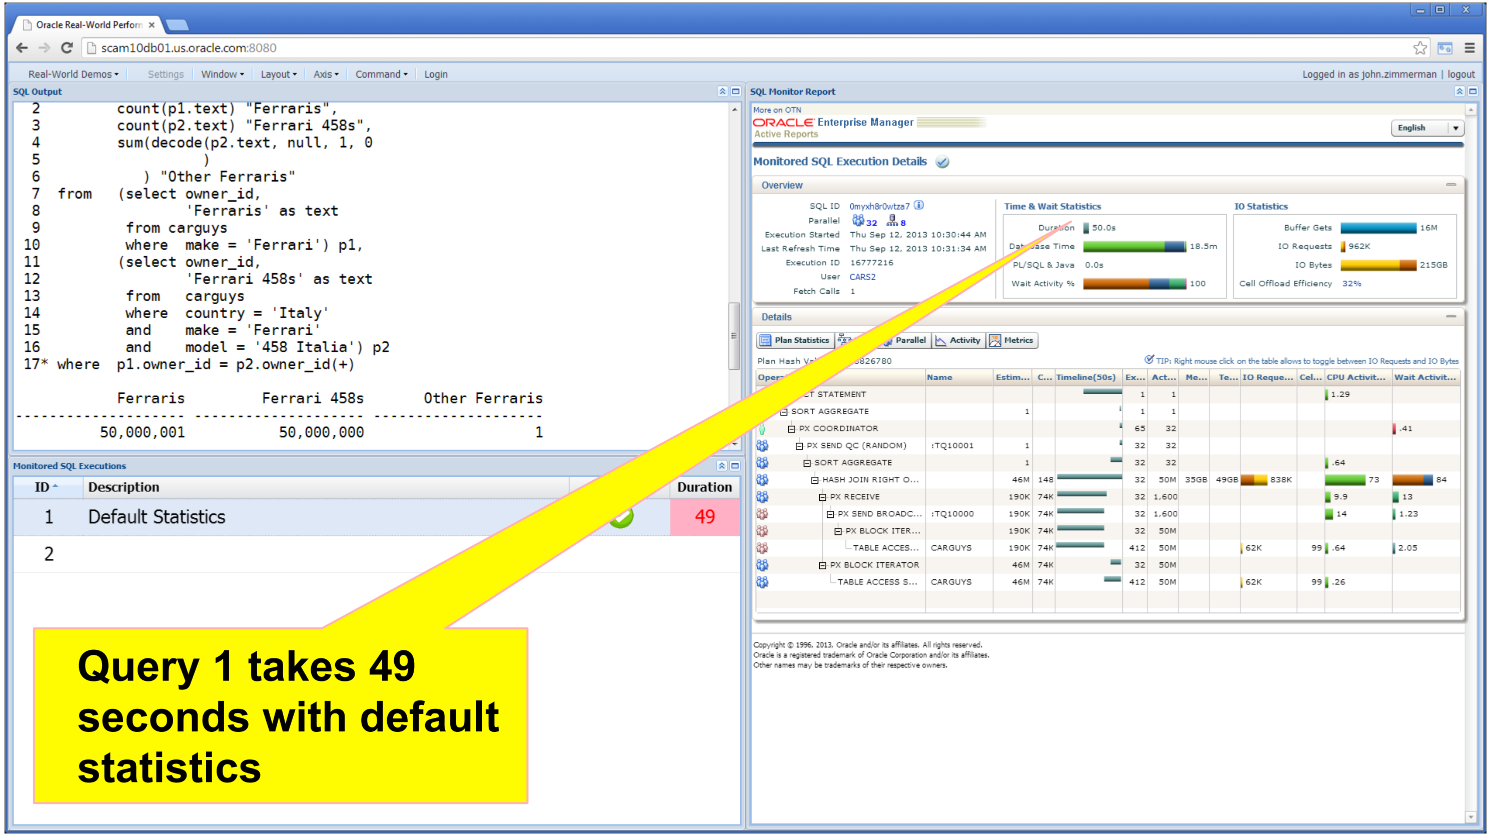

```
        count(p1.text) "Ferraris",
        count(p2.text) "Ferrari 458s",
        sum(decode(p2.text, null, 1, 0
                  )
           ) "Other Ferraris"
 from   (select owner_id,
                'Ferraris' as text
         from carguys
         where  make = 'Ferrari') p1,
        (select owner_id,
                'Ferrari 458s' as text
         from   carguys
         where  country = 'Italy'
         and    make = 'Ferrari'
         and    model = '458 Italia') p2
 where  p1.owner_id = p2.owner_id(+)

        Ferraris          Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1
```

Monitored SQL Executions

ID	Description	Duration
1	Default Statistics	49
2		

SQL Monitor Report

Monitored SQL Execution Details

Overview

SQL ID	0myxh8r0wtza7
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:30:44 AM
Last Refresh Time	Thu Sep 12, 2013 10:31:34 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	50.0s
Database Time	18.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Hash Value 826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(50s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
T STATEMENT					1	1					1.29	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						.41
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.64	
HASH JOIN RIGHT O...		46M	148		32	50M	35GB	49GB	838K		73	84
PX RECEIVE		190K	74K		32	1,600					9.9	13
PX SEND BROADC...	:TQ10000	190K	74K		32	1,600					14	1.23
PX BLOCK ITER...		190K	74K		32	50M						
TABLE ACCES...	CARGUYS	190K	74K		412	50M			62K	99	.64	2.05
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	46M	74K		412	50M			62K	99	.26	

Query 1 takes 49 seconds with default statistics

Default Statistics

REAL-WORLD PERFORMANCE

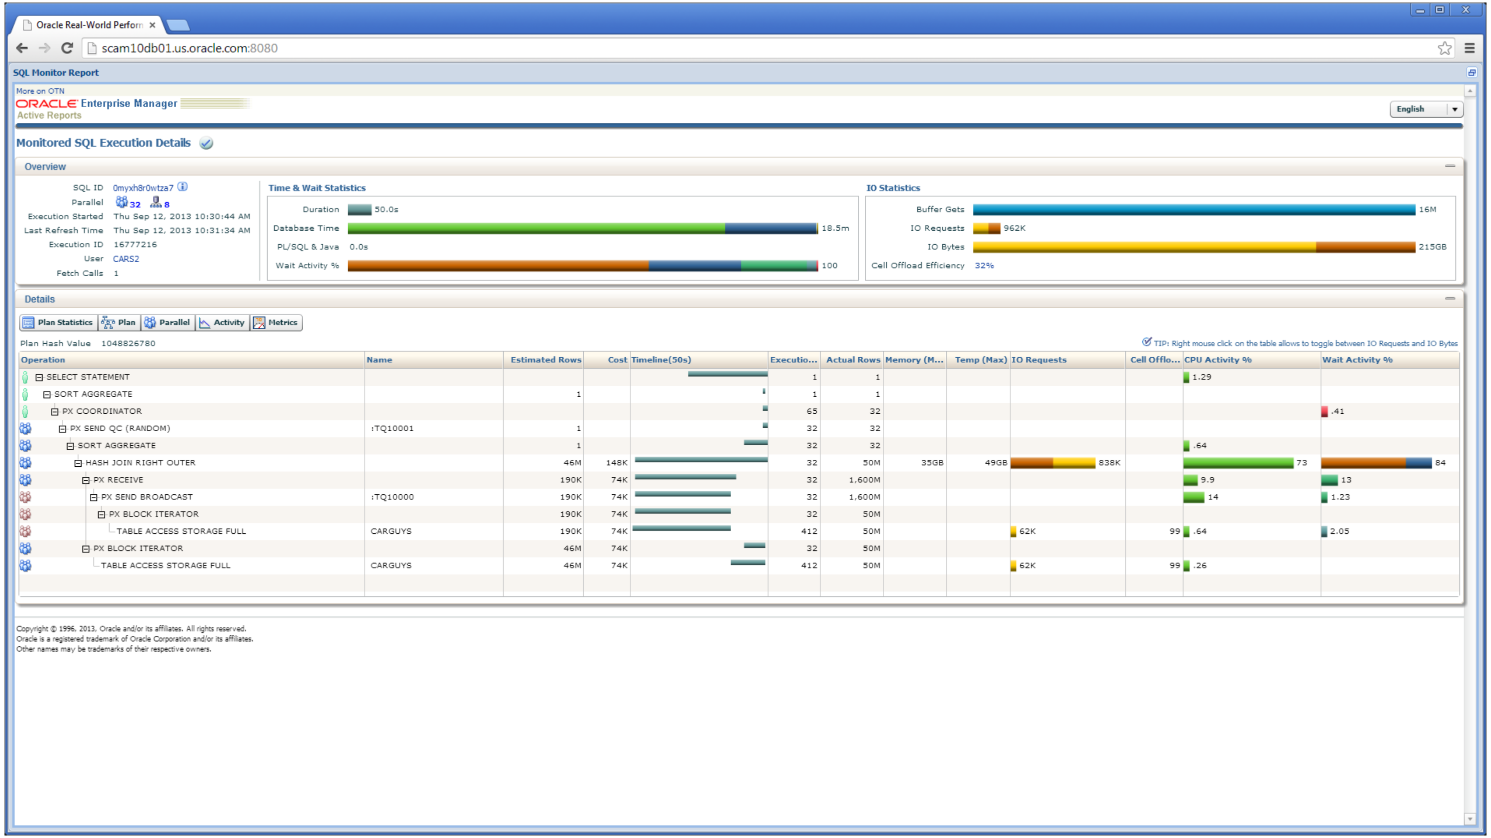

Oracle Real-World Perforn

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	0myxh8r0wtza7
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:30:44 AM
Last Refresh Time	Thu Sep 12, 2013 10:31:34 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	50.0s
Database Time	18.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(50s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					1.29	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						.41
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.64	
HASH JOIN RIGHT OUTER		46M	148K		32	50M	35GB	49GB	838K		73	84
PX RECEIVE		190K	74K		32	1,600M					9.9	13
PX SEND BROADCAST	:TQ10000	190K	74K		32	1,600M					14	1.23
PX BLOCK ITERATOR		190K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	190K	74K		412	50M			62K	99	.64	2.05
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	46M	74K		412	50M			62K	99	.26	

ORACLE
REAL-WORLD PERFORMANCE

Default Statistics

1. Development Findings

- Baseline Performance for Query 1
- Query 1 exceeds target

ORACLE

Initial Optimization Steps—More Predicate Values

More Predicate Values

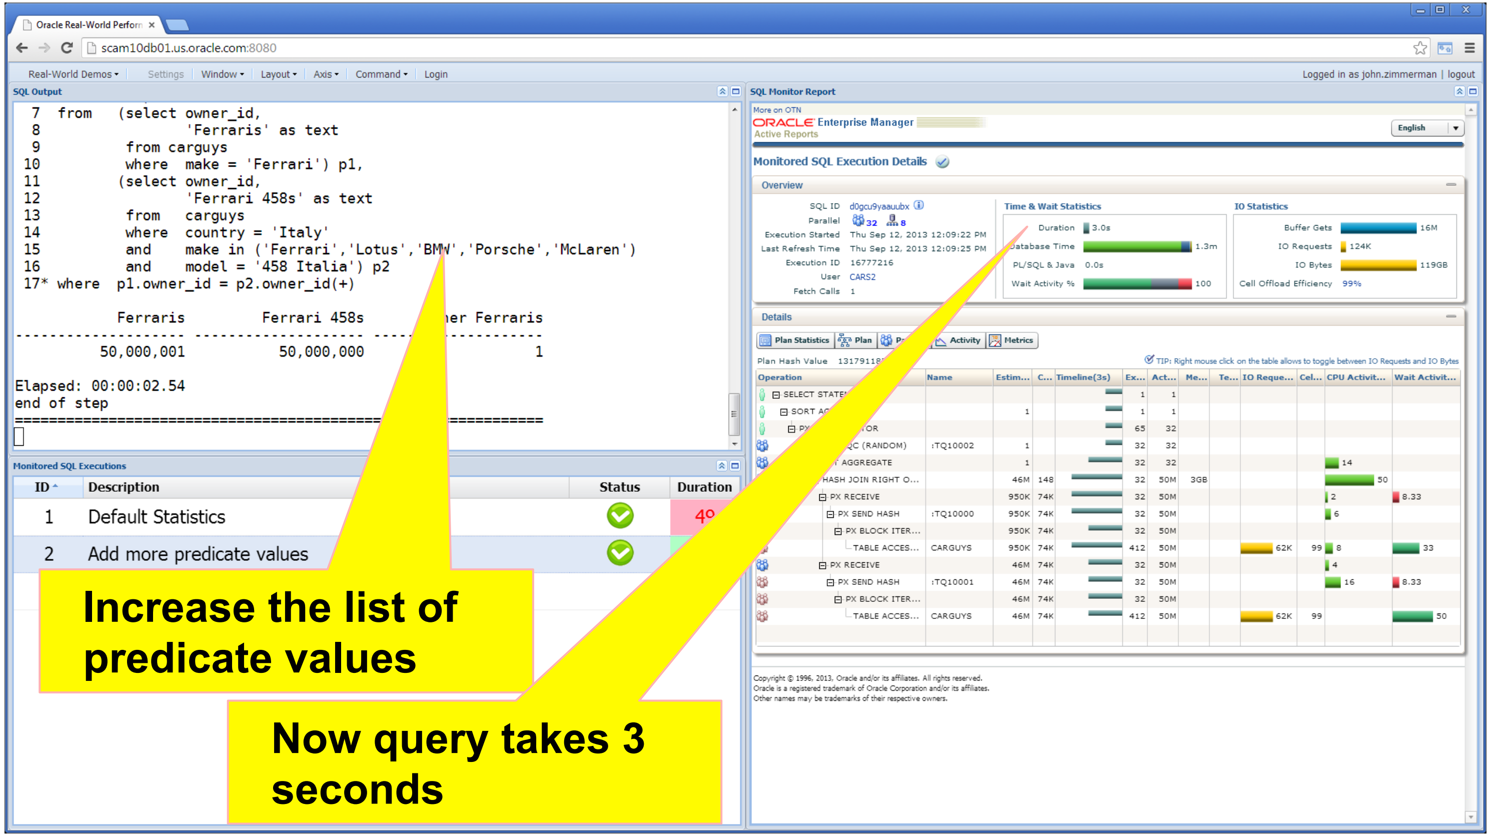

Oracle Real-World Perforn
scam10db01.us.oracle.com:8080
Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout
SQL Output
SQL Monitor Report
Monitored SQL Executions

```
  7  from   (select owner_id,
                   'Ferraris' as text
                from carguys
                where  make = 'Ferrari') p1,
               (select owner_id,
                   'Ferrari 458s' as text
                from   carguys
                where  country = 'Italy'
                and    make in ('Ferrari','Lotus','BMW','Porsche','McLaren')
                and    model = '458 Italia') p2
 17* where  p1.owner_id = p2.owner_id(+)

          Ferraris          Ferrari 458s          er Ferraris
-------------------- -------------------- ----  -------------
          50,000,001           50,000,000                     1

Elapsed: 00:00:02.54
end of step
==================================================================
```

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	d0gcu9yaauubx
Parallel	32 8
Execution Started	Thu Sep 12, 2013 12:09:22 PM
Last Refresh Time	Thu Sep 12, 2013 12:09:25 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.3m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Pa | Activity | Metrics

Plan Hash Value 13179118

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(3s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATE					1	1						
SORT AG		1			1	1						
PX TOR					65	32						
QC (RANDOM)	:TQ10002	1			32	32						
AGGREGATE		1			32	32					14	
HASH JOIN RIGHT O...		46M	148		32	50M	3GB				50	
PX RECEIVE		950K	74K		32	50M					2	8.33
PX SEND HASH	:TQ10000	950K	74K		32	50M					6	
PX BLOCK ITER...		950K	74K		32	50M						
TABLE ACCES...	CARGUYS	950K	74K		412	50M			62K	99	8	33
PX RECEIVE		46M	74K		32	50M					4	
PX SEND HASH	:TQ10001	46M	74K		32	50M					16	8.33
PX BLOCK ITER...		46M	74K		32	50M						
TABLE ACCES...	CARGUYS	46M	74K		412	50M			62K	99		50

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Increase the list of predicate values

Now query takes 3 seconds

More Predicate Values

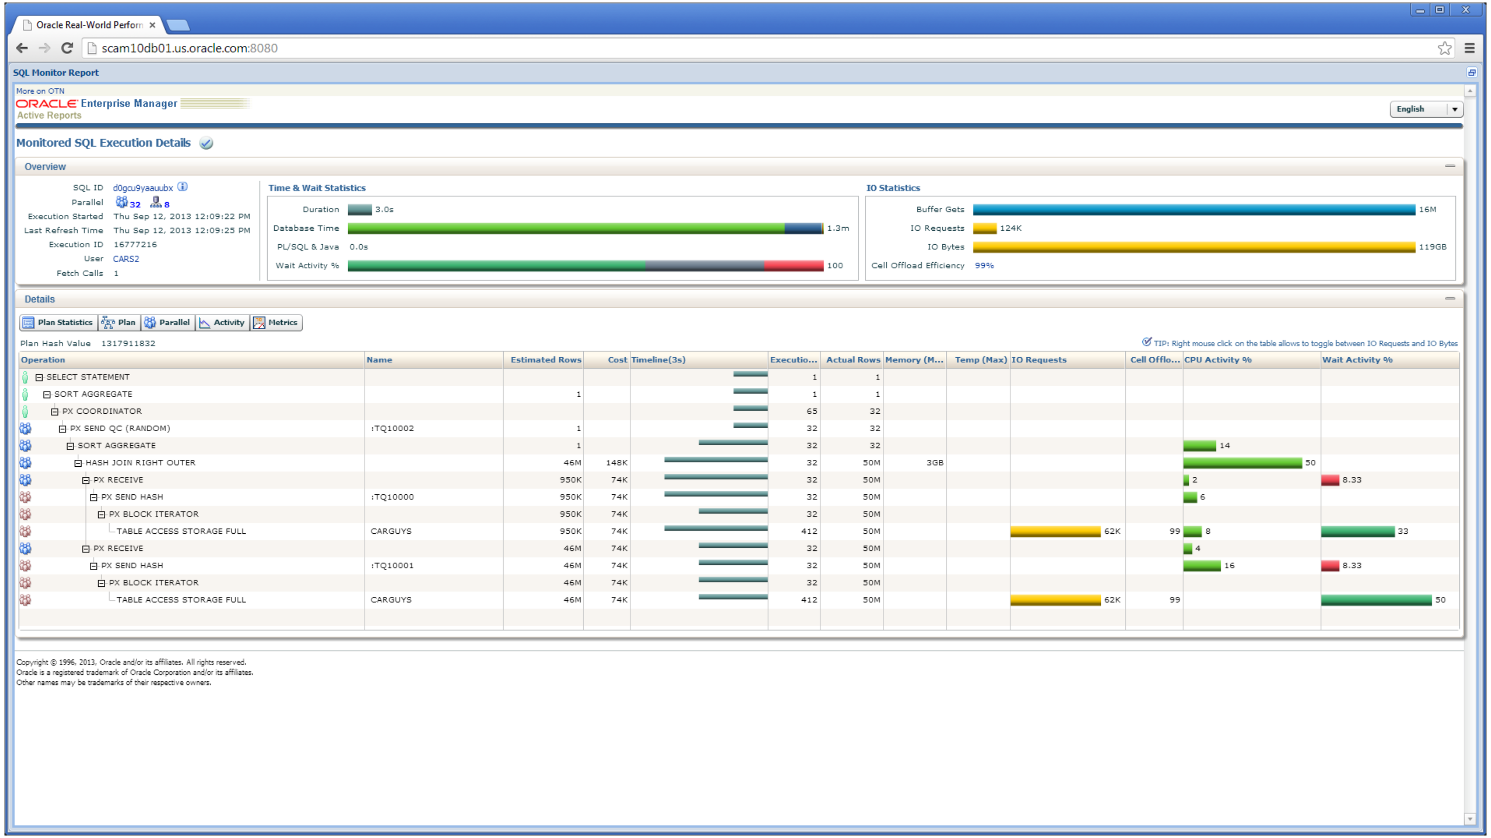

Oracle Real-World Perform ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	d0gcu9yaauubx
Parallel	32 8
Execution Started	Thu Sep 12, 2013 12:09:22 PM
Last Refresh Time	Thu Sep 12, 2013 12:09:25 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.3m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(3s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					14	
HASH JOIN RIGHT OUTER		46M	148K		32	50M	3GB				50	
PX RECEIVE		950K	74K		32	50M					2	8.33
PX SEND HASH	:TQ10000	950K	74K		32	50M					6	
PX BLOCK ITERATOR		950K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	950K	74K		412	50M			62K	99	8	33
PX RECEIVE		46M	74K		32	50M					4	
PX SEND HASH	:TQ10001	46M	74K		32	50M					16	8.33
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	46M	74K		412	50M			62K	99		50

More Predicate Values

2. Development Findings

- Query runs faster just by changing the list of values in the select list
- Plan changed from a broadcast to a hash distribution due to the higher but inaccurate cardinality estimate
- Get correct plan with wrong cardinality estimate—can lead to inconsistent plans and performance

