Club utilisateurs Progress France

Getting the Most Out of Virtualization of Your Progress OpenEdge Environment

Tinco van Engen
Senior Technical Support Engineer
24.06.2014

PUG FRANCE

Agenda

- Virtualization
  - Terms, benefits, vendors, supportability, etc.
- Best practices
  - Disk layout, network, snapshots, etc.
- High Availability
  - Clustering, fault tolerance, backup, etc.
- Replication
  - vSphere SRM

Virtualization

PROGRESS

Virtualization – what is it?

- Way of running multiple OS and applications on a single computer
- Each OS runs within its own Virtual Machine (VM)
  - (virtual) CPU, memory, disk allocation
- Global resource control governed by hypervisor
  - controls the host processor and resources
  - ensures that VMs are isolated from each other

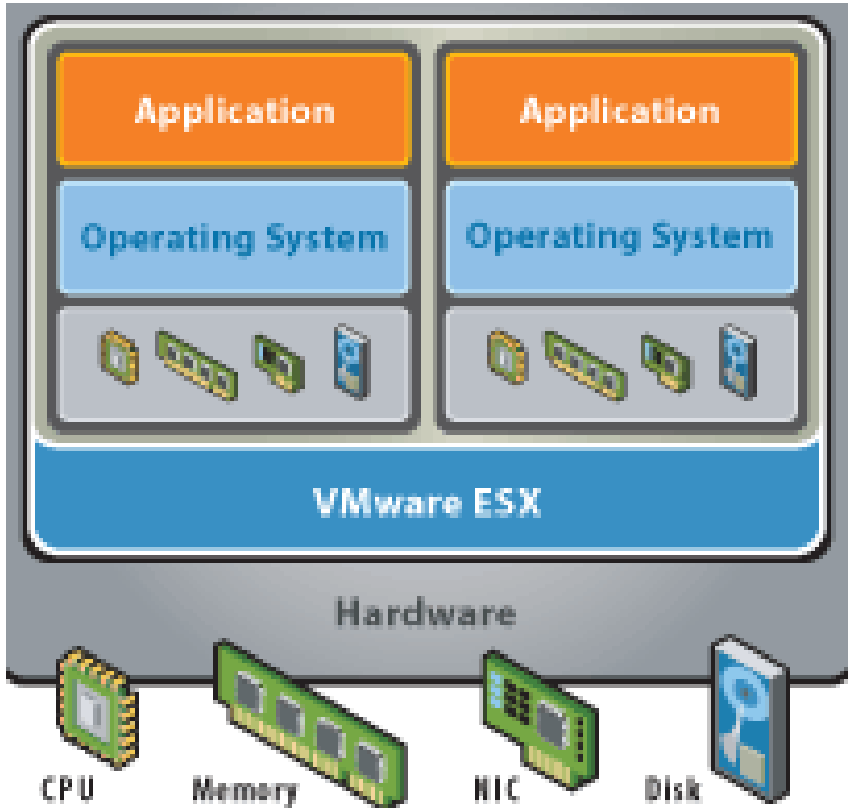

Virtualization – Benefits

- Lets you utilize your hardware more effectively
- Allows you to centrally manage your infrastructure
- Speeds up new deployment
- Supports legacy OS and applications
- Permits encapsulation and isolation
- Reduces overall IT expenses
- … and so on

Virtualization – vendors

- VMware
  - Workstation, ESXi, vSphere
- Microsoft
  - Hyper-V
- Oracle
  - Virtual Box, Solaris zones
- RedHat
  - Xen, RHEV
- IBM
  - LPAR, WPAR

Virtualization – Progress support

- Common questions
  - Does Progress support VMware version X.Y ?
  - Does Progress support VEEAM backup ?

- No – we don’t.
  - The hypervisor vendor does.

- So what does Progress support then?
  - The underlying OS running as VM on the hypervisor

Virtualization – Vocabulary

- Abbreviations
  - VM = virtual machine
  - HA = high availability
  - DR = disaster recovery
  - OE = OpenEdge
  - DRS = distributed resource scheduler
  - LUN = logical unit number
  - SAN = storage area network
  - SRM = site recovery manager
  - FT = fault tolerance
  - vDisk = virtual disk
  - vCPU, vMem = virtual CPU, virtual memory

Best practices

Best practices

- In a nutshell VM is a set of files
  - configuration - .vmx
    - CPU, memory, NIC.. allocation
  - disk(s) - .vmdk
  - RAM - .vmem
  - VM swap file - .vswp
  - log, BIOS …

