Getting more out of your IBM PureFlex system with KVM

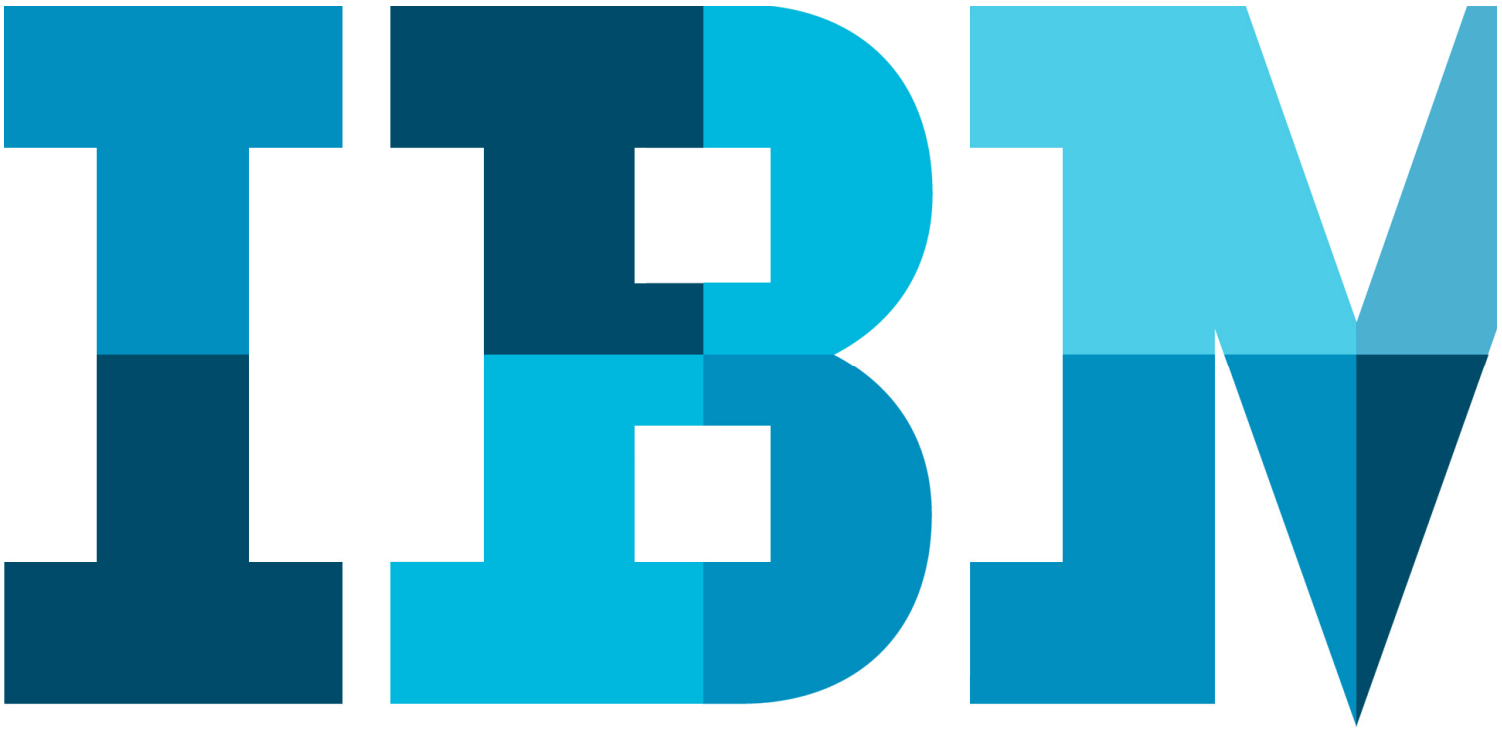

Ashish Billore, IBM VMControl and Cloud Systems Software Ashish.Billore1@in.ibm.com
Kashish Karnick, Brand Manager for IBM PureFlex® systems kkarnick@in.ibm.com
Pradipta Kumar Banerjee, IBM Virtualization Architect bpradipt@in.ibm.com
Vikas Bhardwaj, IBM Solution Architect, Lab Based Services vbhardwaj@in.ibm.com

Contents

Table of Contents

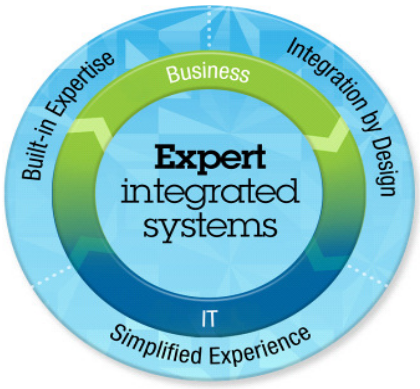

1. Introduction

PureSystems combine the flexibility of a general purpose system, the elasticity of cloud and the simplicity of an appliance. They are integrated by design and come with built in expertise gained from decades of experience to deliver a simplified IT experience. The PureSystems Family is made up of 3 family members:

1. **The PureFlex™ System** The IBM PureFlex System combines compute, storage, networking, virtualisation under a single, unified management console into an infrastructure system, expert at sensing and anticipating resource needs to optimise your infrastructure.
2. **The PureApplication™ System** The IBM PureApplication System is a platform system designed and tuned specifically for transactional web and database applications. This workload-aware, flexible platform is designed to be easy to deploy, customize, safeguard and manage. Whether you operate in a traditional or private cloud environment, this IBM solution can provide you with superior IT economics.
3. **The PureData™ System** As today's big data challenges increase, the demands on data centers have never been greater. PureData System, the newest member of the PureSystems family is optimised exclusively for delivering data services to today's demanding applications with simplicity, speed and lower cost.

The IBM PureFlex System is a subset of PureApplication System, as shown in the figure on the right.

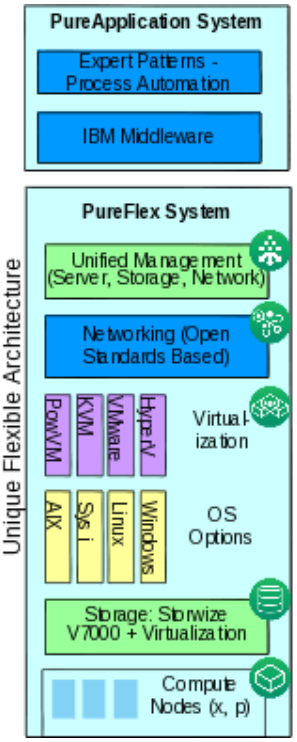

The PureApplication Systems comes with the IBM Middleware stack, which includes the WebSphere® Application Server, Tivoli® Management, DB2®, and Rational® tools.

The real value of the PureApplication System lies in Patterns of Expertise that have process automation built into the system. These capabilities include patterns like Web Application Code, Database Application Code, Data Mart Code, Workload Management, and Metering /License Management.

In fact, these patterns not only allow us to create virtual appliances, which are fine tuned with the help of IBM Expertise, but also with the capability of bringing in expertise from the entire IT community. The IBM PureSystems Center includes 100+ global independent solution vendors (ISVs) who have created appliances on the IBM PureSystems Family.

The first member of the family of PureSystems is the IBM PureFlex System. It is a comprehensive infrastructure system that provides an expert integrated computing system, combining servers, enterprise storage, networking, virtualization, and management into a single infrastructure. Its built-in expertise enables organizations to simply manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management. These systems are ideally suited for customers interested in a system that delivers the simplicity of an integrated solution, but who also want control over tuning middleware and the run-time environment.

