




Customer Experiences With Oracle on Linux on System z

Speaker: David Simpson – IBM
simpson.dave@us.ibm.com


Thursday, March 13, 2014: 08:00 AM - 09:00 AM,
MA, Platinum Ballroom Salon 6
Session Number: **14705**

Twitter -> @IBMANDOracle



Agenda

- Customer Cases
- Oracle 12c Testing
- Oracle Diagnostics
- Q & A / Suggestions



- Running Oracle with Linux on System z since 2009.
- Recent upgrade of z10 -> to zBC12 (less cpu usage observed).
- ext3 file system for the Oracle database files were starting to filling up.
- Used Yast to add new devices and expand the file system online, worked work great.
- Applied OS patch update during a downtime window and rebooted as part of the patching and everything came back up fine.
- Several days later noticed some of the LUNs were not online for the file system that was extended.

Case 1 Cont. (2) - Increasing File System Manual Steps...

- `rescan-scsi-bus.sh` `#scan for any new LUNs`
- `ls -l /dev/dm-*` `#verify device has been added`
- `tail -f /var/log/message` `#another check to verify devices were added ok`

- Update `/etc/multipath.conf` file `#if using aliases update file`
- `Trigger udev -> echo 'add' > /sys/block/<new devname>/uevent` `#add udev rules to bring devices online on next system restart.`
- `/etc/init.d/multipathd restart` `#restart if multipath.conf was updated`

- `fdisk /dev/disk/by-id/dm-uuid-mpath-<UUID>` `#if FCP/SCSI partitioning not required, DASD must partition`

- `pvcreate /dev/disk/by-id/dm-uuid-mpath-<UUID>` `#pvcreate command to make available to add`
- `vgextend <vgname> /dev/disk/by-id/dm-uuid-mpath-<UUID>` `#add to volume group`
- `lvextend <lvname> -l 100%FREE` `#resize logical volume`
- `resize2fs /dev/<vgname>/<lvname>` `#resize file system online`

- **Conclusion** – when adding new LUNs verify /var/log/messages and verify that the UDEV rules needed to bring storage devices online have been successfully created.
- SUSE now has new high performance xfs file system capability for database files so stay tuned on updates with xfs.
- Red Hat 6 customers should use ext4 file systems for performance and extending file systems online.

Customer Experience Case 2:



- Running several E-business Suite, Peoplesoft and other Oracle databases in split tier mode for over 6 years.
- Upgraded from z196's to zEC12 and observed less cpu usage.
- Running Oracle 10gR2, 11gR2 and recently have started using Oracle 12c.
- Only a couple of issues have cropped up over this time.

Customer Case 2 Continued - CMM enabled



In `/var/log/messages` observed the following:

```
Jan 31 14:03:37 zlomsp10 kernel: kernel BUG at mm/page-discard.c:187!
```

```
Jan 31 14:03:37 zlomsp10 kernel: illegal operation: 0001 [#1]
```

```
Jan 31 14:03:37 zlomsp10 kernel: CPU: 0 Not tainted
```

```
Jan 31 14:03:37 zlomsp10 kernel: Process oracle (pid: 16449, task: 000000017cf5ac58, ksp: 000000010a44fa30)
```

15 minutes later (900 seconds) Oracle throws a message to the alert log...

ORA-00494: enqueue [CF] held for too long (more than 900 seconds) by 'inst 1, osid 29091

- CMM (Cooperative Memory Management) has been improved in SLES 11 / Red Hat 6 (mm never accepted upstream Linux code)
- **Solution:** Move to SLES 11/Red Hat 6 or remove '**cmma=yes**' from `zipl.conf` in older Linux distributions.

Customer Case 2 Continued - ASLR



- Linux (all distros) has a feature called ASLR (Address Space Layout Randomization) to help prevent buffer overflow attacks.