Initial Optimization Steps—Increase Degree of Parallelism

Degree of Parallelism

ORACLE
REAL-WORLD PERFORMANCE

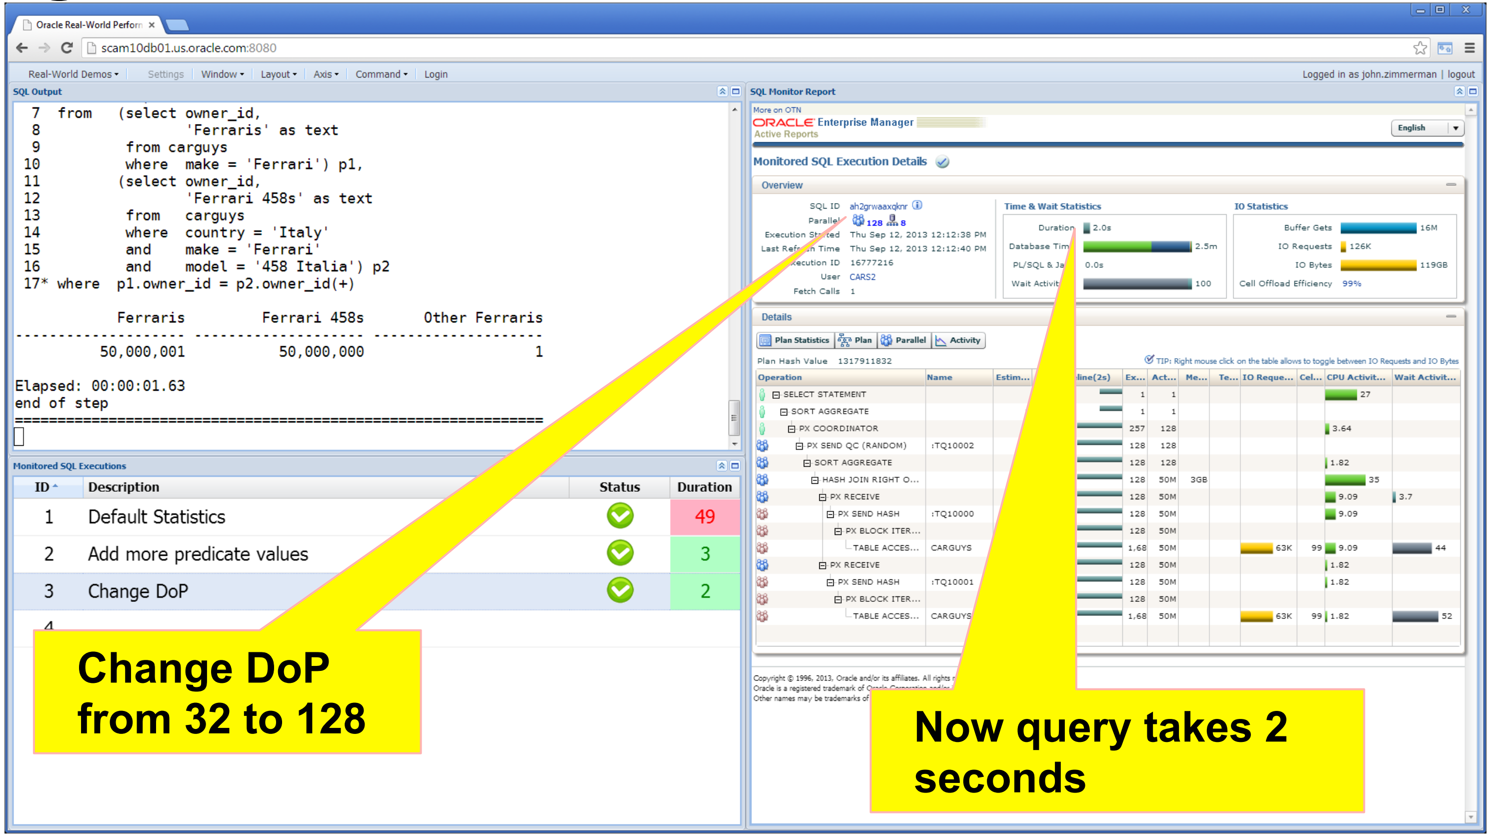

Oracle Real-World Perform ×
scam10db01.us.oracle.com:8080

Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout

SQL Output

```
  7   from    (select owner_id,
  8                   'Ferraris' as text
  9            from carguys
 10            where  make = 'Ferrari') p1,
 11           (select owner_id,
 12                   'Ferrari 458s' as text
 13            from   carguys
 14            where  country = 'Italy'
 15            and    make = 'Ferrari'
 16            and    model = '458 Italia') p2
 17* where   p1.owner_id = p2.owner_id(+)

             Ferraris          Ferrari 458s       Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:01.63
end of step
================================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	ah2grwaaxqknr
Parallel	128 8
Execution Started	Thu Sep 12, 2013 12:12:38 PM
Last Refresh Time	Thu Sep 12, 2013 12:12:40 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Tim	2.5m
PL/SQL & Ja	0.0s
Wait Activit	100

IO Statistics

Buffer Gets	16M
IO Requests	126K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	line(2s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT				1	1					27	
SORT AGGREGATE				1	1						
PX COORDINATOR				257	128					3.64	
PX SEND QC (RANDOM)	:TQ10002			128	128						
SORT AGGREGATE				128	128					1.82	
HASH JOIN RIGHT O...				128	50M	3GB				35	
PX RECEIVE				128	50M					9.09	3.7
PX SEND HASH	:TQ10000			128	50M					9.09	
PX BLOCK ITER...				128	50M						
TABLE ACCES...	CARGUYS			1,68	50M			63K	99	9.09	44
PX RECEIVE				128	50M					1.82	
PX SEND HASH	:TQ10001			128	50M					1.82	
PX BLOCK ITER...				128	50M						
TABLE ACCES...	CARGUYS			1,68	50M			63K	99	1.82	52

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights r
Oracle is a registered trademark of Oracle Corporation and/or
Other names may be trademarks of

Change DoP from 32 to 128

Now query takes 2 seconds

ORACLE

Degree of Parallelism

ORACLE
REAL-WORLD PERFORMANCE

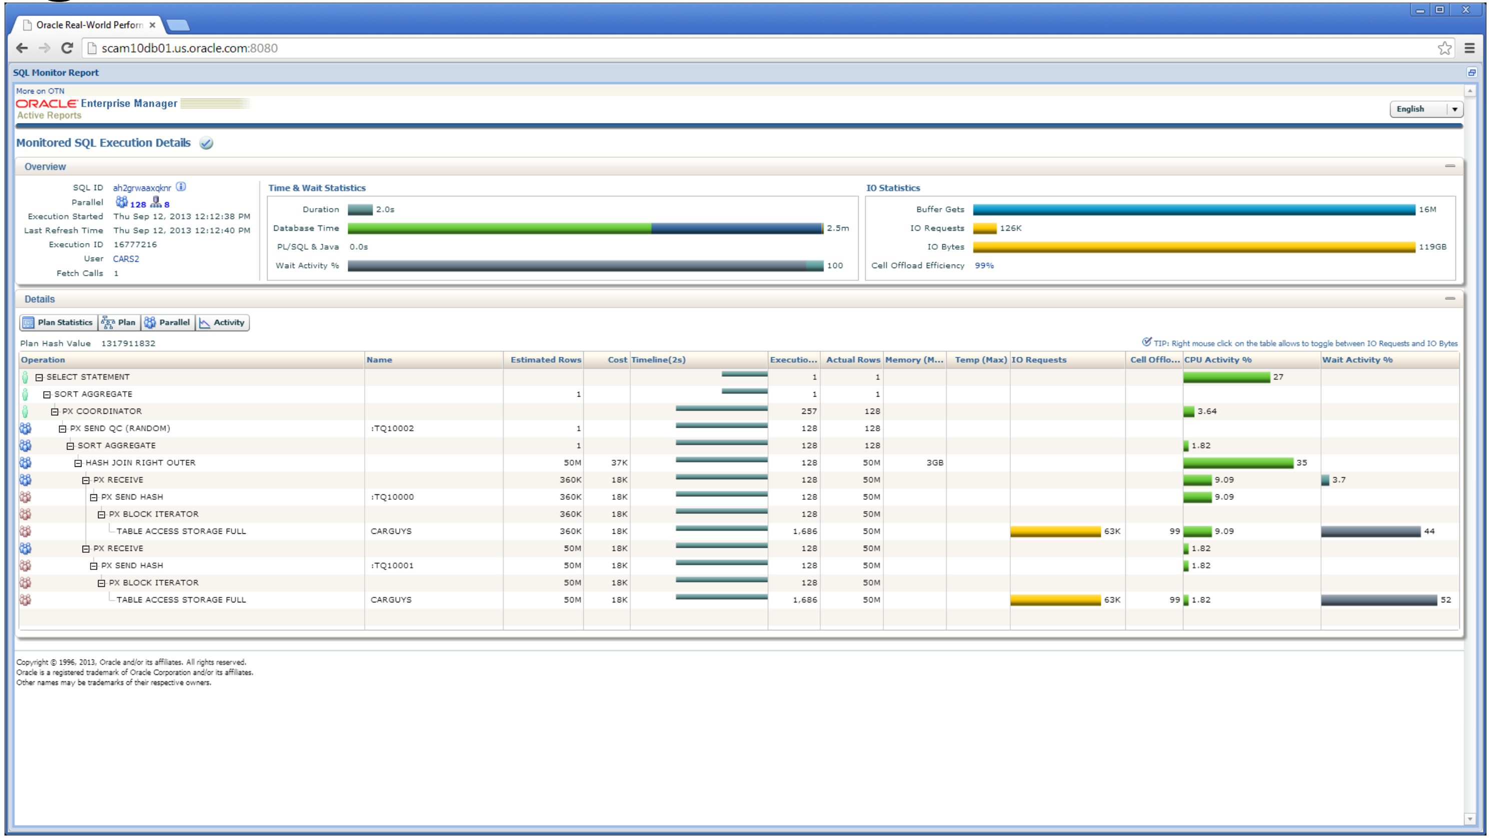

Oracle Real-World Perforr ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	ah2grwaaxqknr
Parallel	128 8
Execution Started	Thu Sep 12, 2013 12:12:38 PM
Last Refresh Time	Thu Sep 12, 2013 12:12:40 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Time	2.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	126K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(2s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					27	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					257	128					3.64	
PX SEND QC (RANDOM)	:TQ10002	1			128	128						
SORT AGGREGATE		1			128	128					1.82	
HASH JOIN RIGHT OUTER		50M	37K		128	50M	3GB				35	
PX RECEIVE		360K	18K		128	50M					9.09	3.7
PX SEND HASH	:TQ10000	360K	18K		128	50M					9.09	
PX BLOCK ITERATOR		360K	18K		128	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	360K	18K		1,686	50M			63K	99	9.09	44
PX RECEIVE		50M	18K		128	50M					1.82	
PX SEND HASH	:TQ10001	50M	18K		128	50M					1.82	
PX BLOCK ITERATOR		50M	18K		128	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	50M	18K		1,686	50M			63K	99	1.82	52

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

REAL-WORLD PERFORMANCE

Degree of Parallelism

3. Development Findings

- Changing DoP from 32 to 128 improves performance and meets the target; 4X more resources yields a 25X performance improvement
- Plan has changed from a broadcast distribution to a hash distribution due to DoP change
- DoP is a resource management technique, not a query tuning tool

Indexes

Indexes

4. Development Findings—Indexes

- Indexes on columns:
  - owner_id
  - country
  - make
  - model
  - country, make, model

Indexes

ORACLE
REAL-WORLD PERFORMANCE

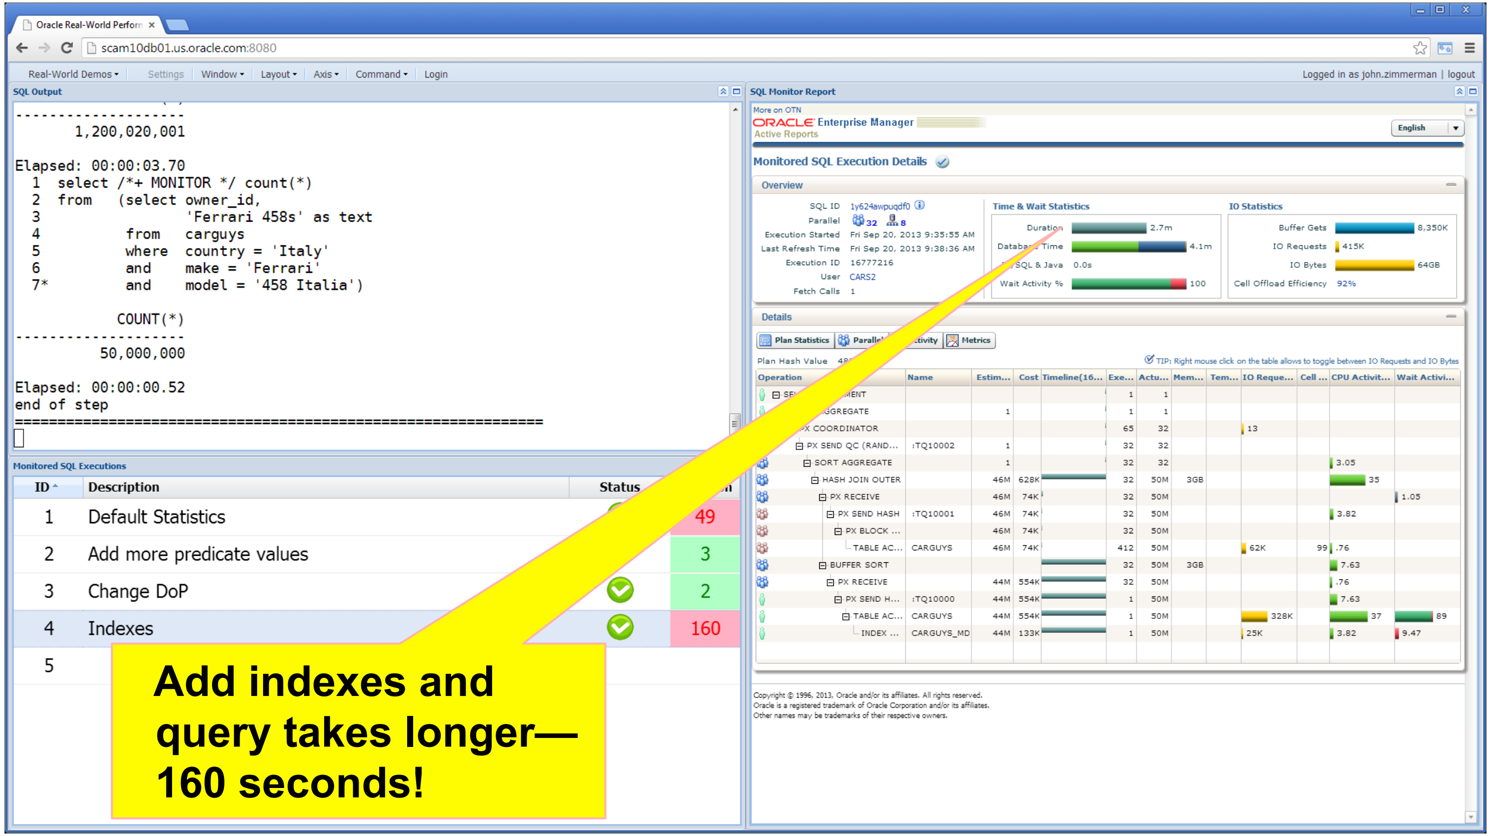

Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout
SQL Output
SQL Monitor Report
More on OTN
ORACLE Enterprise Manager
Active Reports
English

```
--------------------
      1,200,020,001

Elapsed: 00:00:03.70
  1  select /*+ MONITOR */ count(*)
  2  from   (select owner_id,
  3                 'Ferrari 458s' as text
  4          from   carguys
  5          where  country = 'Italy'
  6          and    make = 'Ferrari'
  7*         and    model = '458 Italia')

         COUNT(*)
--------------------
         50,000,000

Elapsed: 00:00:00.52
end of step
==============================================================
```

Monitored SQL Executions

ID	Description	Status	n
1	Default Statistics		49
2	Add more predicate values		3
3	Change DoP	✓	2
4	Indexes	✓	160
5			

Monitored SQL Execution Details

Overview

SQL ID	1y624awpuqdf0
Parallel	32 8
Execution Started	Fri Sep 20, 2013 9:35:55 AM
Last Refresh Time	Fri Sep 20, 2013 9:38:36 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.7m
Databa Time	4.1m
/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	8,350K
IO Requests	415K
IO Bytes	64GB
Cell Offload Efficiency	92%

Details

Plan Statistics | Paralle ctivity | Metrics

Plan Hash Value 48

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	Cost	Timeline(16...	Exe...	Actu...	Mem...	Tem...	IO Reque...	Cell ...	CPU Activit...	Wait Activi...
SE MENT					1	1						
GGREGATE		1			1	1						
X COORDINATOR					65	32			13			
PX SEND QC (RAND...	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					3.05	
HASH JOIN OUTER		46M	628K		32	50M	3GB				35	
PX RECEIVE		46M	74K		32	50M						1.05
PX SEND HASH	:TQ10001	46M	74K		32	50M					3.82	
PX BLOCK ...		46M	74K		32	50M						
TABLE AC...	CARGUYS	46M	74K		412	50M			62K	99	.76	
BUFFER SORT					32	50M	3GB				7.63	
PX RECEIVE		44M	554K		32	50M					.76	
PX SEND H...	:TQ10000	44M	554K		1	50M					7.63	
TABLE AC...	CARGUYS	44M	554K		1	50M			328K		37	89
INDEX ...	CARGUYS_MD	44M	133K		1	50M			25K		3.82	9.47

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Add indexes and query takes longer—160 seconds!

ORACLE

Indexes

ORACLE
REAL-WORLD PERFORMANCE

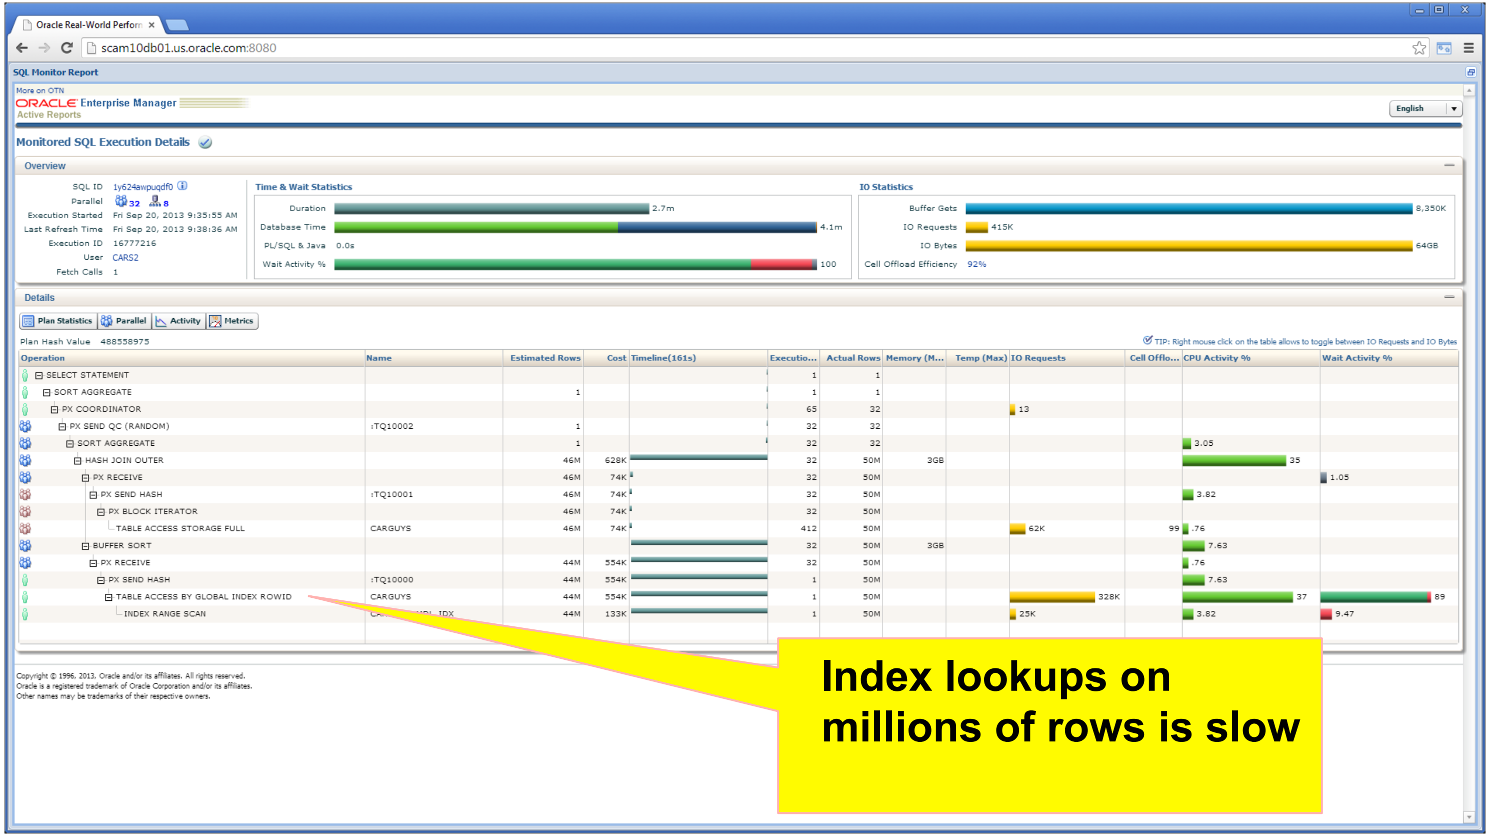

Oracle Real-World Perform ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	1y624awpuqdf0
Parallel	32 8
Execution Started	Fri Sep 20, 2013 9:35:55 AM
Last Refresh Time	Fri Sep 20, 2013 9:38:36 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.7m
Database Time	4.1m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	8,350K
IO Requests	415K
IO Bytes	64GB
Cell Offload Efficiency	92%

Details

Plan Statistics | Parallel | Activity | Metrics

Plan Hash Value 488558975

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(161s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32			13			
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					3.05	
HASH JOIN OUTER		46M	628K		32	50M	3GB				35	
PX RECEIVE		46M	74K		32	50M						1.05
PX SEND HASH	:TQ10001	46M	74K		32	50M					3.82	
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	46M	74K		412	50M			62K	99	.76	
BUFFER SORT					32	50M	3GB				7.63	
PX RECEIVE		44M	554K		32	50M					.76	
PX SEND HASH	:TQ10000	44M	554K		1	50M					7.63	
TABLE ACCESS BY GLOBAL INDEX ROWID	CARGUYS	44M	554K		1	50M			328K		37	89
INDEX RANGE SCAN	CA[illegible]MDL_IDX	44M	133K		1	50M			25K		3.82	9.47

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Index lookups on millions of rows is slow

ORACLE

Indexes

REAL-WORLD PERFORMANCE

4. Development Findings—Indexes

- Not understanding the big/little data challenge
- Indexes are not efficient for operations on a large numbers of rows
- Full table scan is faster with predictable performance

To Index or Not

- Indexing is an OLTP technique for operations on a small number of rows
- A table scan may consume more resources but it will be predictable no matter how many rows are returned
- Indexes impact DML operations
- If I/O bandwidth went from 70MB/sec to 70GB/sec would you change your optimization/execution strategy?

To Index or Not

- Index driven query retrieving 1,000,000 rows
  - Assume the index is cached and the data is not.
    - 1,000,000 random IOPS @ 5ms per I/O
    - This would require 5000 Seconds (or over 1 hour) to Execute
  - How much data could you scan in 5000 Seconds with a fully sized I/O system able to scan 25 GB/Sec ?
    - Over 100 TB !

Histograms

Histograms

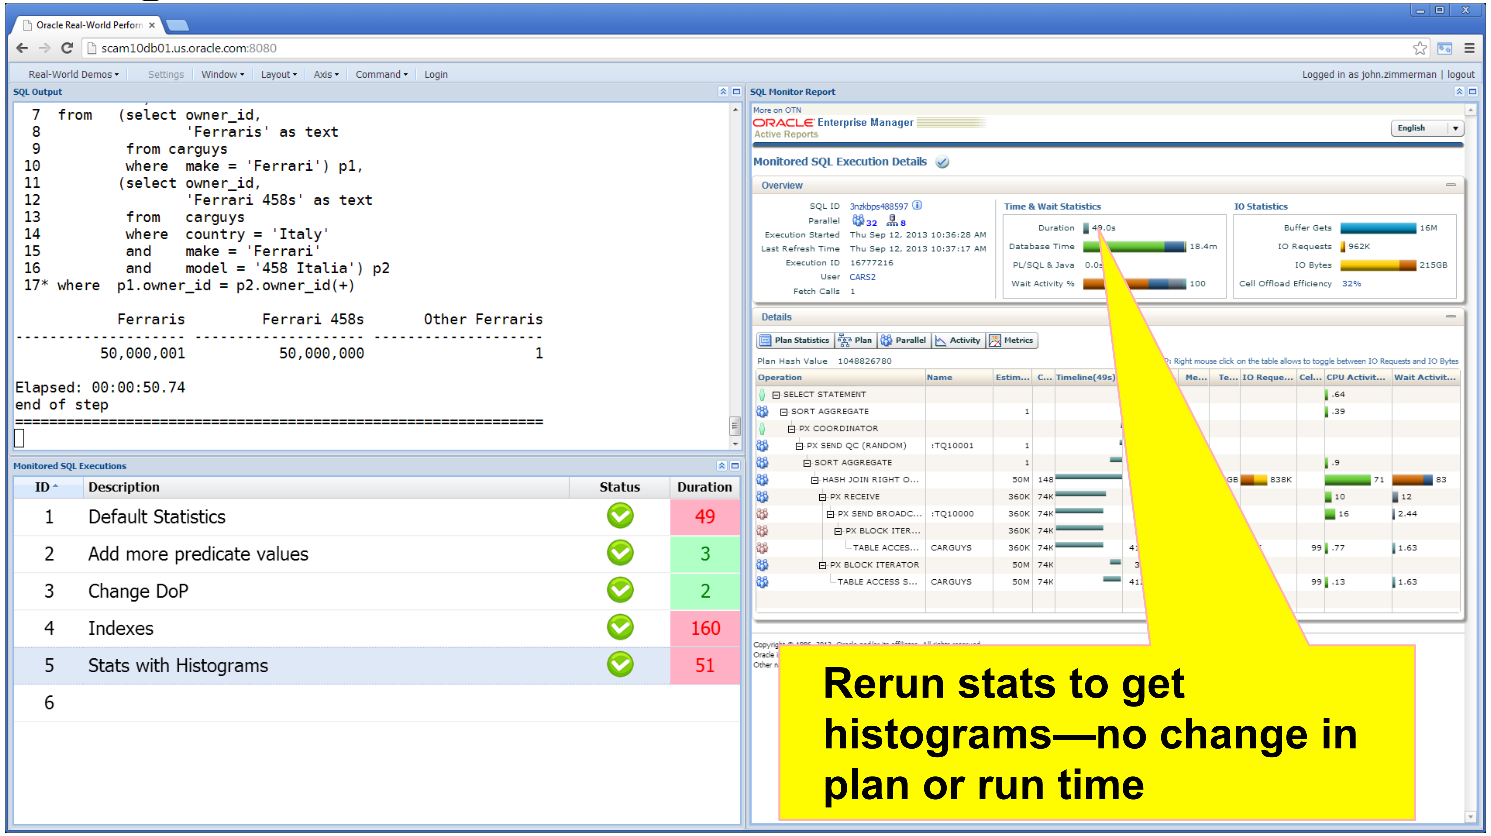

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

Logged in as john.zimmerman | logout

SQL Output

