SOFTWARE. HARDWARE. COMPLETE.

Sun

ORACLE

ORACLE

Linux/Open Source and Cloud computing

Wim Coekaerts
Senior Vice President, Linux and Virtualization Engineering

NIST Definition of Cloud Computing

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

This cloud model promotes availability and is composed of:

5 Essential Characteristics

- **On-demand self-service**
- **Resource pooling**
- **Rapid elasticity**
- **Measured service**
- **Broad network access**

3 Service Models

- **SaaS**
- **PaaS**
- **IaaS**

4 Deployment Models

- **Public Cloud**
- **Private Cloud**
- **Community Cloud**
- **Hybrid Cloud**

Source: NIST Definition of Cloud Computing v15

SaaS, PaaS and IaaS

Software as a Service

Applications delivered as a service to end-users over the Internet

Platform as a Service

App development & deployment platform delivered as a service

Infrastructure as a Service

Server, storage and network hardware and associated software delivered as a service

Evolution of Private and Public Clouds

Private Cloud Evolution

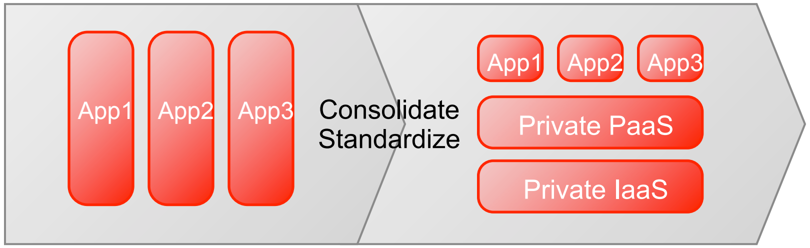

Silo'd

- Physical
- Dedicated
- Static
- Heterogeneous
- HPC / Cluster
- Scale out

Grid/Virtualization

- Virtual
- Shared services
- Dynamic
- Standardized appliances
- Consolidate

Evolution of Private and Public Clouds

Private Cloud Evolution

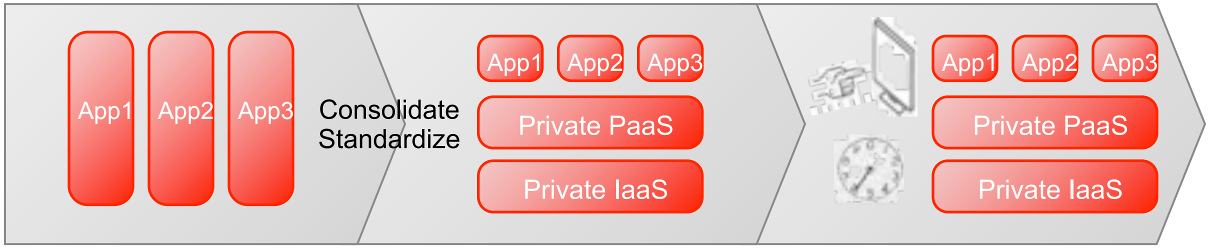

Silo’d

- Physical
- Dedicated
- Static
- Heterogeneous
- HPC / Cluster
- Scale out

Grid/Virtualization

- Virtual
- Shared services
- Dynamic
- Standardized appliances
- Consolidate

Private Cloud

- Self-service
- Policy-based resource mgmt
- Chargeback
- Capacity planning
- Merged Silo/Grid

