

Oracle VM 3.3

Planning and Implementing



Agenda

1. Introduction
2. What are we doing
3. Why did we want to virtualize
4. Building the case for virtualization
5. Hardware sizing for Oracle VM
6. Oracle VM Architecture layout
7. Oracle VM migration plan
8. Findings and experiences
9. Future plans
10. Q & A

Simo Vilmunen

Technical Architect and DBA



Oracle eBS **10+** years

Oracle Databases **15+** years

Implementations

Upgrades

Performance

Uponor Oyj

“Partnering with professionals to create better plumbing, indoor climate and infrastructure solutions”



Uponor Oracle Solutions

Oracle e-Business Suite

12.1.3
11.2.0.4

Implemented
in 2006

Used in 19
countries in
Europe

Distributed
configuration
w/ Demantra

1200 users



Business
criticality

Oracle Business Intelligence

11.1.1.7
11.2.0.3

Implemented
in 2009

Used in 19 countries in Europe

500 users



Business
criticality

Oracle HFM

11.1.2.3
11.2.0.4

Implemented
in 2010

Used across Uponor worldwide

130 users



Business
criticality

Oracle Enterprise Linux 5


Oracle Enterprise Linux 6

Oracle Enterprise Manager 12c


Why did want to virtualize?

Hardware starting point 2014


Physical servers




Manageable devices



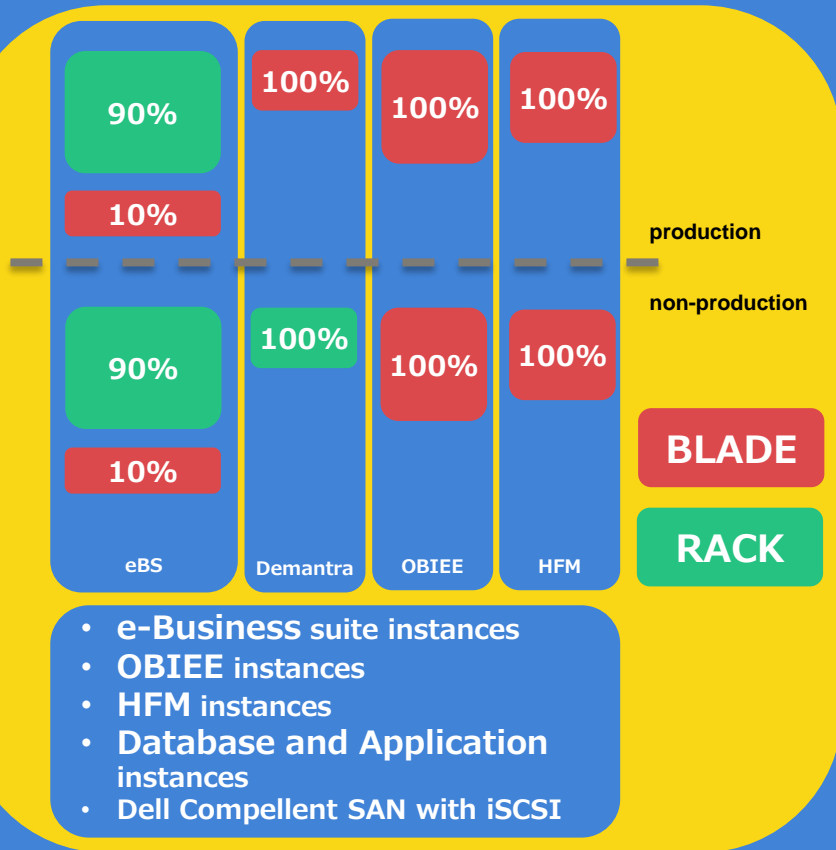
Racks



4 persons



Architecture starting point 2014



Migrated from hosted eBS solution to co-location 2012 with new hardware

Migrated OBIEE to co-location 2012 with new hardware

Migrated HFM to co-location 2014 with new hardware

Dell R410 R510 R610 R810 M620
Intel E5 series v1 and v2 processors

Improvement areas?