```
 from    (select owner_id,
                 'Ferraris' as text
          from carguys
          where  make = 'Ferrari') p1,
         (select owner_id,
                 'Ferrari 458s' as text
          from   carguys
          where  country = 'Italy'
          and    make = 'Ferrari'
          and    model = '458 Italia') p2
 where   p1.owner_id = p2.owner_id(+)

             Ferraris         Ferrari 458s       Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:50.74
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6			

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	3nzkbps488597
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:36:28 AM
Last Refresh Time	Thu Sep 12, 2013 10:37:17 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	49.0s
Database Time	18.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Hash Value 1048826780

: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(49s)		Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT										.64	
SORT AGGREGATE		1								.39	
PX COORDINATOR											
PX SEND QC (RANDOM)	:TQ10001	1									
SORT AGGREGATE		1								.9	
HASH JOIN RIGHT O...		50M	148				GB	838K		71	83
PX RECEIVE		360K	74K							10	12
PX SEND BROADC...	:TQ10000	360K	74K							16	2.44
PX BLOCK ITER...		360K	74K								
TABLE ACCES...	CARGUYS	360K	74K		41				99	.77	1.63
PX BLOCK ITERATOR		50M	74K		3						
TABLE ACCESS S...	CARGUYS	50M	74K		41				99	.13	1.63

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.

Rerun stats to get histograms—no change in plan or run time

Histograms

ORACLE
REAL-WORLD PERFORMANCE

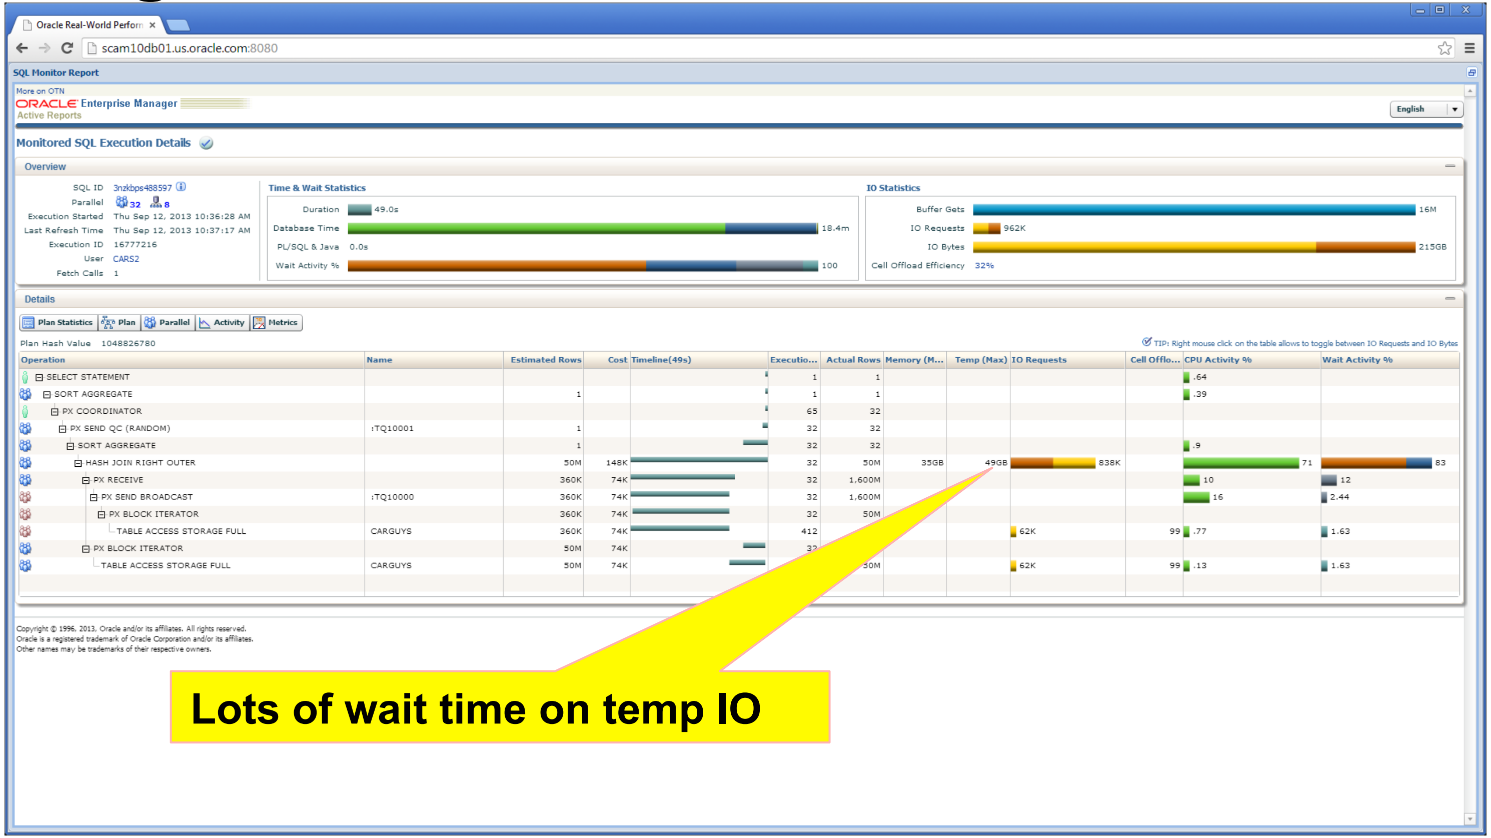
Oracle Real-World Perform ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	3nzkbps488597
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:36:28 AM
Last Refresh Time	Thu Sep 12, 2013 10:37:17 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	49.0s
Database Time	18.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(49s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					.64	
SORT AGGREGATE		1			1	1					.39	
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.9	
HASH JOIN RIGHT OUTER		50M	148K		32	50M	35GB	49GB	838K		71	83
PX RECEIVE		360K	74K		32	1,600M					10	12
PX SEND BROADCAST	:TQ10000	360K	74K		32	1,600M					16	2.44
PX BLOCK ITERATOR		360K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	360K	74K		412				62K	99	.77	1.63
PX BLOCK ITERATOR		50M	74K		32	[illegible]						
TABLE ACCESS STORAGE FULL	CARGUYS	50M	74K		[illegible]	50M			62K	99	.13	1.63

Lots of wait time on temp IO

ORACLE

Histograms

5. Development Findings

- Re-gathered stats to automatically create histograms
  - Frequency histograms on country, make and model columns
- No change in plan—query still exceeds target

Flash Temp

Flash Temp

REAL-WORLD PERFORMANCE

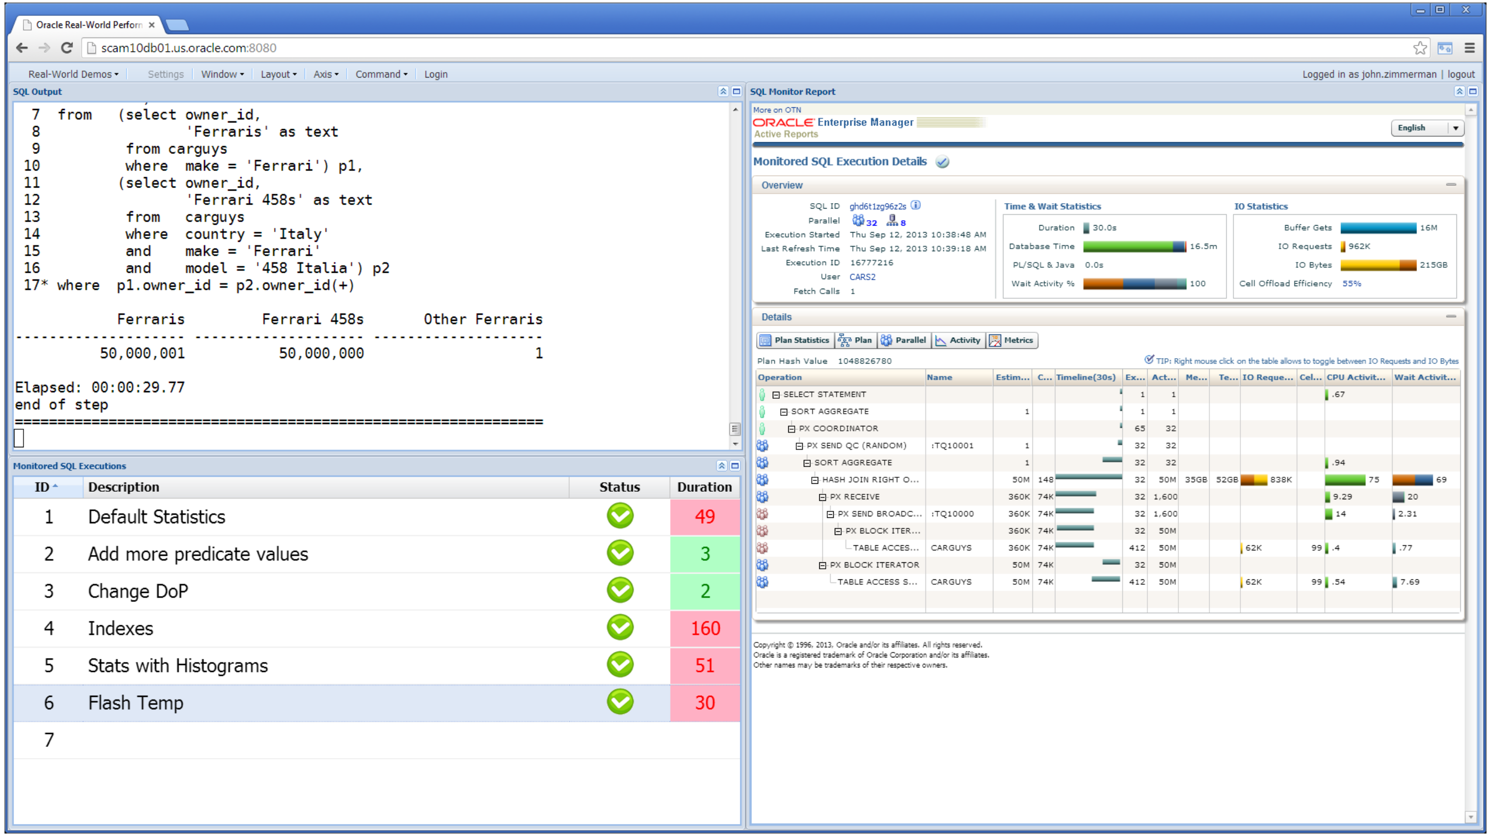

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login

Logged in as john.zimmerman | logout

SQL Output

```
from    (select owner_id,
                'Ferraris' as text
         from carguys
         where  make = 'Ferrari') p1,
        (select owner_id,
                'Ferrari 458s' as text
         from   carguys
         where  country = 'Italy'
         and    make = 'Ferrari'
         and    model = '458 Italia') p2
where   p1.owner_id = p2.owner_id(+)

            Ferraris        Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:29.77
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7			

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	ghd6t1zg96z2s
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:38:48 AM
Last Refresh Time	Thu Sep 12, 2013 10:39:18 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	30.0s
Database Time	16.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	55%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(30s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1					.67	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.94	
HASH JOIN RIGHT O...		50M	148		32	50M	35GB	52GB	838K		75	69
PX RECEIVE		360K	74K		32	1,600					9.29	20
PX SEND BROADC...	:TQ10000	360K	74K		32	1,600					14	2.31
PX BLOCK ITER...		360K	74K		32	50M						
TABLE ACCES...	CARGUYS	360K	74K		412	50M			62K	99	.4	.77
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	50M	74K		412	50M			62K	99	.54	7.69

Flash Temp

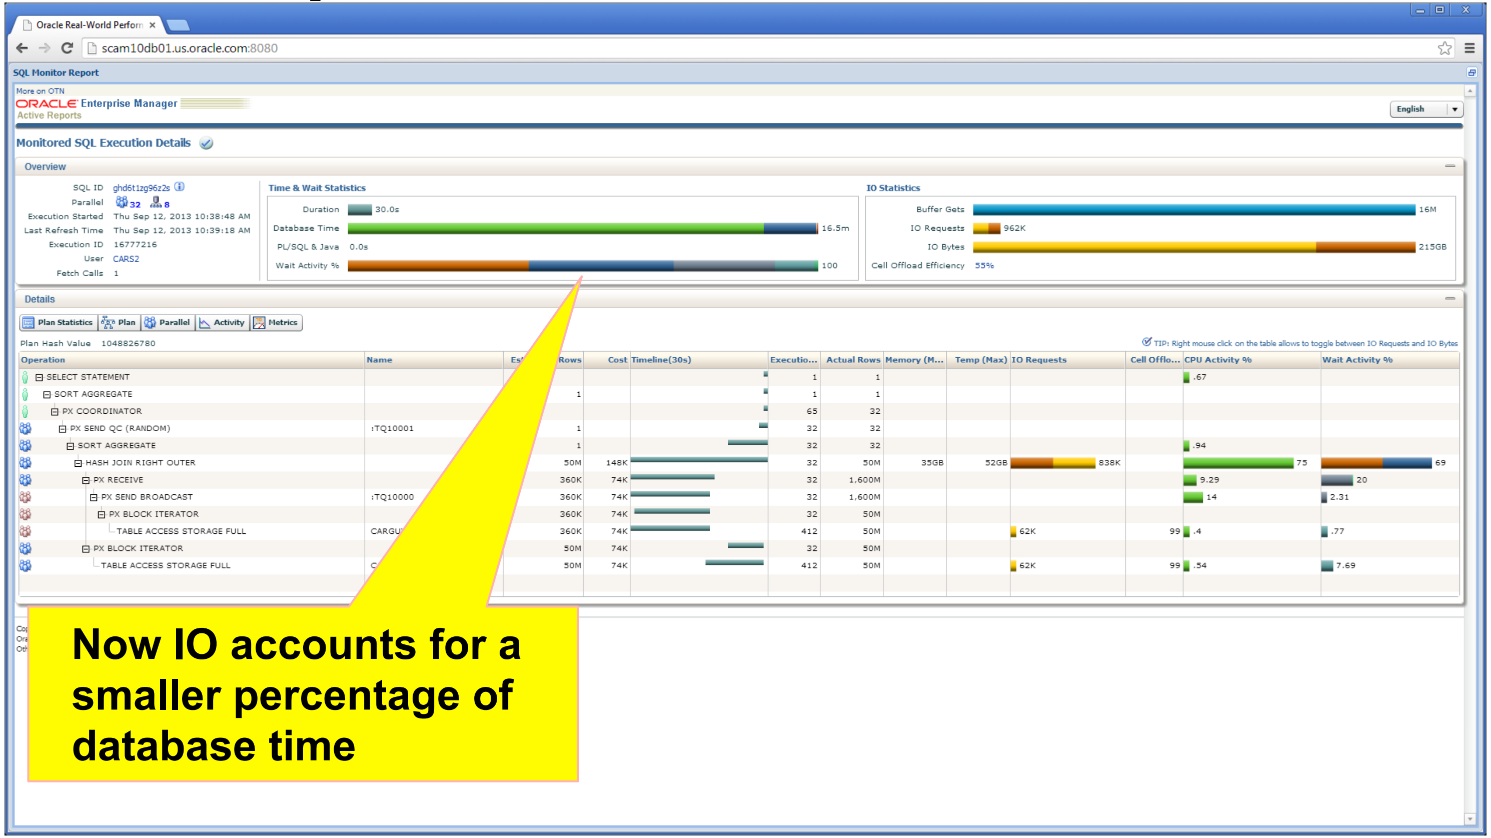

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	ghd6t1zg96z2s
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:38:48 AM
Last Refresh Time	Thu Sep 12, 2013 10:39:18 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	30.0s
Database Time	16.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	55%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Es Rows	Cost	Timeline(30s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					.67	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.94	
HASH JOIN RIGHT OUTER		50M	148K		32	50M	35GB	52GB	838K		75	69
PX RECEIVE		360K	74K		32	1,600M					9.29	20
PX SEND BROADCAST	:TQ10000	360K	74K		32	1,600M					14	2.31
PX BLOCK ITERATOR		360K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGU	360K	74K		412	50M			62K	99	.4	.77
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	C	50M	74K		412	50M			62K	99	.54	7.69

Now IO accounts for a smaller percentage of database time

Flash Temp

6. Development Findings

- Most of the wait time was spent performing IO on temp, so move temp to flash disks
- Improved performance but still does not meet target
- Not a good use of flash
- Incorrect use of tools/products

Manual Memory Parameters

Manual Memory Parameters

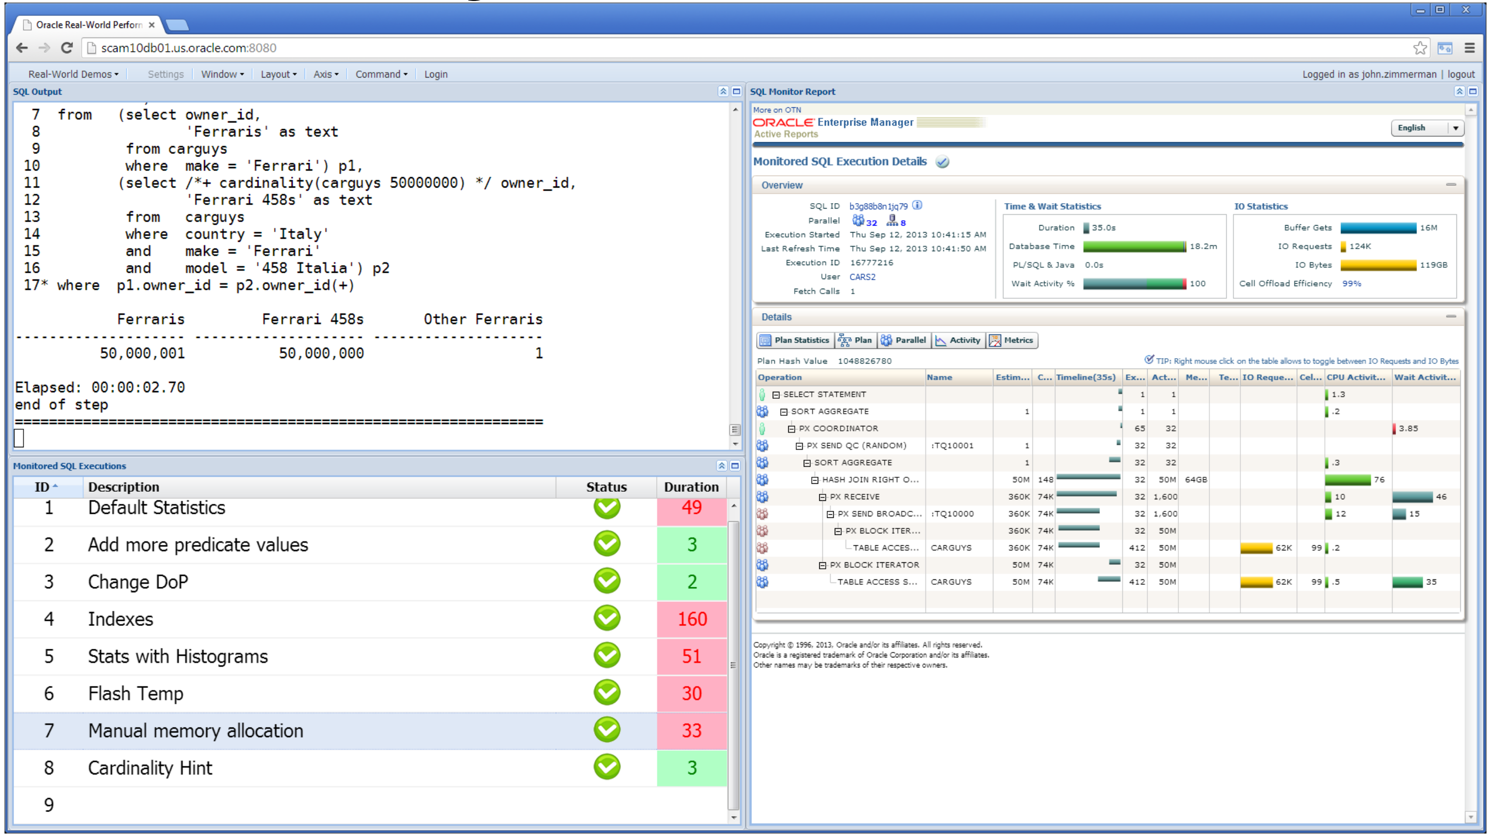
Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos ▾
Settings
Window ▾
Layout ▾
Axis ▾
Command ▾
Login
Logged in as john.zimmerman | logout

SQL Output