IBM invested over $2B in delivering this system, created from the ground up, so as to mitigate IT complexity without compromising the flexibility to tune systems to the tasks that businesses demand. By providing both flexibility and simplicity, IBM PureFlex System provides extraordinary levels of IT control, efficiency, and operating agility that enable businesses to rapidly deploy IT services at a reduced cost. Moreover, the system is built on decades of expertise, enabling deep integration and central management of the comprehensive, open-choice infrastructure system and dramatically cutting down on the skills and training required for managing and deploying the system.

The objective of this paper is to showcase the capabilities of open standards based virtualized infrastructure for achieving a resilient and technologically advanced infrastructure, and delivering a low cost per virtual machine (VM). We accomplish this by leveraging IBM PureSystems as a base platform for

virtualization with various capabilities around server, storage, network, and virtualization management. From an open source perspective, we use Linux Kernel-based Virtual Machine (KVM) hypervisor available as part of Red Hat Enterprise Linux Server.

1.1 Description

An environment where virtual machines can be seamlessly migrated, while maintaining all network settings and securities, can be labeled as a flexible and dynamic virtual environment. We make this statement because today, most virtualization implementations are static in nature, due to the high level of complexity involved in a virtualized IT setup. In fact, this complexity is amplified by scaling up; e.g. managing 40 VMs vs. 400 VMs is a completely different experience. Complexity increases as the number of VMs per applications, VLANs, operating systems and storage LUNs increase. Traditionally, switches cannot see the IP addresses of virtual servers, which have led to many issues in managing virtual machines. These issue include allocating more network ports per physical hosts than needed, hypervisor overhead due to virtual switch, and, most importantly, transferring control of network policies from the network administrator to the server administrator.

These issues have led to static virtual infrastructure because it was complex to manually move virtual machines. Today, task automation is a highly complicated scripting job, which is inflexible to a company's growth needs. Thus, as infrastructure grows and complexity increases, automation using scripts and other tools becomes even more expensive, non-replicable, and complicated. The reality is that there are not that many tasks that need to be done often. Twenty percent of administrator tasks take up to eighty percent of the management time.

The IBM PureFlex system is able to manage server, storage, networking, and virtualization from a single interface. The main component of the IBM PureFlex System is an appliance called the Flex Systems Manager (IBM Flex Systems Manager). The IBM Flex Systems Manager is an x86 compute node that consists of best-of-breed IBM systems management software pre-built in a plug-and-play appliance. Its broad functional capabilities include managing the entire hardware infrastructure within the PureFlex System. But the greatest advantage offered by the PureFlex system is its advanced virtualization capabilities.

1.2 Purpose

The primary purpose of this document is to describe certain scenarios which showcase the potential of IBM PureSystems and Linux Kernel-based virtualization (KVM).

This paper discusses some of the **advanced capabilities** of PureFlex with KVM. More detailed information covering the basic aspects of KVM virtualization management using Flex System Manager™ (IBM Flex Systems Manager), configuration, etc. can be found in the reference material.

We have used IBM PureFlex system running Red Hat Enterprise Linux KVM as the basis for this paper.

This document is not meant to act as a detailed step-by-step setup guide.

1.3 Suggested Reading

The documents listed here provide additional details, including detailed configuration steps.

Item	Document Name	Description	File Location
1.	IBM PureSystems	*Provides an overview of IBM PureSystems*	http://www-03.ibm.com/systems/pureflex/express/index.html
2.	Flex System Manager overview		http://www-03.ibm.com/systems/flex/systems-management/index.html
3.	Flex System Manager VMControl		http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.vim.helps.doc/fsd0_vim_main.html
4.	Storwize V7000		http://www-03.ibm.com/systems/storage/disk/storwize_v7000/index.html
5.	Server system pool	*Describes Server System Pool*	http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_c_learnmore_getting_started_system_pools.html
6.	IBM & KVM Virtualization		http://www.linux-kvm.org/page/Main_Pagehttp://www-03.ibm.com/systems/virtualization/infrastructure/open/
7.	Implementing Systems Management of IBM PureFlex Systems	*Describes IBM Flex Systems Manager functionalities in detail along with config examples*	http://www.redbooks.ibm.com/abstracts/sg248060.html
8.	IBM VMready®	*Describes VMReady configuration examples*	http://www.redbooks.ibm.com/abstracts/sg247985.html
9.	Configuring KVM for VEPA mode		http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaai%2Fvswitch%2Fliaaivswitchlldpad.htm
10.	Automating tasks	*Describes task automation using IBM Flex Systems Manager*	http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.automation.helps.doc/fqm0_t_ea_automating_tasks.html
11.	IBM Flex Systems Manager CLI		http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.8731.doc%2Fusing_the_cli.html

Table 1: Suggested Reading

2. KVM and IBM Flex Systems Manager Resource Pools

Pooling is a resource management term that refers to the grouping together of resources for the purposes of maximizing advantage and/or minimizing risk to the users. A System pool could consist of server system pool, storage system pool, or network system pool. IBM Flex Systems Manager provides two types of pools for KVM, i.e. server system pool and network system pool.

The following sections describe each type of pool in detail.

2.1 Server System Pools

A server system Pool is a logical grouping of hosts with similar characteristics*

Server system pools enable one to group similar hosts. The grouping of similar hosts is one of the building blocks for cloud.

You can use server system pools to do the following:

1. Group hosts with same security policy (Linux iptables, ebtables or selinux rules)
2. Group hosts with same power savings policy (Linux tuned profiles)
3. Group hosts based on the type of connected shared storage (NFS, SAN)
4. Group hosts with similar hardware configuration – high speed CPU, or having SSDs, etc.
5. Group hosts based on the connected storage functionality, for instance:
    1. Hosts connected to a Storwize® V7000 pool using the Easy Tier® function
    2. Hosts connected to a Storwize V7000 pool with mirrored volumes
    3. Hosts connected to a Storwize V7000 pool with thin-provisioned volumes

IBM Flex Systems Manager server system pool provides the following capabilities:

1. Performance hot-spot detection and automatic virtual machine migration across pool members for optimization
2. Automatic virtual machine migration from a host with predictive failure alert to other pool members
3. Automatic placement of virtual machines on pool members when deploying an appliance
4. Automatic virtual machine migration from a specific host to other hosts in the pool for maintenance activities on the specified host

Note that virtual machine (VM) and virtual server (VS) are used synonymously in this paper.

* The characteristics like security policy, power savings policy etc needs to be configured separately, outside of IBM Flex Systems Manager.

2.2 Network System Pool

Network system pools (NSP) simplify and automate network configuration tasks for virtual machines. One can manage the network connections of the pooled network systems to ensure network connectivity across a set of network switches. Following are some of the features provided by NSP:

- Automatic network port configuration (VLANs)
- Automatic MAC address migration
- Automatic Layer 2 profile migration (VLANs)

Restrictions:

- MAC addresses within a network system pool must be unique.
- Link aggregated ports are not supported in a network systems pools environment because of limited Link Aggregation (LAG) VLAN deployment support.

Consequently, you cannot use VEPA (802.1Qbg) with NIC (Network Interface Card) teaming using IBM Flex Systems Manager. You must configure this outside of IBM Flex Systems Manager.

Network system pools, combined with server system pools, provide flexibility and control over how network resources are used. An administrator can:

- Define larger network system pools to allow more efficient use of network resources
- Define logical networks within a network system pool for shaping and isolation purposes
- Define a server system pool with some or all of the servers managed by a network system pool

3. KVM and VM Aware Networking

3.1 VM aware Networking with IBM VMready

IBM VMready provides the ability to configure virtual machine networking at the physical switch layer providing a seamless interface for configuring both physical and virtual server networking. In a KVM environment, VMready ensures network policies follow virtual machines as they migrate from one physical switch port to another using a technology called NMotion. This makes the manual configuration of each individual switch ports to cater for virtual machine networking unnecessary. Once the VM-Aware network is configured by the network administrator, the VMready switch continues to ensure that connectivity and appropriate network policies are enforced for both virtual and physical servers. VMready is available on the PureFlex10Gbe EN4093 chassis switch and other PureFlex switches.