Evolution of Private and Public Clouds

Public Cloud Evolution

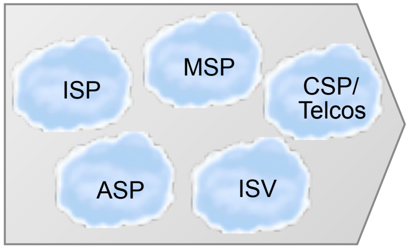
ISP
MSP
CSP/ Telcos
ASP
ISV

Private Cloud Evolution

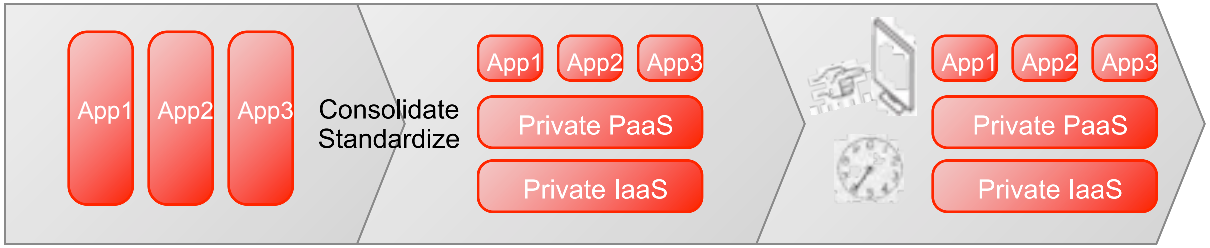
App1
App2
App3
Consolidate
Standardize
App1
App2
App3
Private PaaS
Private IaaS
App1
App2
App3
Private PaaS
Private IaaS

Silo’d

- Physical
- Dedicated
- Static
- Heterogeneous
- HPC / Cluster
- Scale out

Grid/Virtualization

- Virtual
- Shared services
- Dynamic
- Standardized appliances
- Consolidate

Private Cloud

- Self-service
- Policy-based resource mgmt
- Chargeback
- Capacity planning
- Merged Silo/Grid

Evolution of Private and Public Clouds

Public Cloud Evolution

Public Clouds

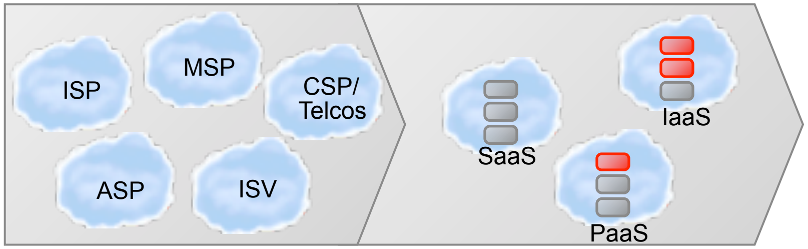

Private Cloud Evolution

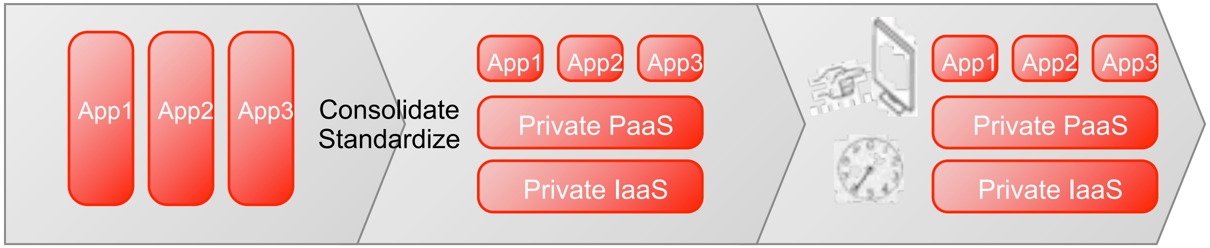

Silo'd

- Physical
- Dedicated
- Static
- Heterogeneous
- HPC / Cluster
- Scale out

Grid/Virtualization

- Virtual
- Shared services
- Dynamic
- Standardized appliances
- Consolidate

Private Cloud

- Self-service
- Policy-based resource mgmt
- Chargeback
- Capacity planning
- Merged Silo/Grid