```
  7  from   (select owner_id,
 8                  'Ferraris' as text
 9            from carguys
10            where  make = 'Ferrari') p1,
11           (select /*+ cardinality(carguys 50000000) */ owner_id,
12                  'Ferrari 458s' as text
13            from   carguys
14            where  country = 'Italy'
15            and    make = 'Ferrari'
16            and    model = '458 Italia') p2
17* where  p1.owner_id = p2.owner_id(+)

            Ferraris          Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:02.70
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	b3g88b8n1jq79
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:41:15 AM
Last Refresh Time	Thu Sep 12, 2013 10:41:50 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	35.0s
Database Time	18.2m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(35s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
⊟ SELECT STATEMENT					1	1					1.3	
⊟ SORT AGGREGATE		1			1	1					.2	
⊟ PX COORDINATOR					65	32						3.85
⊟ PX SEND QC (RANDOM)	:TQ10001	1			32	32						
⊟ SORT AGGREGATE		1			32	32					.3	
⊟ HASH JOIN RIGHT O...		50M	148		32	50M	64GB				76	
⊟ PX RECEIVE		360K	74K		32	1,600					10	46
⊟ PX SEND BROADC...	:TQ10000	360K	74K		32	1,600					12	15
⊟ PX BLOCK ITER...		360K	74K		32	50M						
TABLE ACCES...	CARGUYS	360K	74K		412	50M			62K	99	.2	
⊟ PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	50M	74K		412	50M			62K	99	.5	35

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Manual Memory Parameters

REAL-WORLD PERFORMANCE

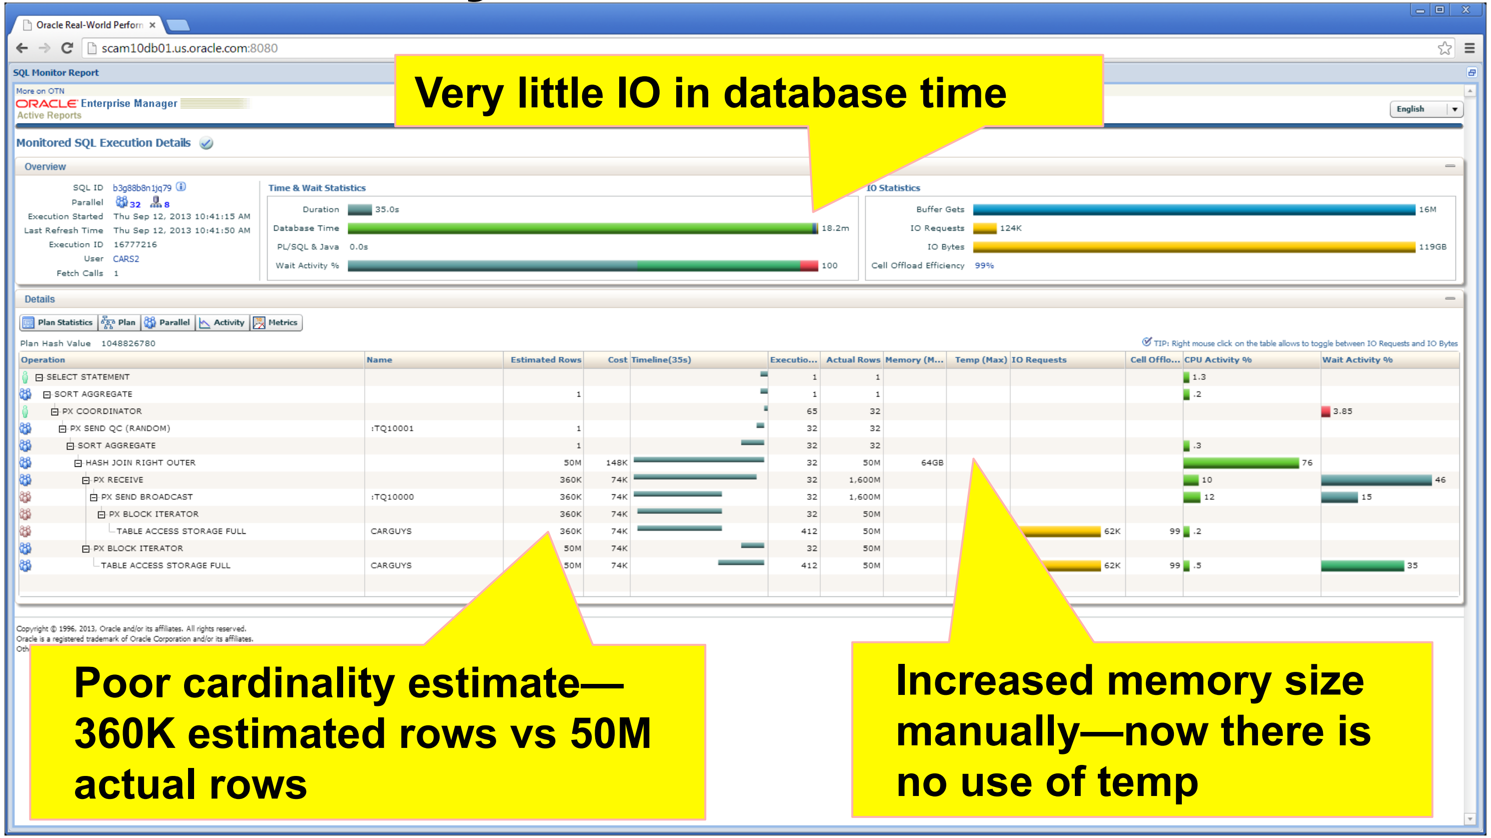
Oracle Real-World Perform
scam10db01.us.oracle.com:8080
SQL Monitor Report
More on OTN
ORACLE Enterprise Manager
Active Reports
English

Very little IO in database time

Monitored SQL Execution Details

Overview

SQL ID	b3g88b8n1jq79
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:41:15 AM
Last Refresh Time	Thu Sep 12, 2013 10:41:50 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	35.0s
Database Time	18.2m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(35s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					1.3	
SORT AGGREGATE		1			1	1					.2	
PX COORDINATOR					65	32						3.85
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.3	
HASH JOIN RIGHT OUTER		50M	148K		32	50M	64GB				76	
PX RECEIVE		360K	74K		32	1,600M					10	46
PX SEND BROADCAST	:TQ10000	360K	74K		32	1,600M					12	15
PX BLOCK ITERATOR		360K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	360K	74K		412	50M			62K	99	.2	
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	50M	74K		412	50M			62K	99	.5	35

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Oth

Poor cardinality estimate—360K estimated rows vs 50M actual rows

Increased memory size manually—now there is no use of temp

Manual Memory Parameters

ORACLE
REAL-WORLD PERFORMANCE

7. Development Findings

- Set sort_area_size and hash_area_size to 2G
- Eliminated temp usage but still did not meet target
- Memory is allocated per parallel server process, which can quickly exceed resources
- Moving to a solution before understanding the problem

ORACLE

Cardinality Estimates

Cardinality Estimates

ORACLE
REAL-WORLD PERFORMANCE

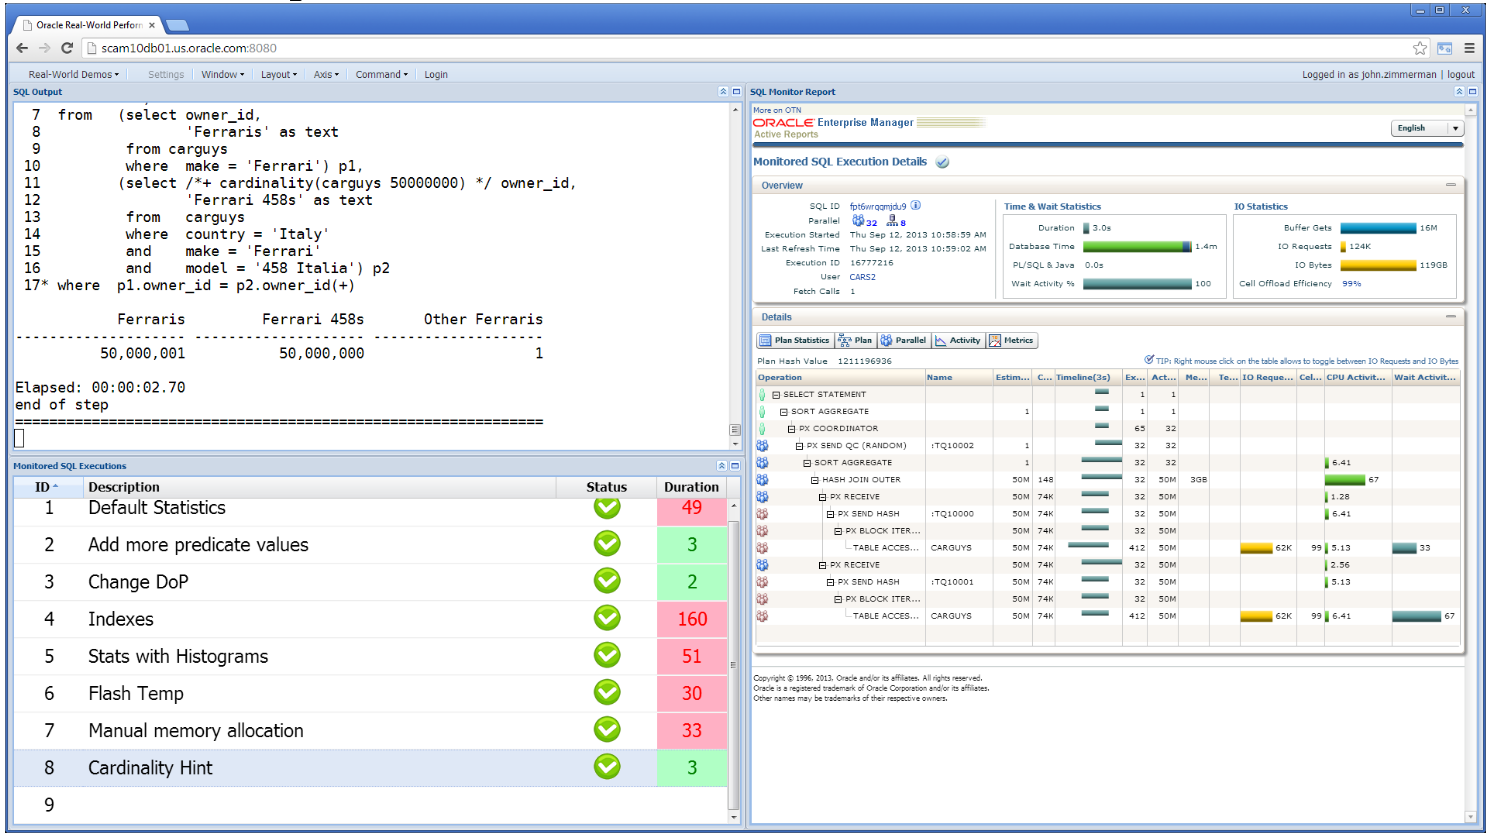
Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos ▾ Settings Window ▾ Layout ▾ Axis ▾ Command ▾ Login
Logged in as john.zimmerman | logout
SQL Output
SQL Monitor Report

```
 7  from   (select owner_id,
                   'Ferraris' as text
             from carguys
             where  make = 'Ferrari') p1,
           (select /*+ cardinality(carguys 50000000) */ owner_id,
                   'Ferrari 458s' as text
             from   carguys
             where  country = 'Italy'
             and    make = 'Ferrari'
             and    model = '458 Italia') p2
 17* where  p1.owner_id = p2.owner_id(+)

            Ferraris         Ferrari 458s       Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:02.70
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
1	Default Statistics	✓	49
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9			

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	fpt6wrqqmjdu9
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:58:59 AM
Last Refresh Time	Thu Sep 12, 2013 10:59:02 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1211196936

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(3s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					6.41	
HASH JOIN OUTER		50M	148		32	50M	3GB				67	
PX RECEIVE		50M	74K		32	50M					1.28	
PX SEND HASH	:TQ10000	50M	74K		32	50M					6.41	
PX BLOCK ITER...		50M	74K		32	50M						
TABLE ACCES...	CARGUYS	50M	74K		412	50M			62K	99	5.13	33
PX RECEIVE		50M	74K		32	50M					2.56	
PX SEND HASH	:TQ10001	50M	74K		32	50M					5.13	
PX BLOCK ITER...		50M	74K		32	50M						
TABLE ACCES...	CARGUYS	50M	74K		412	50M			62K	99	6.41	67

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

| |

Cardinality Estimates

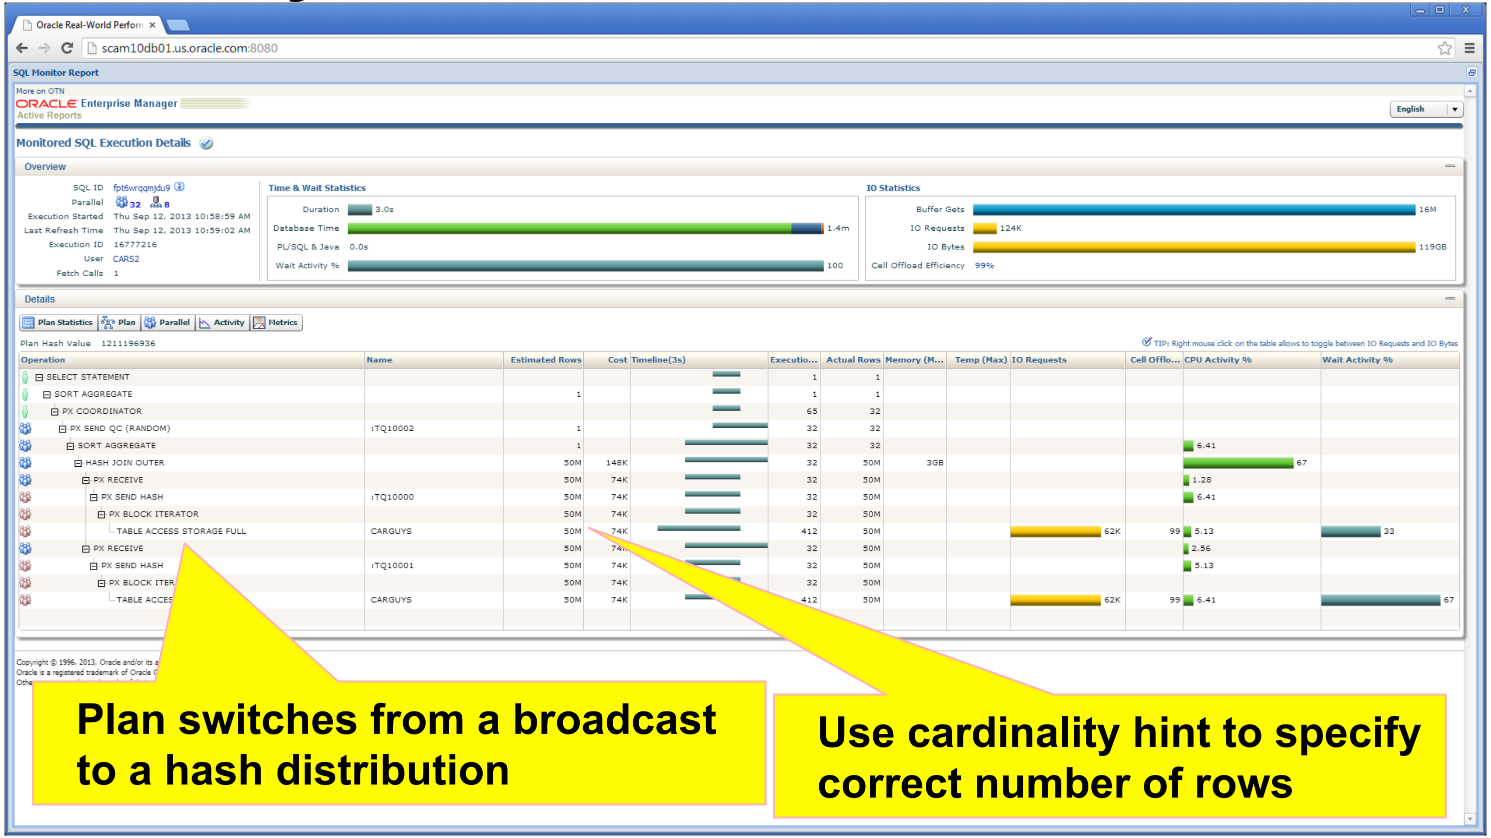

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	fpt6wrqqmjdu9
Parallel	32 8
Execution Started	Thu Sep 12, 2013 10:58:59 AM
Last Refresh Time	Thu Sep 12, 2013 10:59:02 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1211196936

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(3s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					6.41	
HASH JOIN OUTER		50M	148K		32	50M	3GB				67	
PX RECEIVE		50M	74K		32	50M					1.28	
PX SEND HASH	:TQ10000	50M	74K		32	50M					6.41	
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	50M	74K		412	50M			62K	99	5.13	33
PX RECEIVE		50M	74K		32	50M					2.56	
PX SEND HASH	:TQ10001	50M	74K		32	50M					5.13	
PX BLOCK ITER		50M	74K		32	50M						
TABLE ACCES	CARGUYS	50M	74K		412	50M			62K	99	6.41	67

Copyright © 1996, 2013, Oracle and/or its a
Oracle is a registered trademark of Oracle C
Othe

Plan switches from a broadcast to a hash distribution

Use cardinality hint to specify correct number of rows

Cardinality Estimates

8. Cardinality Hint

- SQL Monitor showed poor cardinality estimates
- Cardinality hint gives optimizer the correct number of rows for the table scan
- Plan changed from a broadcast to hash distribution
- Query time now meets target
- Now temp is not an issue

Disable Broadcast Distribution

Disable Broadcast Distribution

ORACLE
REAL-WORLD PERFORMANCE

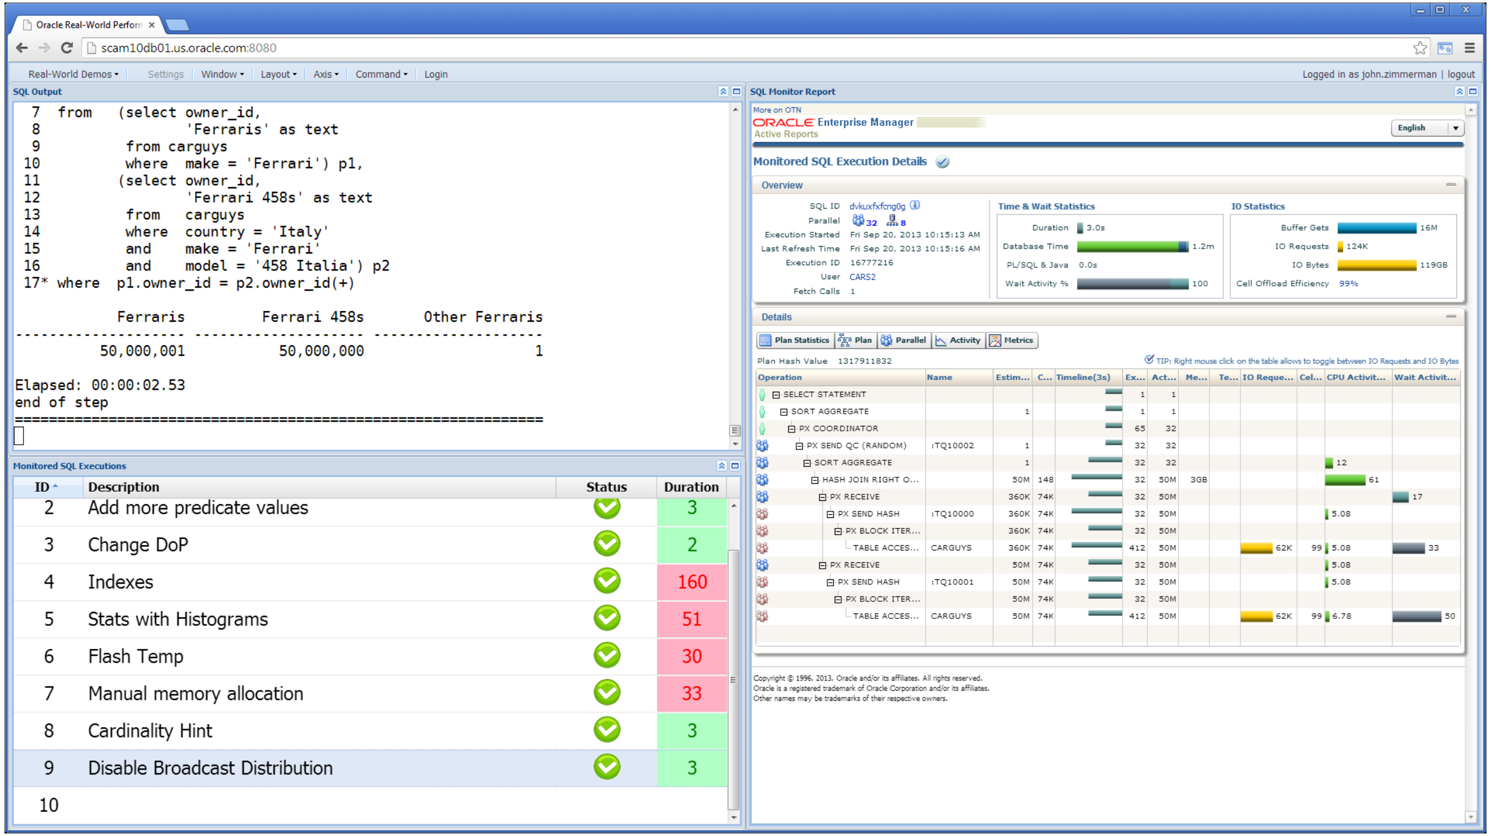

Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout
SQL Output
Monitored SQL Executions
SQL Monitor Report
More on OTN
ORACLE Enterprise Manager
Active Reports
English

```
 7  from   (select owner_id,
                   'Ferraris' as text
 9          from carguys
10          where  make = 'Ferrari') p1,
11         (select owner_id,
12                 'Ferrari 458s' as text
13          from   carguys
14          where  country = 'Italy'
15          and    make = 'Ferrari'
16          and    model = '458 Italia') p2
17* where  p1.owner_id = p2.owner_id(+)