VMready provides enhanced granularity to physical switch network capabilities. It extends the capability of the switch from specifying network parameters at the physical port level down to the virtual machine level.

The following functionalities are available when configuring parameters per virtual machine port:

- VLAN membership
- Traffic shaping and monitoring
- Access Control Lists (ACLs)
- Quality of Service (QoS) attributes

3.2 VM aware Networking with 802.1 Qbg/VEPA (Virtual Ethernet Port Aggregator)

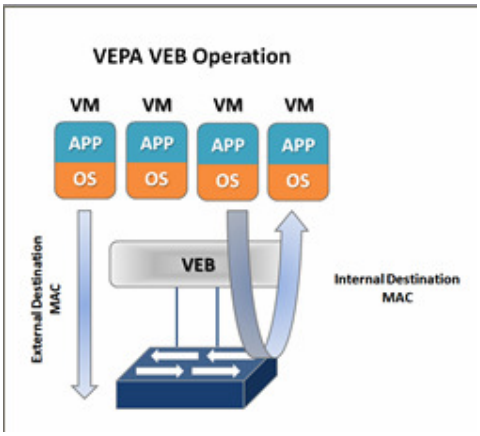

Figure 1: VEPA Overview

Normally, the VMs running on a specific host communicate with each other and to the outside world using a virtual switch or vswitch. The vswitch is, in effect, a Layer 2 switch, running within the hypervisor. Every hypervisor typically has a virtual switch built in.
The virtual switch moves networking into the server realm, bringing with it the need to re-test, re-qualify and re-deploy traditional network based tools and solutions for the virtualized environment

VEPA is an alternate to the traditional vswitch. VEPA simply forces VM traffic to be handled by an external switch. The external network switch in turn provides connectivity between the virtual machines on the same physical server as well as to the rest of the infrastructure.
This allows each VM frame flow to be monitored managed and secured with all of the tools available to the physical switch. This allows for things like flow statistic gathering, ACL enforcement, etc. Additionally this also frees the host resources from network processing.
VEPA in effect moves switching out of the server and back into the physical network and makes all virtual machine traffic visible to the external network switch. By moving virtual machine switching back into the physical network, a VEPA based approach makes existing network tools and processes work consistently across both virtualized and non-virtualized environments as well as across hypervisor technologies.

VEPA requires hypervisor support.

IBM PureFlex 10Gbe EN4093, EN4093R and the CN4093 switches are 802.1Qbg compliant. Coupled with the KVM hypervisor and IBM Flex Systems Manager, these switches provide a seamless way to leverage 802.1 Qbg for VM aware networking.

4. KVM and Storwize V7000

You can leverage Storwize V7000 capabilities in a PureFlex environment to provide for storage features like thin provisioning, mirroring, higher IOPS, replication, backup, disaster recovery etc for KVM.

The following are some scenarios in which you can use KVM and Storwize V7000: Create a KVM VM using a local mirrored volume for backup and recovery

1. Create a KVM VM using a thin-provisioned volume for optimal storage usage
2. Create a KVM VM using a remote replicated volume for disaster recovery
3. Create a KVM VM using an easy-tiered volume when high IOPS are required
4. Create a server system pool with associated easy-tiered enabled storage pool. VMs requiring high-IOPS can be provisioned on this server system pool.
5. Create a server system pool with associated storage pool having thin-provisioned volumes. Less critical VMs can be provisioned on this server system pool.
6. Create a server system pool with associated storage pool having mirrored volumes. VMs requiring backup can be provisioned on this server system pool.

The possibilities are numerous.

5. Automating Tasks

One of the advantages of a systems management solution is the ability to perform some tasks automatically. IBM® Flex System Manager accomplishes this with two features: Scheduler and event automation plans.

5.1 Scheduler

Any task or command that needs to be executed periodically can be scheduled using the Scheduler component.

Following are some examples:

- Run a backup script on KVM host machines to backup VM configuration
- Power off nodes during off-peak hours
- Collect inventory of nodes

5.2 Event Automation Plans

In IBM Flex System Manager, event automation plans are built around events. An event is any occurrence that changes the system or its components in some way. Hardware events, like fan events, are automatically created and sent to IBM Flex System Manager. Other custom events, like exceeding a certain CPU or memory threshold, can be created by setting thresholds.

Using automation, you can have IBM Flex System Manager send an alert when the memory capacity threshold on the server is exceeded. It can send an e-mail when a fan stops running, so that the fan can be replaced before it affects a critical workload. These automated responses to events are called event actions.

Event automation plans handle more complicated tasks than the Scheduler component because event automation plans perform the tasks in response to an event. An event automation plan begins its tasks based on the occurrence of events.

Some examples of automation plans to monitor and respond to typical situations in a PureFlex environment:

- Automatically migrate virtual machines from a host on receiving hardware predictive failure alert from the host
- Automatically migrate virtual machines from a host in case of critical hardware events from the host
- Migrate virtual machines from a specific host to another host if host resource utilization exceeds a specific threshold.
- Send an e-mail to the administrator in case of resource utilization (CPU, memory, or disk) for a specific node crosses a defined threshold.

5.3 Automation by leveraging IBM Flex Systems Manager CLI

IBM Flex Systems Manager provides a very powerful Command Line Interface (CLI) which an administrator can use either standalone or as a scriptable framework to automate management tasks.

Following are some examples with IBM Flex Systems Manager command (smcli <cmd>) and pseudo-code

- *Specify security policy for members of a specific system pool*

  Get the member hosts for a specific system pool and execute specific script/command on the system pool members

```
pool_list = list_system_pools()    # use smcli lssyspool
system_pool = select_the_specific_system_pool(pool_list)
host_list = parse_syspool_output(system_pool)   # use smcli lssyspool -p
for host in host_list:
        # run script containing specific iptables rule sets on all host members
        set_security_policy(iptables_rules_script)
```

- *Specify specific energy savings policy for members of a specific system pool*

  Get the member hosts for a specific system pool and execute specific script/command on the system pool members

```
pool_list = list_system_pools()    # use smcli lssyspool
system_pool = select_the_specific_system_pool(pool_list)
host_list = parse_syspool_output(system_pool)   # use smcli lssyspool -p
for host in host_list:
        # run 'tuned-adm activate powersave|default|performance' on all host members
        set_powersave_energy_policy()
```

- *Specify resource control settings for a specific VM*

  Get the host running the specific VM and execute specific script/command on the host

```
get_resource_control_values_for_vm()
host = get_host(vm)    # use smcli lsvrtsys -l
power_down(vm)
# run script changing the VM domain xml values on the specific host
set_resource_control()
power_up(vm)
```

- *Create an n-tier solution workload from existing set of captured appliances*

```
va_list = list_all_va()  # use smcli lsva -o
va_workload_list = select_va()
```

```
pool_list = list_system_pools()    # use smcli lssyspool
system_pool = select_the_specific_system_pool(pool_list)
for va in va_workload_list:
      deploy_va_list =  deploy_va(va, system_pool)  #  use smcli deployva
create_new_workload(deployed_va_list)
```

6. Working with Server System Pools and Network System Pools

The following section gives a brief overview of server system pools and network system pools handling in IBM Flex Systems Manager. More details can be found in the materials mentioned in the references.