- System z Oracle team has published Note: **1345364.1**

ORA-00445: background process "m001" did not start after 120 seconds

Incident details in: /opt/u01/app/oracle/diag/rdbms/incident/incdir_3721/db1_mmon_7417_i3721.trc

ERROR: Unable to normalize symbol name for the following short stack (at offset 2):

- Verify whether ASLR is in use:

```
# /sbin/sysctl -a | grep randomize
```

```
kernel.randomize_va_space = 1
```

***If param is any value other than 0 then ASLR in use.

- To fix adjust kernel parameter in /etc/sysctl.conf to turn off

```
kernel.randomize_va_space=0 (SUSE & Red Hat)
```

```
kernel.exec-shield=0 (Red Hat only)
```

OR

Change shmmax kernel parameter to be greater than sga_target as well as patch Oracle to latest the latest levels – Oracle BUG.

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
sdcdb                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
```

- 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  
2893
```

```
JIT_TIME_MS  
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  
4148
```

```
JIT_TIME_MS  
182
```

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

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 5000, 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

- **Snapshots** (take snaps before & after test):
 - exec DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT;
- **AWR Report** (@?/rdbms/admin/awrrpt.sql)
 - Workload Repository Report – Top events to start
- **ADDM report** (@?/rdbms/admin/addmrpt.sql)
 - Oracle provided recommendations for further analysis
- **ASH Report** (use particularly for concurrency issues)
 - @?/rdbms/admin/ashrpt.sql

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)**

Remote Diagnostic Agent (RDA) Report – Note: 314422.1

RDA HTML Menu

- [Overview](#)
- [Operating System Setup](#)
- [User Profile](#)
- [Performance](#)
- [Network](#)
- [Oracle Net](#)
- [Oracle Installation](#)
- [RDBMS](#)
 - [RDBMS Memory](#)
 - [RDBMS Log/Trace Files](#)
 - [Backup and Recovery](#)
 - [SQL*Plus/iSQL*Plus](#)
- [IBM WebSphere \(Offline\)](#)
- [J2EE/OC4J](#)
 - [Generic](#)
 - [J2EE Miscellaneous](#)
- [Oracle JDBC](#)
- [Cluster](#)
 - [Hang Analysis](#)
- [ASM](#)
- [Data Guard](#)
- [Enterprise Manager Server](#)
- [Database Control](#)
- [External Data Collection](#)

List of Diagnostic Problems

Using: SHOW PROBLEM -ALL -ORDERBY LASTINC_TIME DSC

From: /opt/oracle/diag/rdbms/edpsprd/edpsprd

Problem ID	Problem Key	Last Incident	Last Incident Time
4	ORA 4031	516429	2013-01-12 12:33:39.529000 -05:00
6	ORA 445	411813	2013-01-08 20:06:34.734000 -05:00
7	ORA 240	381339	2012-12-19 19:59:01.195000 -05:00
5	ORA 600 [15709]	246899	2012-08-25 05:41:55.184000 -04:00
2	ORA 7445 [kggmd5Process()+26]	13410	2011-12-12 18:16:11.498000 -05:00
3	ORA 600 [SKGMHASH]	13209	2011-12-12 11:39:00.697000 -05:00
1	ORA 7445 [kglgob()+8490]	9169	2011-12-06 12:57:10.293000 -05:00