            Ferraris         Ferrari 458s       Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:02.53
end of step
==============================================================
```

ID	Description	Status	Duration
2	Add more predicate values	✓	3
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10			

Monitored SQL Execution Details

Overview

SQL ID	dvkuxfxfcng0g
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:15:13 AM
Last Refresh Time	Fri Sep 20, 2013 10:15:16 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.2m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(3s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					12	
HASH JOIN RIGHT O...		50M	148		32	50M	3GB				61	
PX RECEIVE		360K	74K		32	50M						17
PX SEND HASH	:TQ10000	360K	74K		32	50M					5.08	
PX BLOCK ITER...		360K	74K		32	50M						
TABLE ACCES...	CARGUYS	360K	74K		412	50M			62K	99	5.08	33
PX RECEIVE		50M	74K		32	50M					5.08	
PX SEND HASH	:TQ10001	50M	74K		32	50M					5.08	
PX BLOCK ITER...		50M	74K		32	50M						
TABLE ACCES...	CARGUYS	50M	74K		412	50M			62K	99	6.78	50

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

Disable Broadcast Distribution

REAL-WORLD PERFORMANCE

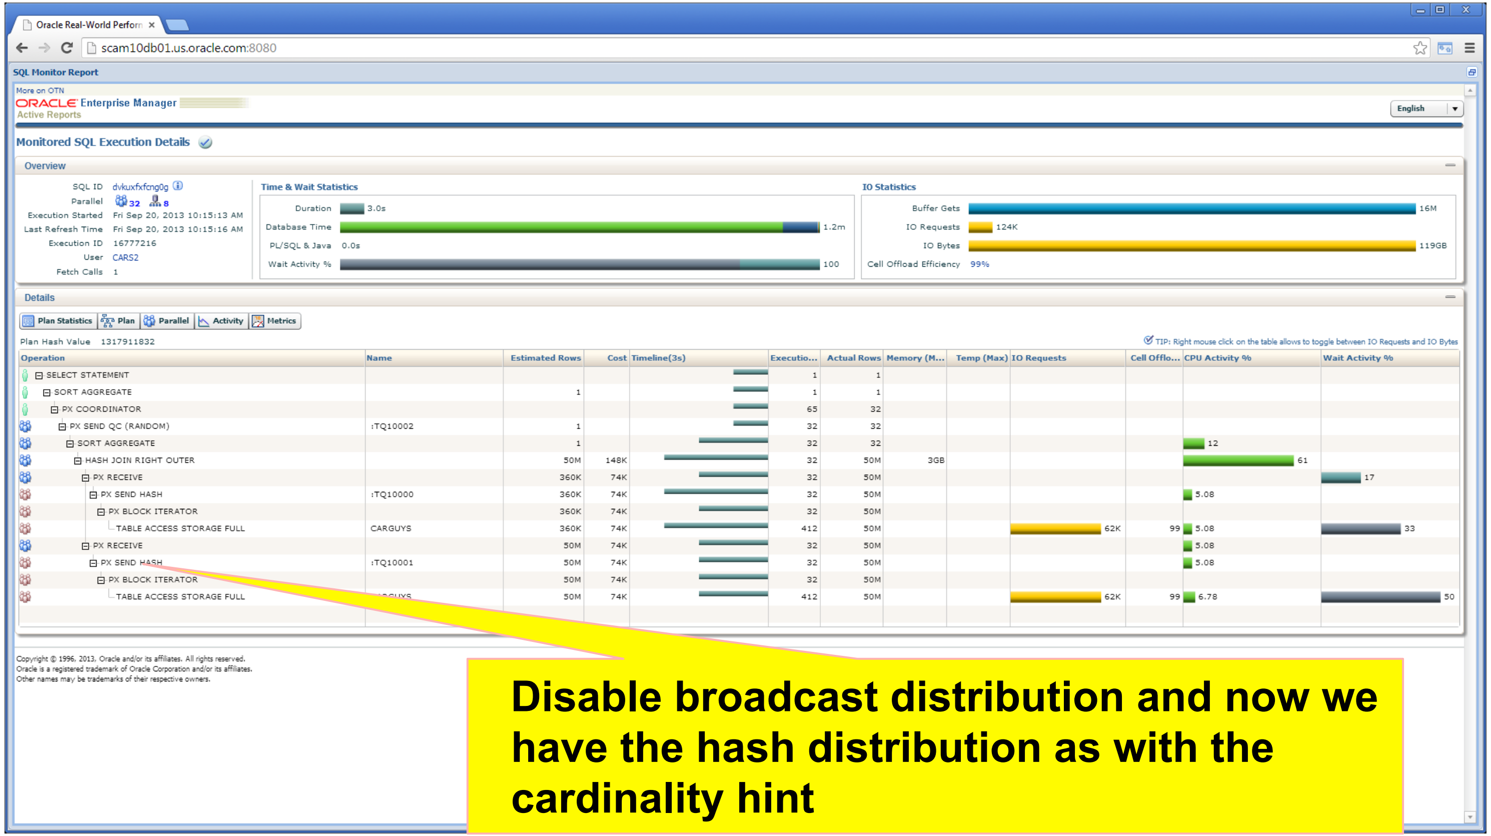

Oracle Real-World Perforn ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	dvkuxfxfcng0g
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:15:13 AM
Last Refresh Time	Fri Sep 20, 2013 10:15:16 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.2m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(3s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					12	
HASH JOIN RIGHT OUTER		50M	148K		32	50M	3GB				61	
PX RECEIVE		360K	74K		32	50M						17
PX SEND HASH	:TQ10000	360K	74K		32	50M					5.08	
PX BLOCK ITERATOR		360K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	360K	74K		412	50M			62K	99	5.08	33
PX RECEIVE		50M	74K		32	50M					5.08	
PX SEND HASH	:TQ10001	50M	74K		32	50M					5.08	
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	[illegible]GUYS	50M	74K		412	50M			62K	99	6.78	50

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Disable broadcast distribution and now we have the hash distribution as with the cardinality hint

Disable Broadcast Distribution

9. Development Findings

- Googling reveals a hidden parameter to disable broadcast distribution
- Plan and run times are similar to cardinality hint, meeting target
- Moving to a solution before understanding the problem

Second Query with Broadcast Distribution Disabled

Query 2: Broadcast Distribution Disabled

ORACLE
REAL-WORLD PERFORMANCE

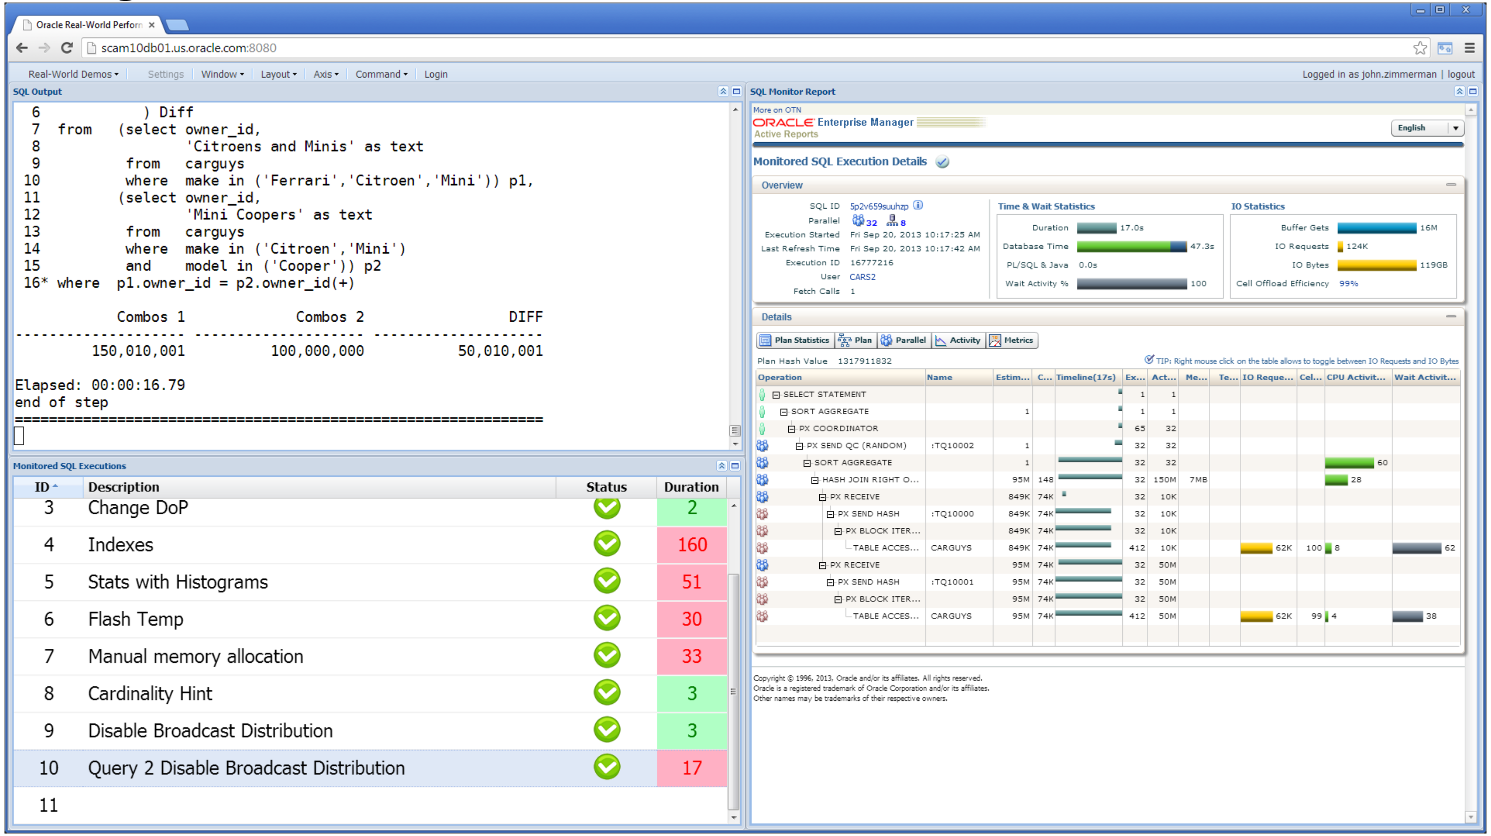

Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout

SQL Output

```
  6             ) Diff
  7  from    (select owner_id,
  8                  'Citroens and Minis' as text
  9           from   carguys
 10           where  make in ('Ferrari','Citroen','Mini')) p1,
 11          (select owner_id,
 12                  'Mini Coopers' as text
 13           from   carguys
 14           where  make in ('Citroen','Mini')
 15           and    model in ('Cooper')) p2
 16* where  p1.owner_id = p2.owner_id(+)

           Combos 1             Combos 2                 DIFF
-------------------- -------------------- --------------------
         150,010,001          100,000,000           50,010,001

Elapsed: 00:00:16.79
end of step
================================================================
```

Monitored SQL Executions

ID ▴	Description	Status	Duration
3	Change DoP	✓	2
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	5p2v659suuhzp
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:17:25 AM
Last Refresh Time	Fri Sep 20, 2013 10:17:42 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	17.0s
Database Time	47.3s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(17s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					60	
HASH JOIN RIGHT O...		95M	148		32	150M	7MB				28	
PX RECEIVE		849K	74K		32	10K						
PX SEND HASH	:TQ10000	849K	74K		32	10K						
PX BLOCK ITER...		849K	74K		32	10K						
TABLE ACCES...	CARGUYS	849K	74K		412	10K			62K	100	8	62
PX RECEIVE		95M	74K		32	50M						
PX SEND HASH	:TQ10001	95M	74K		32	50M						
PX BLOCK ITER...		95M	74K		32	50M						
TABLE ACCES...	CARGUYS	95M	74K		412	50M			62K	99	4	38

ORACLE

Query 2: Broadcast Distribution Disabled

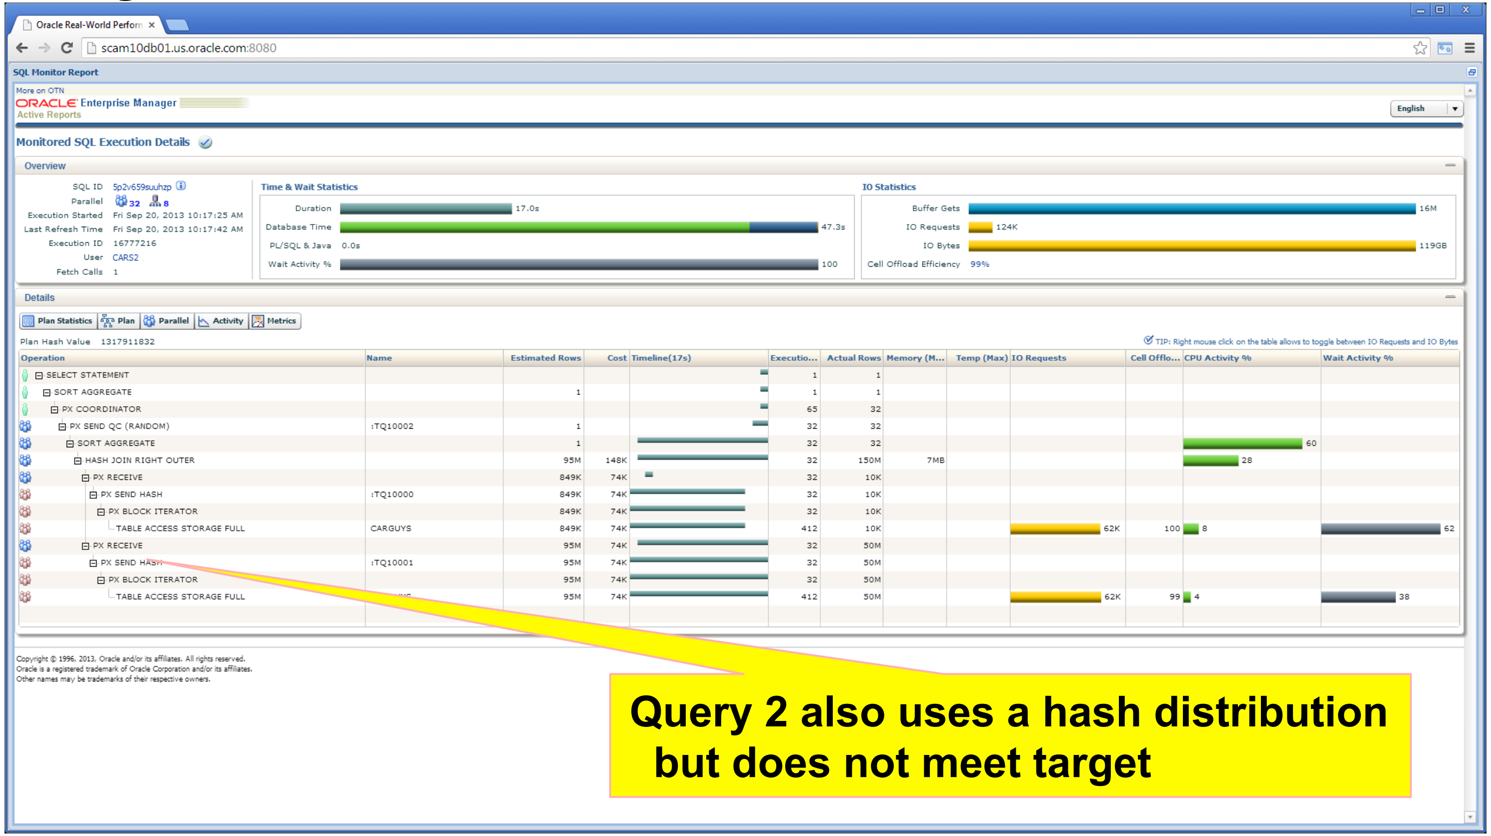

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	5p2v659suuhzp
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:17:25 AM
Last Refresh Time	Fri Sep 20, 2013 10:17:42 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	17.0s
Database Time	47.3s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1317911832

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(17s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					60	
HASH JOIN RIGHT OUTER		95M	148K		32	150M	7MB				28	
PX RECEIVE		849K	74K		32	10K						
PX SEND HASH	:TQ10000	849K	74K		32	10K						
PX BLOCK ITERATOR		849K	74K		32	10K						
TABLE ACCESS STORAGE FULL	CARGUYS	849K	74K		412	10K			62K	100	8	62
PX RECEIVE		95M	74K		32	50M						
PX SEND HASH	:TQ10001	95M	74K		32	50M						
PX BLOCK ITERATOR		95M	74K		32	50M						
TABLE ACCESS STORAGE FULL	[illegible]	95M	74K		412	50M			62K	99	4	38

Query 2 also uses a hash distribution but does not meet target

Query 2: Broadcast Distribution Disabled

10. Development Findings

- Plan uses a hash distribution
- Exceeds target

Second Query with Broadcast Distribution Enabled

ORACLE
REAL-WORLD PERFORMANCE

Query 2: Broadcast Distribution Enabled

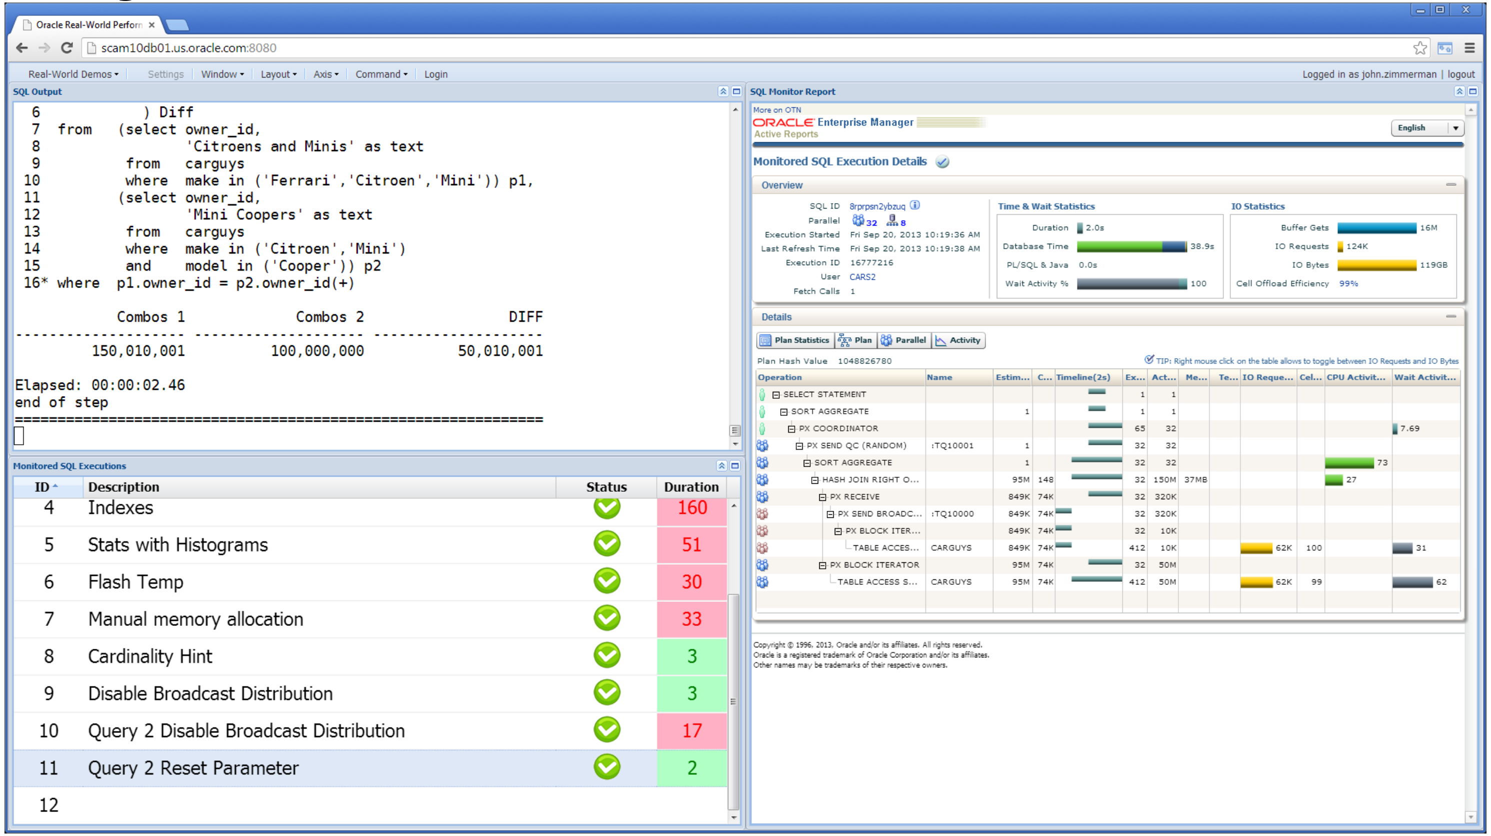

Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos
Settings
Window
Layout
Axis
Command
Login
Logged in as john.zimmerman | logout
SQL Output

```
  6            ) Diff
  7  from    (select owner_id,
                     'Citroens and Minis' as text
  9           from   carguys
 10           where  make in ('Ferrari','Citroen','Mini')) p1,
 11          (select owner_id,
 12                  'Mini Coopers' as text
 13           from   carguys
 14           where  make in ('Citroen','Mini')
 15           and    model in ('Cooper')) p2
 16* where   p1.owner_id = p2.owner_id(+)

            Combos 1             Combos 2                 DIFF
-------------------- -------------------- --------------------
         150,010,001          100,000,000           50,010,001

Elapsed: 00:00:02.46
end of step
==============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
4	Indexes	✓	160
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12			

SQL Monitor Report

Monitored SQL Execution Details

Overview

SQL ID	8rprpsn2ybzuq
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:19:36 AM
Last Refresh Time	Fri Sep 20, 2013 10:19:38 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Time	38.9s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(2s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						7.69
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					73	
HASH JOIN RIGHT O...		95M	148		32	150M	37MB				27	
PX RECEIVE		849K	74K		32	320K						
PX SEND BROADC...	:TQ10000	849K	74K		32	320K						
PX BLOCK ITER...		849K	74K		32	10K						
TABLE ACCES...	CARGUYS	849K	74K		412	10K			62K	100		31
PX BLOCK ITERATOR		95M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	95M	74K		412	50M			62K	99		62

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Query 2: Broadcast Distribution Enabled

