# Experiences Implementing Oracle Solutions in a Linux on IBM System z Environment

## SHARE Session #9205

Denver, CO August 27, 2009

**Gaylan Braselton, IBM**
gbrasel@us.ibm.com

**Dave Jones, V/Soft**
dave@vsoft-software.com

# User Experiences with Oracle on zLinux

**Dave Jones**
- *V/Soft Software*
- *Houston, TX*
- *dave@vsoft-software.com*

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**

AIX*
DB2*
HiperSockets
IBM*
IBM logo*
Lotus*
System p
System z
System z9
System z10
WebSphere*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

**The following are trademarks or registered trademarks of other companies.**

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows, Windows NT and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. More information on Oracle trademarks can be found at www.oracle.com/html/copyright.html.

* All other products may be trademarks or registered trademarks of their respective companies.

**Notes**:

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

# Agenda

- Client A: Large Canadian provincial government
  - Environment
  - System architecture
  - Results
- Client B: Transzap
  - Environment
  - System architecture
  - Results
- Best practices
- Conclusions

# Large Canadian provincial government

Provide IT services for 125+ provincial government agencies

Many on-line service visible to the public

- Unemployment benefits
- Heath care
- Educational benefits

# Large Canadian provincial government

Initial needs :

- Must solve many issues with the intermediate platform
- Many operation systems
- Many versions
- Unsupported software
- Unsatisfactory DR
- Fast growing (unprecedented growth)
- Understaffed
- Need a flexible solution with rapid deployment

Obvious answer:

- Use z/VM and Linux as a virtual guest

# Large Canadian provincial government

## Environment

Existing mainframe site

1 z9 EC for Linux on z/VM

5 z890s, 2 z800s, 1 G5 for z/OS workload

450+ physical servers (750+ logical) (HP, SUN, pSeries, …)

# Large Canadian provincial government

System architecture, hardware:

- 1 z9 EC mainframe with 5 IFLs (~ 2750 mips)
- 8 LPARs
  - Oracle/DB
  - WAS (2)
  - Open Source
  - Domino (2)
  - Service Zone
  - Lab Zone
- 40+ internal networks

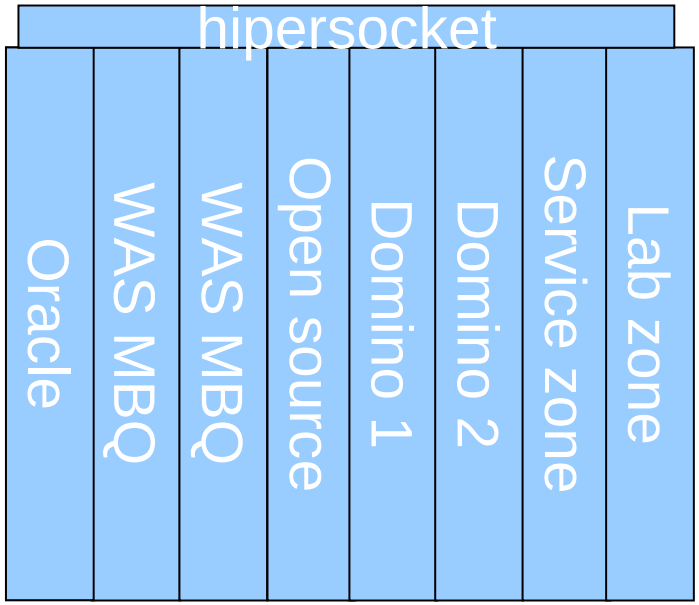

# Large Canadian provincial government

System architecture, software:

- Software
- Novell SLES (versions 9 & 10)
- z/VM v. 5.3 RSU 802
- Oracle/DB (versions 10g & 10gR1)
- Velocity Software Performance Tools
- CA products (Automation, Scheduler)

# Large Canadian provincial government

Networks in a box:

OSA devices

- Traditional connectivity from mainframe to physical switches

HiperSockets

- Inter and Intra LPAR connectivity

Guest LANs

- Connect virtual machines on virtual networks within an LPAR

VSWITCHes

- Connect guest LANs to physical switches using real OSA devices

40+ VSWITCHes on 8 LPARs

# Large Canadian provincial government

Resource Sharing

Some Linux file systems (/boot, /etc, /bin, /opt, others) are shared r/o between Linux guests on same z/VM host.