Evolution of Private and Public Clouds

Public Cloud Evolution ▶ Public Clouds ▶ Hybrid

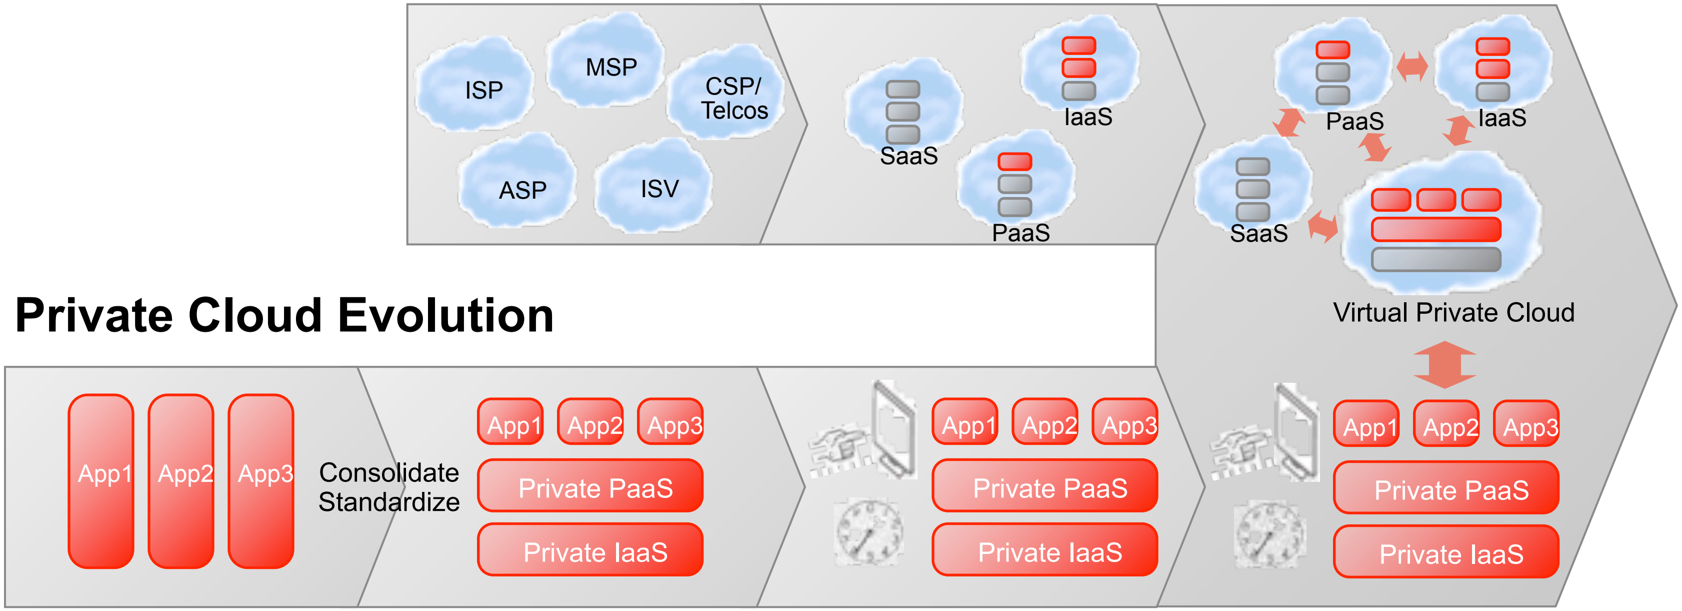

Private Cloud Evolution

Silo'd

- Physical
- Dedicated
- Static
- Heterogeneous
- HPC / Cluster
- Scale out

▶ Grid/Virtualization

- Virtual
- Shared services
- Dynamic
- Standardized appliances
- Consolidate

▶ Private Cloud

- Self-service
- Policy-based resource mgmt
- Chargeback
- Capacity planning
- Merged Silo/Grid

▶ Hybrid

- Federation with public clouds
- Interoperability
- Cloud bursting

Linux Evolution

General Purpose OS

Linux runs on small/medium/large systems and continues to evolve. More CPU’s, more memory, higher performance...

Embedded Linux

Focus on Embedded market, handheld, tablets but there’s similarities with cloud / virtualization. Special Linux kernel build and custom Linux OS distributions

Linux-based Appliances

Minimal install, easy to package and repackage Linux distribution to deploy apps on for virtualization and cloud deployments. Based on general purpose Linux but with ideas from embedded world

Our efforts on Linux

- Work on Linux scalability - as a hardware company we have very large x86 systems 128-160 threads 1-2TB memory. NUMA awareness, high speed network, flash disk, tracing, fault management, filesystems (btrfs) ... products like ExaData and ExaLogic

- **-All-** our code changes go into the mainline tree (Linus's tree). Our Oracle Linux source trees are also publicly available including changelogs :

http://oss.oracle.com/git/?p=linux-2.6-unbreakable.git

- Create packages to make packaging of Linux-based assemblies or appliances easy and reproducible. This includes easy configuration of application installation on top as well. This work comes in Oracle VM Templates with Oracle Linux (freely downloadable including source).