[machines03] ll

Name	Size	Provisioned Size	Type
ll.vmx.lck	0.00 KB		File
ll-1235b827.hlog	0.07 KB		File
ll.vmsd	0.55 KB		File
ll.vmxf	2.96 KB		File
ll.vmx	3.67 KB		File
ll.vmx~	3.67 KB		File
ll.nvram	8.48 KB		Non-volatile memory file
ll-Snapshot63.vmsn	31.30 KB		Snapshot file
vmware-32.log	118.93 KB		Virtual Machine log file
vmware-31.log	119.81 KB		Virtual Machine log file
vmware-30.log	120.12 KB		Virtual Machine log file
vmware-28.log	133.34 KB		Virtual Machine log file
vmware-29.log	325.28 KB		Virtual Machine log file
vmware-27.log	623.68 KB		Virtual Machine log file
vmware.log	739.59 KB		Virtual Machine log file
ll-ctk.vmdk	5,765.19 KB		File
ll-000001-ctk.vmdk	5,765.19 KB		File
vmx-ll-305510439-2.vswp	94,208.00 KB		File
ll-000001.vmdk	2,689,024.00 KB	94,448,640.00 KB	Virtual Disk
ll-1235b827.vswp	3,145,728.00 KB		File
ll.vmdk	94,449,660.00 K		Virtual Disk

- Possible to scale up or down resource allocation after VM creation
- Always install VMTools, Synthetic drivers for your VM

Best practices – VM CPU allocation

- More vCPUs do not always guarantee best performance
- Remember that your VM is not the only VM on the host!

- Example
  - Dual quad with HT enabled
  - 16 cores available
  - 8 vCPUs assigned for my database VM
  - Other VMs using 10 cores at the moment
  - My VM has to wait for a time slice till 8 cores are free

- Make sure that DBA has an access to host performance stats

Best practices – VM Memory allocation

- Generally more memory – better
  - Larger –B/-B2 makes your database faster
  - Do not go over memory allocated for the VM
- ESX host creates a .vswp file equal to the amount of allocated physical memory
  - Not entirely true if reservation is in play
  - 64GB memory allocation creates 64GB .vswp file
  - For 10 VMs that's 640GB disk space
- Unnecessary memory allocation can lead to disk space issues
  - Hard to track

Best practices – CPU & Memory allocation

- CPU/Memory “hot plug”
  - if supported by the host OS
- Rule of thumb
  - start with less resources
  - scale up for performance if needed

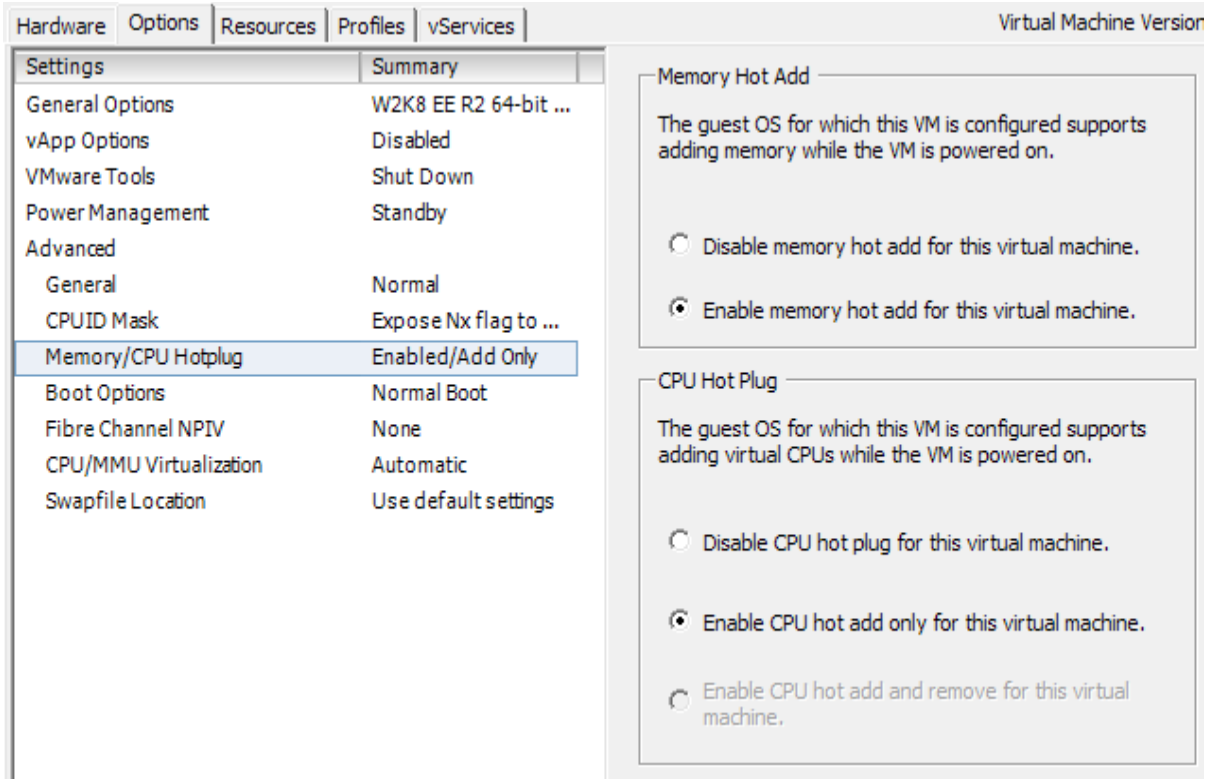

Best practices – VM Disk allocation

- Considerations when creating a virtual disk
  - How the disk is created
  - When the space is allocated

