
ORACLE®

Session id: 31951

Configuring, Tuning, and Using Oracle9i in a Linux Environment

Edward Whalen
Principal Consultant
Bradmark Professional Services

ORACLE

Introduction

- Principal Consultant, Bradmark Professional Services
- Database Performance Consultant
- Author
  - Teach Yourself Oracle9i in 21 Days
    - SAMS – 2003
  - Oracle Performance Tuning
    - Addison Wesley – 2002
  - Teach Yourself Oracle8 in 21 Days
    - SAMS – 1997
  - Oracle Performance Tuning and Optimization
    - SAMS – 1997

Introduction (cont.)

- SQL Server Books
  - 4 from MS Press
- Experience
  - Oracle performance tuning
  - SQL Server performance tuning
  - Oracle RAC configurations (over 12)
    - Mostly Linux
  - Oracle Apps, Peoplesoft, JD Edwards tuning
  - Application tuning

Overview

- A brief introduction to Linux
- Configuring Linux Systems for Oracle
- Oracle RAC on Linux
- Tuning Linux
- Monitoring Linux
- Issues and Gotchas
- Conclusions

Introduction to Linux

- What is Linux?
- Is Linux the same as UNIX?
- Is it a real OS?
- Do software vendors support it?

What is Linux?

- Linux is an operating system with the following attributes
  - Multi-user
  - Multi-tasking
  - Virtual Memory system
  - X-Windows
  - Compatible with the IEEE POSIX.1 standard
  - Non-proprietary source code
  - Source code is available

Is Linux the same as UNIX?

- The command set is similar
- They work the same
- They look the same
- Linux cannot use the UNIX trademark
- Linux is essentially UNIX

Is Linux the same as UNIX?

- It looks like a duck
- It quacks like a duck
- It walks like a duck

- But it’s a Penguin

Is Linux the same as UNIX?

- It looks like a duck
- It walks like a duck
- It quacks like a duck

- But it’s a Penguin

Is Linux a real OS?

- Is Linux a real OS?
- Yes, absolutely
- Linux is used in mission critical applications
- Linux is used around the world
- There are many versions of Linux available

Do Software Vendors Support Linux?

- Software on Linux
  - Oracle
    - Not only do they support it, they run their business on it
  - Oracle Applications (Oracle Financials)
  - Sybase
  - Informix
  - DB/2
  - Netscape
  - Corel

Do Hardware Vendors Support Linux?

- Hardware that Linux runs on
  - Intel i386
    - Dell, Compaq, IBM, HP
  - Intel iA64
  - Sun Sparc
  - PPC
  - Alpha
  - M68000

Linux Flavors (alphabetical)

- Red Hat
- SuSe
- SCO/Caldera
- Corel
- Debian
- Mandrake
- Slackware
- About 100 more

A Brief History of Linux

- 1969 – AT&T Bell Labs - Ken Thompson & Dennis Ritchie started development using a PDP-7
- 1971 – 1st edition is released
- 1973 – 4th edition rewritten in “C”
- 1975 – 6th edition released outside of Bell Labs
- 1980 – Microsoft releases Xenix

A Brief History of Linux (cont.)

- 1982 – AT&T's Unix System Group (USG) releases System III
- 1983 – UNIX System V, the first supported UNIX is released
- 1984 – University of California Berkely releases 4.2BSD
- 1989 – UNIX system V release 4 is released, unifying System V, BSD & Xenix
- 1991 – USL (Unix Systems Lab) is spun off

A Brief History of Linux (cont.)

- 1991 – Linus Torvalds begins work on Linux at the University of Helsinki in Finland and the first version was released
- 1993 – Novell buys USL
- 1993 – Novell transfers rights to UNIX trademark and specification to X/Open
- 1999 – Linux version 2.2 kernel released
- 2001 – Linux version 2.4 kernel released

Why Linux?