Our efforts on Linux

- btrfs development efforts continue
  - helps in providing an easy way to take snapshots and upgrade /
  - for local filesystem based VM disks reflink clones

- ocfs2
  - reflink provides us the ability to clone VMs on the fly across a cluster instead of just being able to do local dependent snapshots
  - dlmfs is a very convenient solution for implementing simple cluster based lock infrastructure for VM HA (we use it in Oracle VM)

- Make Linux into a better virtualized OS - projects like zcache and ramster. cooperative guests are much more interesting than each VM on its own

Open Source Virtualization

Hypervisors

True hypervisor concept - Xen most well known version of open source hypervisor used in many production systems today.

OS with hypervisor characteristics

Ability to run VMs on top of a host kernel, like KVM or Oracle VM VirtualBox.

Container style

Form of virtualization where you isolate processes and applications from each other on top of a single OS kernel. Open VZ, LXC

Our efforts on Xen and Oracle VM VirtualBox

- Contrary to popular believe, the Linux tree has full support for Xen. For Xen guest as well as Xen privileged (dom0) domains. We worked with Citrix and the Xen community to get the code submitted and in acceptable form for Linus.

  Linux 2.6.39 rc6 :

  config XEN_DOM0 (including pciback, netback and blockback coming)

  config XEN

  config XEN_PVHVM

  config XEN_MAX_DOMAIN_MEMORY

  config XEN_SAVE_RESTORE

- **-All-** our code changes go into the mainline tree (Xen.org Xen tree).

- Oracle VM VirtualBox and Oracle VM Server (Xen-based) are both complete open source Virtualization products. Free download source + binaries and production use.

Our efforts on Xen and Oracle VM VirtualBox

- VirtualBox development is actively continuing
  - new releases every few months even debian includes 4.0 .debs
  - used on server side in VDI - Cloudy desktops and we can run Linux with multimedia (RCA support) not just windows, remote USB, remote audio (happily run Ubuntu VMs)
- Data Integrity integration into VMs (WIP)
  - work to eliminate silent data corruption
  - data corruption unnoticed no errors or warnings until its too late
  - logical block checksum is not enough, used at READ time, too late
  - requires end-to-end integrity checking app -> disk device to **prevent** bad data from being written
  - Cern study, 207 : write known data pattern to more than 3000 nodes over a 5 week period had 8.7TB worth of corrupt data

Open Standards - or the need for Open Standards is important

Packaging format for appliances (OVF)

OVF is a good start to have a common format of describing a set of VMs that work together. Assemblies with configuration metadata. In DMTF

Cloud API

Many clouds but other than de facto standards no true standard yet. Oracle submitted a proposal to DMTF in 2010

Application and OS Configuration

Above OVF

- startup scripts (first boot)
- cleanup scripts (clean virtual disks, swap space, temp files, reset to first boot state)

Helping open standards along

- We provide support for OVF in both Oracle VM VirtualBox(first to implement support) and Oracle VM server
- Work with the DMTF to suggest enhancements and propose additions to the OVF standard to fit complex application configuration
- Use OVF in our own assembly builder products (Oracle Virtual Assembly Builder) and support for OVF in our management products
- Oracle Cloud API submission and feedback forums
- Enhance Oracle products such as MySQL, and many others to be easily deployed in a cloud environment (optimized, preconfigured MySQL/Oracle Linux template on Oracle VM)
- Storage Connect

Oracle VM Templates

Rapid Application Deployment

Save days or weeks in installation and configuration time

ORACLE E-Delivery

Download from Oracle

- Pre-built, pre-configured VM
- Complete app, middlewar DB installation
- Complete MySQL, Siebe Database 11g, Enter Manager…