- 3 types of disk provisioning
  - Thick lazy zeroed
    - vDisk not zeroed upfront
  - Thick eager zeroed
    - Entire vDisk zeroed out before becoming accessible
  - Thin provision
    - Instant access, allocation on demand

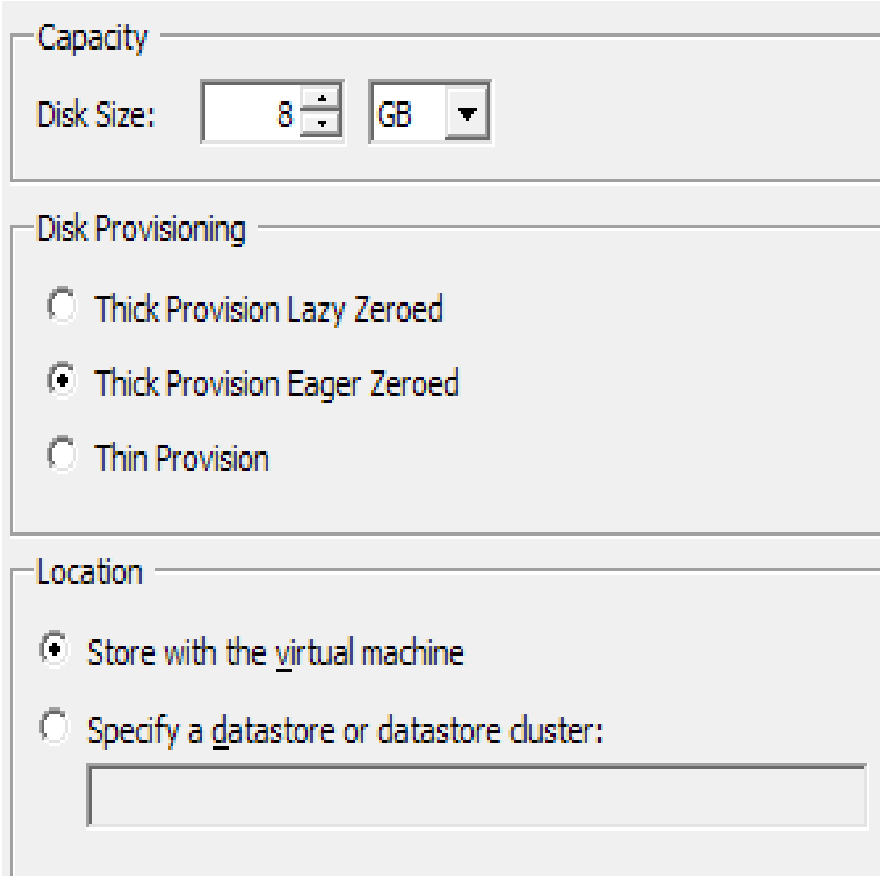

- Database VM does not belong on thinly provisioned vDisk

Best practices – VM Disk allocation

- RAW device mapping (RDM)
  - vDisk in native OS format directly on SAN
  - benefits from SAN snapshots and replication
  - native SAN optimization
  - configure with care
    - use virtual compatibility as opposite to physical
    - LUN has to have same LUN ID across all the hosts
    - test before going live!

Best practices – NIC configuration

- Several network adapters available
  - usage depends on the guest OS
  - vmxnet2, vmxnet3, e1000, e1000e …
- Usually selected by default on VM creation
  - vmxnet3 – paravirtualized NIC
  - e1000e – emulation of Intel Gb Ethernet Controller
- Whenever possible use vmxnet3
  - Less CPU intensive
  - Gives better throughput