Oracle database software (executables) are shared r/o as well

- Insures that all Linux guests are indeed running at exactly the same release and patch level

Custom cloning tool developed for client.

All cloning done from a master, hardened Linux guest in the security zone

# Large Canadian provincial government

Resource sharing, hipersockets:

One hipersocket connects all 8 LPARs
Internal network only.
Used for administrative purposes.
Applications include the cloner, telnet, RSCS (file transfer and message queues).
Secure memory-to-memory transfer.

# Large Canadian provincial government

Results:

Over 200 Oracle databases up and running
On 150+ Linux guests
80 WAS servers running
Confirmed our expectation that 2 Linux administrators can support all virtual Linux servers.
*100:1 ratio of Linux virtual machines to administrator*
2 z/VM systems programmers supporting 8 LPARs: (could support many more)
- z/VM support will be integrated into z/OS group

# Large Canadian provincial government

Results:

Overall, the cost of the software and hardware is reduced by 30%.

- Every extra instance will help to reduce the cost.

Oracle Project

- Cloned 53 Oracle Databases: Took 4 days with a team of 4 people

Original break-even time: 36 months
Now, under 24 months!

# Transzap

Provides Software as a Service (SaaS) to the oil/gas industry
Simplifies financial and operational processes by replacing paper-based tasks with digital data and workflow.
Provides one of the energy industry's largest Internet-based data exchanges:

- 4200 companies
- 44,000 users

# Transzap

Original systems ran on a distributed infrastructure,
Linear increases in hardware and software costs as business grew
Technical configurations of hardware grew more complex
Difficult to manage effectively
Unplanned outages and time to repair a problem

# Transzap

Environment

Numerous rack mounted Intel-based servers
Windows and Sun Solaris
Oracle 10g
In-house developed applications
Less than 100 employees

No history of mainframe usage, no mainframe “culture”

# Transzap

Environment

Z9 BC

- 4 IFL engines
- 15104 MB total storage
  - 13696 MB main,
  - 1408MB xstor
- 2 LPARs, one production, one test

DS6800 Enterprise Storage 18 TB

TS3400 Tape Library

# Transzap

Environment

z/VM 5.4 at RSU 0802
SLES10 SP1 Linux
Oracle 10g
1.2TB database size

- 6 Oracle servers
- Others for test, DR, development, etc.

Beginning WebSphere deployment

# Transzap

Results

Unplanned outages reduced in number and duration.
New Oracle guests can be created to meet growing business needs in under 2 hours via FLASHCOPY on DS6800
Oracle developers use same tools and methodologies as before
Oracle backups done by RMAN to a Linux virtual machine NFS server, and then that server is backed up by Veritas to an Intel server

# Best Practices

Networking:

Performance data collection using private VSWITCHes

A TCPIP stack with multiple guest LANs and VLANs collects data for the Velocity SNMP data collection.

Manage multiple networks from a single TCPMAINT

The VSWITCHes are defined without real devices.

Production VLANs

Membership in the VSWITCH and VLAN is RACF protected.

# Best Practices

Systems:

Golden images (z/VM & Linux)

- Install and test once, clone many times

Cloning engine

- Buy or write one

Shared resources

- z/VM is all about sharing of resources

# Best Practices

Systems:

Where possible, employ FLASHCOPY or vendor equivalent technology

Use emulated 3390 (ECKD) devices for Linux file system and Oracle code

- Much easier to manager with tools like DIRMAINT
- Allows FLASCOPY technology to work

Use FCP (SCSI) disks to hold Oracle databases

- Better performance than LVM of a number of (smaller) 3390 DASD devices

# Conclusions

- Oracle database on zLinux is an excellent solution
  - Runs well in a zLinux virtualization environment
  - Oracle support for this platform is good
- z/VM is all about sharing resources
  - Clone guests from tested, hardened master Linux images

## User Experiences

# Questions?

# Contact information

**Dave Jones**
**V/Soft Software**
**Houston, TX**

**dave@vsoft-software.com**
**281.578.7544**