Customer Experiences With Oracle on Linux on System z

Speaker: ***David Simpson – IBM***

simpson.dave@us.ibm.com

Thursday, August 7, 2014: 08:30 AM - 09:30 AM, DLLCC, Room 304

Session Number: ***14705***

Twitter -> ***@IBMandOracle***

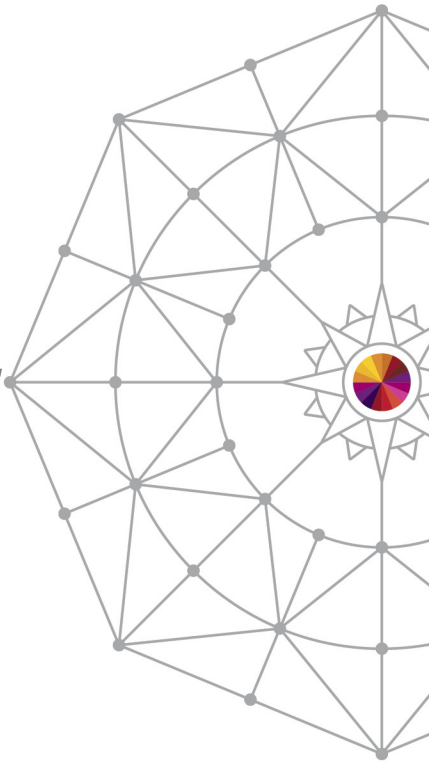

#SHAREorg

SHARE in Pittsburgh 2014

Copyright and Trademark Information

- For IBM – can be found at http://www.ibm.com/legal/us/en/copytrade.shtml
- For Oracle – can be found at http://www.oracle.com/us/legal/index.html

Agenda

- Recent Customer Cases
  - Oracle RAT Replay Test Customer Case
  - Oracle with Golden Gate Customer Case
- Oracle 12c Testing & New Features
- Q & A

Customer Case 1:

- Running Oracle with Linux on System z since 2010.
- Moving new workloads to Linux on System z.
- Encountered some unique performance challenges with one new Application.
- New to Oracle Real Application Clusters (RAC) and Oracle Real Application Testing (RAT),

Current z196 Linux Layout

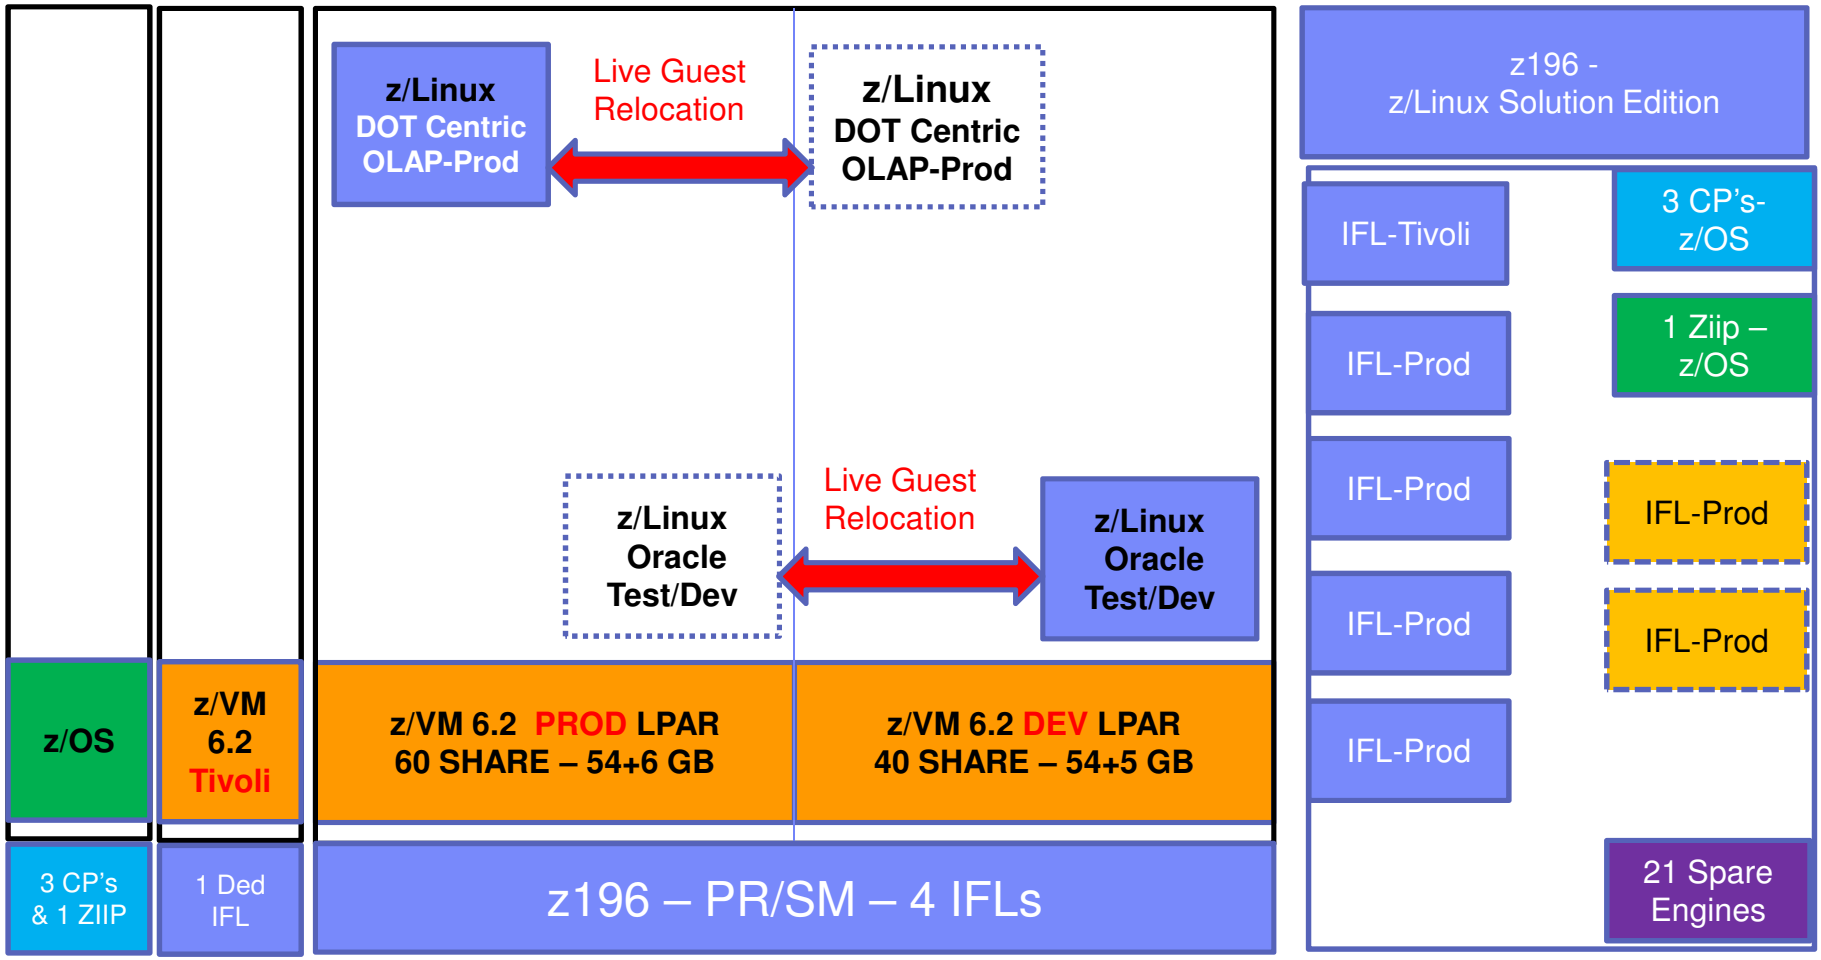

** z196 – 2 books with 21 additional IFLs that could be "activated" with microcode

The Problem….