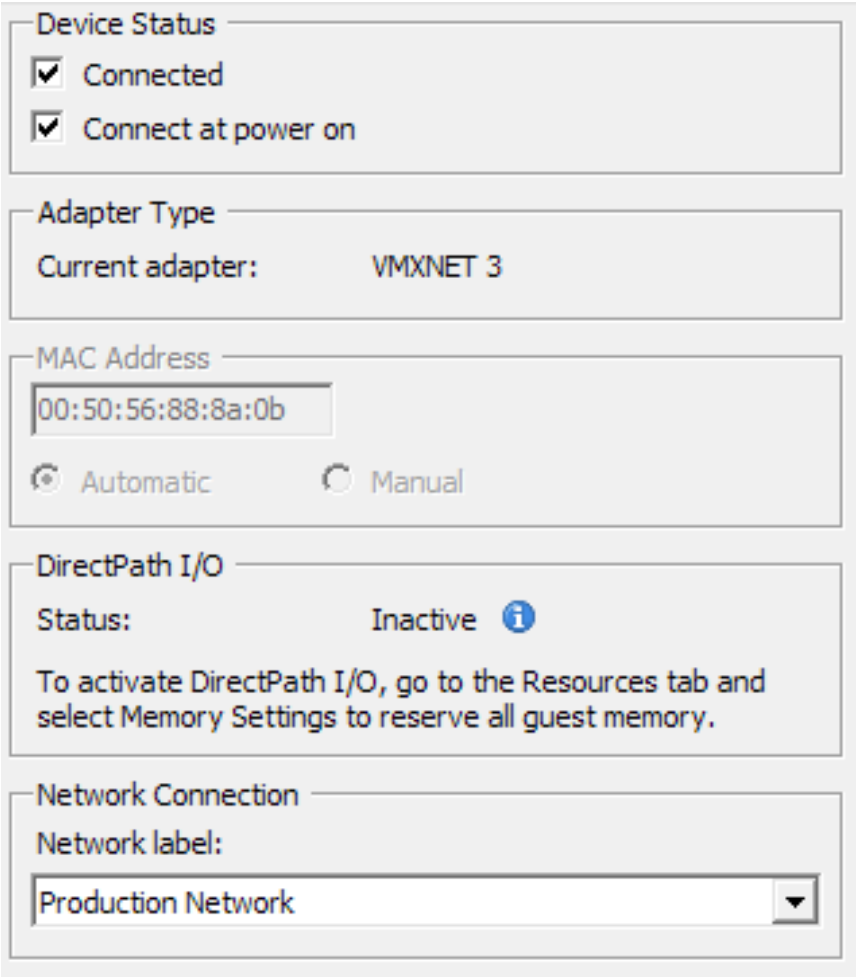

Best practices – NIC configuration

- VMDirectPath I/O “passthrough”
  - in case of “network intensive” applications
    - hundreds of AppServer clients
    - hundreds of Client-Server connections
- Has to be enabled on the device level first
  - then it becomes available for the vNIC
- While it improves performance, there are limitations
  - complicates HA
  - up to 6 devices

Best practices – Snapshots

- Copy of the virtual machine's disk file (VMDK) at a given point in time
  - offline and online
- Great when installing OS, app patch or a new version
- Snapshots are NOT backups!
- Not for prolonged use
  - Can and will cause performance issues
- Careful when taking it while having a database running
  - use a quiet point
  - **verify** that a quiet point's been enabled
  - … unless you like playing Russian roulette

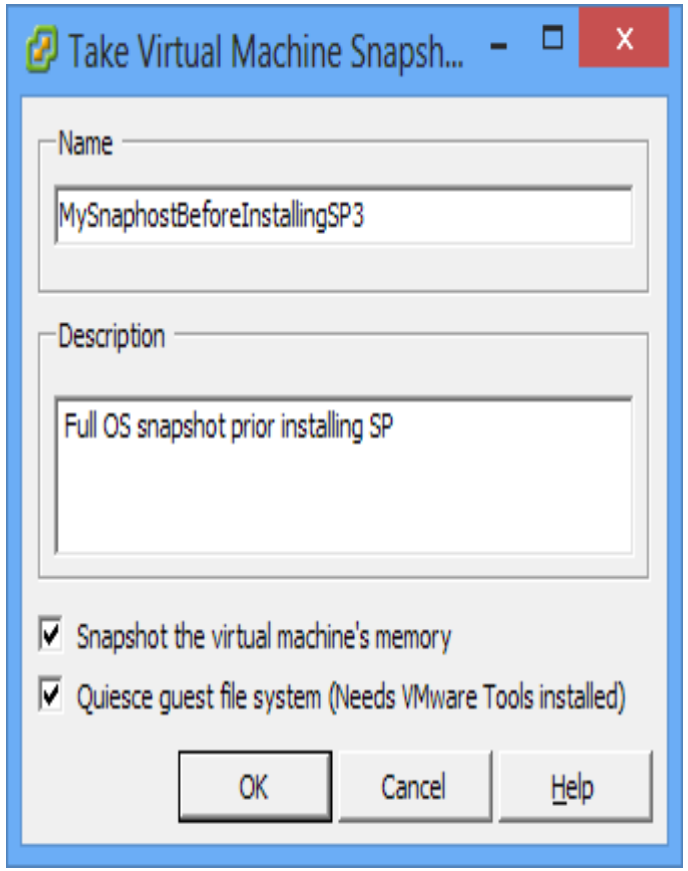

Best practices – Snapshots

- Taking snapshot “willy-nilly” could cause DBDOWN with
  - bkioRead:Insufficient disk space during .... (9450)
  - SYSTEM ERROR: read wrong dbkey at offset ... (9445)

- VMware provides hooks
  - requires VMTools
  - pre-freeze-script & post-thaw-script

