



EURO RSCG CIRCLE



Migrating From Unix to Linux – An Oracle's Journey

By Yomi Odedeyi, O.C.P.
yomi.odedeyi@eurorscg.com
January 2003

Roadmap

- Part 1: Justifications
- Part 2: Preparation
- Part 3: Making the Move
- Part 4: Lessons Learned



DBA Studio



EURO RSCG CIRCLE



Housekeeping Issues

- I am no authority, just a humble dba speaking from my own personal experience. I reserve the right to be wrong.
- This is an interactive session — feel free to stop me and ask questions at any time. I reserve the right not to know the answer.
- There will be no Euro RSCG freebies afterwards. This is all you're gonna get!!

EURO RSCG CIRCLE



Part 1

Justifications

EURO RSCG CIRCLE



What Does Circle Do?

- A Subsidiary of EURO RSCG Worldwide — advertising services, corporate communications and marketing services conglomerate.
- EURO RSCG Circle's offerings are e-CRM consulting, creative development, digital marketing services and technology integration
- Some Major customers:
 - IBM
 - Intel
 - MSNBC
 - Schlumberger
 - Bermuda Dept of Tourism
 - Symantec
 - Heinz

EURO RSCG CIRCLE



Why Did We Migrate?

- Client acceptance — Our “Solaris only” solution was becoming increasingly difficult to sell to new customers
- Future compatibility
 - Circle is committed to Oracle, and Oracle is moving internally to Linux
- Lower on-going costs
- Reuse existing Intel hardware
- Transfer Existing Skillsets (or not) 😊

EURO RSCG CIRCLE



Typical Solaris Hardware



- ***Sun Enterprise 3500 Server***
- An 8-way multiprocessing server with dynamic reconfiguration capabilities
- Base Price ~ \$140,000 — this gets you:
 - Operating System, Four UltraSparc 900Mhz CPUs, 8GB RAM, 2 x 18GB SCSI Drives with Controller Hardware, No database storage.
 - Add another \$15K to \$25K for RAID database storage depending how much storage you need.
- Don't forget tape drives, uninterruptible power supply, management software.

EURO RSCG CIRCLE



Typical Linux Hardware



- ***Dell PowerEdge 6600 Enterprise Server***

- An 4-way multiprocessing server with dynamic reconfiguration capabilities
- Base Price ~ \$20,000 — this gets you:
 - RedHat Linux Operating System, Four Intel Xeon 1.5Ghz CPUs, 8GB RAM, 2 x 18GB SCSI Drives with Controller Hardware, No database storage.
 - Add another \$15K to \$25K for RAID database storage depending how much storage you need.
- Don't forget tape drives, uninterruptible power supply, management software.



Oracle Software Cost

Oracle 9i Enterprise Edition

- Perpetual License Base price ~ \$40K per processor
- Plus support cost ~ \$6K per processor per annum

Total for a 4 cpu server is around \$150,000 less applicable discounts depending on whether you're buying direct or through a reseller, etc.

EURO RSCG CIRCLE



Cost Difference is Dramatic

- Note: That was NOT an apples-to-apples comparison. But: *almost* any way you do the math, one solution is much cheaper than the other.