Testing new workload with Linux on System z, higher than expected CPU observed in testing.

Source non System z – 4 cpu cores no RAC

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		3,060		49.80	
log file sync	308,417	1,746	6	28.43	Commit
direct path read	1,016,061	853	1	13.89	User I/O
SQL*Net more data to client	1,953,042	282	0	4.60	Network
db file sequential read	39,396	118	3	1.92	User I/O

Host CPU (CPUs: 4 Cores: 4 Sockets: 4)

Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
		17.4	7.0		75.7

Target System z – 1 cpu core w/ RAC

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
gc cr block 2-way	462,628	2,246	5	8.94	Cluster
gc buffer busy acquire	122,076	1,172	10	4.67	Cluster
DB CPU		982		3.91	
gc current block 2-way	156,316	923	6	3.67	Cluster
log file sync	70,396	769	11	3.06	Commit

Host CPU (CPUs: 1 Cores: 1 Sockets: 1)

Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
6.00	15.54	66.1	15.7	0.5	13.8

Perform Oracle Health Check

Asked to do an Oracle Health Check:

- Apply Oracle 11.2.0.4 with PSU 2 (11.2.0.4.2) to grid and then DB HOME (*Standard Answer*)
- set **filesystemio_options='setall'** if db files on a file system (non RAC)
- For highly transactional Oracle RAC databases consider using 4K block size
  - Cannot use database replay in this case
  - **4K block region– small indexes and sequences**
- For Oracle RAC databases – suggest creating a "service" for distinct workloads
- Try Oracle sequence caching, pre fetch 1000+ sequences (noorder) on each RAC Cluster node
- Tune SQL & Investigate index compression (pctfree 0) for queries with block contention.

SQL Ordered by Gets of AWR Report:

Buffer Gets	Executions	Gets per Exec	%Total	Elapsed Time (s)	%CPU	%IO	SQL Id	SQL Module	SQL Text
745,786	4,994	149.34	4.65	65.96	1.8	0	az33m61ym46y4	JDBC Thin Client	SELECT NULL AS table_cat, o.ow...
722,860	20	36,143.00	4.51	30.78	2.8	0	2yqr2axb9cxub	JDBC Thin Client	select * from (select detecto...
472,485	13	36,345.00	2.95	12.80	4.4	0	bd6amr4mdfq3y	JDBC Thin Client	select * from (select detecto...

After Getting Some Data….Some Initial Recommendations…

(1) Turn REORDER (VM) off for Large Virtual Memory Oracle Linux guests.

- Reorder can delay the guest for approximately 1seconds per 8g.
- Recommend turn REORDER off for any Oracle guest 8G or greater.
- Upgrade to z/VM 6.3 (which has re-order turned off automatically)

(2) Oracle Support Recommended -> turn off ASLR (System z Linux kernel parameter – Address Space

- ASLR Linux feature is enabled and it is recommended to disable it by setting kernel.randomize_va_space = 0.
- For it add/modify this parameter in /etc/sysctl.conf kernel.randomize_va_space=0

(3) Increase virtual memory from 8GB to 10GB on each guest

- Observed some Linux swap when running on just one Linux guest

(4) Implement HugePages for better memory management.

- kernel parameter change, and re-start Oracle instances)
- Reduces the # of pages the kernel must manage and makes the system more efficient

(5) Increase Linux guests from 1 virtual CP to 2 virtual CPs

- Increase share so it can use 2 IFL's if it needs the resources.
- Objective is to get the workload to run; then scale back after we get the a successful run.

(6) Run the workload in a non RAC environment to tune, then implement with Oracle RAC

- Once we get the workload to run, implement the desired 2-node environment to provide the high availability.

Resolution Steps

Adjusting # of Virtual IFLs from 1 -> 2 reduced Oracle concurrency and overall cpu load

Linux on System z – 2 virtual cpu non RAC

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		4,811		67.43	
log file sync	322,112	548	2	7.68	Commit
direct path read	271,759	354	1	4.96	User I/O
db file scattered read	24,153	152	6	2.13	User I/O
latch: shared pool	774	89	115	1.25	Concurrency

Host CPU (CPUs: 2 Cores: 1 Sockets: 1)

Load Average Begin	Load Average End	%User	%System	%WIO	%Idle
0.71	3.88	45.9	3.8	5.2	49.1

Oracle Parm with 1 Virtual IFL

Parameter	Session Value	Instance Value
_spin_count	1	1
cpu_count	1	1
_lm_lms_spin	FALSE	FALSE
_mutex_spin_count	255	255

Oracle Parm with 2 Virtual IFL's

Parameter	Session Value	Instance Value
_spin_count	2000	2000
cpu_count	2	2
_lm_lms_spin	FALSE	FALSE
_mutex_spin_count	255	255

- An Oracle latch helps prevent two processes from simultaneously updating the same area of the SGA. **Oracle spin_count parameter is based on # of cpus. _spin_count will wait x# of cpu cycles per Oracle process that has to wait for a latch that is busy.**

Recommended Testing Methodologies

- Oracle Real Application Testing (RAT)
  - Compare DB Time's good stat to compare
  - LOBs sometimes an Issue.
  - Licensed product (may be able to get trial)
  - Potential downtime (to get good backup)
- Batch Job or SQL comparison scripts
  - Use user wall clock time to compare
  - Can test generate typical production load?
- Mercury Load runner or IBM Rational (Application Load Generators)
  - Licensable products
  - Scripting necessary (sequences)
  - Can scripts generate/mimic production workload?