6.1 Server System Pool Creation

You can create KVM server system pools by going to the IBM Flex Systems Manager VMControl main window and clicking **Server system pools**.

Figure 2: IBM Flex Systems Manager VMControl Home Page

The above figure (Fig 2) shows the VMControl plugin home page and the various options available. Select “System Pools” option to work with System Pools in VMControl.

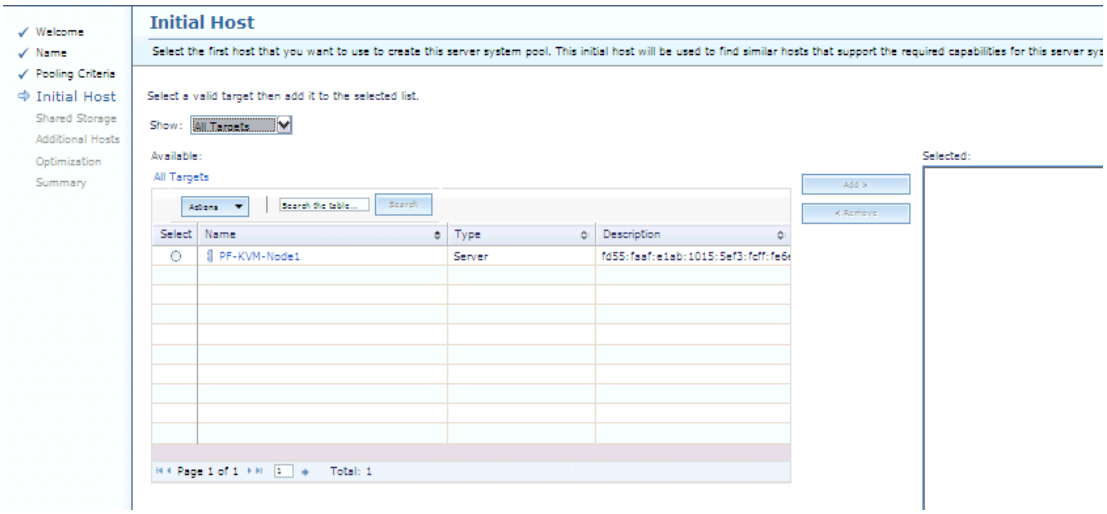

Figure 3: IBM Flex Systems Manager VMControl Server System Pool Creation

As shown in Figure 3 above, the Server System Pool creation wizard will guide you through steps and details needed to create a Server System Pool. Select the initial host to base you Server System Pool on, subsequent members of the Server System Pool should be similar to the initial host.

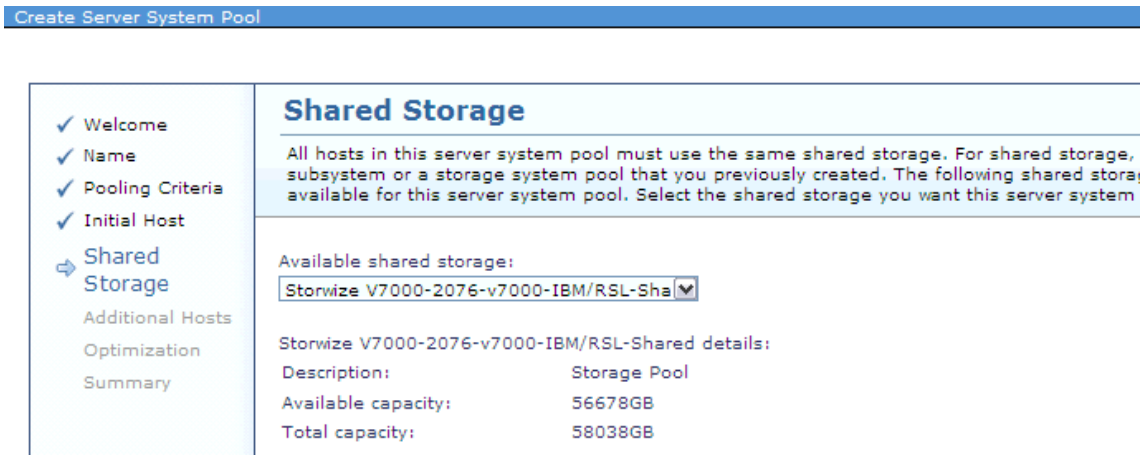

Figure 4: Assign Storage Pool to the Server System Pool

Once the initial Host is selected for the Server System Pool, as shown in Figure 4, select the “Shared Storage” available for this Server System Pool. This storage must be shared by all the member hosts of the Server System Pool.

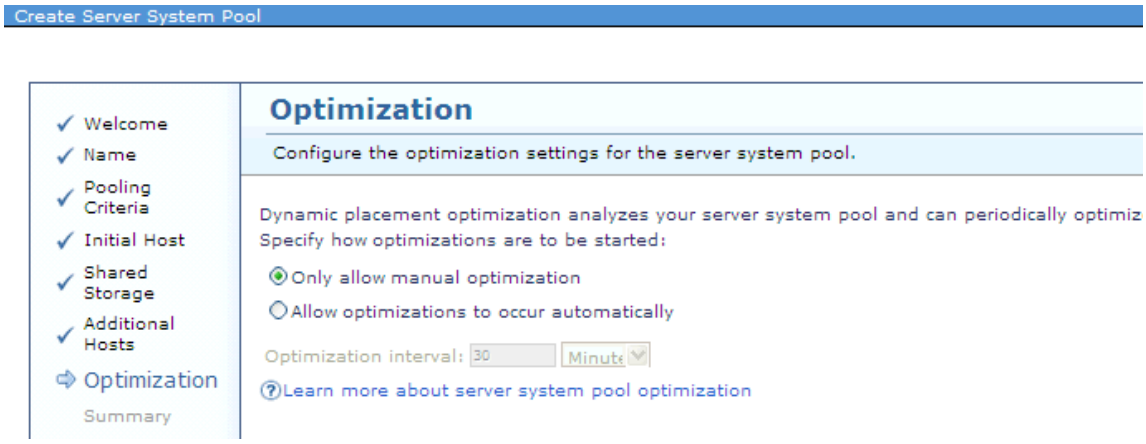

Figure 5: Server System Pool Optimization Settings

Specify the “Optimization” options for the Server System Pool. You can choose to have manual optimization or automatic optimization for your Server System Pool with specified “Optimization Intervals”, as shown in Figure 5.

6.2 Create VEPA Logical Network Profile and Attach to Network System Pool

Details on configuring IEEE 802.1 Qbg can be found here
http://www.redbooks.ibm.com/abstracts/sg247985.html