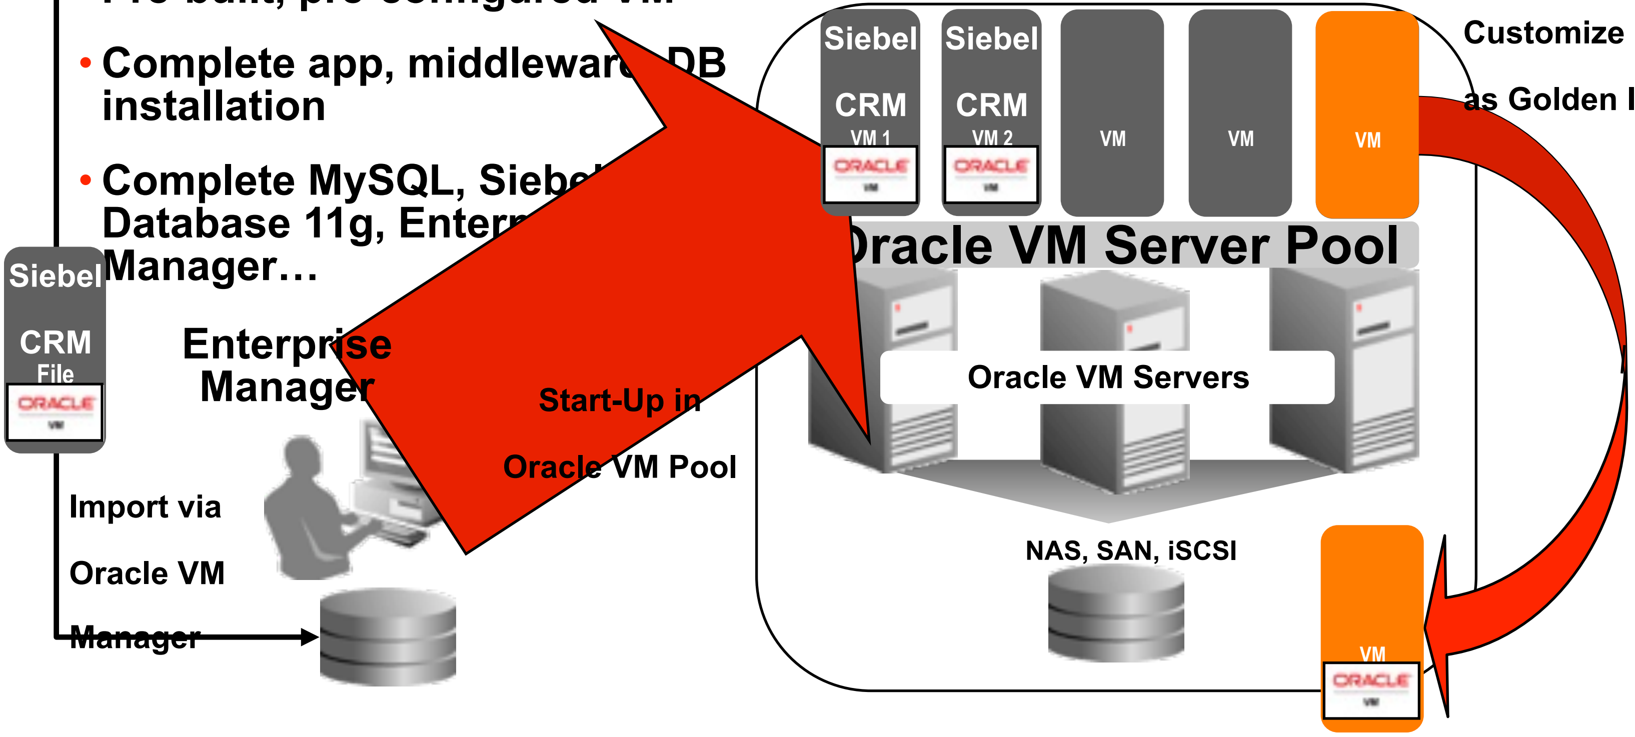

Oracle Virtual Assembly Builder

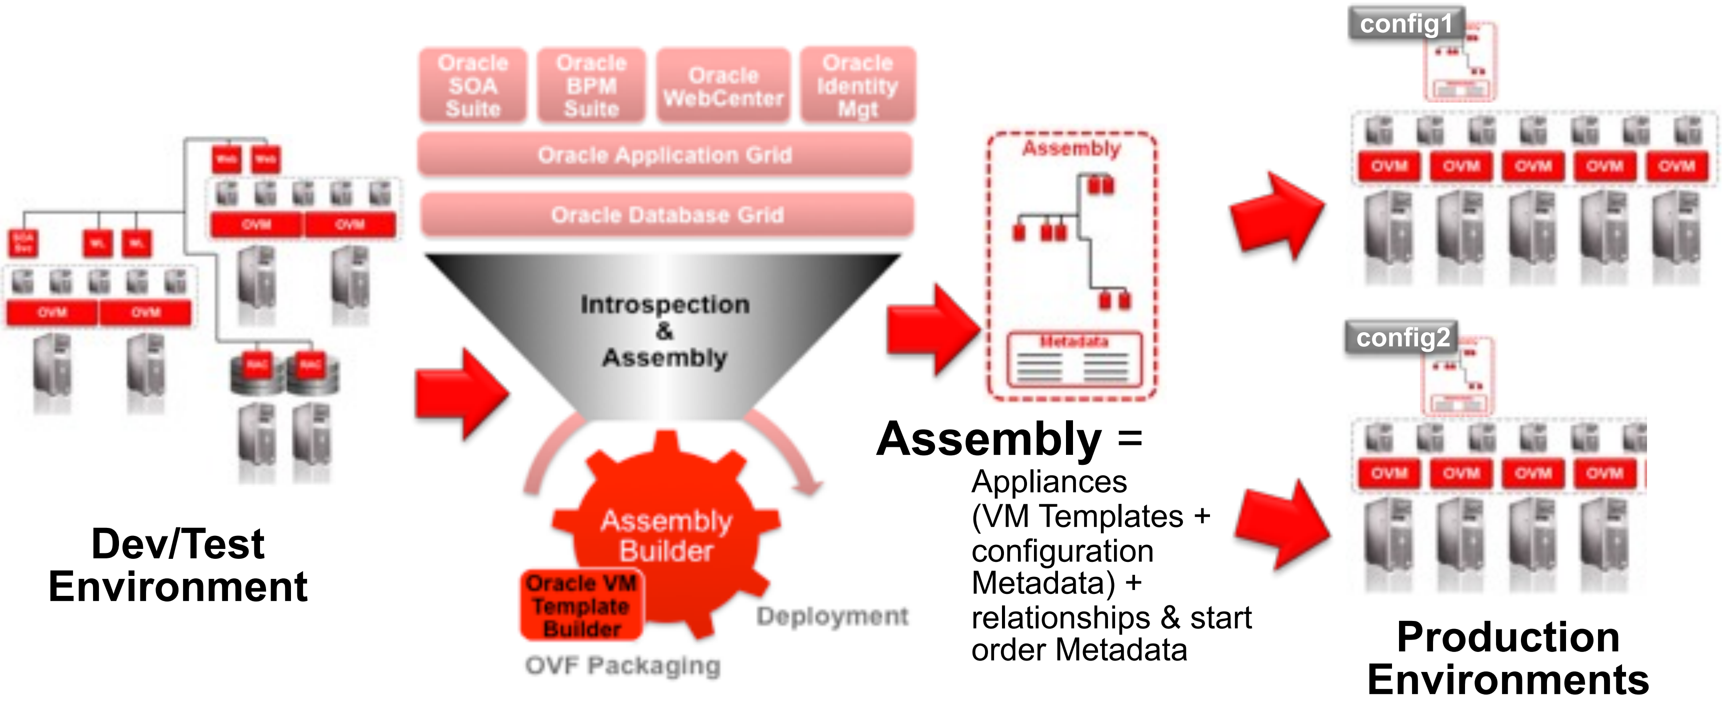

Dev/Test Environment

Assembly =
Appliances
(VM Templates +
configuration
Metadata) +
relationships & start
order Metadata

Production Environments

- **Package** up complex structure from dev/test and reconstitute in production
- **Minimize setup time and risk** of hard-to-debug configuration errors
- **Easily replicate** in production with minor variations
- Each production instance has **well-contained configuration** parameters for flexibility

Case Study

Oracle’s Software Development Cloud - “Devops”

Nearly Ten Years Development / Use

- Internal **hardware resource management** application leveraging existing development automation as a ‘private cloud API’
- A platform allows us to move more users & hardware resources into a **self-service ‘cloud’***
- A **self-service reimaging & reboot** portal for user