Cost is NOT the full story

■ Solaris

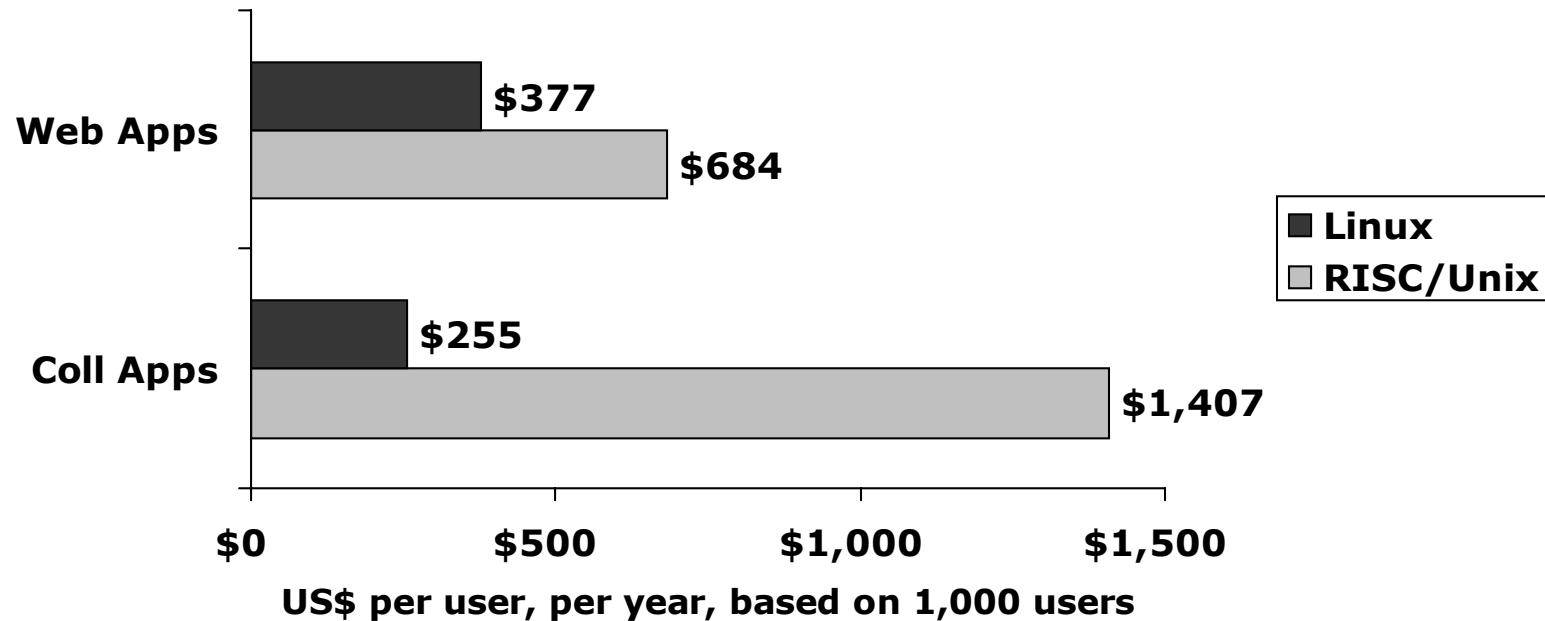
- Offers unmatched number-crunching ability
- Economies of scale for really large applications
- Downside: single point of failure

■ Linux Intel

- Single 4-way server does not match the number-crunching capacity of the 8-way Sun box
- If you were to cluster enough Intel boxes together to match the Sun server, the Cluster software license would almost eat up the hardware cost difference.



However Linux is Proven to have Lower TCO



- “Linux saves enterprise customers 45% to 80% in TCO over RISC/Unix”
- Source: IDC 2001

EURO RSCG CIRCLE



Part 2

Preparation

EURO RSCG CIRCLE



Size Up Your Databases

- Size up Oracle databases to be migrated
 - take a tape and measure:
 - (a) dimensions of the server room
 - (b) dimensions of the server boxesthen (a) minus (b) = Oracle database size
- Use a select such as this instead:

```
select a.tablespace_name name, b.tablespace_name dummy,
sum(b.bytes)/count( distinct a.file_id||'.'||a.block_id ) bytes,
sum(b.bytes)/count( distinct a.file_id||'.'||a.block_id ) -
sum(a.bytes)/count( distinct b.file_id ) used,
sum(a.bytes)/count( distinct b.file_id ) free,
100 * ( (sum(b.bytes)/count( distinct a.file_id||'.'||a.block_id )) -
(sum(a.bytes)/count( distinct b.file_id )) ) /
(sum(b.bytes)/count( distinct a.file_id||'.'||a.block_id )) pct_used,
to_char(sysdate, 'fmMonth dd, yyyy hh24:mi:ss') as Today
from sys.dba_free_space a, sys.dba_data_files b
where a.tablespace_name = b.tablespace_name
group by a.tablespace_name, b.tablespace_name;
```