- TCO (Total Cost of Ownership)
  - Standards based
  - Multivendor support
  - Ease of use
  - Good performance
- Industry support
  - Oracle
  - Dell, IBM, HP, ….
  - Third party support for tools (including Bradmark)

Operating system concepts

- Linux is based on a modular non-microkernel architecture
  - All device drivers share the same memory with the kernel
  - Device drivers are modular and loadable/unloadable
- Linux is multi-tasking
- Linux is a virtual-memory OS
- Most kernel parameters can be modified on the fly

Overview of the Linux Structure

- Each user logs into their own environment with a username and password
- User programs cannot modify sensitive system information (hence less crashes)
- Each user may be given permission to access directories other than their own

Configuring Linux System for Oracle

- Linux is configured by setting parameters in the /etc/conf directory
- Pseudo files look like files
  - Read from them
    - To view parameter settings
  - Cat data to them
    - To set parameter settings
- Parameter settings can be put in /etc/rc.local

```
echo 2147483649 > /proc/sys/kernel/shmmax
echo 1024 65000 > /proc/sys/net/ipv4/ip_local_port_range
```

Configuring the Linux System for Oracle

- Parameter settings can be put in /etc/sysctl.conf

```
# Disables packet forwarding
net.ipv4.ip_forward = 0
# Enables source route verification
net.ipv4.conf.default.rp_filter = 1
# Enables the magic-sysrq key
kernel.sysrq = 1
kernel.sem=250 32000 100 128
kernel.shmmax=2147483648
kernel.shmmni=4096
fs.file-max=409200
```

Configuring the Linux System for Oracle

- Do you need to relink the Linux Kernel for Oracle?
  - No, use /proc
- Do you ever need to relink the Linux Kernel?
  - Yes, under certain conditions
    - Addition of new device drivers
    - Changing some parameters
      - SCSI timeout
    - Creating a non-modular kernel

Oracle RAC on Linux

- Linux is an excellent platform for Oracle RAC
- Shared Disk Subsystem is required
- High performance interconnect is required
- Many vendors offer Linux RAC solutions
  - Dell
  - IBM
  - HPQ

Tuning Oracle on Linux

- Choose the right kernel

Kernel	# CPUs	RAM
Standard	1	4 GB
SMP	Multi	4 GB
Enterprise	Multi	> 4 GB

Tuning Oracle on Linux

- Proper sizing is the key to good performance
  - CPU resources
  - Memory
  - I/O Subsystem
    - Use a sufficient number of disk drives
      - No more than 125 IOPS per disk
    - Use high performance disks
    - Done overtax the I/O subsystem

Tuning the Linux Kernel

- Done via /proc directory structure
  - /proc is a virtual file system
  - Provides an instantaneous view of the operation of the system
    - /proc/meminfo, /proc/mounts, /proc/partitions
    - Can be viewed with cat, more, less
  - Can be used to configure kernel parameters
    - /proc/sys
    - Can be set with the echo command

Monitoring Linux

- Manual
  - You can look at the state of the system via files in the /proc virtual file system

```
[oracle@localhost 2500w]$ cat /proc/meminfo
        total:    used:    free:  shared: buffers:  cached:
Mem:  4022009856 4016812032  5197824 220852224 57868288
   3519877120
Swap: 4342824960        0 4342824960
MemTotal:      3927744 kB
MemFree:          5076 kB
..
..
LowTotal:       847616 kB
LowFree:          3048 kB
SwapTotal:     4241040 kB
SwapFree:      4241040 kB
BigPagesFree:        0 kB
```

Monitoring Linux

- Tools
  - Top
  - Memstat
  - Vmstat
  - Iostat
  - Free
  - Third Party (such as NORAD Surveillance)

