

# Getting the Most Out of Virtualization of Your Progress OpenEdge Environment

Libor Laubacher  
Principal Technical Support Engineer  
8.10.2013

**PROGRESS**  
**EXCHANGE** 2013  
DISCOVER. DEVELOP. DELIVER.

# Agenda

---

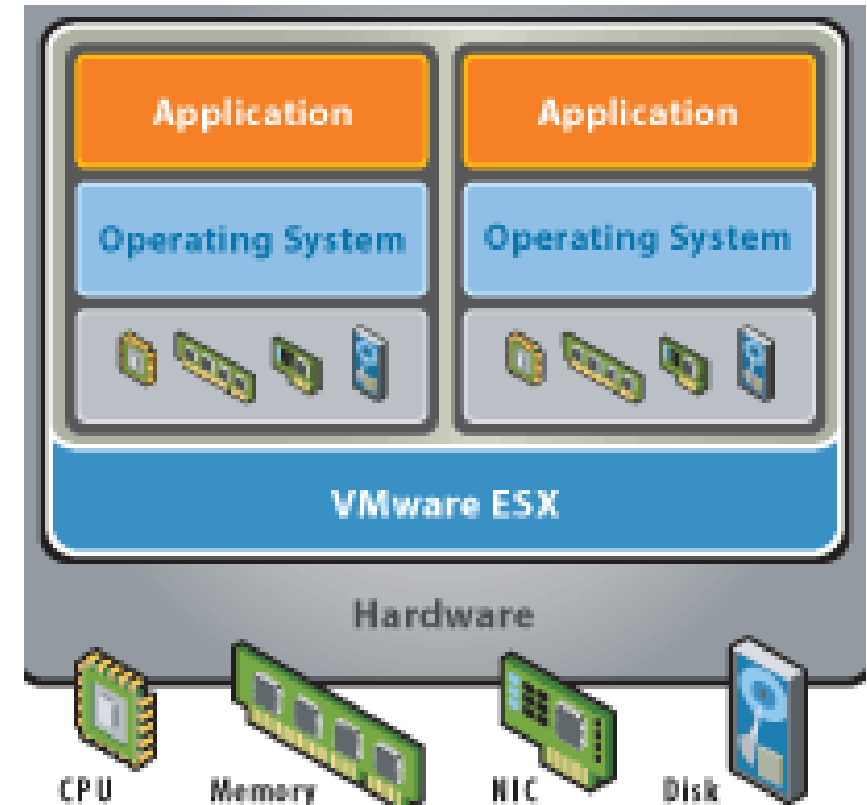
- Virtualization
  - Terms, benefits, vendors, supportability, etc.
- Best practices
  - Disk layout, network, snapshots, etc.
- High Availability
  - Clustering, fault tolerance, backup, etc.
- Replication
  - vSphere SRM

Virtualization



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

[machines03] II

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.vmx	2.96 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