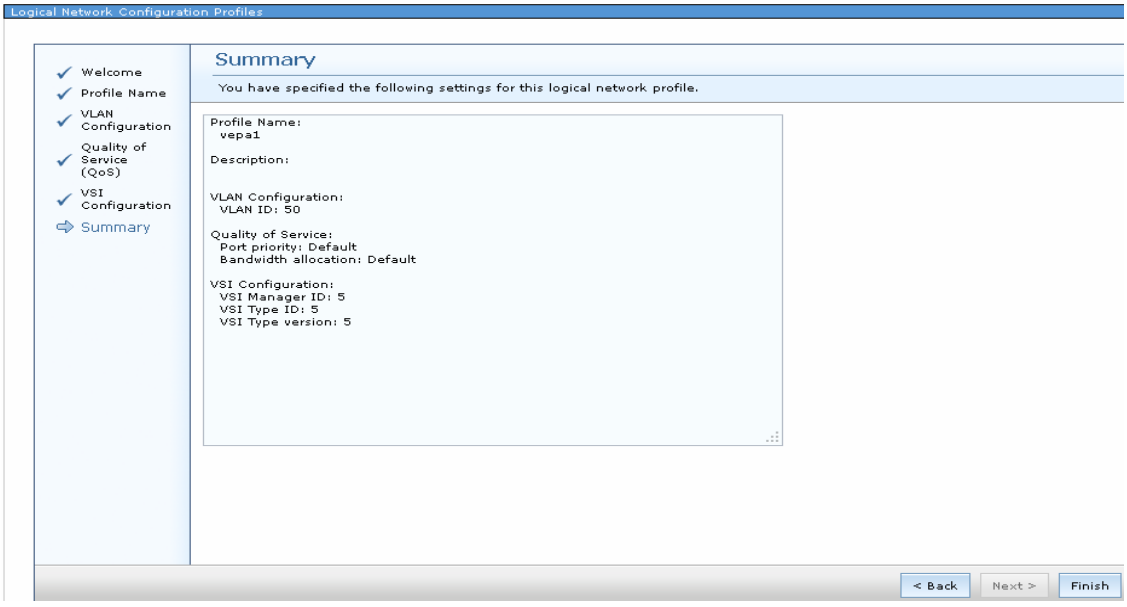

Figure 6: Sample VEPA based Logical Network Profile

Figure 6 above shows the summary of a Sample VEPA based Logical Network Profile. A Logical Network Profile can be configured using the wizard shown above.

Add logical network profiles to the network system pool:

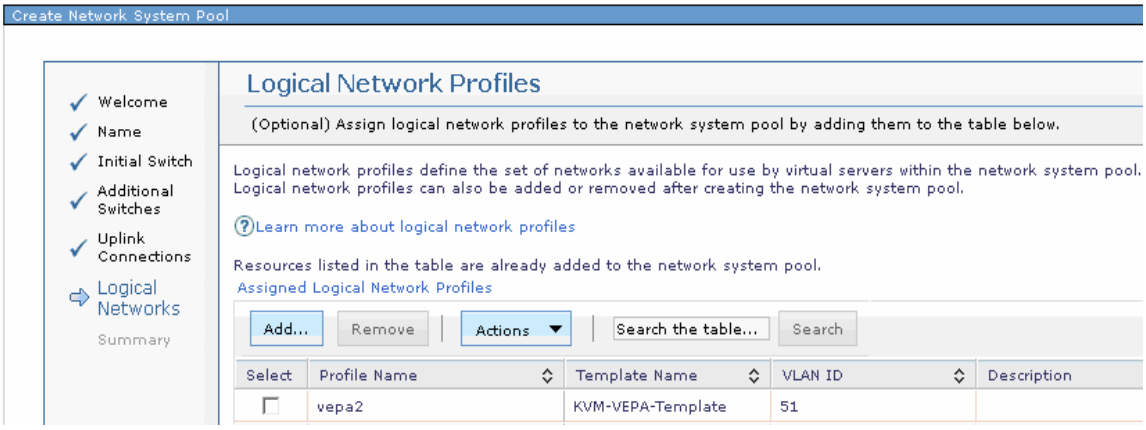

Figure 7: Make Logical Network Profile Part of Network System Pool

Figure 7 above shows how a previously defined logical network profile (as in Figure 6) can be associated with a Network System Pool, using the Network System Pool creation wizard.

6.3 Create VM with IEEE 802.1 Qbg support

Navigate to the Virtual Server and Members view and select the host and right click **System Configuration -> Create Virtual Server.**

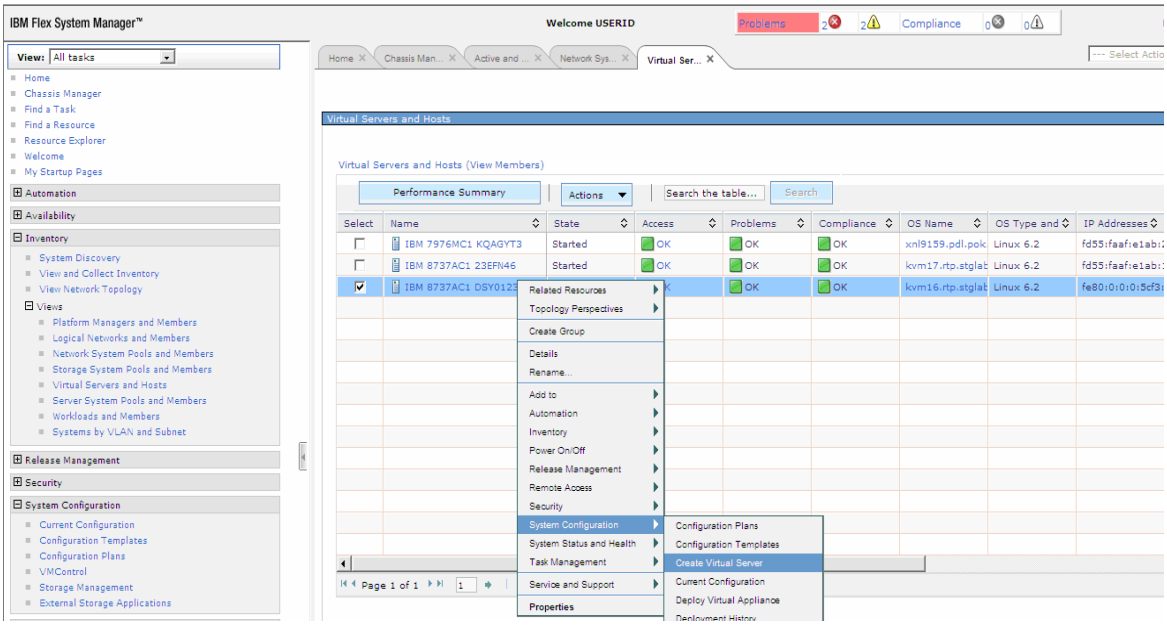

Figure 8: Create Virtual Server Wizard

Use the wizard to create a virtual server and provide the required information.
As shown in Figure 9 below, the Virtual Server creation wizard allows you to select the VEPA logical network profile created earlier.