Summarized Errors

Current CPU Hogs / Top 15 by CPU Time

F	S	UID	PID	PPID	C	PRI	NI	ADDR	SZ	WCHAN	STIME	TTY	TIME	CMD
0	R	oracle	23639	1	65	79	0	-	21093142	stext	13:15	?	04:59:23	ora_j000_edpsprd
0	R	oracle	24814	1	47	78	0	-	21089063	stext	16:13	?	02:12:07	oracleedpsprd (LOCAL=NO)
0	S	oracle	17293	1	7	75	0	-	21088031	sk_wai	Jan14	?	02:02:05	oracleedpsprd (LOCAL=NO)
0	S	oracle	31422	1	8	75	0	-	21088013	sk_wai	Jan14	?	01:45:42	oracleedpsprd (LOCAL=NO)
0	S	oracle	1879	1	3	75	0	-	21090269	sk_wai	Jan13	?	01:42:19	oracleedpsprd (LOCAL=NO)
0	S	oracle	29474	1	3	75	0	-	21092455	semtim	Jan13	?	01:39:25	ora_dbw0_edpsprd
0	S	oracle	29478	1	2	75	0	-	21090149	semtim	Jan13	?	01:26:40	ora_dbw1_edpsprd
0	S	oracle	29482	1	1	75	0	-	21095330	semtim	Jan13	?	00:54:31	ora_lgwr_edpsprd
0	R	oracle	1349	1	54	85	0	-	21097455	stext	20:00	?	00:28:37	oracleedpsprd (LOCAL=NO)
4	S	root	27853	1	0	79	0	-	43180	rt_sig	Jan13	?	00:24:34	/opt/tivoli/tsm/StorageAgent/bin/dsmsta
0	S	oracle	7960	7933	0	75	0	-	230979	futex	Jan13	?	00:19:24	/opt/oracle/product/11.2.0.3/db/jdk/bin/java
0	R	oracle	16863	1	13	75	0	-	21089235	stext	18:43	?	00:17:18	oracleedpsprd (LOCAL=NO)
0	S	oracle	16879	1	13	75	0	-	21089235	sk_wai	18:43	?	00:17:14	oracleedpsprd (LOCAL=NO)
0	S	oracle	16855	1	13	75	0	-	21089235	sk_wai	18:43	?	00:16:59	oracleedpsprd (LOCAL=NO)
0	S	oracle	16897	1	13	75	0	-	21089235	sk_wai	18:43	?	00:16:50	oracleedpsprd (LOCAL=NO)

[Back to top](#)

Root CPU Hogs / Top 5 by CPU Time

F	S	UID	PID	PPID	C	PRI	NI	ADDR	SZ	WCHAN	STIME	TTY	TIME	CMD
4	S	root	27853	1	0	79	0	-	43180	rt_sig	Jan13	?	00:24:34	/opt/tivoli/tsm/StorageAgent/bin/dsmsta
5	S	root	25436	1	0	-40	-	-	34880	futex	Jan13	?	00:05:56	/sbin/multipathd
4	S	root	21726	20943	0	76	0	-	797	select	13:03	pts/2	00:02:34	top
4	S	root	27841	1	0	75	0	-	17482	compat	Jan13	?	00:02:17	/opt/tivoli/tsm/client/ba/bin/dsmc sched
1	S	root	24	1	0	70	-5	-	0	worker	Jan13	?	00:00:29	[events/0]