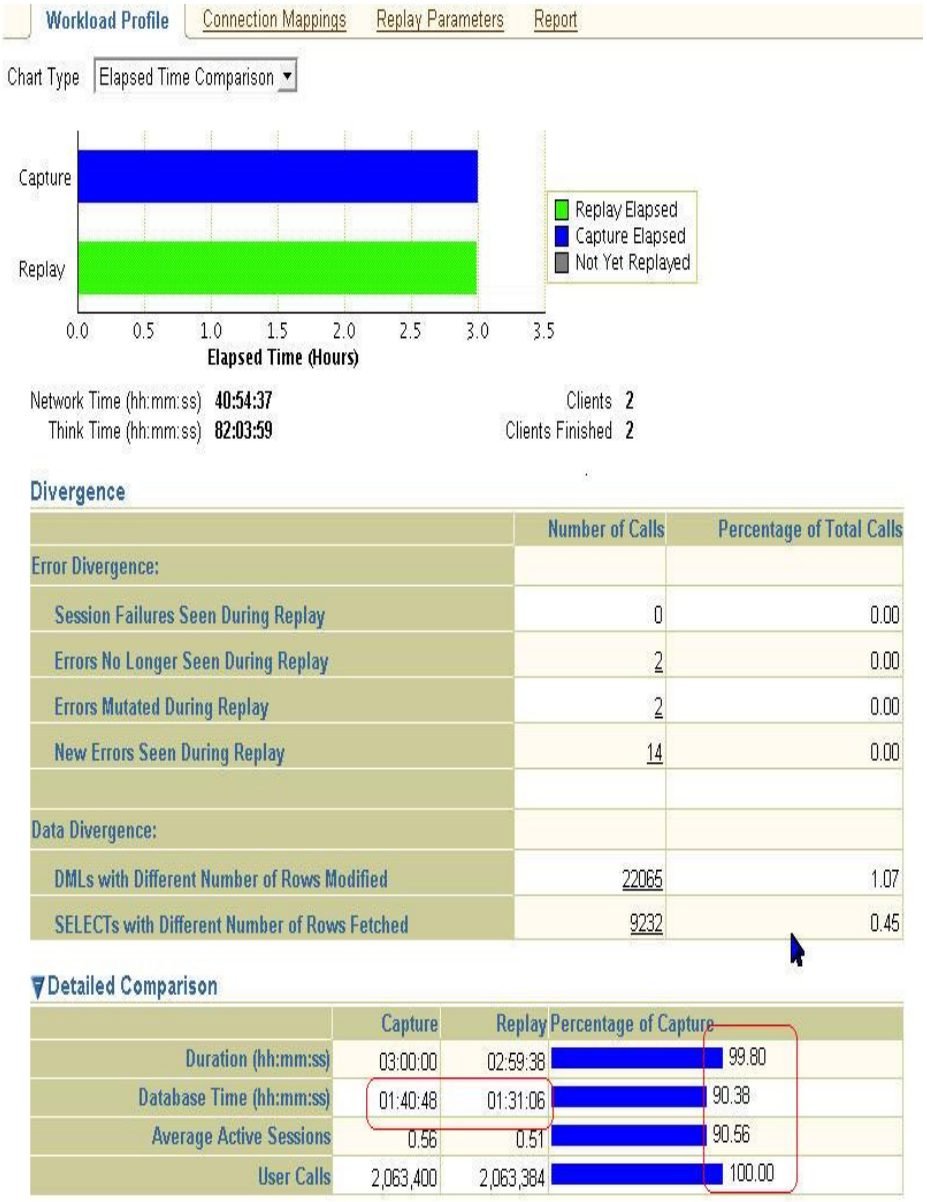

Oracle Real Application Testing

- Helps resolve the Issue of having the Application Team setup (multiple hours) and test each test run.
- Allows the capability to “capture” a production or test workload and replay in test with various configurations.

Replay Information

Information	Replay	Capture
Name	REPLAY-ATMSUATO-20140606142341	CAPTURE-OPENTMS-WINUAT-20140515
Status	COMPLETED	COMPLETED
Database Name	ATMSUATO	OPENTMS
Database Version	11.2.0.3.0	11.2.0.3.0
Start Time	15-05-14 14:21:06	15-05-14 14:21:37
End Time	15-05-14 15:43:37	15-05-14 15:36:31
Duration	1 hour 22 minutes 31 seconds	1 hour 14 minutes 54 seconds

Test 1 – Get Oracle RAT Working

Test 1 Initial Test Configuration:

- Set up RAT (Real Application Testing) workload client drivers on another Linux server (e.g. RMAN server)
- Increase 1 -> 2 virtual cpus on the Oracle (non RAC) Linux guest
- Applied Oracle RAT patch (13947480) to 11.2.0.3 target system to get it working

Replay Statistics

Statistic	Replay	Capture
DB Time	6988.917 seconds	36918.593 seconds
Average Active Sessions	1.38	8.22
User calls	16901532	16901520

Summary: System z Linux DB Time was 5.28 times (DB Time) better than original test environment.

Test 2 – Resource Tuning

Test 2 - Performance Tuning Single Node Test

- logoff non essential non-prod guests
- Striped database file system
- dedicated IFLs to the Linux guest (to help isolate/make more reproducible in Shared environment)

Replay Statistics

Statistic	Replay	Capture
DB Time	9020.154 seconds	36918.593 seconds
Average Active Sessions	1.82	8.22
User calls	16901532	16901520

Summary: System z Linux DB Time was 4.09 times better than current environment (slightly slower with tuning ☹).

Test 3 – Test workload with Oracle RAC

Test 3 RAC environment,

- configure IFL's to be dedicated to each guest (rac1 & rac2)
- Applied RAT patch to 11.2.0.3 Oracle RAC guests
- Set up and ran the Replay of the workload on the RAC guests
- Oracle workload balanced to both nodes in the cluster

Replay Statistics

Statistic	Replay	Capture
DB Time	9344.161 seconds	36918.593 seconds
Average Active Sessions	1.65	8.22
User calls	16901532	16901520

Summary: System z DB Time was **3.95 times** better than current environment.

Test 4 – Test Oracle RAC with DB Service for Contention

Test 4 Oracle RAC environment with DB service

- see if localizing workload to one node will help (i.e. address contention issues running RAC)

Replay Statistics

Statistic	Replay	Capture
DB Time	11416.267 seconds	36918.593 seconds
Average Active Sessions	2.14	8.22
User calls	16901532	16901520

Summary: System z DB Time is **3.23 times** better than test environment. Slight degradation from load balanced test (due top concurrency and spin_count NOT Oracle RAC).

Oracle RAT Testing Summary

- Testing showed having a minimum 2 Virtual IFL's reduced concurrency.
- Over allocating resources/dedicating IFLs did not help this particular workload.
- Using Oracle service to funnel transactions to one member of RAC node did not help for this workload.
- Oracle 11.2.0.3 -> 11.2.0.4 RAT replay test did not complete.

Test	DB Time (seconds)	x Improvement
Baseline (non System z)	36919	
Single Node (2 Virtuals)	6989	5.28x
Single Node (2 dedicated IFLs / disk striping)	9020	4.09x
RAC 2 node (2 dedicated IFLs each node – rac1 & rac2)	9344	3.95x
RAC 2 node (with Service)	11416	3.23x

Customer Experience Case 2:

- Customer asked us to size an application running IBM WebSphere (WAS), Oracle Database with Oracle GoldenGate replication (help desk & reporting)
- AWR for Oracle , WAS (heap & garbage collector size) to size memory.
- Linux sar data gathered for peak workload days.
- Gathered existing source system server model and number of cores (processor speed not given initially)
- Sizing tools used a Java light and database workload model.
- Conducted a POC to validate the configuration.