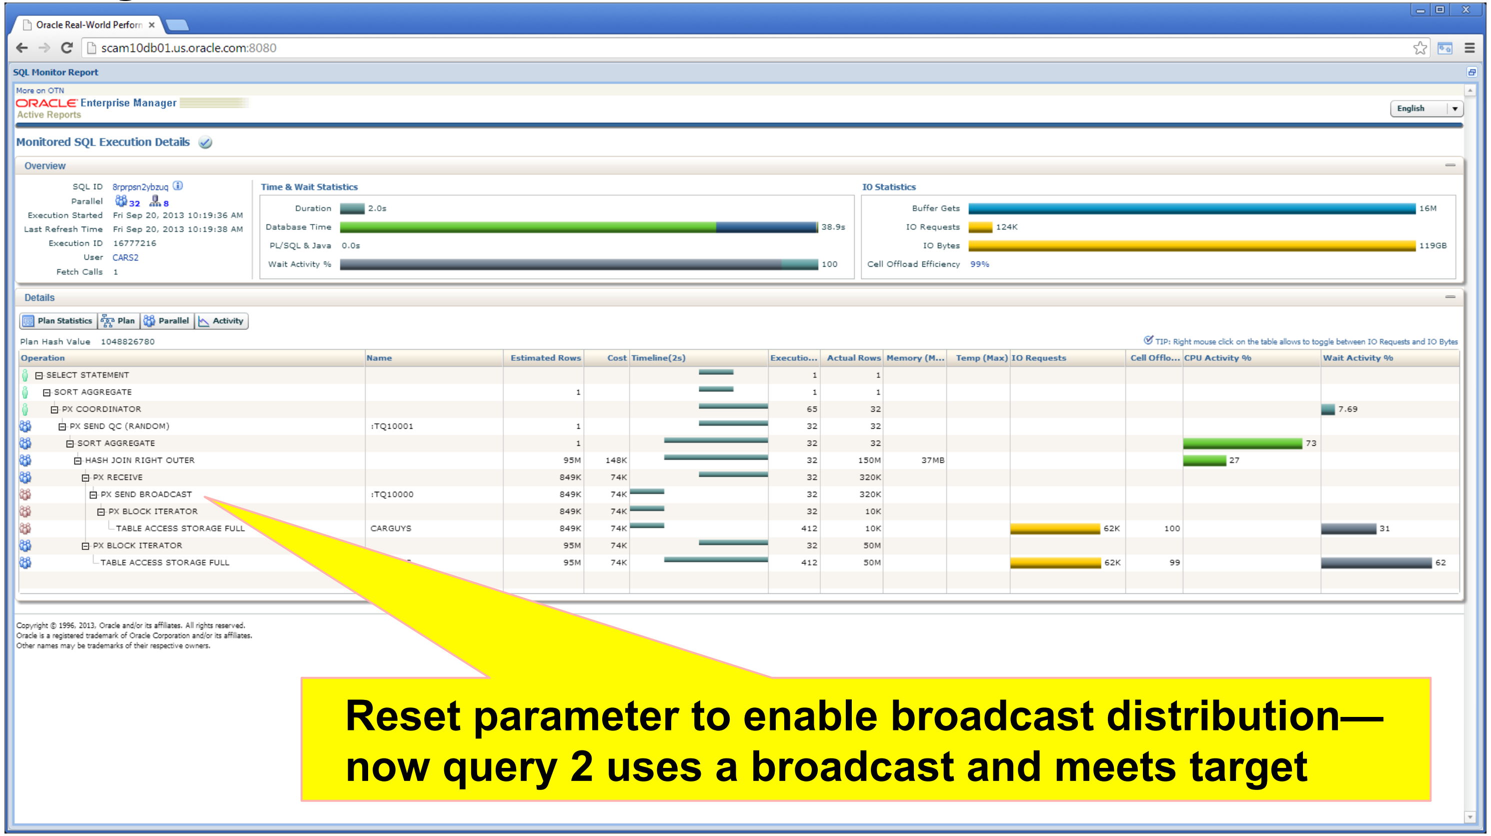

Oracle Real-World Perforn ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	8rprpsn2ybzuq
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:19:36 AM
Last Refresh Time	Fri Sep 20, 2013 10:19:38 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Time	38.9s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(2s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						7.69
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					73	
HASH JOIN RIGHT OUTER		95M	148K		32	150M	37MB				27	
PX RECEIVE		849K	74K		32	320K						
PX SEND BROADCAST	:TQ10000	849K	74K		32	320K						
PX BLOCK ITERATOR		849K	74K		32	10K						
TABLE ACCESS STORAGE FULL	CARGUYS	849K	74K		412	10K			62K	100		31
PX BLOCK ITERATOR		95M	74K		32	50M						
TABLE ACCESS STORAGE FULL		95M	74K		412	50M			62K	99		62

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Reset parameter to enable broadcast distribution—now query 2 uses a broadcast and meets target

ORACLE

Query 2: Broadcast Distribution Enabled

ORACLE
REAL-WORLD PERFORMANCE

11. Development Findings

- Reset _parallel_broadcast_enabled
- Plan now uses a broadcast distribution
- Meets target
- Should not change system parameters to tune one query

ORACLE

Extended Stats

Extended Stats

ORACLE
REAL-WORLD PERFORMANCE

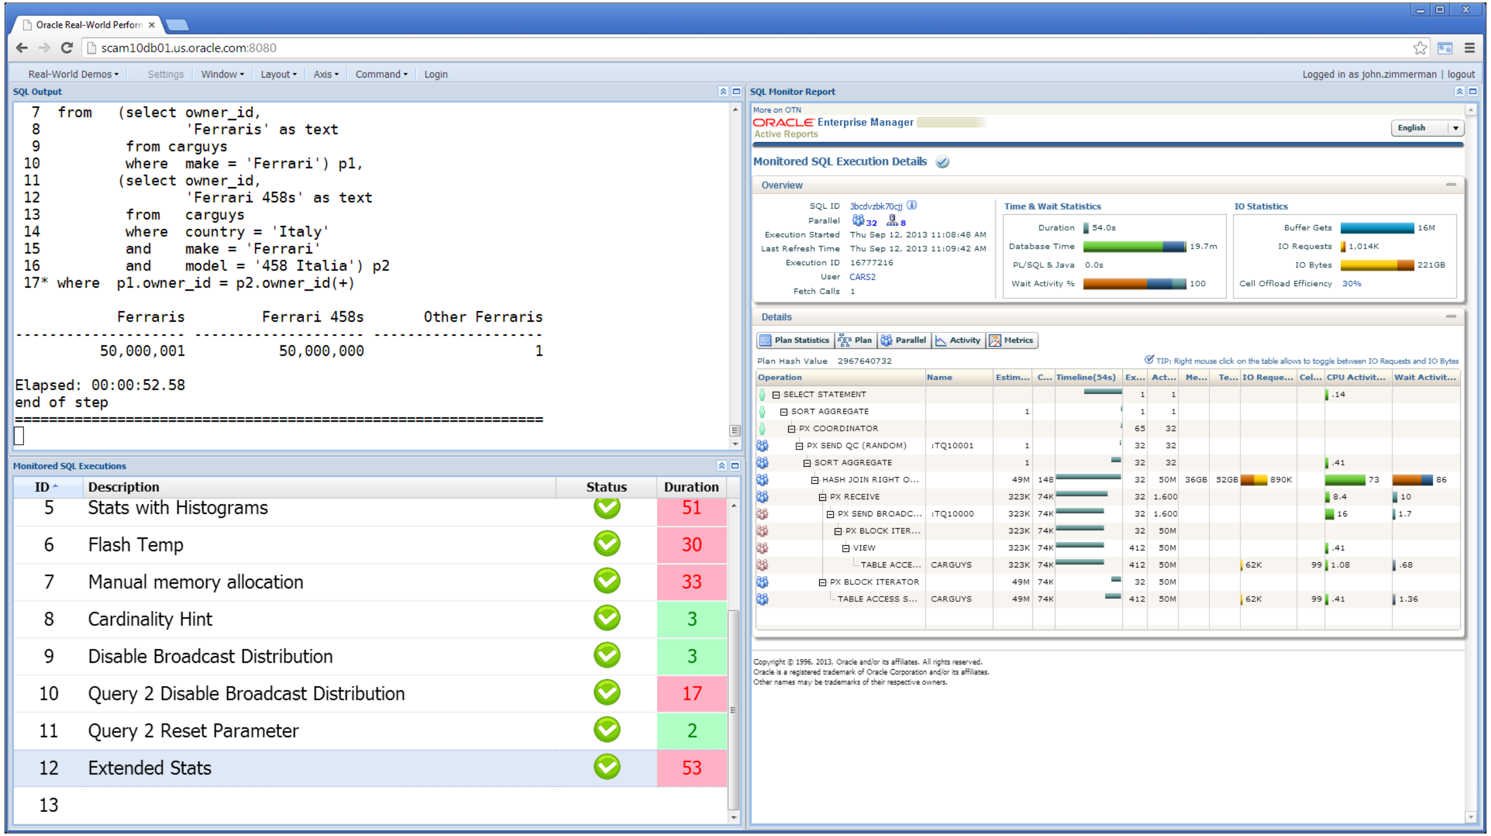
Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos
Settings
Window
Layout
Axis
Command
Login
Logged in as john.zimmerman | logout

SQL Output

```
 7  from   (select owner_id,
                   'Ferraris' as text
 9          from carguys
10          where  make = 'Ferrari') p1,
11         (select owner_id,
12                 'Ferrari 458s' as text
13          from   carguys
14          where  country = 'Italy'
15          and    make = 'Ferrari'
16          and    model = '458 Italia') p2
17* where  p1.owner_id = p2.owner_id(+)

            Ferraris          Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:52.58
end of step
==============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
5	Stats with Histograms	✓	51
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12	Extended Stats	✓	53
13			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	3bcdvzbk70cjj
Parallel	32 8
Execution Started	Thu Sep 12, 2013 11:08:48 AM
Last Refresh Time	Thu Sep 12, 2013 11:09:42 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	54.0s
Database Time	19.7m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	1,014K
IO Bytes	221GB
Cell Offload Efficiency	30%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 2967640732

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(54s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1					.14	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.41	
HASH JOIN RIGHT O...		49M	148		32	50M	36GB	52GB	890K		73	86
PX RECEIVE		323K	74K		32	1,600					8.4	10
PX SEND BROADC...	:TQ10000	323K	74K		32	1,600					16	1.7
PX BLOCK ITER...		323K	74K		32	50M						
VIEW		323K	74K		412	50M					.41	
TABLE ACCE...	CARGUYS	323K	74K		412	50M			62K	99	1.08	.68
PX BLOCK ITERATOR		49M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	49M	74K		412	50M			62K	99	.41	1.36

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

Extended Stats

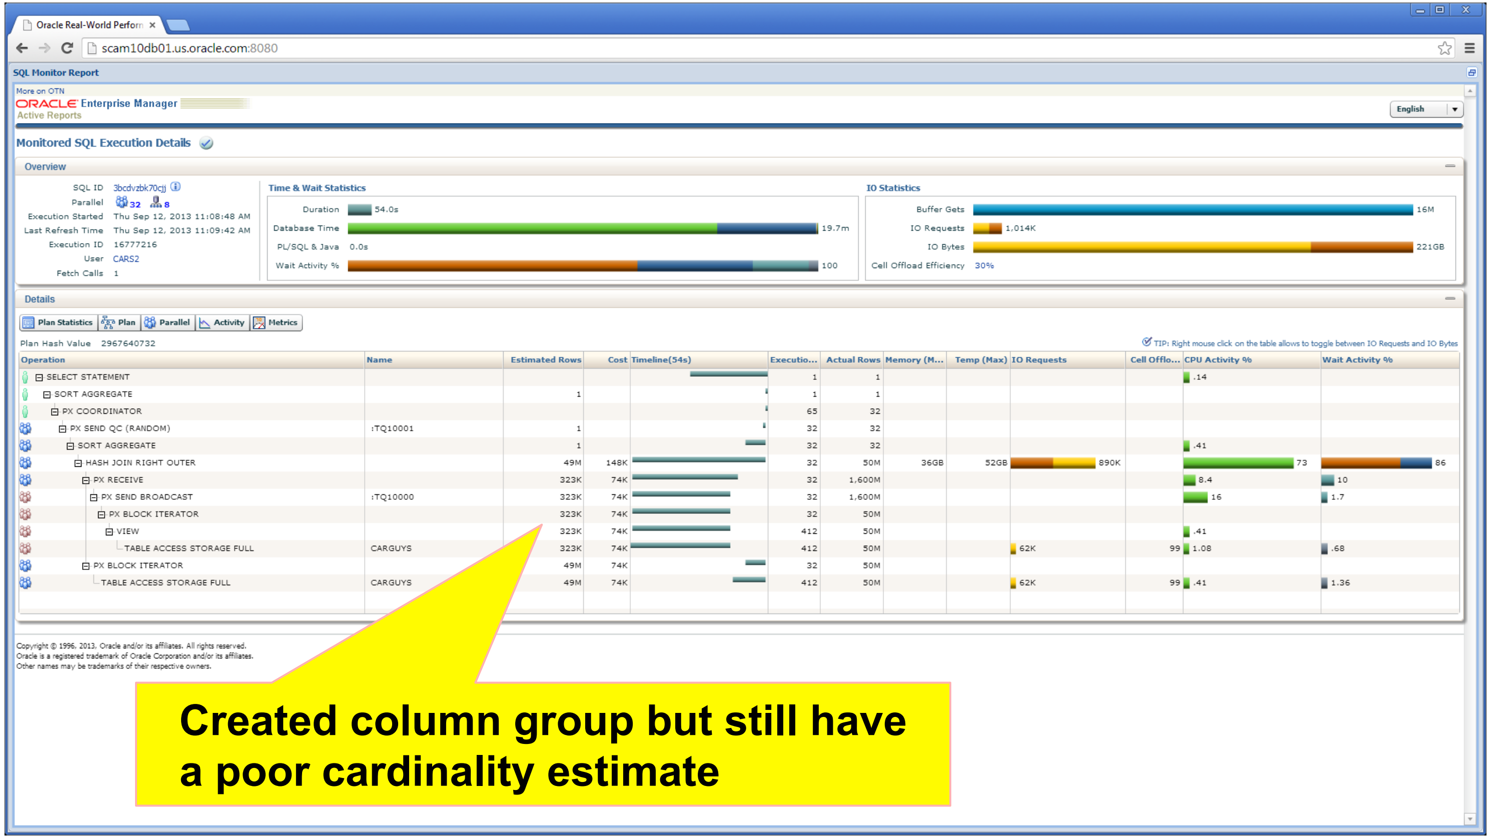

Oracle Real-World Perform ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	3bcdvzbk70cjj
Parallel	32 8
Execution Started	Thu Sep 12, 2013 11:08:48 AM
Last Refresh Time	Thu Sep 12, 2013 11:09:42 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	54.0s
Database Time	19.7m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	1,014K
IO Bytes	221GB
Cell Offload Efficiency	30%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 2967640732

✓ TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(54s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1					.14	
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.41	
HASH JOIN RIGHT OUTER		49M	148K		32	50M	36GB	52GB	890K		73	86
PX RECEIVE		323K	74K		32	1,600M					8.4	10
PX SEND BROADCAST	:TQ10000	323K	74K		32	1,600M					16	1.7
PX BLOCK ITERATOR		323K	74K		32	50M						
VIEW		323K	74K		412	50M					.41	
TABLE ACCESS STORAGE FULL	CARGUYS	323K	74K		412	50M			62K	99	1.08	.68
PX BLOCK ITERATOR		49M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	49M	74K		412	50M			62K	99	.41	1.36

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Created column group but still have a poor cardinality estimate

Extended Stats

12. Development Findings

- High correlation between Country, Make and Model columns
- Created column group
- Query still exceeds target
- Still have poor cardinality estimate

Histogram on Column Group

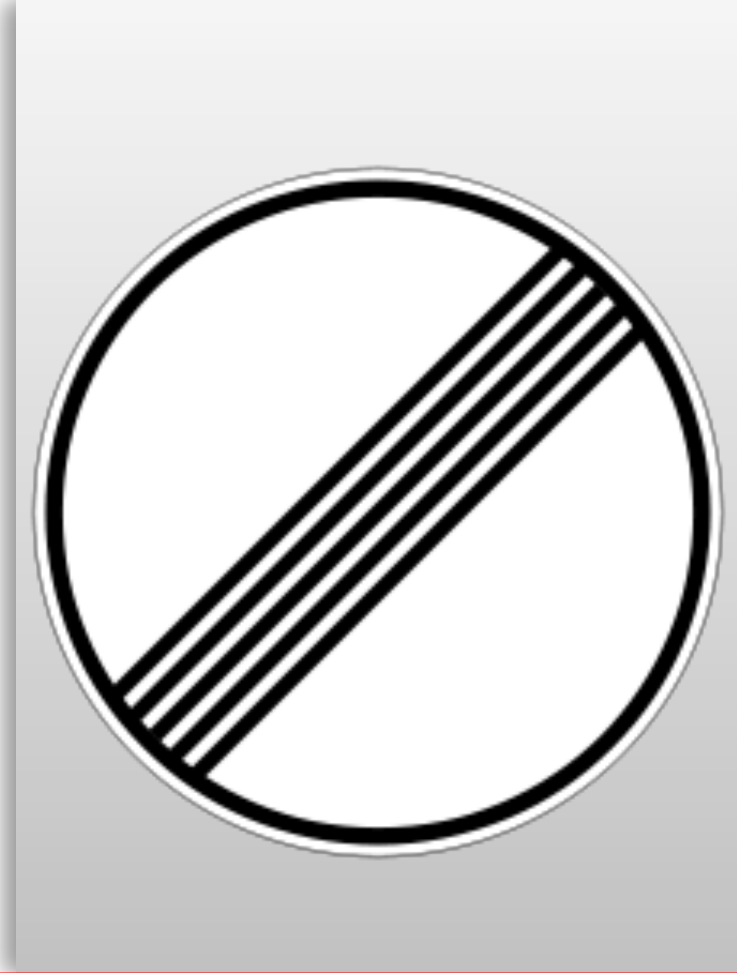

Histogram on Column Group

ORACLE
REAL-WORLD PERFORMANCE

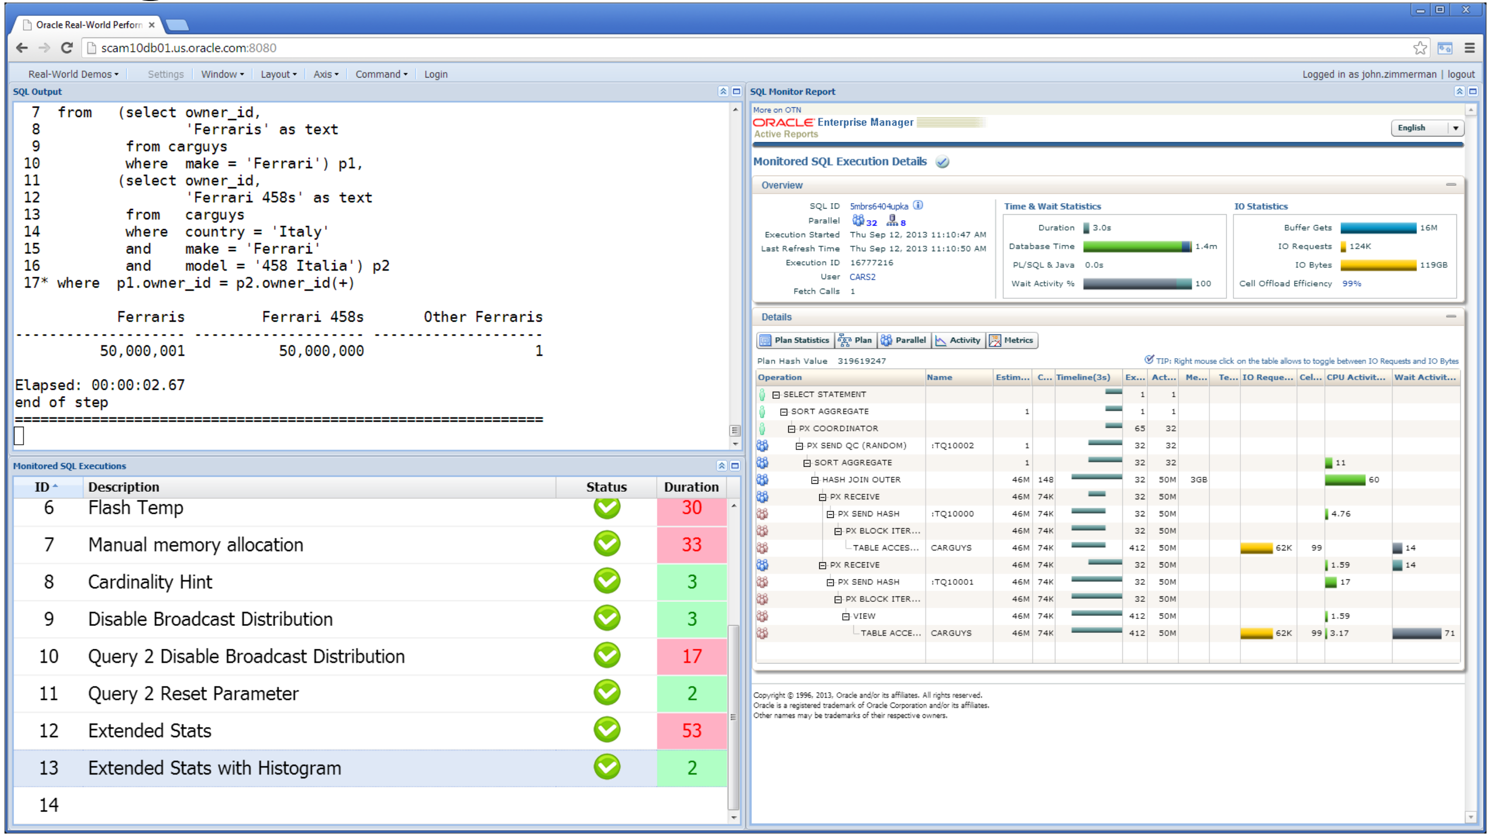

Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos
Settings
Window
Layout
Axis
Command
Login
Logged in as john.zimmerman | logout

SQL Output

```
from   (select owner_id,
               'Ferraris' as text
        from carguys
        where  make = 'Ferrari') p1,
       (select owner_id,
               'Ferrari 458s' as text
        from   carguys
        where  country = 'Italy'
        and    make = 'Ferrari'
        and    model = '458 Italia') p2
where  p1.owner_id = p2.owner_id(+)

          Ferraris          Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:02.67
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
6	Flash Temp	✓	30
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12	Extended Stats	✓	53
13	Extended Stats with Histogram	✓	2
14			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	5mbrs6404upka
Parallel	32 8
Execution Started	Thu Sep 12, 2013 11:10:47 AM
Last Refresh Time	Thu Sep 12, 2013 11:10:50 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 319619247

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(3s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					11	
HASH JOIN OUTER		46M	148		32	50M	3GB				60	
PX RECEIVE		46M	74K		32	50M						
PX SEND HASH	:TQ10000	46M	74K		32	50M					4.76	
PX BLOCK ITER...		46M	74K		32	50M						
TABLE ACCES...	CARGUYS	46M	74K		412	50M			62K	99		14
PX RECEIVE		46M	74K		32	50M					1.59	14
PX SEND HASH	:TQ10001	46M	74K		32	50M					17	
PX BLOCK ITER...		46M	74K		32	50M						
VIEW		46M	74K		412	50M					1.59	
TABLE ACCE...	CARGUYS	46M	74K		412	50M			62K	99	3.17	71