Resource allocation

Manageability

Cost

Improved future proof architecture

Environment lifecycle

Better ways to support business

Building case for virtualization

Why choose Oracle VM

Licensing and Support

- Option for soft or hard partitioning
- Oracle VM **licensing and support**
- Support for Oracle products when using Oracle VM
- Less vendors

Consideration factors

- Experience with Oracle VM
- Features
- Stability
- Dell hardware compatibility
- Snapshot and backup capabilities

Compatibility

- Oracle products under Oracle Enterprise Linux or Windows
- Future versions compatibility

+ **Cost**

+ **Compatibility**

+ **Support**

Planning Phase

Initial Planning

- How are you currently utilizing your hardware
- How many instances are you running
- Future requirements
- Licenses
- Plan your initial Oracle VM architecture (Oracle VM server pools)

Analyse and calculate

- Total count of CPU and RAM
- Use **www.spec.org** to get estimate on current hardware capability
- Spec CINT2006 benchmarks for processor per hardware vendor
- Calculate total score for your hardware per Oracle VM server pool

Server	Processor	Amount	Spec CINT2006 processor score	Total score	RAM
MT 1	E5640	1	119	119	16
MT 2	E52609	2	112	224	16
DB 1	E5640	2	119	238	64
DB 2	E5-2680	2	166	332	64
DB 3	E5-2650v2	2	341,5	683	64

Server pool 1 {

Server pool 2 {

Licensing with Oracle VM (with Intel processors)

License with cores

All Oracle VM servers in the server pool

$\text{Servers} * \text{CPUs} * \text{Cores} * 0.5$

License with NUP (named user plus)

All Oracle VM servers in the server pool

Database: $\text{Servers} * \text{CPUs} * \text{Cores} * 0.5 * 25$
(NUP minimum)

IAS: $\text{Servers} * \text{CPU amount} * \text{Cores} * 0.5 * 10$
(NUP minimum)

Enable hyper threading to maximize amount of vCPUs

License only per assigned vCPUs when using hard partitioning

(1cpu/4c hyper threading enabled and 8 vCPU available, assign 4vCPU and license only one core)

Our setup

Oracle VM **3.3.2**

four server pools

test application servers

test database servers

production application servers

administration tools

Dell M630

X12

2x Intel E5-2667 v3 (3.2 GHz, 8c/CPU)

Intel x520 for 10GbE iSCSI

Intel x520 for 1GbE LAN

320GB RAM

Dell Compellent SAN with 10GbE iSCSI (multipath)

Architecture after design phase

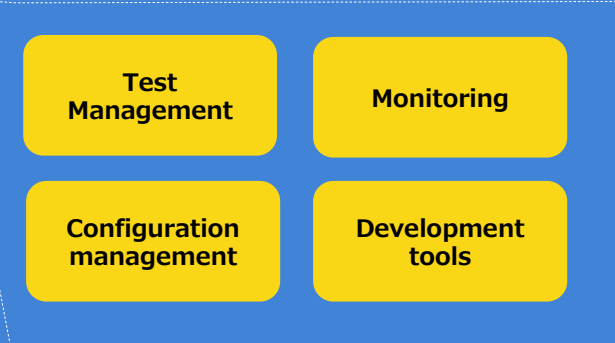
Identical hardware

- Improved fault tolerance for failures
- Easier manageability
- Scalability possibilities



Administration pool

- Consolidate monitoring and administration services
- Reused hardware



OVM Manager

EM 12c



Production databases

Architecture decisions

Naming

1. Templates
2. Servers
3. Identical naming for volumes in OVM and SAN

Network functions

1. Server Management
2. Cluster Heartbeat
3. Live Migration
4. Storage network
5. VM Network

Volumes

1. iSCSI volumes as physical disks
2. OVS repository for templates, ISOs and assemblies (500GB)
3. Storage pool size for VMs (12GB)

Tags

1. Grouping VMs with tags per instance name, application, purpose
2. Easier manageability for operations

Templates

1. Operating systems with pre-installed binaries (OEL5, OEL6, Windows Server 2008R2..)
2. Applications (Oracle Database, Oracle eBS Apps, Demantra, Weblogic..)
3. Oracle has preinstalled VM templates available – Use them for starting point!