Script to Check Database Size

- You can use the online version of the select at <http://www.conspectech.ca/tips/archives/000008.html> which gives you a nice looking report like this:

Tablespace Name	Current Total Bytes	Auto-Extend Max Bytes	Current Used Bytes	Current Free Bytes	% Used Current	% Used Max Bytes
DATA_01	1,468,006,400	2,097,152,000	751,828,992	716,177,408	51.21	35.85
DRSYS	88,080,384	0	8,192	88,072,192	0.01	#####
INDEX_01	1,468,006,400	2,097,152,000	1,372,192,768	95,813,632	93.47	65.43
LOB_01	157,286,400	0	8,192	157,278,208	0.01	#####
RBS	541,065,216	0	121,643,008	419,422,208	22.48	#####
SYSTEM	501,219,328	0	271,720,448	229,498,880	54.21	#####
TEMP	262,144,000	0	152,838,144	109,305,856	58.30	#####
sum	4,485,808,128	4,194,304,000	2,670,239,744	1,815,568,384		



More on Sizing

- In sizing your database, don't forget to account for space used by:
 - ORACLE_HOME directory
 - Online redo logs
 - Control files (these can be large if RMAN info is stored inside)
 - Init and parameter files
 - archive logs kept on disk
 - online backups kept on disk
 - udump, bdump, cdump and trace files



Reorg Your Databases

- This is a good time to reorganize the database:
 - Delete duplicate or inaccurate data
 - Archive redundant data (e.g. transaction records > 1 year)
 - Consider table partitioning
 - Drop unused indexes, etc
- Agree with business owners and users on rules to govern deletion/archiving.
 - Watch out for “oh s**t” factor
 - Have good backups and be ready to rollback to status quo ante



Last but not least

- Verify your server config
`/usr/platform/sun4u/sbin/prtdiag`
 - Verify server usage stats
 - Check OS and Oracle error logs
 - Check and list your cron jobs
 - Develop a project timeline and identify who will play each role
 - Communicate, communicate, communicate
- Use all of above information as inputs in determining your Linux hardware and storage needs



Part 3

Making the Move

EURO RSCG CIRCLE



KISS

- Start with a KISS.
 - Do a dry run on a few test tables
 - Don't try to bite off too much at once
- Stay on same Oracle release/version, consider upgrading only after successfully changing OS
- Verify availability/certification of your Oracle version on your particular flavor of Linux



Certification Matrix

Certification Matrix: Oracle Server - Enterprise Edition on Intel Based Server LINUX

Server Certifications

Linux OS	Oracle Version	Status	Install Issue
SuSE SLES7	9.2 (9i)	Certified	None
SuSE SLES7	9.0.1 (9i)	Certified	None
SuSE 7.2	9.0.1 (9i)	Certified	None
SuSE 7.1	9.0.1 (9i)	Certified	None
Red Hat 7.1	9.0.1 (9i)	Certified	None
Intel Caldera Open Unix 8 with LKP 8.0	9.0.1 (9i)	Certified	None
SuSE SLES7	8.1.7 (8i)	Certified	<u>Yes</u>
SuSE 7.2	8.1.7 (8i)	Certified	<u>Yes</u>
SuSE 7.1	8.1.7 (8i)	Certified	None
SuSE 7.0	8.1.7 (8i)	Certified	None
Red Hat 7.1	8.1.7 (8i)	Certified	None
Red Hat 7.0	8.1.7 (8i)	Certified	<u>Yes</u>
Red Hat 6.2 EE	8.1.7 (8i)	Certified	None
Red Hat 6.2	8.1.7 (8i)	Certified	None
Miracle Linux 2.0	8.1.7 (8i)	Certified	None
Miracle Linux 1.0	8.1.7 (8i)	Certified	None
Intel Caldera Open UNIX 8 with LKP 8.0	8.1.7 (8i)	Certified	None

Copyright (c) 1995,2000 Oracle Corporation. All Rights Reserved.