# 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

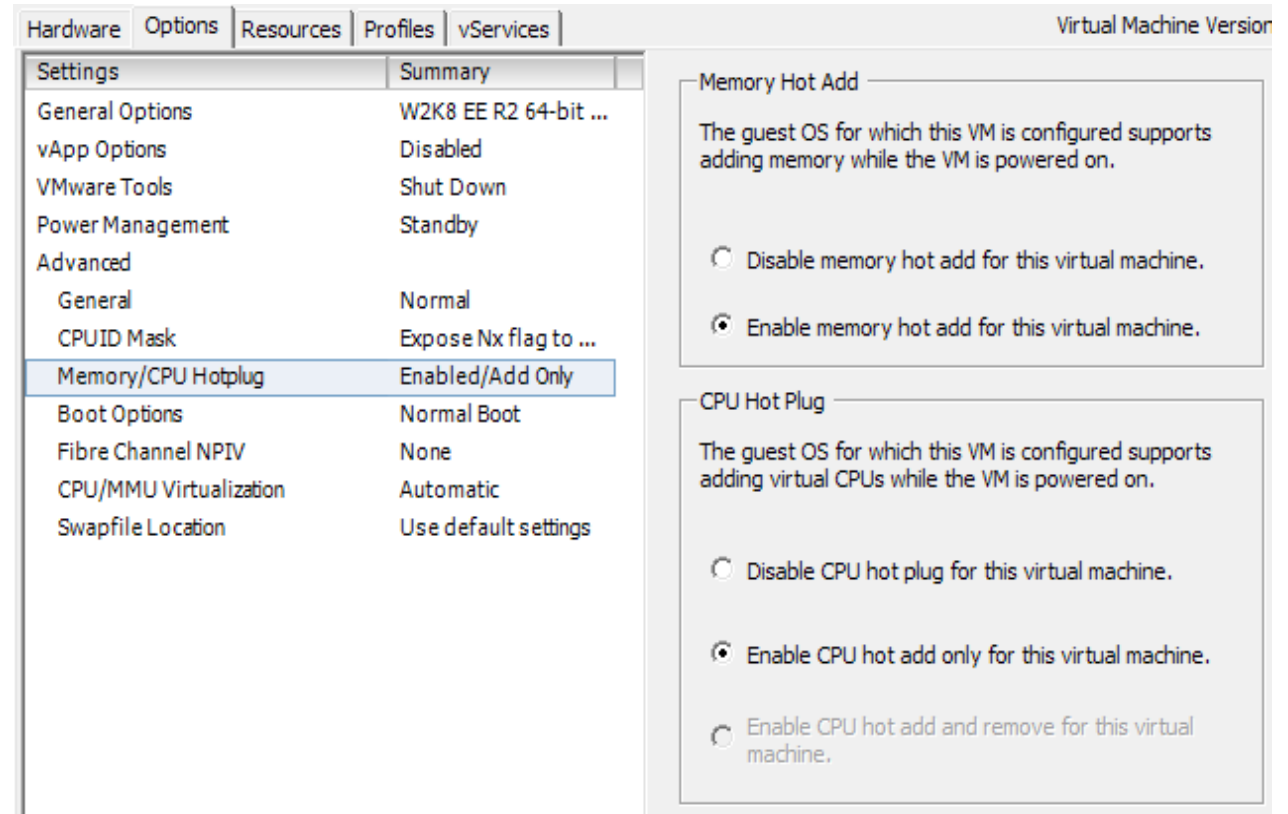
# Best Practices – VM Memory Allocation

---

- 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

The screenshot shows a configuration window for a virtual disk. It is divided into three sections: Capacity, Disk Provisioning, and Location. In the Capacity section, the Disk Size is set to 8 GB. In the Disk Provisioning section, the 'Thick Provision Eager Zeroed' option is selected. In the Location section, the 'Store with the virtual machine' option is selected.

Capacity

Disk Size: 8 GB

Disk Provisioning

- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision

Location

- Store with the virtual machine
- Specify a datastore or datastore cluster:

# Best Practices – VM Disk allocationA

---

- 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

The screenshot shows the configuration window for a network adapter. It is divided into several sections:

- Device Status:** Contains two checked options:  Connected and  Connect at power on.
- Adapter Type:** Shows 'Current adapter:' as VMXNET 3.
- MAC Address:** A text box contains '00:50:56:88:8a:0b'. Below it are two radio buttons:  Automatic and  Manual.
- DirectPath I/O:** Shows 'Status:' as Inactive with an information icon. Below it is a note: 'To activate DirectPath I/O, go to the Resources tab and select Memory Settings to reserve all guest memory.'
- Network Connection:** A dropdown menu is set to 'Production Network'.



# Best Practices – NIC Configuration

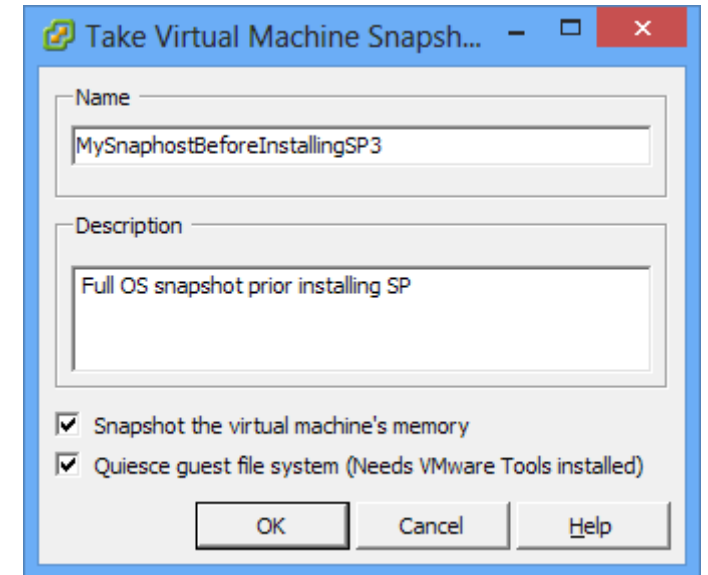
---

- 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

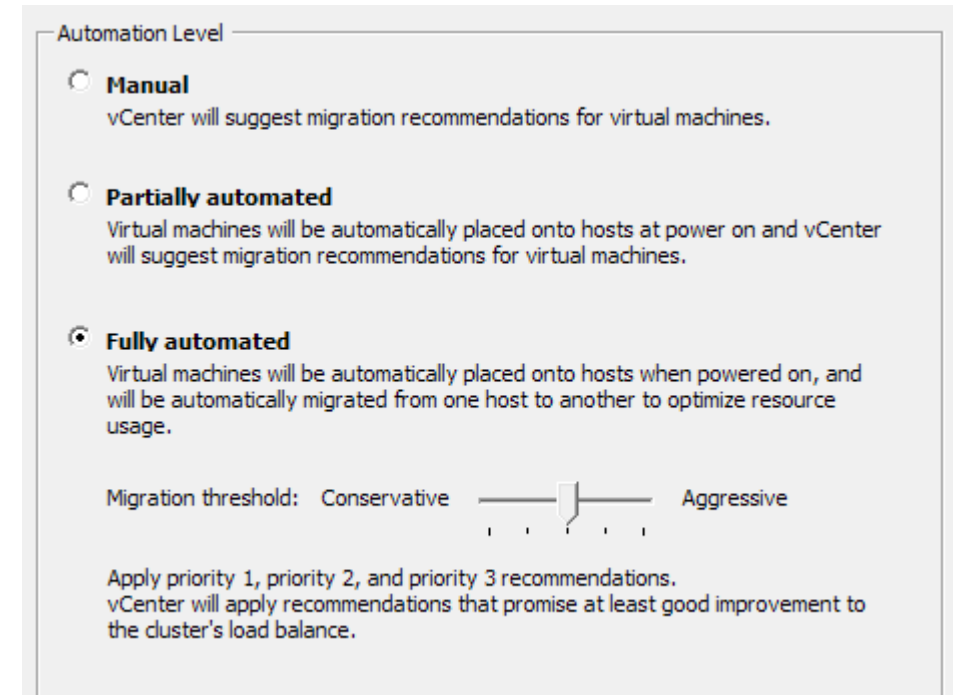
---

- 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 subs
  - 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

A large orange rectangular box with the text "SRM Demo" centered inside in white font.

SRM Demo



Summary



# 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**