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

Histogram on Column Group

ORACLE
REAL-WORLD PERFORMANCE

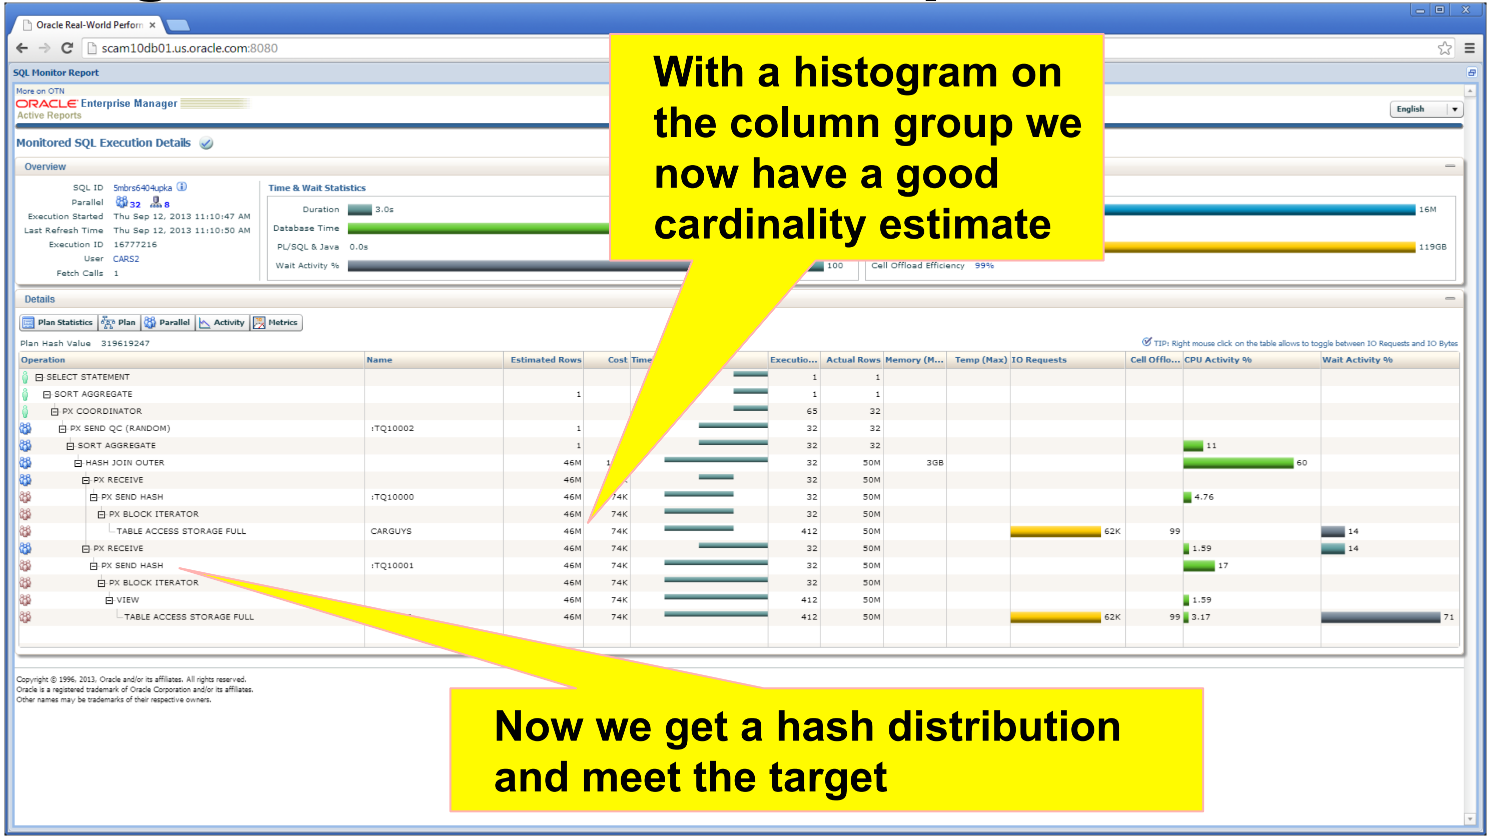

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	5mbrs6404upka
Parallel	32 8
Execution Started	Thu Sep 12, 2013 11:10:47 AM
Last Refresh Time	Thu Sep 12, 2013 11:10:50 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	
PL/SQL & Java	0.0s
Wait Activity %	100

16M

119GB

Cell Offload Efficiency 99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 319619247

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Time	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					11	
HASH JOIN OUTER		46M	1		32	50M	3GB				60	
PX RECEIVE		46M			32	50M						
PX SEND HASH	:TQ10000	46M	74K		32	50M					4.76	
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	46M	74K		412	50M			62K	99		14
PX RECEIVE		46M	74K		32	50M					1.59	14
PX SEND HASH	:TQ10001	46M	74K		32	50M					17	
PX BLOCK ITERATOR		46M	74K		32	50M						
VIEW		46M	74K		412	50M					1.59	
TABLE ACCESS STORAGE FULL		46M	74K		412	50M			62K	99	3.17	71

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

With a histogram on the column group we now have a good cardinality estimate

Now we get a hash distribution and meet the target

ORACLE

Histogram on Column Group

ORACLE
REAL-WORLD PERFORMANCE

13. Development Findings

- Re-gathered stats after running the query with the column groups
- Frequency Histogram on the column group
- Accurate cardinality estimates
- Optimizer now uses a hash distribution

ORACLE

Second Query with Histogram Column Group

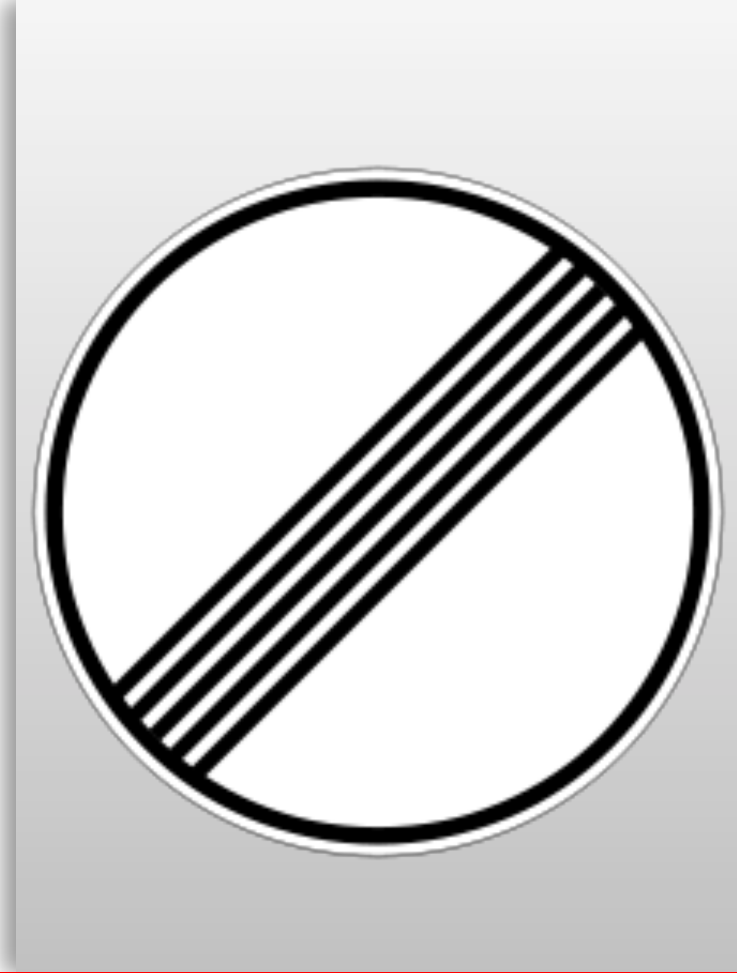

Query 2: Histogram Column Group

ORACLE
REAL-WORLD PERFORMANCE

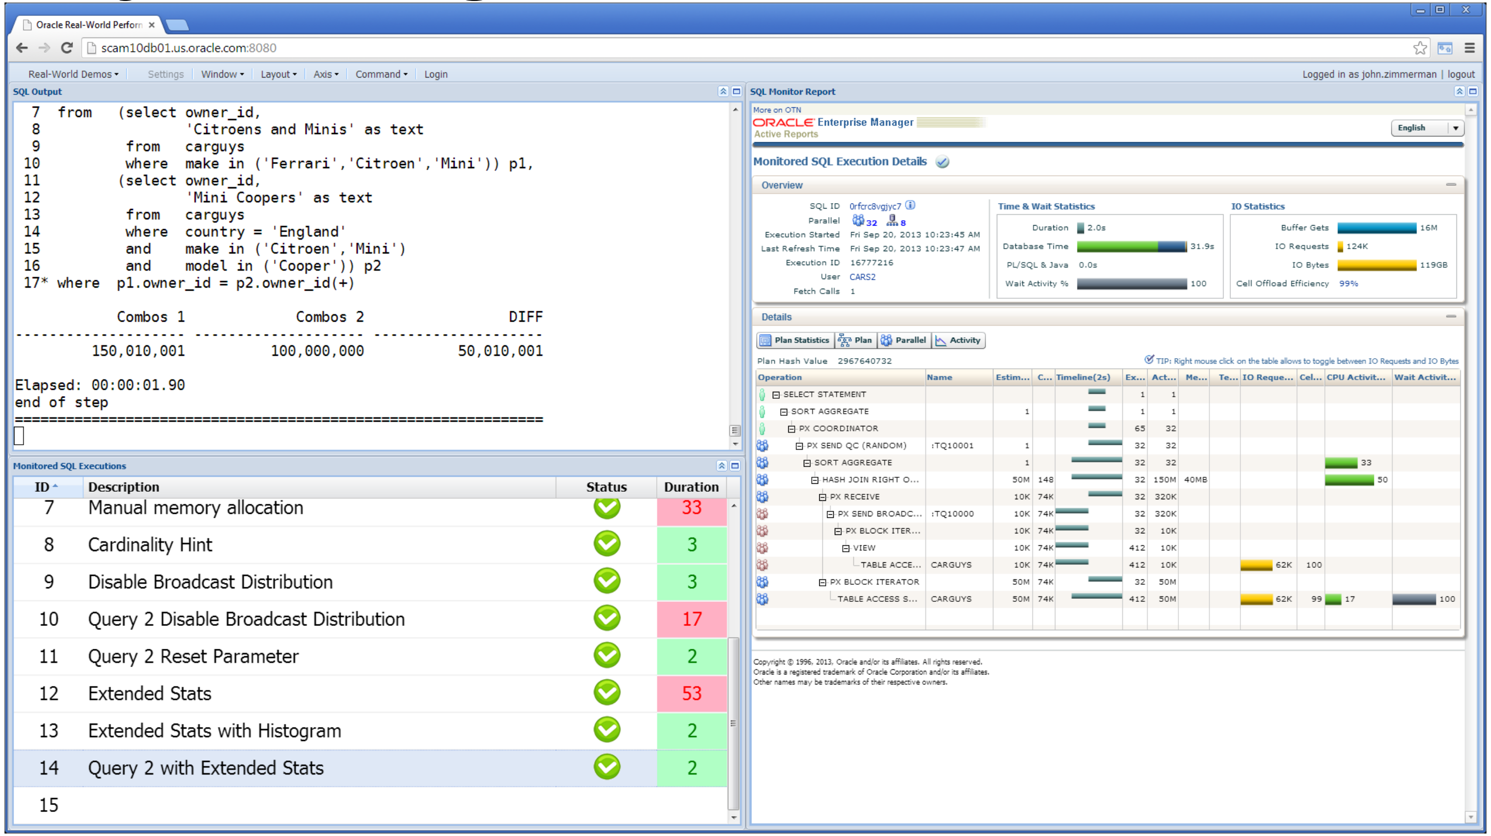
Oracle Real-World Perform
scam10db01.us.oracle.com:8080
Real-World Demos
Settings
Window
Layout
Axis
Command
Login
Logged in as john.zimmerman | logout
SQL Output
SQL Monitor Report

```
from   (select owner_id,
               'Citroens and Minis' as text
        from   carguys
        where  make in ('Ferrari','Citroen','Mini')) p1,
       (select owner_id,
               'Mini Coopers' as text
        from   carguys
        where  country = 'England'
        and    make in ('Citroen','Mini')
        and    model in ('Cooper')) p2
where  p1.owner_id = p2.owner_id(+)

           Combos 1             Combos 2                 DIFF
-------------------- -------------------- --------------------
         150,010,001          100,000,000           50,010,001

Elapsed: 00:00:01.90
end of step
==============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
7	Manual memory allocation	✓	33
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12	Extended Stats	✓	53
13	Extended Stats with Histogram	✓	2
14	Query 2 with Extended Stats	✓	2
15			

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	0rfcrc8vgjyc7
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:23:45 AM
Last Refresh Time	Fri Sep 20, 2013 10:23:47 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Time	31.9s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity

Plan Hash Value 2967640732

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(2s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					33	
HASH JOIN RIGHT O...		50M	148		32	150M	40MB				50	
PX RECEIVE		10K	74K		32	320K						
PX SEND BROADC...	:TQ10000	10K	74K		32	320K						
PX BLOCK ITER...		10K	74K		32	10K						
VIEW		10K	74K		412	10K						
TABLE ACCE...	CARGUYS	10K	74K		412	10K			62K	100		
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	50M	74K		412	50M			62K	99	17	100

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

ORACLE

Query 2: Histogram Column Group

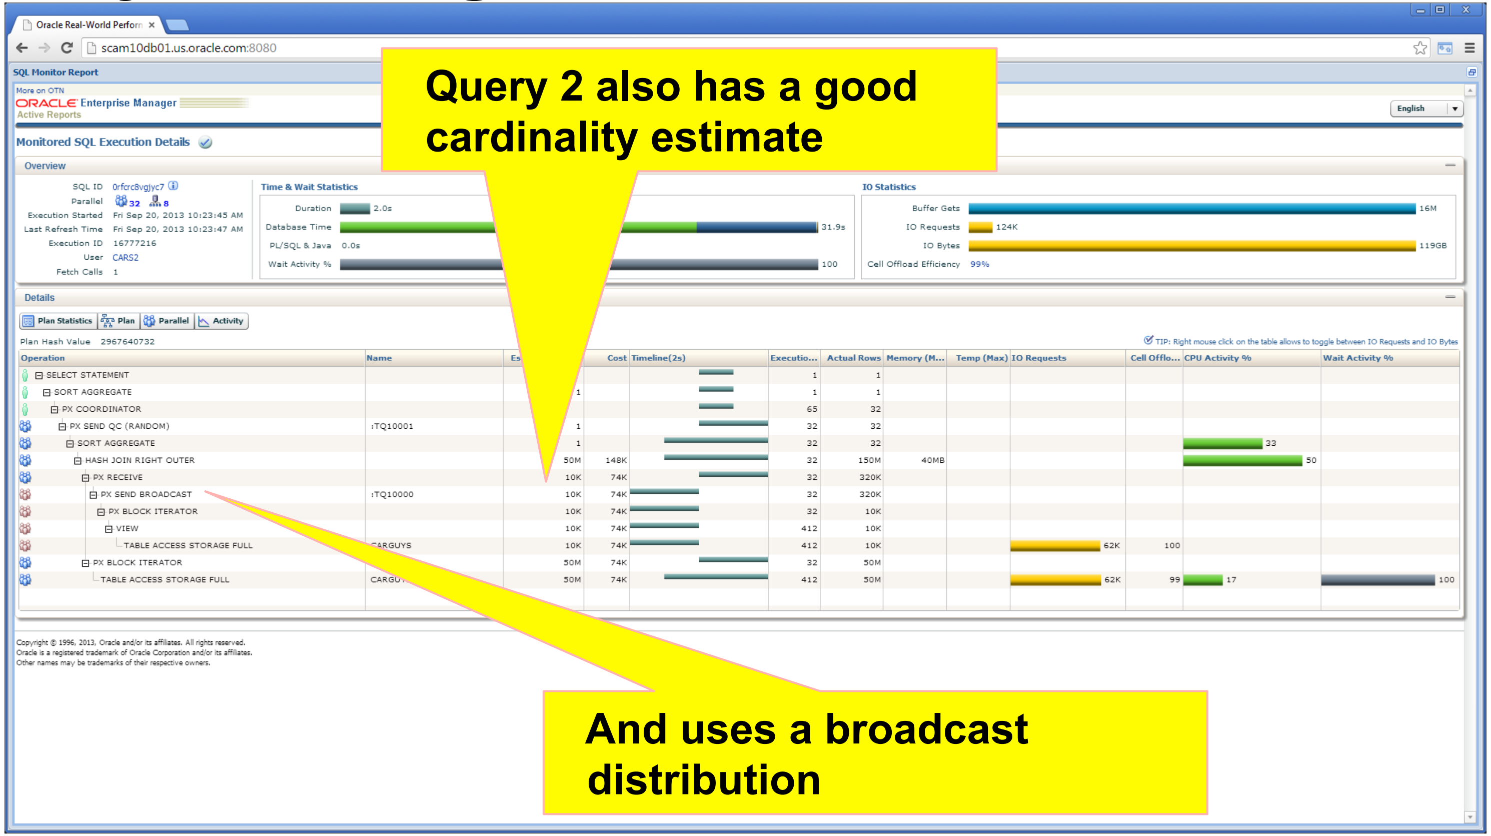

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN

ORACLE Enterprise Manager

Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	0rfcrc8vgjyc7
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:23:45 AM
Last Refresh Time	Fri Sep 20, 2013 10:23:47 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	2.0s
Database Time	31.9s
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity

Plan Hash Value 2967640732

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Es	Cost	Timeline(2s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					33	
HASH JOIN RIGHT OUTER		50M	148K		32	150M	40MB				50	
PX RECEIVE		10K	74K		32	320K						
PX SEND BROADCAST	:TQ10000	10K	74K		32	320K						
PX BLOCK ITERATOR		10K	74K		32	10K						
VIEW		10K	74K		412	10K						
TABLE ACCESS STORAGE FULL	CARGUYS	10K	74K		412	10K			62K	100		
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGU	50M	74K		412	50M			62K	99	17	100

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Query 2 also has a good cardinality estimate

And uses a broadcast distribution

Query 2: Histogram Column Group

14. Development Findings

- Accurate cardinality estimates
- Optimizer uses a broadcast distribution on second query

Histogram on Column Groups

Now we have the correct solution!

- Both queries have good cardinality estimates
- Correct plans
- Meet targets

Auto Column Group Creation: Seed Column Usage

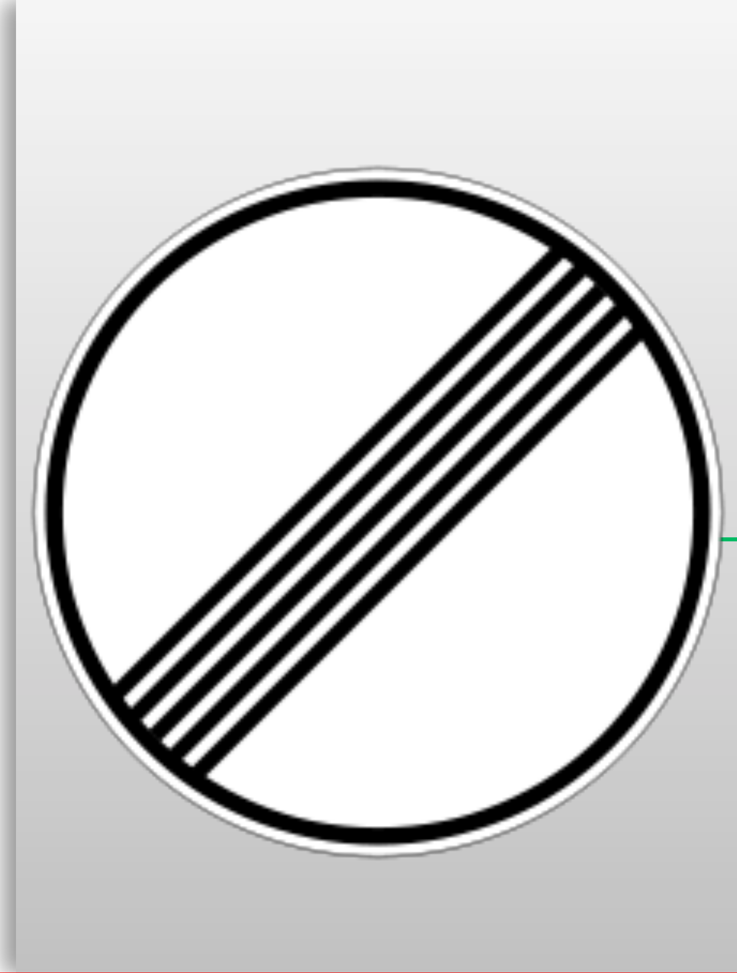

Auto Column Group Creation

REAL-WORLD PERFORMANCE

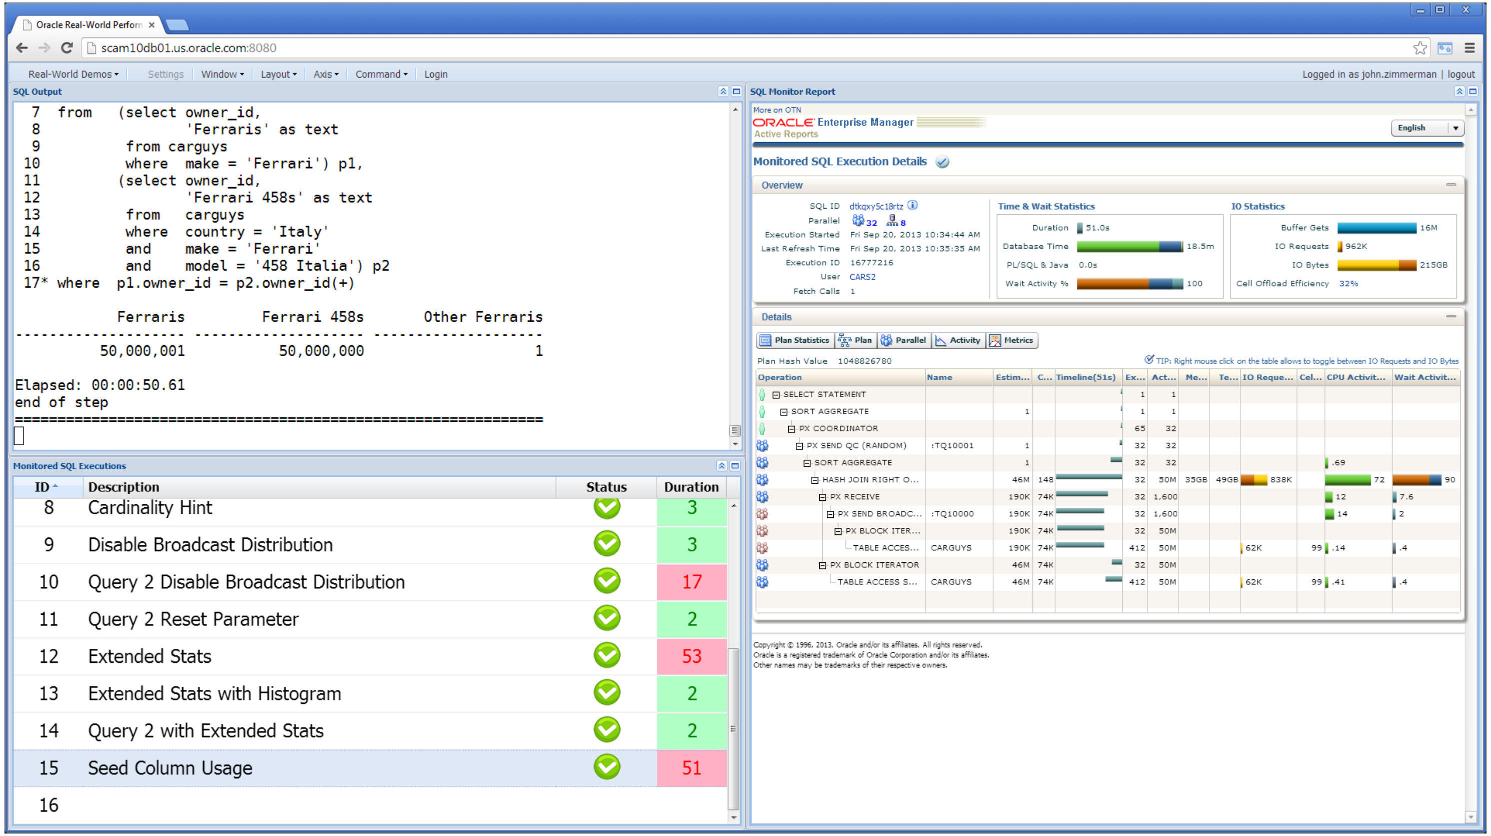

Oracle Real-World Perform

scam10db01.us.oracle.com:8080

Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login — Logged in as john.zimmerman | logout

SQL Output