- General solution for any tool using snapshots under the hood
  - pre-freeze-script: proquiet db –C enable
  - post-thaw-script: proquiet db –C disable

Best practices – Snapshots

- Independent vDisk
  - not affected by snapshots
- Non-persistent
  - content of non-persistent vDisk is discarded on power off
  - do NOT place your database on non-persistent vDisk
  - application/client temporary files
    - dbi, lbi, rcd, srt ….
- Persistent
  - any static part of your application
  - database backed up by online probkup

Best practices – DRS

- Distributed resource scheduler
  - optimizes workload with available resources
    - based on CPU, memory & storage load of a host
    - live migration to a less utilized host
  - resource prioritization per VM (application)
  - isolation based on business
    - resource pools
    - production, QA, development, testing, etc.
  - affinity rules
    - where and how VMs can run
    - both Application server VM and database VM have to start
    - OE Replication source and replication target VMs always on different hosts
    - at least one failover cluster node have to be on a different host than the rest

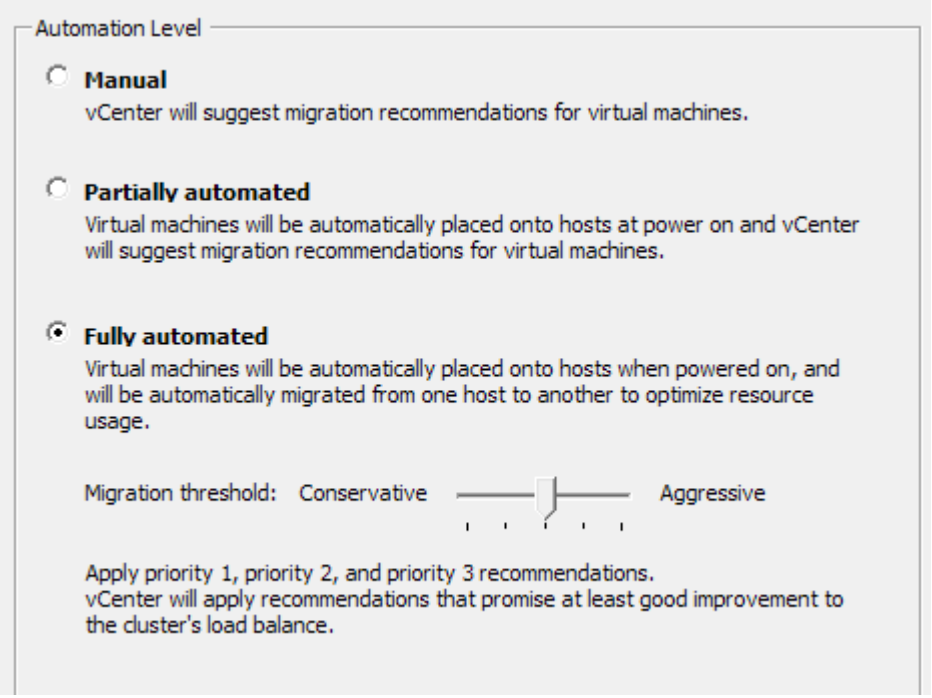

High availability

High availability

- Progress HA/DR solutions
  - Failover clusters
  - OE Replication
  - (NameServer) load balancing

- VMware provides their own on the VM level
  - vMotion
  - Storage vMotion
  - Fault tolerance
  - Cluster
    - HA
    - DRS

High availability

- VMware High Availability features can enhance resilience and uptime of OpenEdge processes
  - Database
  - AdminServer
  - OE Management & Explorer
  - AppServer & WebSpeed brokers
  - OE Application itself

- Let’s drill down

High availability - vMotion

- vMotion
  - migration of a VM between 2 different hosts
  - cold
    - offline
  - live
    - online