Oracle VM installation

1. **Download Oracle VM Manager**
2. **Install Oracle VM Manager to dedicated server**
3. **Download Oracle VM Server for x86**
4. **Install Oracle VM Server for x86 to all servers in server pool**
5. **Discover Oracle VM Servers from Oracle VM Manager**

Installation checklist

Oracle VM Server

- **Enable hyper threading**
- **Disable power saving**
- **Number of dom0 (Management domain) vCPUs**
Since 3.2.2 default 20 /boot/menu/grub.lst
dom0_vcpus_pin dom0_max_vcpus=16
Change this to match your CPU configuration
- **iSCSI configuration**
/etc/iscsid/iscsi.conf
node.session.cmds_max
node.session.queue_depth
- **sysctl parameters**
/etc/sysctl.conf
vm.swappiness
vm.dirty_ratio
vm.dirty_background_ratio

Oracle VM Manager configuration

- **Server pool repository (OVS) setup**
- **Set MTU 9216 (enable jumbo frames) for iSCSI network**
- **Use of:**
 - **Distributed Resource Scheduling (DRS)**
 - **Distributed Power Management (DPM)**

Documentation

- **Oracle VM 3: 10GbE Network Performance Tuning**
 - <http://www.oracle.com/technetwork/server-storage/vm/dom3-10gbe-perf-1906632.pdf>
- **Tuning Virtual Memory**
 - https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Performance_Tuning_Guide/s-memory-tunables.html