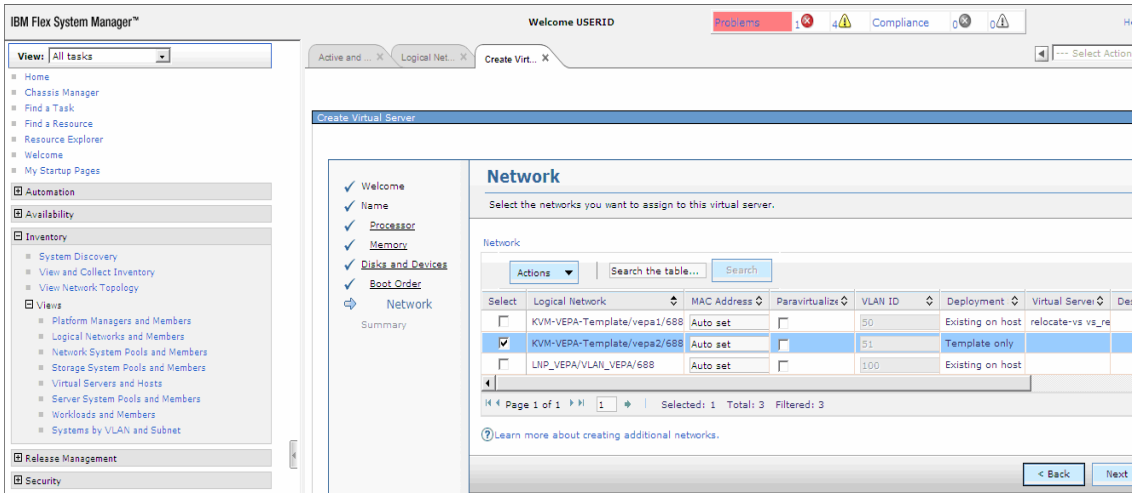

Figure 9: Select VEPA profile for the VM

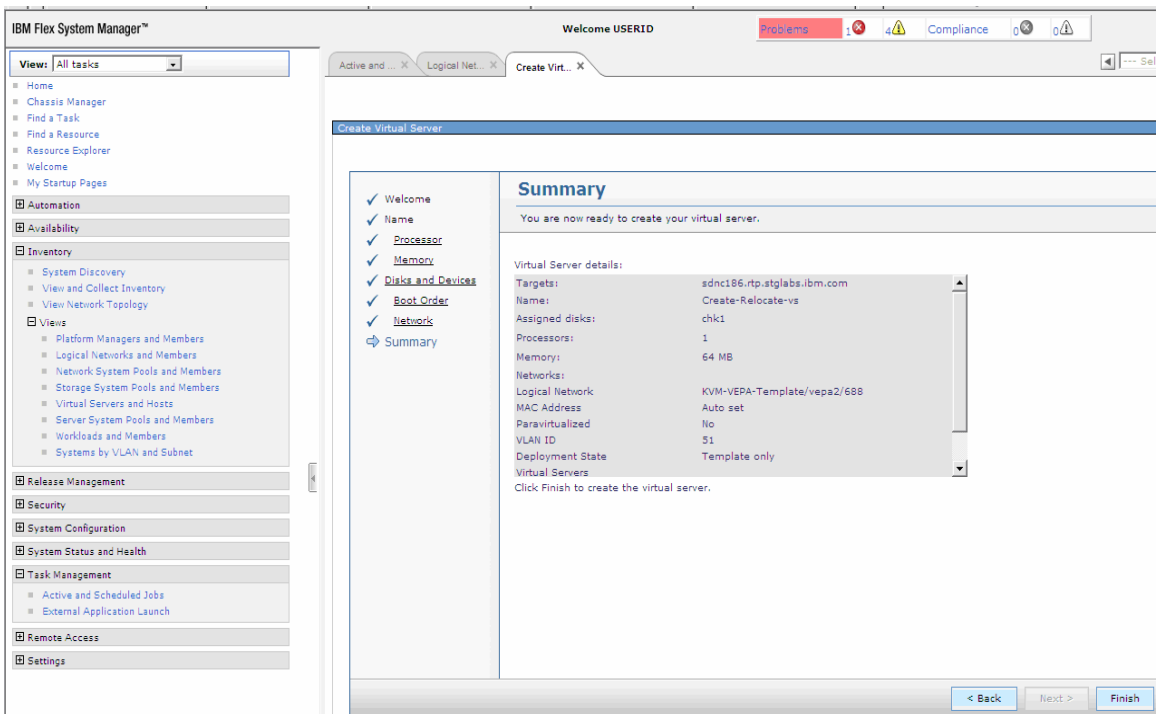

Figure 10: Create VM Summary Screen

Figure 10 above lists the summary of options selected for the virtual server in the Virtual Server creation wizard for reference.

```
<interface type='direct'>
  <mac address='06:f2:19:00:00:04'/>
  <source dev='eth1.51' mode='vepa'/>
  <virtualport type='802.1Qbg'>
    <parameters managerid='5' typeid='5' typeidversion='5'
```

Figure 11: VM Description File Containing VEPA Attributes

As shown in Figure 11 above, the VEPA attributes are associated with Virtual Server created using the virtual server creation wizard above (figure 9).

6.4 Sample VM Relocation with 802.1 Qbg and Preservation of all Network Policies

Select the virtual machine and right click **Availability -> Relocate**.

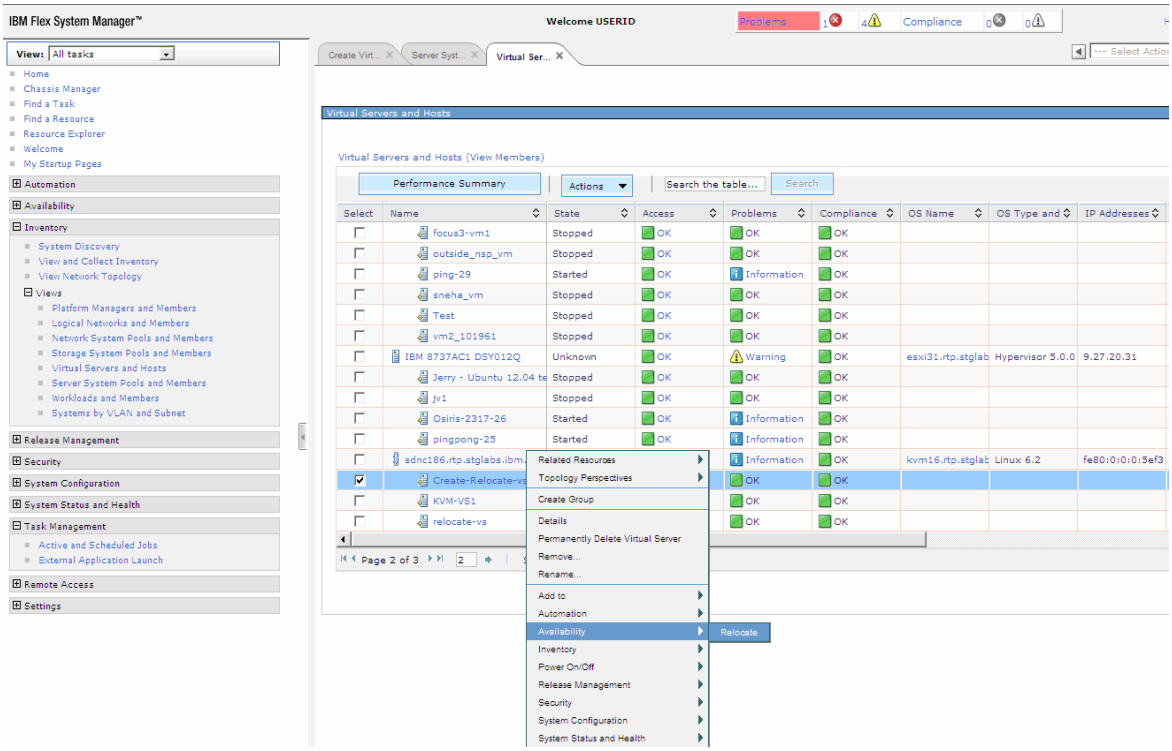

Figure 12: Relocate VM Wizard

Invoke Virtual Server relocation wizard as show in Figure 12 above.