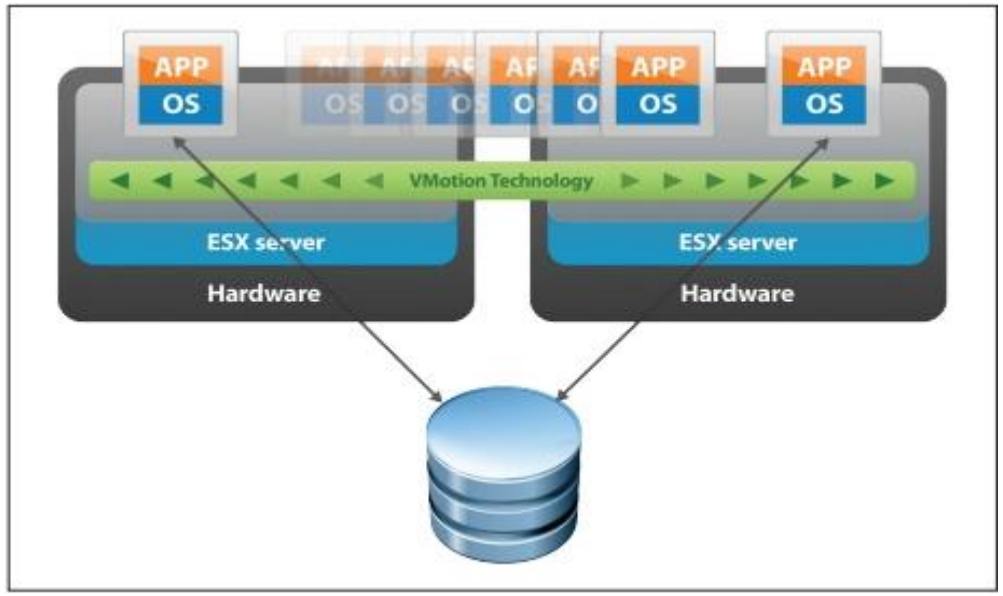

- Live vMotion
  - quick way of offloading a VM from a busy host (while VM is powered on)
  - can be automated on ESX cluster level to balance server utilization
  - minimum or no business disruption

High availability – Storage vMotion

- Storage vMotion
  - enables live migration of virtual disks on the fly
  - way of offloading an online VM from a busy disk subsystem
  - performance considerations

- Cannot prevent VM or ESX host failure
  - it will bring the VM up and running
  - there will be a business disruption

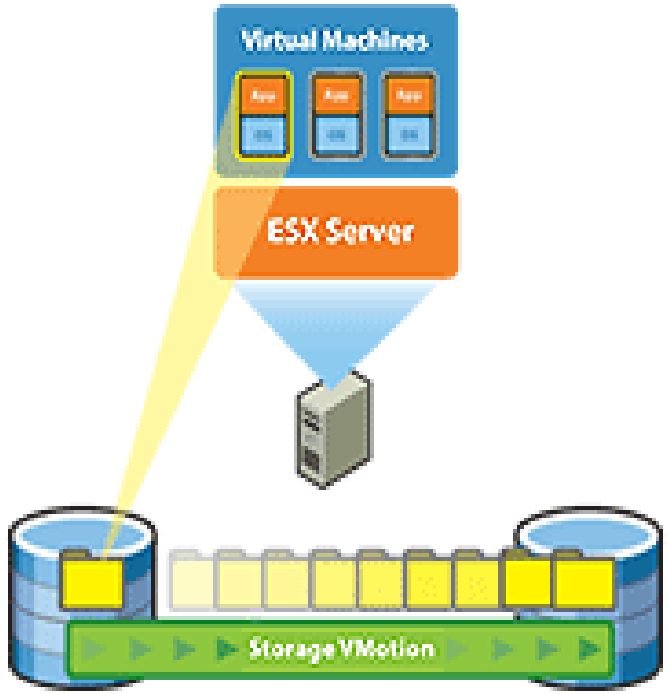

High availability – Fault tolerance

- Fault tolerance
  - not a load balancing solution
  - protects VM against ESX host failure
  - prevents un-planned downtime
  - requires 2 ESX hosts
  - dedicated network
  - synchronous replication

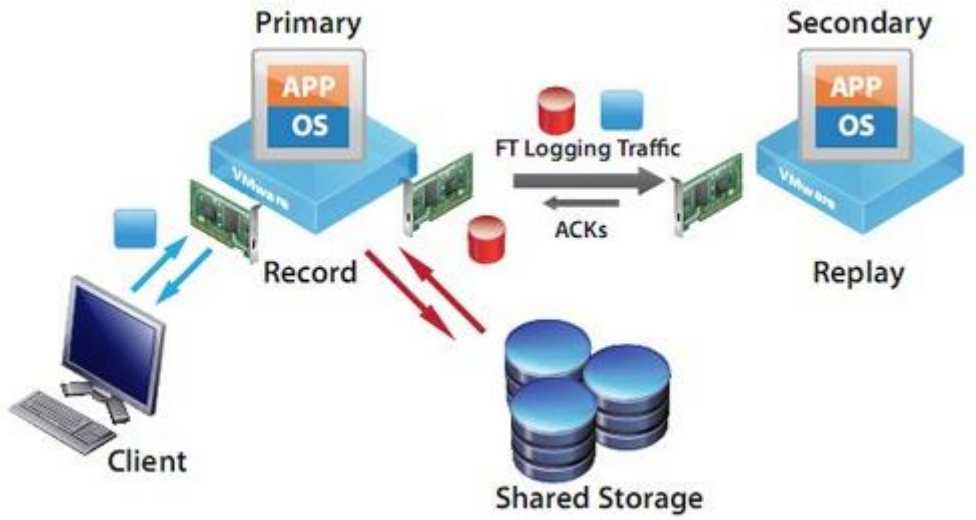

High availability – Fault tolerance

- Considerations
  - additional CPU/Memory allocation might be required
  - fast network
  - best suited for:
    - Webserver/JSE VM
    - OE Management/Explorer VM
    - AppServer VM
    - VM hosting your application
  - always test and evaluate results