EURO RSCG CIRCLE



The Nitty Gritty 1

- Reverse-engineer source database — see online version of the script for doing this at http://www.conspectech.com/pro_tips/archives/000363.php
- Note where your sequence numbers are at: select sequence_name, last_number from user_sequences;
- Export data from source tables — use one file for each table instead of all tables in a single export file.
- Some suggested export parameters
 - DIRECT=TRUE
 - COMPRESS=YES
 - BUFFER=52428800



The Nitty Gritty 2

- Prepare target Linux box with same mount points as source server /u01, /opt, /raid, etc — more on next slide
- Install Oracle software on Linux box
 - For Oracle 8i only, see Metalink Note ID 148733.1 Oracle Installer will not start on Redhat 7.1 due to glibc issue. Download and install the workaround for this from Metalink.
- Adjust OS params as needed
 - LD_ASSUME_KERNEL=2.2.5
 - Ensure the \$ORACLE_HOME is not in your /etc/ld.so.conf
- Create new instance with Locally Managed Tablespaces
- Install client software on web and app servers if applicable



OLTP Mount Points

- A proper production OLTP database requires the following mount points:
 - 1. Oracle software installation and dba home directory – no raid
 - 2. online redo logs and controlfiles – multiplexed, no raid
 - 3. data tablespaces – raid 5 (unless cost is no object)
 - 4. index tablespaces – raid 5 (unless cost is no object)
 - 5. archive logs and exports – no raid
 - 6. temporary datafiles – no raid, if disk fails simply recreate
 - 7. rollback tablespace – raid 0+1 this is a critical volume that Oracle cannot multiplex so we let the OS do it
- To balance cost and performance, use lots of smaller mirrored disks. The only way you could lose data is if you have simultaneous multiple disk failures. The current state of disk technology makes this possibility very remote.
- See the following doc for detailed explanation of this methodology
http://technet.oracle.com/deploy/availability/pdf/oow2000_sane.pdf

EURO RSCG CIRCLE



The Nitty Gritty 3

- Edit reverse-engineer script to remove old storage clauses; add table partitioning clauses if needed
- Pre-create users on new Linux instance
- Run reverse-engineer script to create objects on new Linux instance
- Disable all constraints
- Import data into new Linux instance with "IGNORE=YES"
- Enable your constraints
- Restart sequences at appropriate numbers



The Nitty Gritty 4

- Recreate your cron jobs
- FTP your scripts and admin programs to new box
- Test application connectivity
- Test data integrity
- Pat yourself on the back!



Part 4

Lessons Learned

EURO RSCG CIRCLE



Gothcas

- As we used to say in my Tech Support days “RTFM”
- Linux != Solaris — there are many annoying little differences in things
 - You will need to install an ftp server
 - Setup authorized_keys file for ssh access
 - Install your fav utils — gzip, gpg
- Exceed works differently on Solaris than on Linux.
 - Linux supports color terminals, so the screen highlights in an irritating way on B/W terminals



Research, research, research

- Have access to Metalink
- Refer to <http://support.redhat.com>
- Refer to the ASK TOM site:
<http://asktom.oracle.com/pls/ask/f?p=4950:1>
- Relax. Others have done it successfully!!





EURO RSCG CIRCLE



That's a wrap!

Any other questions?

Additional Resources

- *Presentation on Developing and Deploying on Linux*
http://www.oracle.com/features/9i/index.html?t1db_sp_endlessdomore.html
- RISC versus Linux
http://www.redhat.com/pdf/seminars/RISC-UNIX_to_RHL.pdf
- Feel free to consult my personal knowledge base at http://www.conspectech.com/pro_tips which I am in the process of migrating to <http://www.conspectech.ca/tips>.
- Red Hat Linux Administrator's Handbook – Mohammed J. Kabir. ISBN: 0764547976