Migration Plan

Migration plan

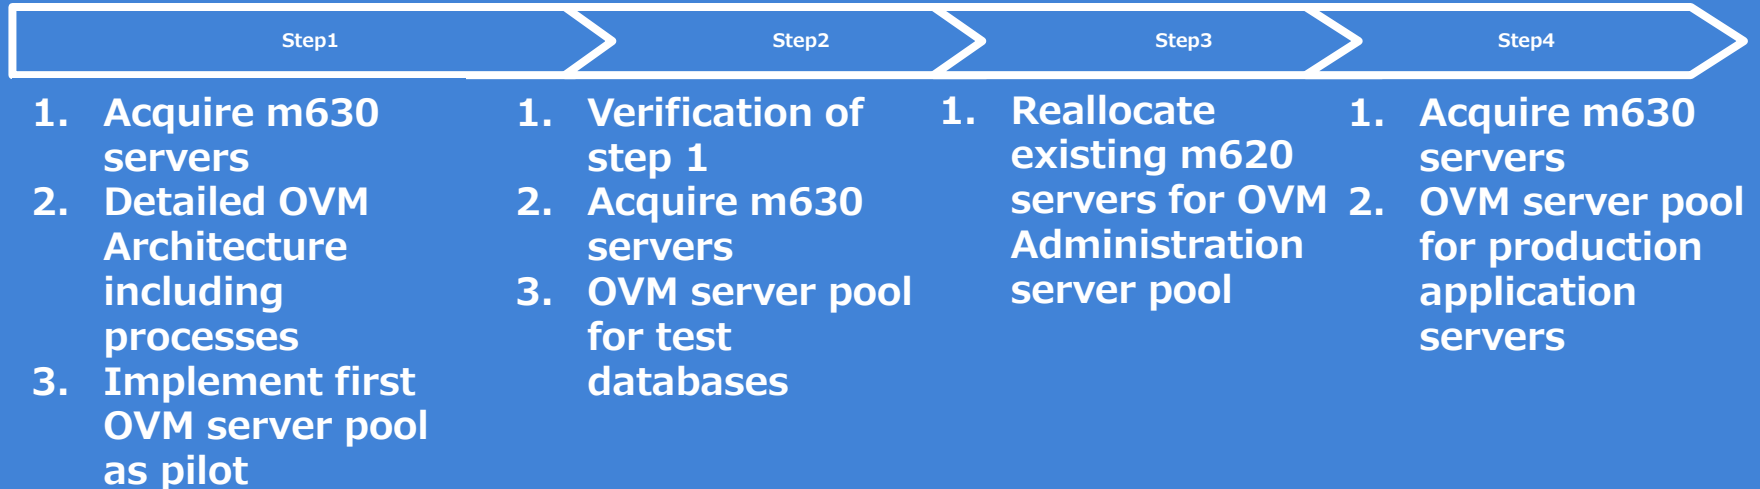
Step 1: test application servers

Step 2: test databases

Step 3: Administration and monitoring tools

Step 4: production application servers

8 months



Findings and experiences

+ Installation and Oracle VM administration is straightforward

! Plan your template usage and naming convention early

! Pay attention to network configuration

Latency issues with network cards

With iSCSI default configuration is not sufficient

! Obviously.. **TEST properly**

Test all features implemented (failover, live migration and performance)

Implementation findings

OVM Manager

- Best server functionality
- OVS Repository
- Refreshing storage takes time
- Good for management **but..**

Oracle EM12c (Oracle VM Infrastructure)

- Enable plugins:
 - Oracle Virtualization
 - Oracle Virtual Infrastructure
- VM Guest name taken from initial discovery

OVS & VM guest

- Starting up VMs is quick
- OS volume sizing (12GB)
- Configure hugepages for database VMs
- One issue with rescanning LUNs which required rolling server pool bounce
- One VM crash due to wrong SGA configuration

Performance tests

FIO

1. I/O Benchmark tool to get good baseline and produce detailed statistics
2. Read and write tests with different block size and parallel sessions
3. Baselines with different hardware (Dell m620 vs m630 & Broadcom vs Intel)

Tests

1. iSCSI, sysctl.conf and I/O scheduler parameters
2. File system performance ext3 vs ext4
3. Operating system differences OEL5 vs OEL6

Main performance findings


OVM Server x86

- Improved performance compared to older generation hardware running Oracle Enterprise Linux 5 or 6
- Configuration required for optimal iSCSI performance
- **Intel X520 NIC outperforms** Broadcom Netextreme II on OVM Server with iSCSI
- Enable jumbo frames on all supported devices (SAN, switches, iSCSI NIC, OVS)

VM guest performance

- Latency and throughput issues with guests running OEL 5 or OEL6 using PVM drivers
- Latency increasing with large block reads and multiple parallel sessions

Breaking point performance tests showed



OEL 5 and ext3

VM performance averaging 50% compared to Oracle VM server (ovs) with large block reads

Small block read and write tests perform closer to Oracle VM server (ovs)

xenblk-front parameter for VM from support improved overall performance (values 64-4096)

Disk partition alignment to reduce I/O

What next

I

Implement production application server pool with hard partitioning

II

Efficient and systematic use of Oracle Enterprise Manager 12c for OVM


III

Standardize environment provisioning and lifecycle management


IV

Create custom templates with Oracle Virtual Assembly Builder

Where we are currently



90% virtualized



-30% physical servers
+20% total servers

Q & A

References

Oracle VM 3: 10GbE Network Performance Tuning

<http://www.oracle.com/technetwork/server-storage/vm/vm3-10gbe-perf-1998932.pdf>

Tuning Virtual Memory

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/5/html/Performance_Tuning_Guide/s-memory-tunables.html

Optimising Oracle VM 3.x Oracle Linux 5 guest I/O write performance through guest/host partition boundary alignment (Doc ID 1632576.1)

Aligning Partitions to Maximize Storage Performance

<http://www.oracle.com/technetwork/server-storage/asm-unified-storage/documentation/partitionalign-111512-1979960.pdf>

FIO: <http://freecode.com/projects/fio>

FIO scripts by Kyle Hailey

https://github.com/khailey/fio_scripts/blob/master/README.md