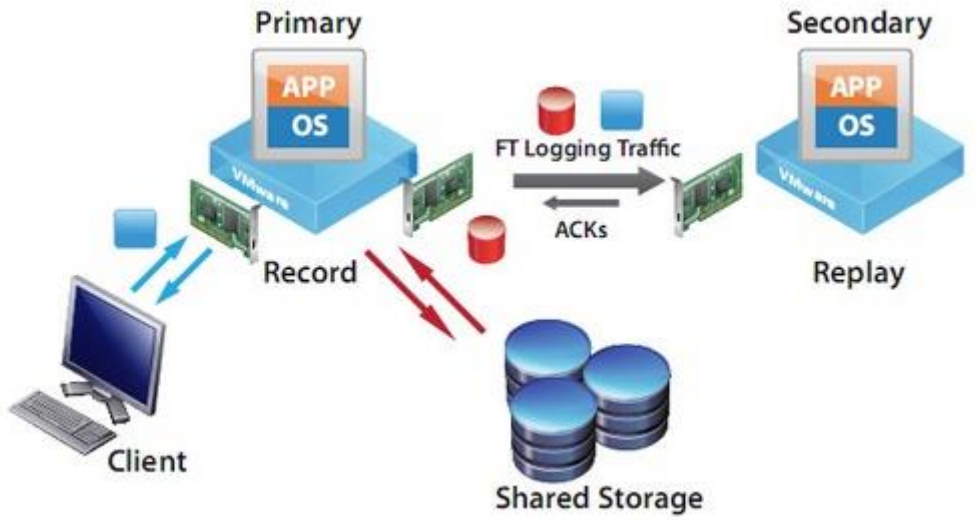

Replication

Replication

- vSphere SRM
  - SRM – Site Recovery Manager
  - provides replication of powered on VM to a secondary site
  - provides tools for failure testing
    - encapsulation
  - can replicate all vDisks or some
  - does not require 2 SANs
  - requires
    - 2 vCenters
    - extra appliances

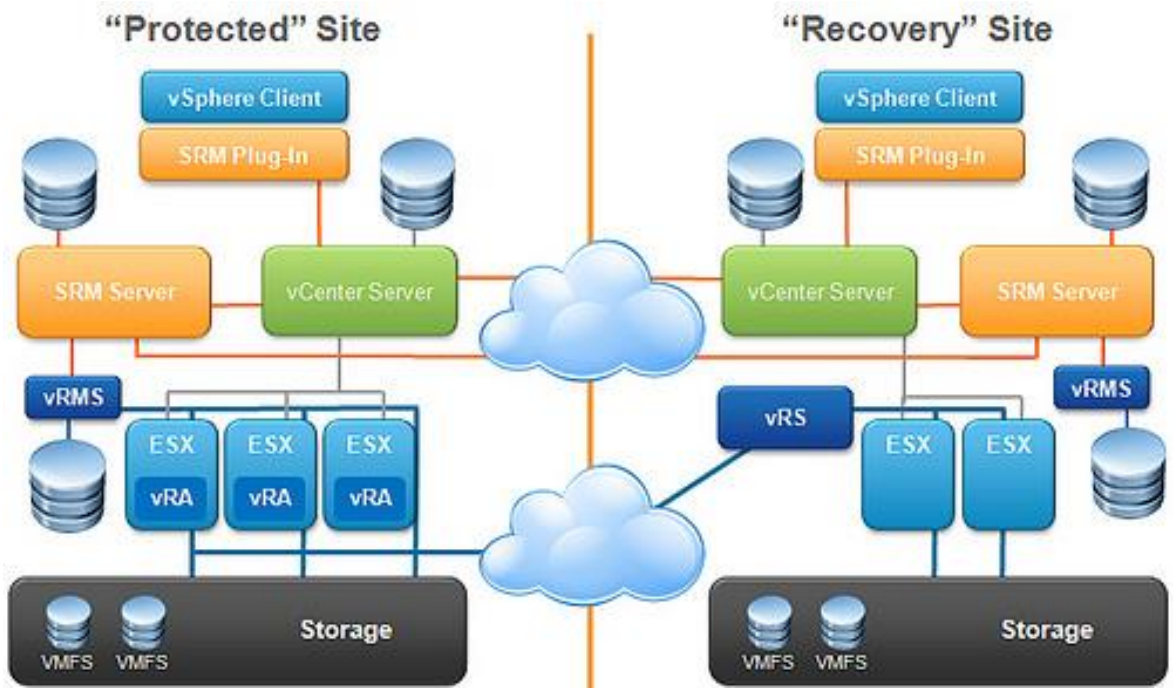

Replication

- vSphere replication
  - not “online”
  - minimum RPO is 15 minutes
  - done by using vDisk deltas
    - similar to after imaging
  - use case: AppServer, WebSpeed and application VMs