Original Sizing Estimate:

Current System Configuration:

Vendor	Server Hardware Description and Lookup Identification	# OEM Servers		Default Values		Workload Assignment	
		Enter #	Result	90.0%	65.0%	No.	Description
Fujitsu	SPARC T4-2 (3U) SPARC T4 2.85GHz (2ch/16co)	4.00	4.00	90.0%	25.0%	5	Java Lite
Fujitsu	SPARC Enterprise M9000-64 SPARC64 VII 2.52GHz 6MB (32ch/128co)	1.00	1.00	90.0%	25.0%	6	Database

Current System to IFLs with Incorrect Model & Server Utilizations

Processor	Feature	MSU	Capacity Rating	Utilization for Case 1				Utilization for Case 2			
				< Complementary		Peaks	Concurrent >	< Complementary		Peaks	Concurrent >
				0%	40.0%	70.0%	100%	0%	40.0%	70.0%	100%
2827-7xx I12	12W IFL		17,448	173%	247%	303%	359%	48%	69%	85%	100%
2827-7xx I13	13W IFL		18,733	161%	230%	282%	334%	45%	64%	79%	93%
2827-7xx I14	14W IFL		19,998	151%	216%	265%	313%	42%	60%	74%	87%
2827-7xx I15	15W IFL		21,245	142%	203%	249%	295%	40%	57%	70%	82%
2827-7xx I16	16W IFL		22,472	134%	192%	235%	279%	38%	54%	66%	78%

Rough Sizing Initially Estimated 14-16 IFLs …

Source: Oracle Federal Acquisition Price List

POC Results Summary

Test Configuration	App Server IFLs	DB Server IFLs	Total IFLs	Ack. Txn Rate	App Server CPU %	DB Server CPU %	IFL Usage	Acks/IFL
4 App Servers to 1 DB	24	12	36	226	84.4	98.0	32.0	7.06
4 App Servers to 1 DB	32	18	50	275	78.0	82.0	39.7	6.92
4 App Servers to 1 DB	40	20	60	310	67.0	83.0	43.4	7.14
2 App Servers to 1 DB	12	8	20	147	92.0	78.0	17.3	8.51
2 App Servers to 1 DB	20	12	32	226	80.0	91.0	26.9	8.40
2 App Servers to 1 DB	40	20	60	376	65.4	84.4	43.0	8.74

- **Alternate topologies (2 and 4 nodes) were used to show vertical scalability with Linux on System z.**
- **Throughput per IFL remained relatively constant (8.x) in the alternative topologies**
- **Workload mix used for all test cases were the same including batch scheduling jobs, replication etc.**
- **No lag in replication was observed unlike in current production**
- **Application ported with NO Code changes or issues.**

Application Showed Linear Growth as Capacity Increased

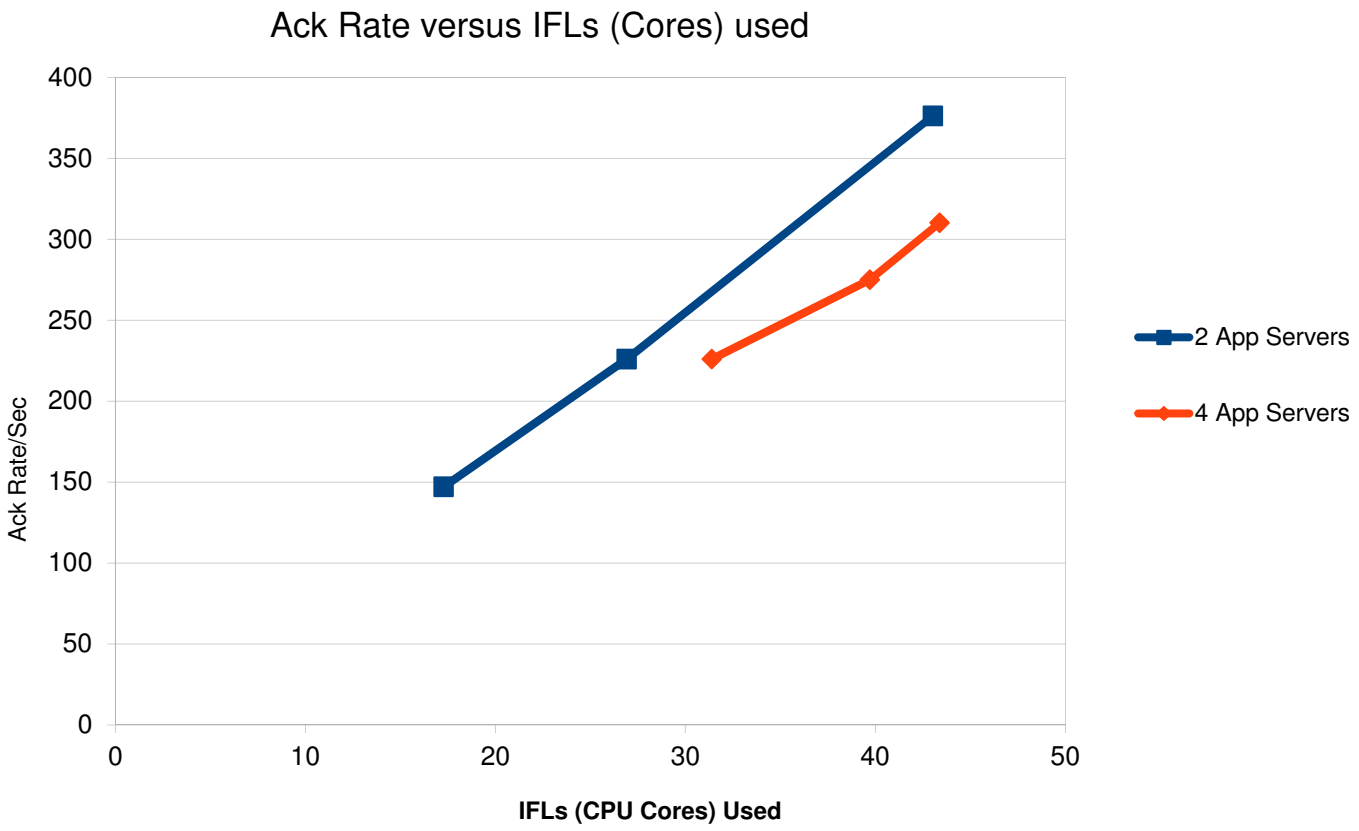

- **Transaction rate increased proportionally to net increase in number of IFLs (capacity)**

New Sizing Tool Calculations:

Current System Configuration:

Vendor	Server Hardware Description and Lookup Identification	Enter #	Result	90.0%	65.0%	No.	Description
Fujitsu	SPARC Enterprise T5440 (4U) UltraSPARC T2+ 1.4GHz 4MB (4ch/32co)	4.00	4.00	90.0%	50.0%	3	Java Heavy
Fujitsu	SPARC Enterprise M9000-64 SPARC64 VII+ 3.0GHz 12MB (32ch/128co)	1.00	1.00	90.0%	35.0%	7	Mixed High

Current System to IFLs with Correct Model, CPU types, and Utilizations:

Processor	Feature	MSU	Capacity Rating	Utilization for Case 1 < Complementary 0%	Peaks 40.0%	Peaks 70.0%	Concurrent > 100%	Utilization for Case 2 < Complementary 0%	Peaks 40.0%	Peaks 70.0%	Concurrent > 100%
2827-7xx I27	27W IFL		34,973	112%	147%	173%	200%	50%	70%	86%	101%
2827-7xx I28	28W IFL		36,039	108%	143%	168%	194%	48%	68%	83%	98%
2827-7xx I29	29W IFL		37,095	105%	139%	164%	188%	47%	66%	81%	95%
2827-7xx I30	30W IFL		38,140	102%	135%	159%	183%	46%	64%	78%	93%
2827-7xx I31	31W IFL		39,175	100%	131%	155%	178%	44%	63%	76%	90%
2827-7xx I32	32W IFL		40,199	97%	128%	151%	174%	43%	61%	74%	88%
2827-7xx I33	33W IFL		41,213	95%	125%	147%	170%	42%	60%	73%	86%
2827-7xx I34	34W IFL		42,217	93%	122%	144%	166%	41%	58%	71%	84%

New estimate 32 IFLs when data recalculated based on correct workload

Source: Oracle Federal Acquisition Price List

Customer Case 2 Summary

- Large reduction in software costs with new environment.

- Enhanced scalability to support:
  - New Features / functionality
  - Increase in user demand
  - Supported cloud consumption model – have resources available when needed.

- When sizing existing workloads it's important to:
  - gather accurate system utilization percentages
  - Gather **both** make and model of server (Ghz)
  - Apply the **correct workload model** (Oracle Golden Gate mixed workload
  - Understand concurrent peaks of all systems being sized.
  - Consider growth of database in both size & increased transactions and new functionality.
  - Understand storage being utilized (ECKD DASD, FCP SCSI, Flash)
  - Always validate the model with a POC / and performance testing.

Image source: maggieofmadison.com

12c Database Cloud Architectures

IBM

Database Cloud

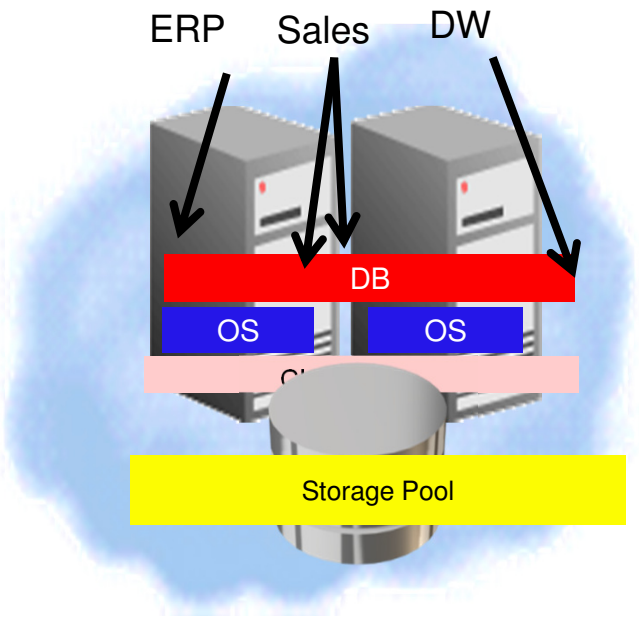

11gR2 Schema Cloud Model

- Multiple schemas share same database
  - Flexible Database Services
  - Fine grain service level elasticity
  - Use with most databases

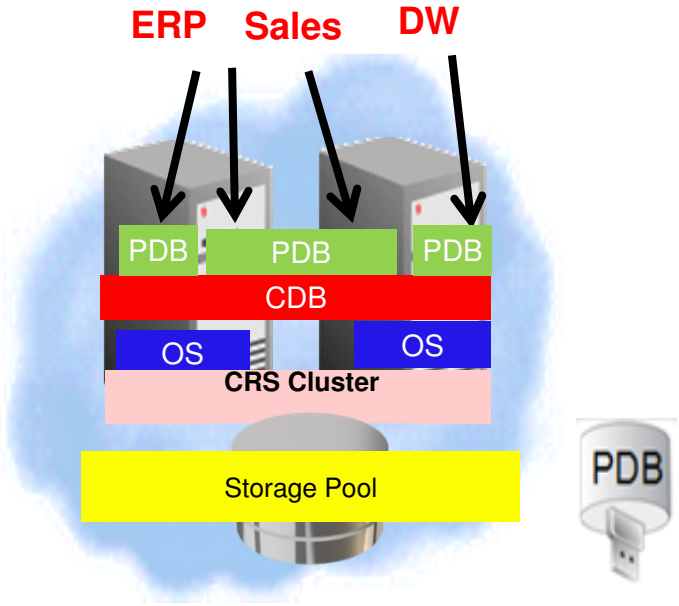

12c Multitenant Pluggable DB

- Share container database
- Integrates with Real Application Clusters

Limitations with Schema Based Cloud Provisioning

- Schema Consolidation of many distinct databases into a single database provides a SaaS Cloud model for databases.
- Some applications such as E-business Suite have many interdependencies, OE, AR, PO -> APPS, trying to put multiple schemas can get complex.
- Security limitations – some applications/users may have access that can effect other applications – i.e. ***Delete Any Table***
- RMAN Point in Time Recovery – can be problematic if one application needs a restore.
- Cloning One application for a test – typically means cloning the whole DB unless Datapump is used.

Cloud Control 12c with Multitenant Databases

- EM 12c (with 12.1.0.3+ Database Plug-in supports Pluggable Databases)
- Can use 12.1.0.2 cloud control with the 12.1.0.3 DB plugin.