Current Metrics

✓ Avg. new **VM reservatio** per day: **~50**

✓ Avg. **self-service reboot** per day: **~25**

✓ Avg. **self-service reimag** per day: **~100**

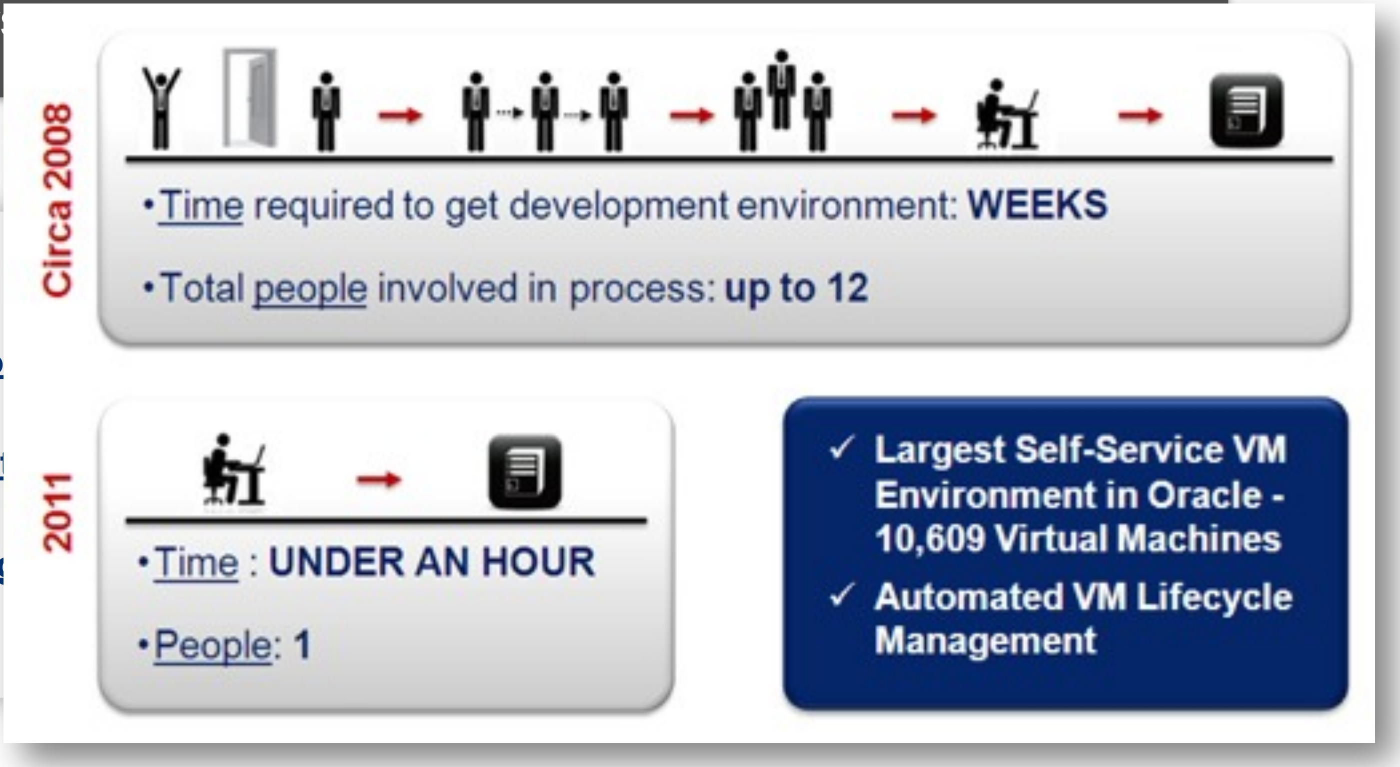

Circa 2008

- Time required to get development environment: **WEEKS**
- Total people involved in process: **up to 12**

2011

- Time : **UNDER AN HOUR**
- People: **1**

✓ **Largest Self-Service VM Environment in Oracle - 10,609 Virtual Machines**

✓ **Automated VM Lifecycle Management**

*That is not limited to a single: OS imaging technology, hypervisor, or cloud API

What’s next

- For developers
  - it’s important to work together and create tools based on open standards, participate and implement open standards
  - Interoperability should be in the design
  - Small footprint is ‘in’, disk and memory size matters

- For businesses, to make clouds work
  - it’s more than just virtualization
  - think about full stack deployment and how all components interact
  - save on sysadmin and application admin costs