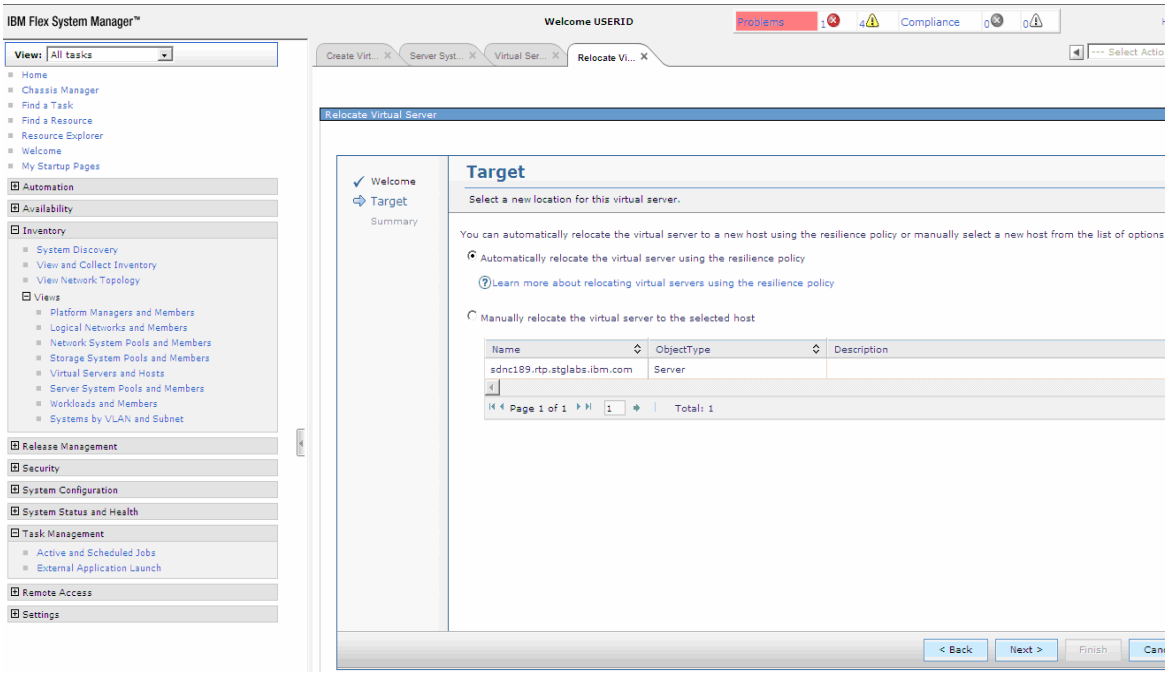

Figure 13: Automatic or Manual Selection of Relocation Target

In the virtual server relocation wizard, specify the target for relocation, this can be either automatic or chosen manually.

After relocation, you can check the VEPA setting for the relocated VM on the target host.

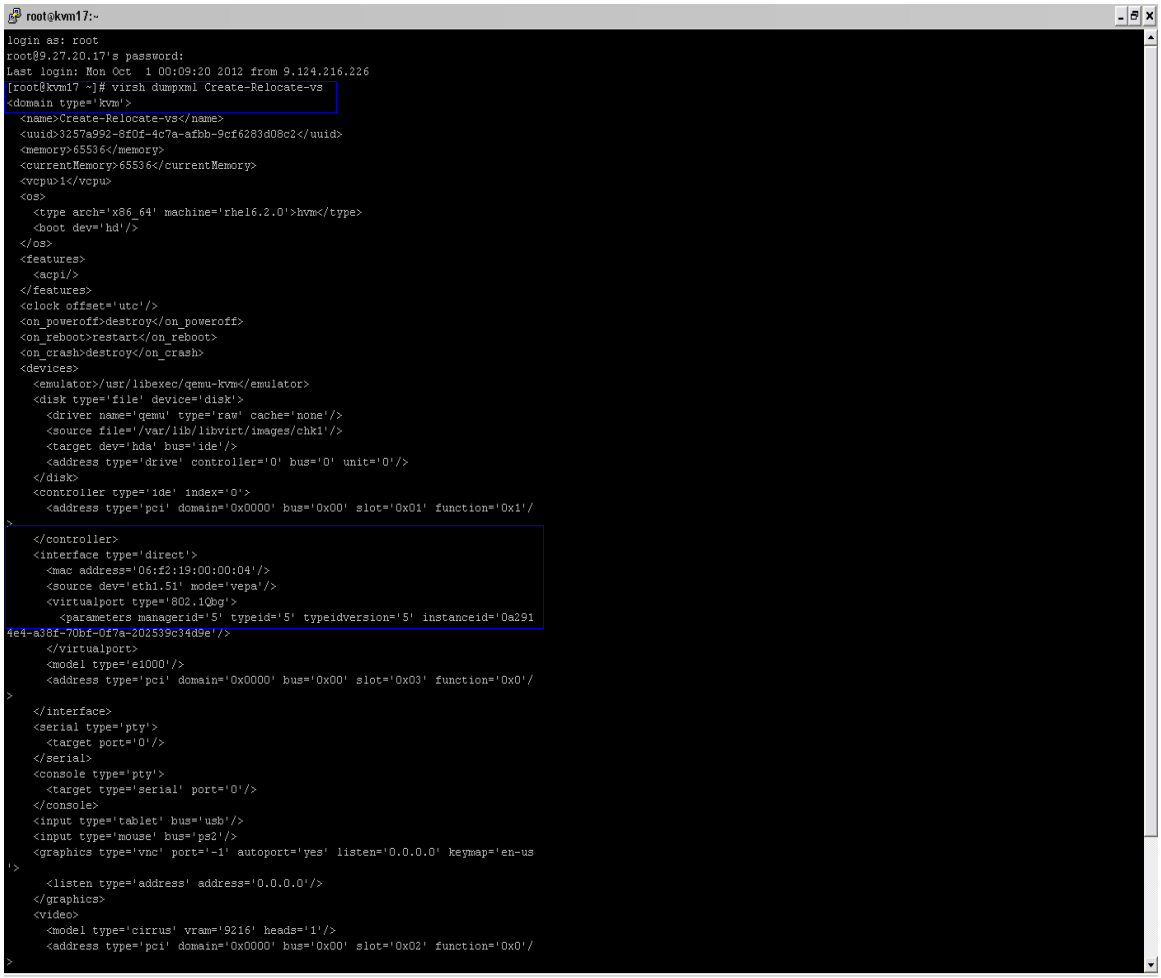

Figure 14: VM Settings after Relocation

Figure 14 above shows the virtual server settings being preserved after relocation to the new host.

7. Capture and Deployment of Virtual Machine Appliances

Using IBM Flex Systems Manager VMControl one can perform the following tasks when working with virtual appliances.

- Import a virtual appliance
- Deploy a virtual appliance to create a virtual machine
- Capture a virtual machine and create appliance or golden image for rapid VM creation
- Create a workload consisting of deployed virtual machine
- Enable resiliency policy for the workloads to ensure automatic relocation in case of any host error

7.1 Import an Existing Virtual Appliance

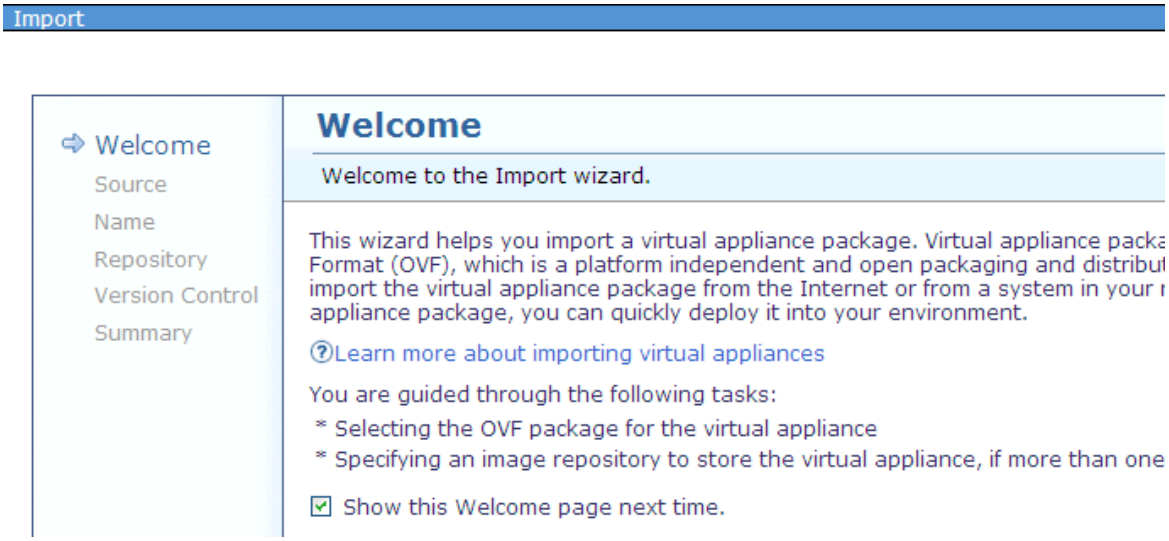

Figure 15: Import Wizard

In order to import a virtual appliance a simple import wizard of IBM Flex Systems Manager can be used. This wizard takes all the required inputs which are required to add the new virtual appliance to IBM Flex Systems Manager.