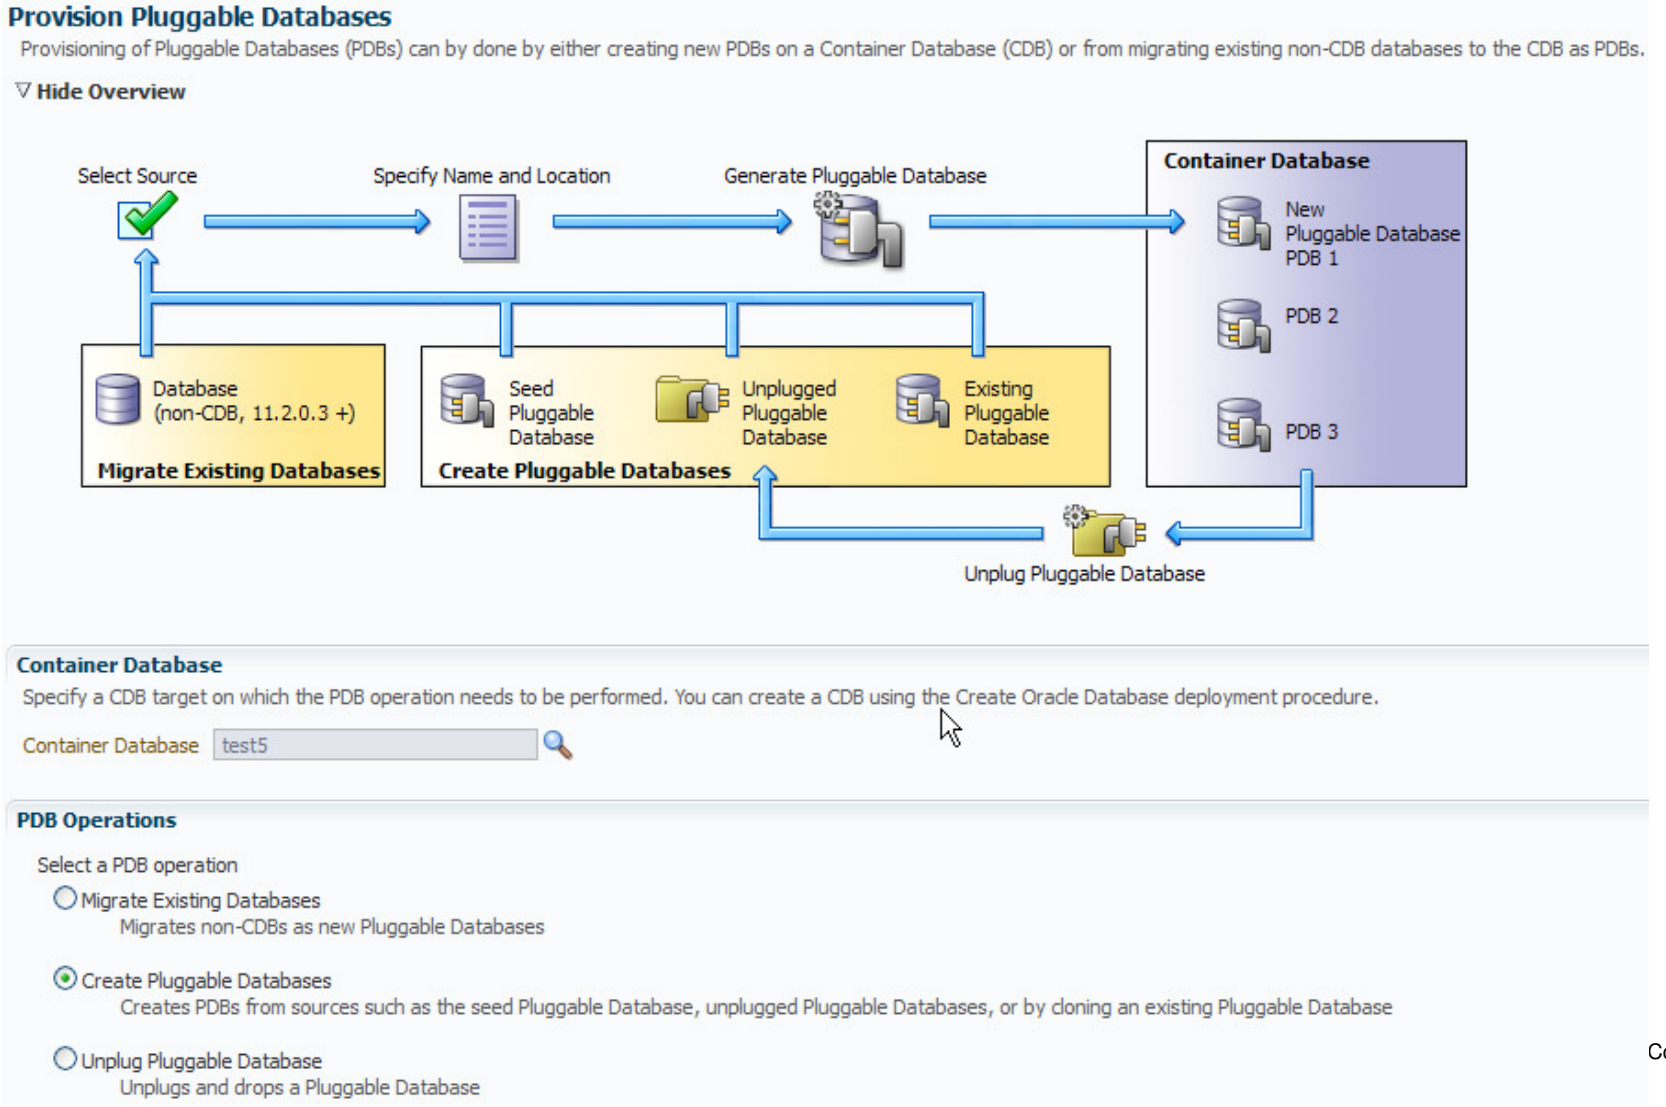

Performance of Multitenant Databases

IBM

- Performance of a pluggable and a non-pluggable database about the same.

DB Name & Test 10 Minutes	Total Trans (swingbench)	User avg utilization %	Sys avg utilization %	I/O wait avg %
test5_1_pdb + test5_2_pdb	3,658,385	88.7	5.2	4.8
test3 + test4 (non pluggable)	3,647,106	89.9	5.4	2.8

Benefits of Multitenant Databases for SAAS Cloud

- Consolidating many databases into an Oracle 12c multitenant architecture provides a new **SaaS** Cloud model option for database consolidation.

- Multitenant provides the following benefits:
  - Previous schema consolidation issues (RMAN, Security) have been addressed.
  - One set of "infrastructure" i.e. Oracle redo logs, UNDO.
  - Cloning a database is quicker and easier.
  - Unplug/Plug database options provides lower downtime for database patching.

- Things to consider for grouping Multitenant databases:
  - Downtime of projects all contained in the CDB i.e. Oracle patching
  - Character sets
  - Memory/cpu and process parameters

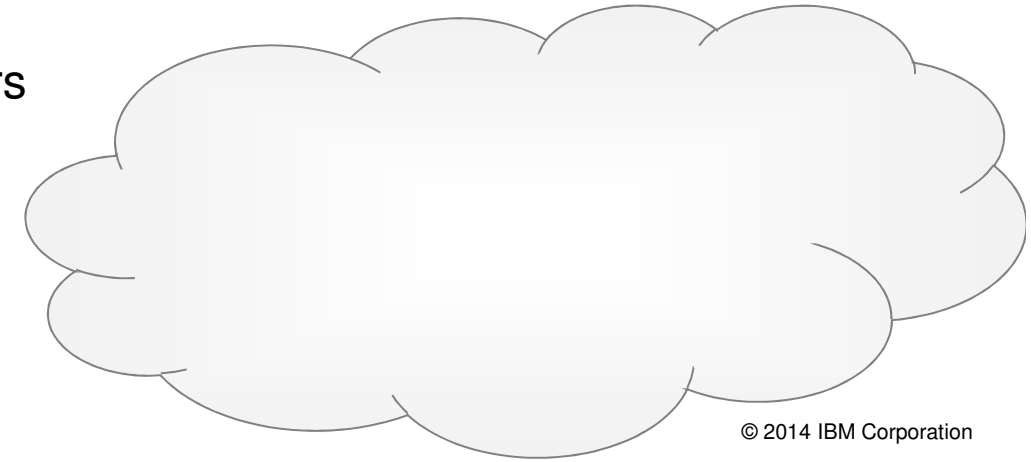

Oracle Testing 11.2.0.4 -> 12.1.0.1 - CPU Intensive Test

18.9% improvement in response time (cpu intensive test)