```
 from   (select owner_id,
                'Ferraris' as text
         from carguys
         where  make = 'Ferrari') p1,
        (select owner_id,
                'Ferrari 458s' as text
         from   carguys
         where  country = 'Italy'
         and    make = 'Ferrari'
         and    model = '458 Italia') p2
 where  p1.owner_id = p2.owner_id(+)

           Ferraris          Ferrari 458s       Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:50.61
end of step
==============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12	Extended Stats	✓	53
13	Extended Stats with Histogram	✓	2
14	Query 2 with Extended Stats	✓	2
15	Seed Column Usage	✓	51
16			

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports
English

Monitored SQL Execution Details

Overview

SQL ID	dtkqxy5c18rtz
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:34:44 AM
Last Refresh Time	Fri Sep 20, 2013 10:35:35 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	51.0s
Database Time	18.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(51s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.69	
HASH JOIN RIGHT O...		46M	148		32	50M	35GB	49GB	838K		72	90
PX RECEIVE		190K	74K		32	1,600					12	7.6
PX SEND BROADC...	:TQ10000	190K	74K		32	1,600					14	2
PX BLOCK ITER...		190K	74K		32	50M						
TABLE ACCES...	CARGUYS	190K	74K		412	50M			62K	99	.14	.4
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS S...	CARGUYS	46M	74K		412	50M			62K	99	.41	.4

Auto Column Group Creation

ORACLE
REAL-WORLD PERFORMANCE

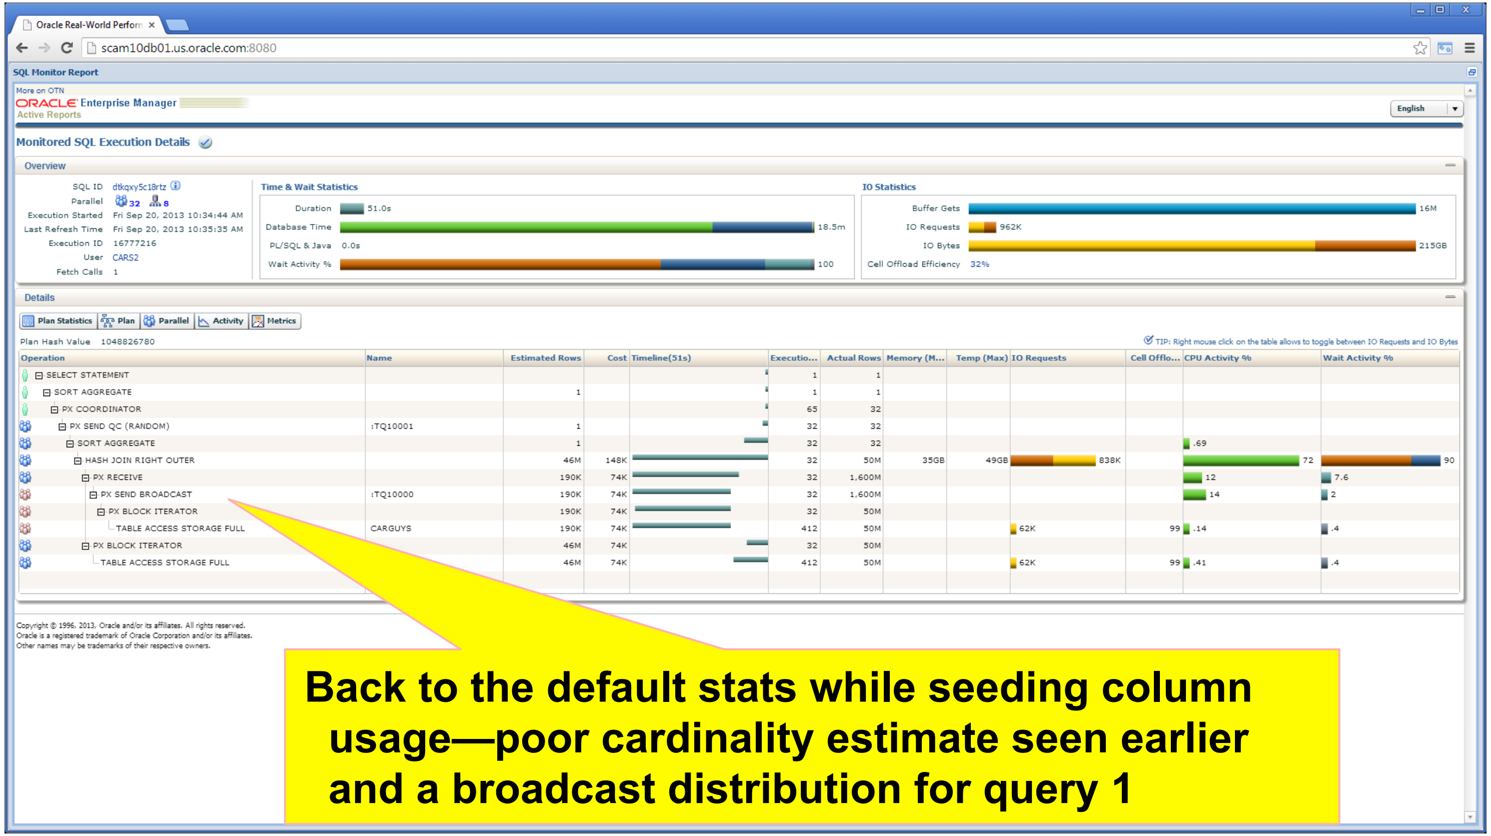

Oracle Real-World Perform
scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

English

Monitored SQL Execution Details

Overview

SQL ID	dtkqxy5c18rtz
Parallel	32 8
Execution Started	Fri Sep 20, 2013 10:34:44 AM
Last Refresh Time	Fri Sep 20, 2013 10:35:35 AM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	51.0s
Database Time	18.5m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	962K
IO Bytes	215GB
Cell Offload Efficiency	32%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 1048826780

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated Rows	Cost	Timeline(51s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10001	1			32	32						
SORT AGGREGATE		1			32	32					.69	
HASH JOIN RIGHT OUTER		46M	148K		32	50M	35GB	49GB	838K		72	90
PX RECEIVE		190K	74K		32	1,600M					12	7.6
PX SEND BROADCAST	:TQ10000	190K	74K		32	1,600M					14	2
PX BLOCK ITERATOR		190K	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	190K	74K		412	50M			62K	99	.14	.4
PX BLOCK ITERATOR		46M	74K		32	50M						
TABLE ACCESS STORAGE FULL		46M	74K		412	50M			62K	99	.41	.4

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Back to the default stats while seeding column usage—poor cardinality estimate seen earlier and a broadcast distribution for query 1

ORACLE

Auto Column Group Creation: Seed Column Usage

15. Development Findings

- Start with default statistics
- Execute dbms_stats.seed_col_usage to monitor column usage
- Run query

Auto Column Group Creation: Create Extended Stats

Auto Column Group Creation

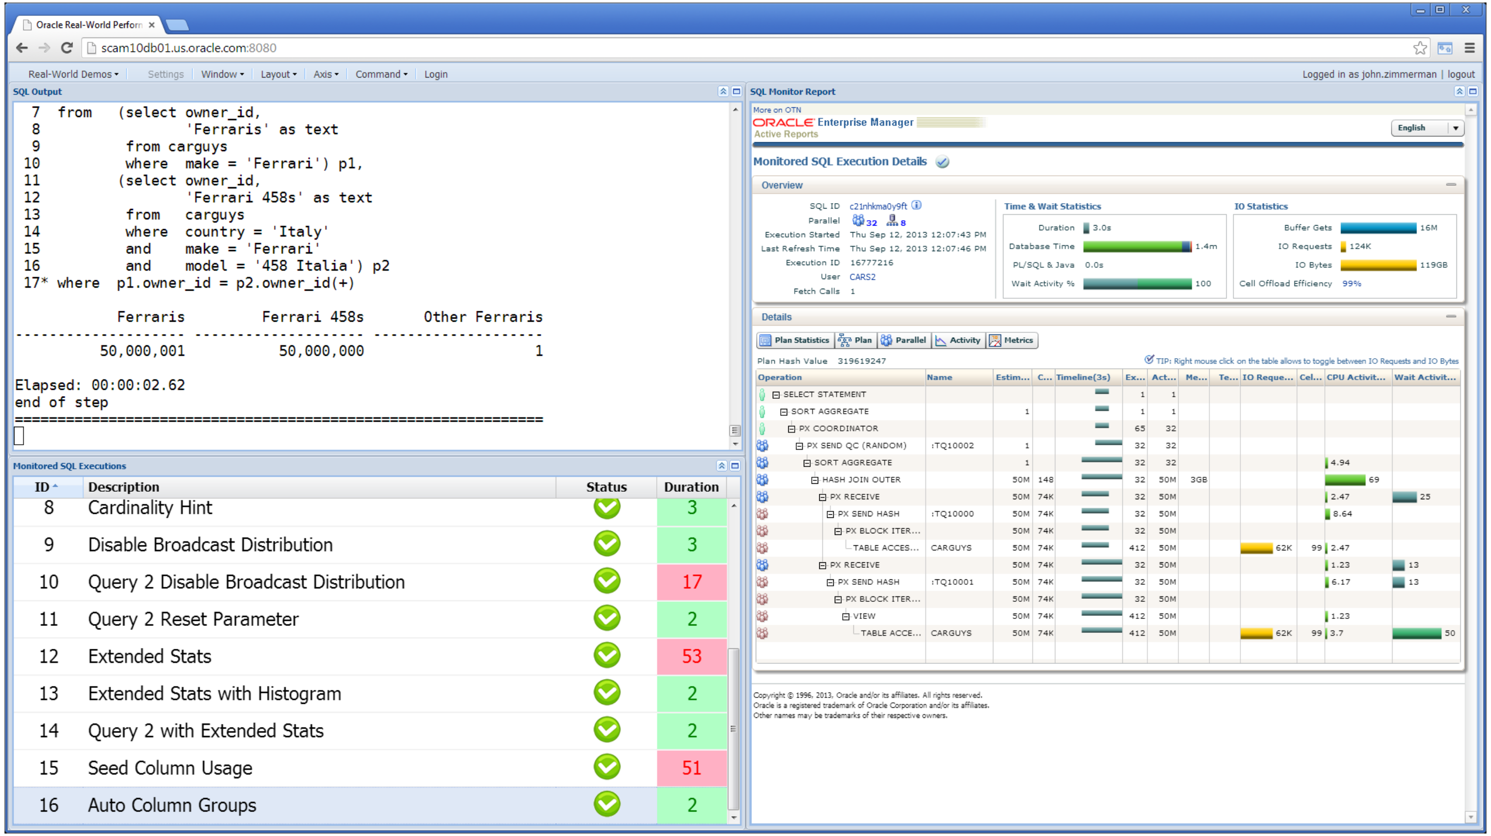
Oracle Real-World Perforn
scam10db01.us.oracle.com:8080
Real-World Demos ▾ | Settings | Window ▾ | Layout ▾ | Axis ▾ | Command ▾ | Login
Logged in as john.zimmerman | logout

SQL Output

```
 7  from   (select owner_id,
                   'Ferraris' as text
 9          from carguys
10          where  make = 'Ferrari') p1,
11         (select owner_id,
12                 'Ferrari 458s' as text
13          from   carguys
14          where  country = 'Italy'
15          and    make = 'Ferrari'
16          and    model = '458 Italia') p2
17* where  p1.owner_id = p2.owner_id(+)

           Ferraris         Ferrari 458s        Other Ferraris
-------------------- -------------------- --------------------
          50,000,001           50,000,000                    1

Elapsed: 00:00:02.62
end of step
============================================================
```

Monitored SQL Executions

ID	Description	Status	Duration
8	Cardinality Hint	✓	3
9	Disable Broadcast Distribution	✓	3
10	Query 2 Disable Broadcast Distribution	✓	17
11	Query 2 Reset Parameter	✓	2
12	Extended Stats	✓	53
13	Extended Stats with Histogram	✓	2
14	Query 2 with Extended Stats	✓	2
15	Seed Column Usage	✓	51
16	Auto Column Groups	✓	2

SQL Monitor Report

Monitored SQL Execution Details

Overview

SQL ID	c21nhkma0y9ft
Parallel	32 8
Execution Started	Thu Sep 12, 2013 12:07:43 PM
Last Refresh Time	Thu Sep 12, 2013 12:07:46 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Hash Value 319619247

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estim...	C...	Timeline(3s)	Ex...	Act...	Me...	Te...	IO Reque...	Cel...	CPU Activit...	Wait Activit...
SELECT STATEMENT					1	1						
SORT AGGREGATE		1			1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002	1			32	32						
SORT AGGREGATE		1			32	32					4.94	
HASH JOIN OUTER		50M	148		32	50M	3GB				69	
PX RECEIVE		50M	74K		32	50M					2.47	25
PX SEND HASH	:TQ10000	50M	74K		32	50M					8.64	
PX BLOCK ITER...		50M	74K		32	50M						
TABLE ACCES...	CARGUYS	50M	74K		412	50M			62K	99	2.47	
PX RECEIVE		50M	74K		32	50M					1.23	13
PX SEND HASH	:TQ10001	50M	74K		32	50M					6.17	13
PX BLOCK ITER...		50M	74K		32	50M						
VIEW		50M	74K		412	50M					1.23	
TABLE ACCE...	CARGUYS	50M	74K		412	50M			62K	99	3.7	50

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

Auto Column Group Creation

ORACLE
REAL-WORLD PERFORMANCE

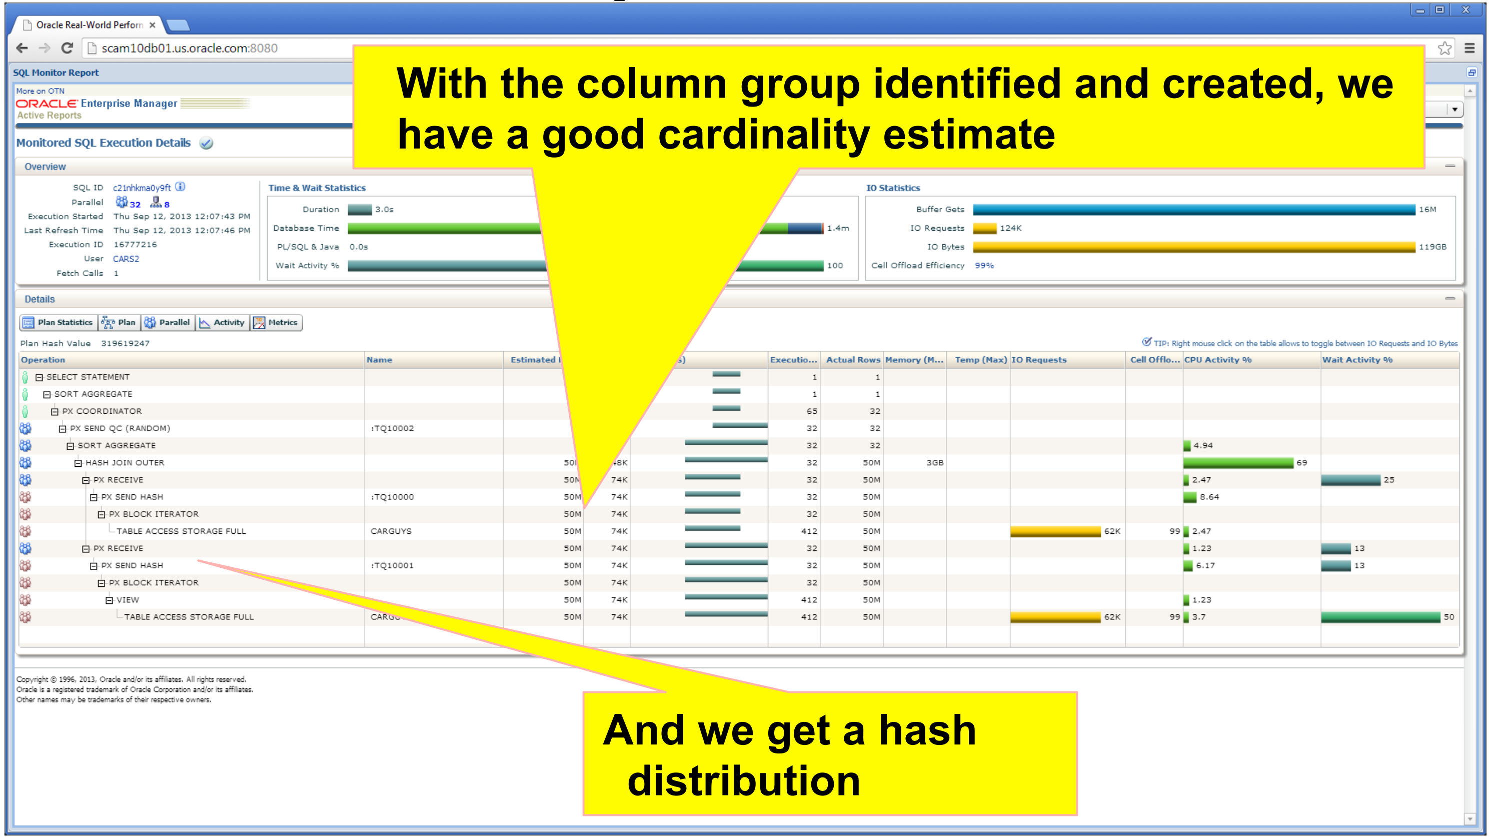

Oracle Real-World Perforn ×

scam10db01.us.oracle.com:8080

SQL Monitor Report

More on OTN
ORACLE Enterprise Manager
Active Reports

Monitored SQL Execution Details

Overview

SQL ID	c21nhkma0y9ft
Parallel	32 8
Execution Started	Thu Sep 12, 2013 12:07:43 PM
Last Refresh Time	Thu Sep 12, 2013 12:07:46 PM
Execution ID	16777216
User	CARS2
Fetch Calls	1

Time & Wait Statistics

Duration	3.0s
Database Time	1.4m
PL/SQL & Java	0.0s
Wait Activity %	100

IO Statistics

Buffer Gets	16M
IO Requests	124K
IO Bytes	119GB
Cell Offload Efficiency	99%

Details

Plan Statistics | Plan | Parallel | Activity | Metrics

Plan Hash Value 319619247

TIP: Right mouse click on the table allows to toggle between IO Requests and IO Bytes

Operation	Name	Estimated		s)	Executio...	Actual Rows	Memory (M...	Temp (Max)	IO Requests	Cell Offlo...	CPU Activity %	Wait Activity %
SELECT STATEMENT					1	1						
SORT AGGREGATE					1	1						
PX COORDINATOR					65	32						
PX SEND QC (RANDOM)	:TQ10002				32	32						
SORT AGGREGATE					32	32					4.94	
HASH JOIN OUTER		50	8K		32	50M	3GB				69	
PX RECEIVE		50M	74K		32	50M					2.47	25
PX SEND HASH	:TQ10000	50M	74K		32	50M					8.64	
PX BLOCK ITERATOR		50M	74K		32	50M						
TABLE ACCESS STORAGE FULL	CARGUYS	50M	74K		412	50M			62K	99	2.47	
PX RECEIVE		50M	74K		32	50M					1.23	13
PX SEND HASH	:TQ10001	50M	74K		32	50M					6.17	13
PX BLOCK ITERATOR		50M	74K		32	50M						
VIEW		50M	74K		412	50M					1.23	
TABLE ACCESS STORAGE FULL	CARG	50M	74K		412	50M			62K	99	3.7	50

Copyright © 1996, 2013, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.

With the column group identified and created, we have a good cardinality estimate

And we get a hash distribution

ORACLE

Auto Column Group Creation: Create Extended Stats

16. Development Findings

- dbms_stats.report_col_usage shows column groups identified during Seed Column Usage
- dbms_stats.create_extended_stats creates column groups identified
- Automatically identifies usage of Country, Make and Model columns together and creates column group

ORACLE
REAL-WORLD PERFORMANCE

Auto Column Group Creation: Create Extended Stats

16. Development Findings

- Regather stats
- Automatically creates Histogram on the column group
- Query meets target

ORACLE

What Did We Learn ?

What Did We Learn

- Moving to a solution before determining the root cause
- Not understanding the big/little data problem
  - Incorrect use of indexes
- Incorrect use of the product
  - Use DoP to “tune” the query
  - Use of flash cache for TEMP
  - Manual Memory Parameters
- System wide changes to fix a single SQL
  - Disable broadcast distribution

Performance Improvement Techniques

Which one are you?

- A Hacker?
- Performance Engineer?

Performance Improvement Techniques

REAL-WORLD PERFORMANCE

The Hacker

- Tuning by Google
- Tuning by pattern matching or word association
- Tuning by what worked well on another system
- Tuning by Folklore

ORACLE
REAL-WORLD PERFORMANCE

Performance Improvement Techniques

The Performance Engineer

- Understands Performance is all about work done in a period of time and systematically learns where the time is spent before making any recommendations
- Tuning is seen as a never ending process to systematically locate where the time is spent and making changes that reduces where the time is spent most
- Able to clearly articulate the problem and recommend the appropriate solution

Hardware and Software

ORACLE®

Engineered to Work Together

ORACLE®