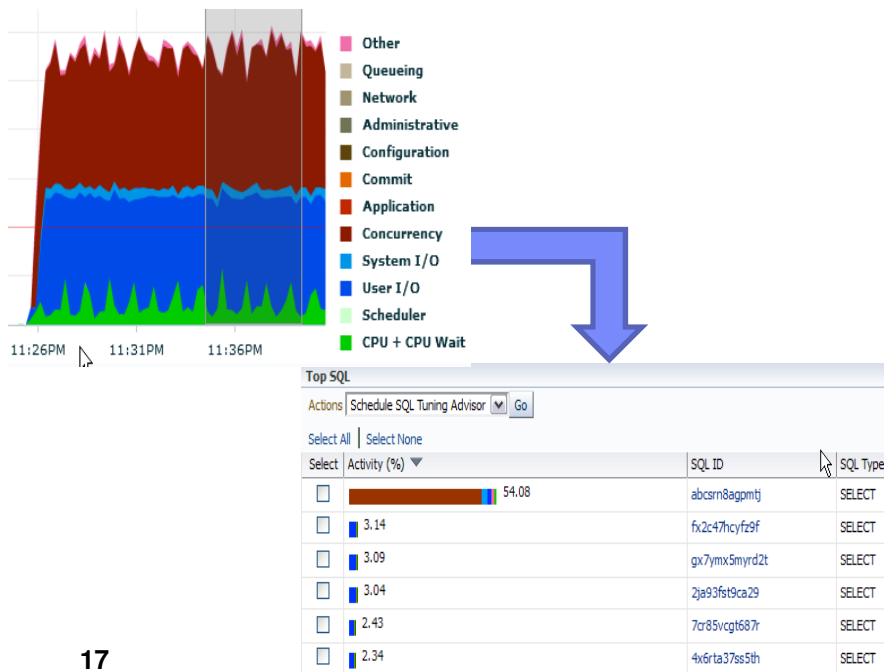
Performance Reports

Oracle Enterprise Manager – Linux on System z

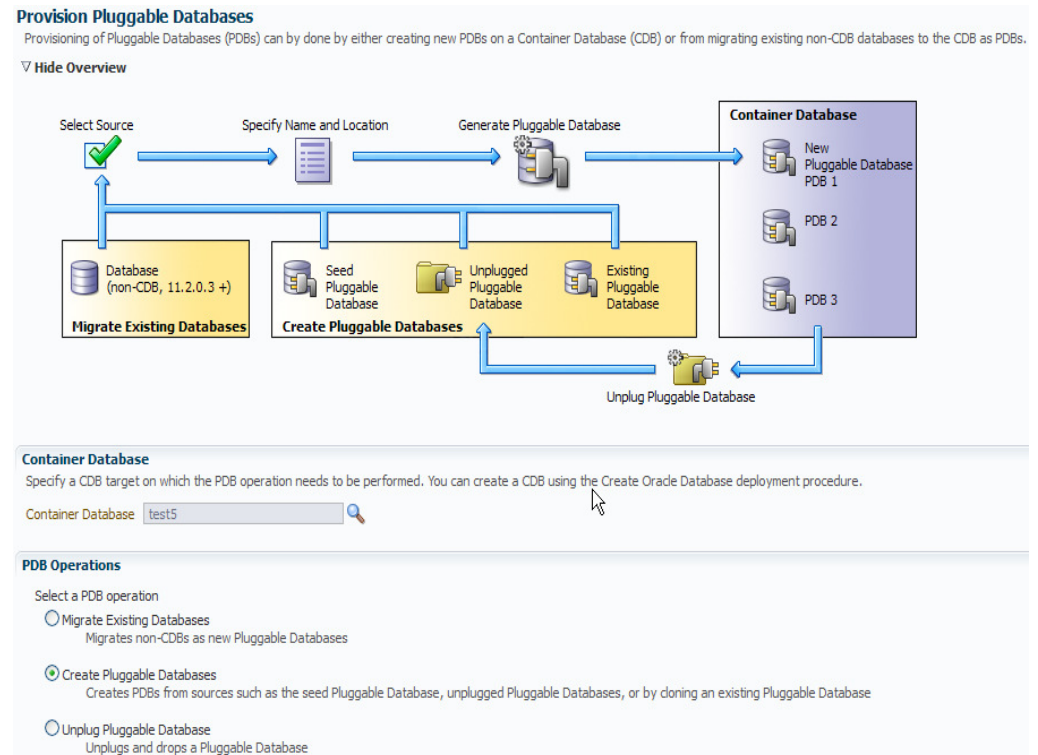


- Oracle performance problems usually come in 2 flavours :
 - Database wide problems (problems which affect all users and sql queries)
 - Query specific problems (problems which affect only one, or few, queries)
- Oracle Enterprise Manager 12c has drill down capabilities system and query issues.
- SQL Tuning advisor provides recommendations for SQL query specific issues.