Oracle 11.2.0.4

Running Parallel Processes: 32
real 0m12.01s
user 0m0.20s
sys 0m0.13s

Running Parallel Processes: 64
real 0m23.84s
user 0m0.40s
sys 0m0.26s

```
procs -----------memory---------- ---swap-- -----io---- --system-- -----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa st
 0  0      0 64919572 202576 1475116    0    0  8070    73    0   28  1  1 96  2  0
 0  0      0 64919476 202576 1475120    0    0     0    19    0 4419  0  0 100  0  0
32  0      0 64659544 202596 1475388    0    0   188   101    0 5914 55  1 44  0  0
32  0      0 64659172 202596 1475404    0    0     0    12    0 4567 100  0  0  0  0
32  0      0 64659172 202612 1475404    0    0     0   151    0 4536 100  0  0  0  0
25  0      0 64713216 202616 1475396    0    0    21    51    0 4618 100  0  0  0  0
64  0      0 64398020 202628 1475868    0    0   171   180    0 6679 93  2  6  0  0
64  0      0 64398020 202628 1475868    0    0     0   100    0 4754 100  0  0  0  0
64  0      0 64398020 202636 1475868    0    0    21   201    0 4757 100  0  0  0  0
64  0      0 64398020 202636 1475868    0    0     0    12    0 4746 100  0  0  0  0
64  0      0 64396484 202648 1475868    0    0     4    37    0 4749 100  0  0  0  0
64  0      0 64396500 202652 1475864    0    0    21    32    0 4769 100  0  0  0  0
64  0      0 64396500 202660 1475868    0    0    21    17    0 4748 100  0  0  0  0
29  0      0 64674340 202664 1475840    0    0     0    19    0 4967 100  0  0  0  0
 0  0      0 64909796 202672 1475680    0    0    21    29    0 4767 34  0 66  0  0
 0  0      0 64910676 202676 1475680    0    0     0    45    0 4571  0  0 100  0  0
```

Oracle 12.1.0.1

Running Parallel Processes: 32
real 0m10.12s
user 0m0.16s
sys 0m0.14s

Running Parallel Processes: 64
real 0m20.05s
user 0m0.34s
sys 0m0.27s

```
procs -----------memory---------- ---swap-- -----io---- --system-- -----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa st
 0  0      0 64820020 202224 1632084    0    0  8090    73    0   27  1  1 96  2  0
 0  0      0 64819800 202224 1632088    0    0    43    12    0 4368  0  0 100  0  0
32  0      0 64571376 202248 1632328    0    0   107   116    0 5899 56  1 43  0  0
32  0      0 64570896 202248 1632364    0    0    43    16    0 4618 100  0  0  0  0
28  0      0 64600612 202272 1632364    0    0    21   156    0 4729 100  0  0  0  0
64  0      0 64319352 202296 1632280    0    0   192   247    0 7806 94  2  5  0  0
64  0      0 64317628 202304 1632816    0    0    43    33    0 4744 100  0  0  0  0
64  0      0 64317212 202312 1632816    0    0    21   204    0 4745 100  0  0  0  0
64  0      0 64317260 202320 1632820    0    0    21    35    0 4705 100  0  0  0  0
64  0      0 64316640 202324 1632820    0    0    43    37    0 4735 100  0  0  0  0
64  0      0 64317012 202332 1632820    0    0    21    29    0 4695 100  0  0  0  0
55  0      0 64395324 202332 1632816    0    0    43    43    0 4864 100  0  0  0  0
 0  0      0 64812836 202340 1632632    0    0    43    29    0 4988 45  0 55  0  0
 0  0      0 64812852 202344 1632636    0    0    21    47    0 4351  0  0 100  0  0
```

11.2.0.4 -> 12.1.0.1 - I/O Test

- Test: With Oracle I/O Calibrate (high I/O)
- Not much change between releases (for this particular I/O test)

Oracle 11.2.0.4

max_iops = 332989
latency = 0
max_mbps = 3109

Oracle 12.1.0.1

max_iops = 333576
latency = 0
max_mbps = 3116

```
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
          12.56    0.00   36.50   41.64    1.92    7.39

Device:         rrqm/s   wrqm/s     r/s     w/s    rsec/s    wsec/s avgrq-sz avgqu-sz   await  svctm  %util
sdz               0.00     0.00 3029.33    0.00 24234.67      0.00     8.00    20.84    6.89   0.32  98.00
sdba              0.00     0.00 3033.33    0.00 24266.67      0.00     8.00    14.70    4.89   0.31  94.00
sdcb              0.00     0.00 2995.00    0.00 23986.67      0.00     8.01    53.64   17.74   0.33  99.67
sdem              0.00     0.00 3033.00    0.00 24264.00      0.00     8.00    23.24    7.68   0.33 100.00
dm-17             0.00     0.00 12113.67    0.00 96909.33     0.00     8.00   113.11    9.31   0.08 100.67
```

Oracle 12c – JIT Compiler Improvements

- New in Oracle 12.1.0.1 – JIT Compiler for Java Stored Procedures versus interpreted.

Oracle 11.2.0.4

alter session set java_jit_enabled=true;
ERROR: ORA-02097: parameter cannot be modified because specified value is invalid

```
var time_compiled NUMBER;
var time_interpreted NUMBER;
exec :time_compiled := factorial(20);

alter session set java_jit_enabled=false;
exec :time_interpreted := factorial(20);
```

INTERP_TIME_MS	JIT_TIME_MS
2893	**2856**

Oracle 12.1.0.1

```
alter session set java_jit_enabled=true;
-- Force compile
select dbms_java.compile_method
('JITDemo', 'factorial', '(J)J') from dual;

var time_compiled NUMBER;
var time_interpreted NUMBER;
exec :time_compiled := factorial(20);

alter session set java_jit_enabled=false;
exec :time_interpreted := factorial(20);
```

INTERP_TIME_MS	JIT_TIME_MS
4148	**182**

Help Desk Reporting System Test 11.2.0.3 -> 12.1.0.1

	Oracle 11.2.0.4	12.1.0.1
1000 Concurrent Reports: Report Time (mm:ss)	43:44.96	37:16.91

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
latch: cache buffers chains	759,643	1,010,447	1330	94.45	Concurrency
DB CPU		18,109		1.69	
cursor: pin S	63,045	3,029	48	0.28	Concurrency
latch free	1,589	884	556	0.08	Other
library cache: mutex X	460	44	95	0.00	Concurrency

- DB with High Concurrency / Hot Data blocks
- **17.3 % Improvement** from 11.2.0.3 -> 12.1.0.1

Help Desk System Test Summary

11.2 (1000 users) - as is	11.2 (tuned query - 1000 user)	Oracle 12c (no other changes)	Ora 12c with 100% large pages
4477.58 (01:14:37.58)	2624.96 (00:43:44.96) + 41% better	2501.87 (00:41:41.87) + 4.7% better	2236.91 (00:37:16.91) + 14.8%

- Embedded Function Call was being called multiple times per row…
  - Modified Index to be more compact – PCTFREE 0 (columns were non update)
  - Consolidated 2 sql checks to 1 to help reduce contention
  - Reduced cpu consumption by 41% with just these 2 changes