Monitoring Linux with Top

```
 11:41am  up 32 min,  2 users,  load average: 1.40, 0.98, 0.59
89 processes: 87 sleeping, 2 running, 0 zombie, 0 stopped
CPU0 states: 13.1% user, 33.3% system,  0.0% nice, 53.0% idle
CPU1 states:  0.0% user,  6.4% system,  0.0% nice, 93.0% idle
CPU2 states:  0.0% user,  0.4% system,  0.0% nice, 99.1% idle
CPU3 states:  0.1% user,  1.1% system,  0.0% nice, 98.2% idle
Mem:  3927744K av, 3922648K used,    5096K free,  216432K shrd,   59680K buff
Swap: 4241040K av,       0K used, 4241040K free                3453976K cached

  PID USER     PRI  NI  SIZE  RSS SHARE STAT %CPU %MEM   TIME COMMAND
 1515 oracle    17   0 23344  22M 17404 R    46.7  0.5   1:22 oracle
   12 root      15   0     0    0     0 SW    5.7  0.0   0:36 bdflush
 1522 oracle    15   0  1072 1072   840 R     0.5  0.0   0:00 top
    1 root      15   0   512  512   444 S     0.0  0.0   0:04 init
    2 root      15   0     0    0     0 SW    0.0  0.0   0:00 keventd
    3 root      15   0     0    0     0 SW    0.0  0.0   0:00 keventd
    4 root      15   0     0    0     0 SW    0.0  0.0   0:00 keventd
    5 root      15   0     0    0     0 SW    0.0  0.0   0:00 keventd
    6 root      34  19     0    0     0 SWN   0.0  0.0   0:00 ksoftirqd_CPU0
    7 root      34  19     0    0     0 SWN   0.0  0.0   0:00 ksoftirqd_CPU1
    8 root      34  19     0    0     0 SWN   0.0  0.0   0:00 ksoftirqd_CPU2
    9 root      34  19     0    0     0 SWN   0.0  0.0   0:00 ksoftirqd_CPU3
   10 root      15   0     0    0     0 SW    0.0  0.0   0:22 kswapd
   11 root      16   0     0    0     0 SW    0.0  0.0   0:00 kreclaimd
   13 root      15   0     0    0     0 SW    0.0  0.0   0:00 kupdated
```

Monitoring Linux with vmstat

```
[root@localhost root]# vmstat 1 5
   procs                  memory    swap        io     system         cpu
 r  b  w   swpd   free   buff  cache  si  so    bi    bo   in    cs  us  sy  id
 0  1  0      0   4424  74348 3413052   0   0    18  2377  112   166   7   4  89
 0  1  0      0   4428  74480 3412780   0   0     4 129280 1084  1168   3  15  81
 0  1  0      0   4428  74608 3412624   0   0     4 129024 1073  1122   3  14  83
 1  0  0      0   4428  74724 3412512   0   0     0 111904 1077  1021   6  10  85
 1  0  0      0   4444  74808 3412364   0   0     4 81156 1071   749   1  11  88
```

Monitoring Linux with iostat

```
[root@localhost root]# iostat
Linux 2.4.9-e.3custom (localhost.localdomain)   11/07/2002

avg-cpu:  %user   %nice    %sys   %idle
          6.75    0.00    4.76   88.49

Device:            tps   Blk_read/s   Blk_wrtn/s   Blk_read   Blk_wrtn
dev8-0           12.98       135.95        42.39     295014      91984
dev8-1           22.11         1.52       755.65       3304    1639776
dev8-2          119.74         0.85     21274.60       1840   46166088
dev8-3           94.09         0.80     19332.26       1726   41951192
dev8-4           67.53         0.64     13839.33       1384   30031480
dev8-5           48.48         0.51      9884.14       1096   21448672
```

Issues and Gotchas

- Sufficient Temp Space
  - /tmp must be at least 400M
  - This might be part of /
  - Configure sufficiently large or the install will fail
- Sufficient Swap Space
  - Make swap at least 1GB
- Development Environment
  - Oracle will install, but will fail to link without the Linux development system

Conclusions

- Linux is a stable and robust platform
- Many vendors are behind Oracle on Linux
- It is easy to work with
- Performs well
- RAC on Linux is great
  - We have implemented several into production
  - RAC works well and is stable on Linux


QUESTIONS
ANSWERS

More Questions?

Come by our booth

Send an email

ewhalen@bradmark.com

Call us

(800) 275-2723

Reminder –
please complete the OracleWorld session survey

Thank you.

Next Steps….

- Relevant web sites to visit for more information
  - www.bradmark.com
  - www.redhat.com
    www.dell.com


ORACLE®