EM 12c - performance drilldowns:



EM 12c - provisioning workflows:



Linux Monitoring Tools



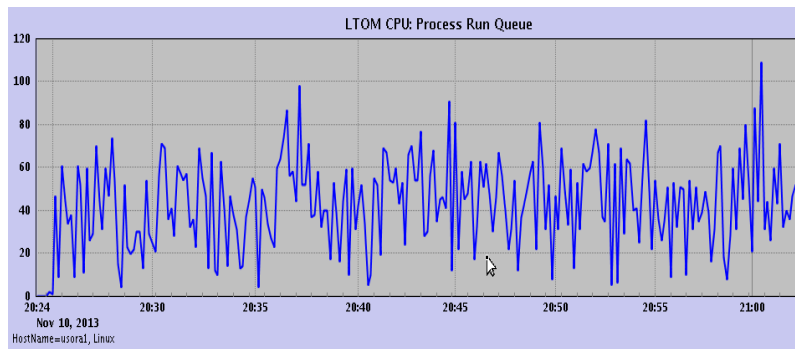
Collect Linux monitoring data for historical analysis

- Standalone performance collection - sysstat rpm (sadc, sar, vmstat)
- More information -> <http://linuxmain.blogspot.com/2011/12/gathering-performance-data-with-sysstat.html>

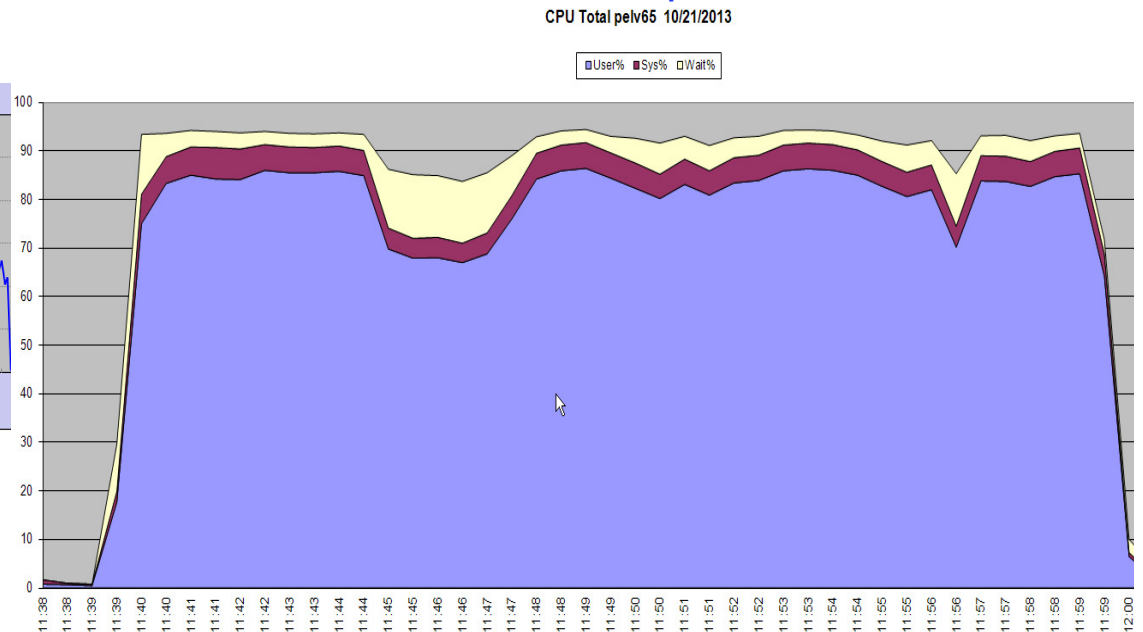
OSWatcher, Linux nmon, and other z/VM tools can be integrated to providing graphing capabilities.

- ✓OSW Document Doc ID 301137.1, OSWg Document Doc ID 461053.1,
- ✓LTOM Document Doc ID - 352363.1


Oracle LTOM Graphs:



Linux nmon report:



Open Microphone (Suggestions / Questions)





Customer Experiences With Oracle on Linux on System z

Speaker: David Simpson – IBM
simpson.dave@us.ibm.com

Thursday, March 13, 2014: 08:00 AM - 09:00 AM,
MA, Platinum Ballroom Salon 6
Session Number: **14705**

Twitter -> @IBMANDOracle