- Linux Large Pages recommended for Large databases with SGA > 10GB with many user sessions, for both memory stability and performance (10%)

11.2.0.4 -> 12.1.0.1 Banking Transaction Test

- Team started with **200** banking transactions per second (tps) on 2 IFLs
- **15.3% improvement** with 12c from 9685 to 11676 banking transactions per second, after upgrade to Oracle 12c (no other changes made).
- Reduced "concurrency" observed from 11.2.0.4 to 12.1.0.1

- Target tps was 10,000 tps -> achieved 12739 with 12 IFLs

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Wait Avg(ms)	% DB time	Wait Class
DB CPU		5705.5		33.2	
latch: cache buffers chains	70,118	2437.8	35	14.2	Concurrency
library cache: mutex X	111,292	1879.3	17	10.9	Concurrency
cursor: pin S	70,871	1441.8	20	8.4	Concurrency
buffer busy waits	11,514	192.2	17	1.1	Concurrency
db file sequential read	21,684	183.9	8	1.1	User I/O
latch free	2,394	82	34	.5	Other
db file scattered read	19,562	58.6	3	.3	User I/O
enq: TX - index contention	3,123	48	15	.3	Concurrency
cursor: pin S wait on X	870	5.9	7	.0	Concurrency

11.2.0.4 -> 12.1.0.1 Banking Transaction Test

IBM

- Team started with **200** banking transactions per second (tps) on 2 IFLs
- **15.3% improvement** with 12c from 9685 to 11676 banking transactions per second, after upgrade to Oracle 12c (no other changes made).
- Reduced "concurrency" observed from 11.2.0.4 to 12.1.0.1
- **Target tps was 10,000 tps -> achieved 12739 with 12 IFLs**

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Wait Avg(ms)	% DB time	Wait Class
DB CPU		5705.5		33.2	
latch: cache buffers chains	70,118	2437.8	35	14.2	Concurrency
library cache: mutex X	111,292	1879.3	17	10.9	Concurrency
cursor: pin S	70,871	1441.8	20	8.4	Concurrency
buffer busy waits	11,514	192.2	17	1.1	Concurrency
db file sequential read	21,684	183.9	8	1.1	User I/O
latch free	2,394	82	34	.5	Other
db file scattered read	19,562	58.6	3	.3	User I/O
enq: TX - index contention	3,123	48	15	.3	Concurrency
cursor: pin S wait on X	870	5.9	7	.0	Concurrency

Oracle FIPS 140-2 Compliance & Linux on System z (s390x)

Oracle Advanced Security provides TDE (transparent data encryption) - encrypt data tablespaces for protection against data at rest and encrypt SQL*Net database network traffic.

RSA has completed its FIPS 140-2 validation of RSA BSAFE Crypto-C Micro Edition 4.0.1.

RSA BSAFE® Crypto-C Micro Edition (Hardware Version: Software Version: 4.0.1)
http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp2047.pdf

RSA BSAFE® Crypto-C Micro Edition (Software Version: 4.0.1)
http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp2056.pdf
http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp2097.pdf
*Actual version of this module used by Oracle may be newer in order in case of defects (no new features/functionality will be introduced).

Oracle version 12.1.0.2 . This includes FIPS-140-2 validated libraries as standard – more info -> http://docs.oracle.com/database/121/DBSEG/asoappe.htm#DBSEG1041 .

Patch for 11.2.0.4 is still being tested.

IBM & Oracle are currently conducting tests to measure overhead of enabling Crypto, main overhead is in startup as system checks are made.

Performance Triage & Resolution

Gather Diagnostic Data

Start with *MOS ID 1121043.1* for How-to & Best Practices

- SQL Tuning
  - Trace files
  - SQLT output (MOS ID: 215187.1)
  - Trace Analyzer (MOS ID: 224270.1)
  - AWR Report (MOS ID: 748642.1)
  - 11g SQL Monitor Report
  - AWR SQL Report (awrsqrpt.sql)
- PL/SQL Tuning
  - Product logs
  - PL/SQL Profiler (MOS ID: 808005.1)
- Reports Tracing (MOS ID: 111311.1)

- Database Tuning
  - AWR Report (MOS ID: 748642.1)
  - ADDM report (MOS ID: 250655.1)
  - Active Session History (ASH)
- Forms Tuning
  - Forms Tracing (MOS ID: 373548.1)
  - FRD Log (MOS ID: 445166.1)
  - Generic note (MOS ID: 438652.1)
- Middletier Tuning
  - JVM Logs
  - JVM Sizing/Tuning (MOS ID: 362851.1,278868.1)
- OS - OSWatcher (MOS ID: 301137.1)

Source: **Getting Optimal Performance from Oracle E-Business Suite - Samer Barakat OOW 2013**

References

- White Papers / Presentations:
  - Oracle Database on Linux on System z - Disk I/O Connectivity Study
  - Oracle Real Application Clusters on Linux on IBM System z: Set up and network performance tuning
  - Performance of an Oracle 10g R2 Database Import Environment
  - Using the Linux cpuplugd Daemon to manage CPU and memory resources from z/VM Linux guests
  - Oracle Database Auditing: Performance Guidelines
  - Analyzing BI Oracle Workloads Performance Tuning Results – Real Customer Examples
  - Oracle RAC Networking Alternatives on Linux on System z and Red Hat 6 Oracle DB Support

Redbooks:

sg248159 - 12c Experiences http://www.redbooks.ibm.com/redpieces/abstracts/sg248159.html?Open

sg248104 - 11gR2 Experiences http://www.redbooks.ibm.com/abstracts/sg248104.html?Open

Cookbook for Linux & z/VM:

http://www.redbooks.ibm.com/redpieces/abstracts/sg248147.html?Open

zSeries Oracle SIG Conference Presentations:

http://zseriesoraclesig.org

- **SHARE Conference presentations:**

Information Sources

- http://www.oracle.com/ibm
  - Oracle IBM Partner Relationship
- http://otn.oracle.com
  - Oracle Select “Downloads”
- http://www.vm.ibm.com/perf/tips
  - General z/VM Performance & Tuning Tips, Capacity planning
- https://support.oracle.com
  - Oracle Support Webpage (My Oracle Support)
- http://www-124.ibm.com/developerworks/oss/linux390/index.shtml
  - Lot’s of information on Linux for zSeries, IBM DeveloperWorks
- http://www-128.ibm.com/developerworks/linux/linux390/perf/index.html
  - Hints and Tips for tuning Linux on System z
- http://www.zseriesoraclesig.org
  - Special Interest Group of Oracle users on the mainframe (z/OS and Linux)
- http://www.mail-archive.com/linux-390%40vm.marist.edu/
  - Marist List Server
- http://www.ibm.com/redbooks
  - SG24-7573-00 Using Oracle Solutions on Linux on System z
  - SG24-7634-00 Experiences with Oracle Solutions on Linux for IBM System z
  - REDP-4788-00 Installing Oracle 11gR2 RAC on Linux on System z

Experiences With Oracle on Linux on System z Customer Panel

Questions