- Storage replication
  - online
  - based on EMC SRDF technology
  - disk level replication
  - use case: Database VM

vSphere SRM Replication

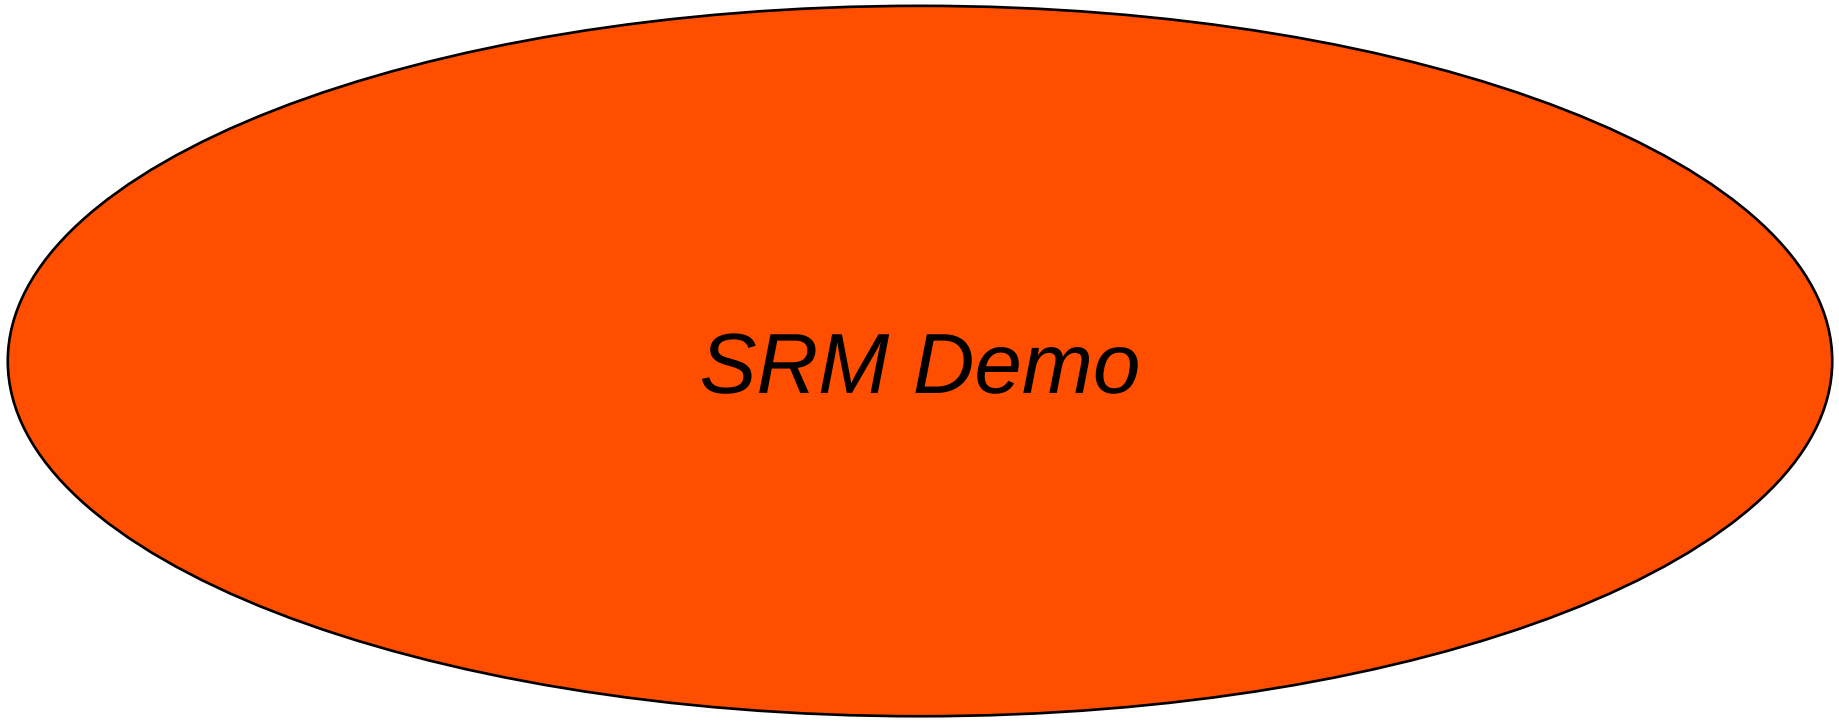

Summary

PROGRESS

Summary

- Virtualization is here
  - whether you like it or not
- Excellent QA/testing capabilities
- Rapid deployment of new servers
- DR/HA solution out of the box
- Application isolation
- Extends the life of legacy apps

- Not a “free lunch” universal solution
  - hypervisor still has and will have a performance overhead
- Sometimes real (physical) hardware is better
  - YMMV, test!

PROGRESS