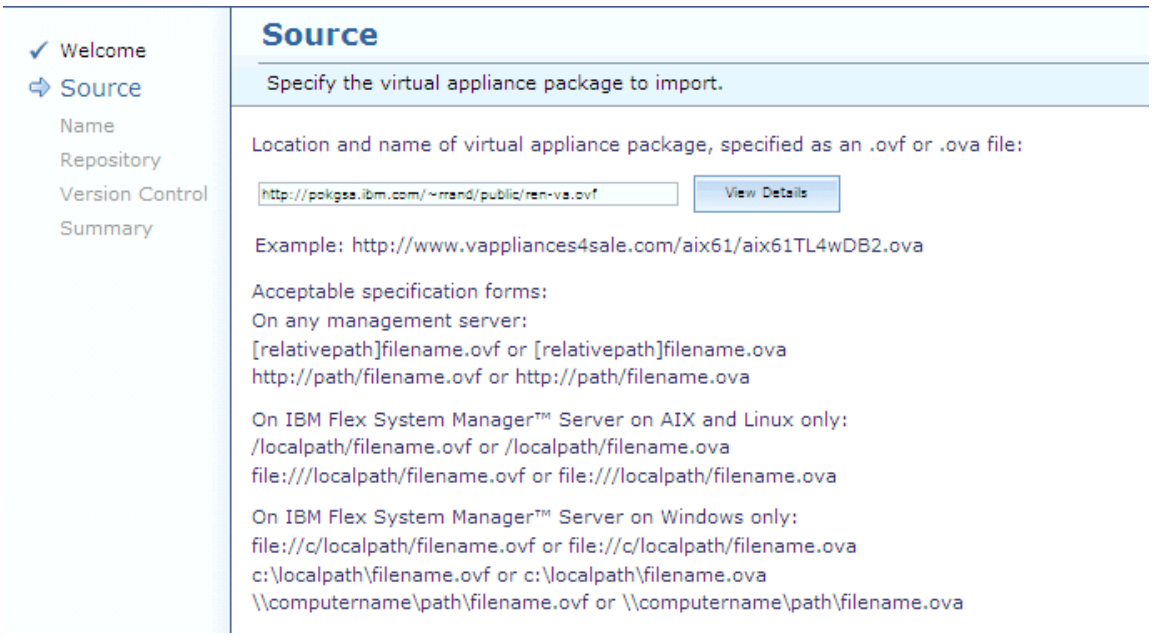

Figure 16: Specify the Source OVF file for the Virtual Appliance

Specify the location and name of the virtual appliance package as shown in the Figure 16. Relative path of the .ovf file or .ova file on a management server or location on the IBM Flex Systems Manager server are acceptable.

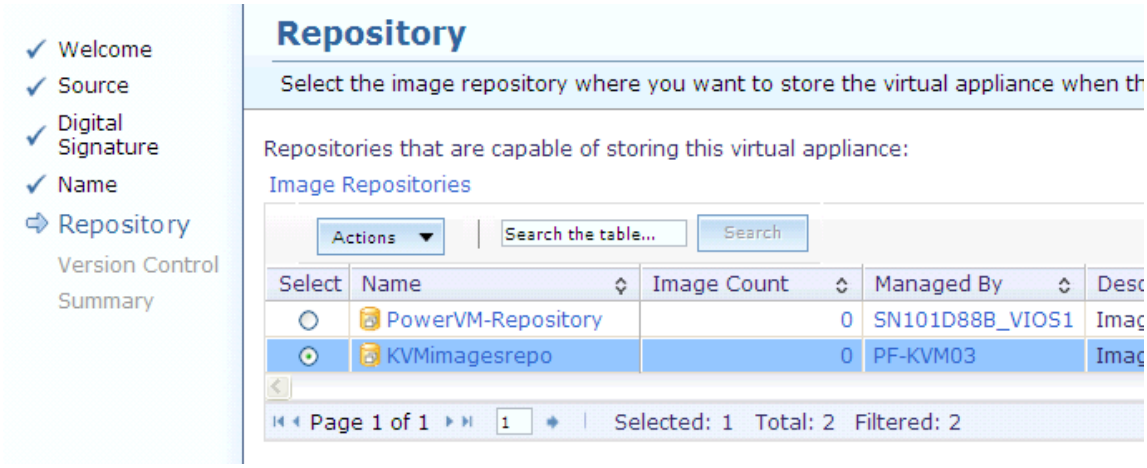

Figure 17: Select the Image Repository to Store the Imported Appliance

Specify the image repository where the virtual appliance will get stored. Import wizard lists the repositories that are capable of storing the virtual appliance.

7.2 Deploy a Virtual Appliance

You can deploy virtual appliances in the Linux KVM virtualization environment on IBM IBM Flex Systems Manager VMControl to new or existing virtual servers, or to server system pools.

Additionally, you can specify the storage to be used for VM disks.

Deploy virtual appliance wizard takes all the required inputs in order to deploy a virtual appliance. As shown in Figure 18 one can select the Virtual Appliance to be deployed.

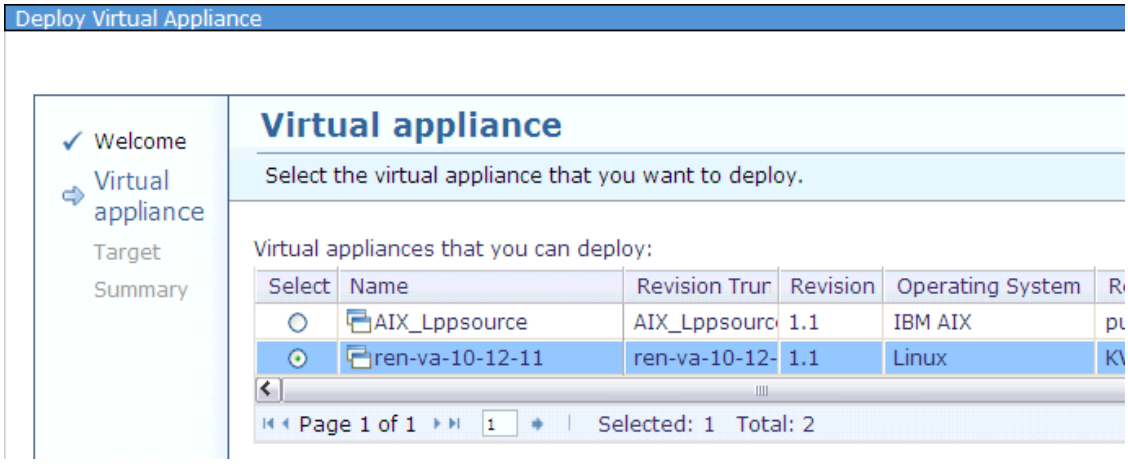

Figure 18: Deploy Virtual Appliance Wizard

Specify the target server or server system pool where the new virtual appliance should be deployed.

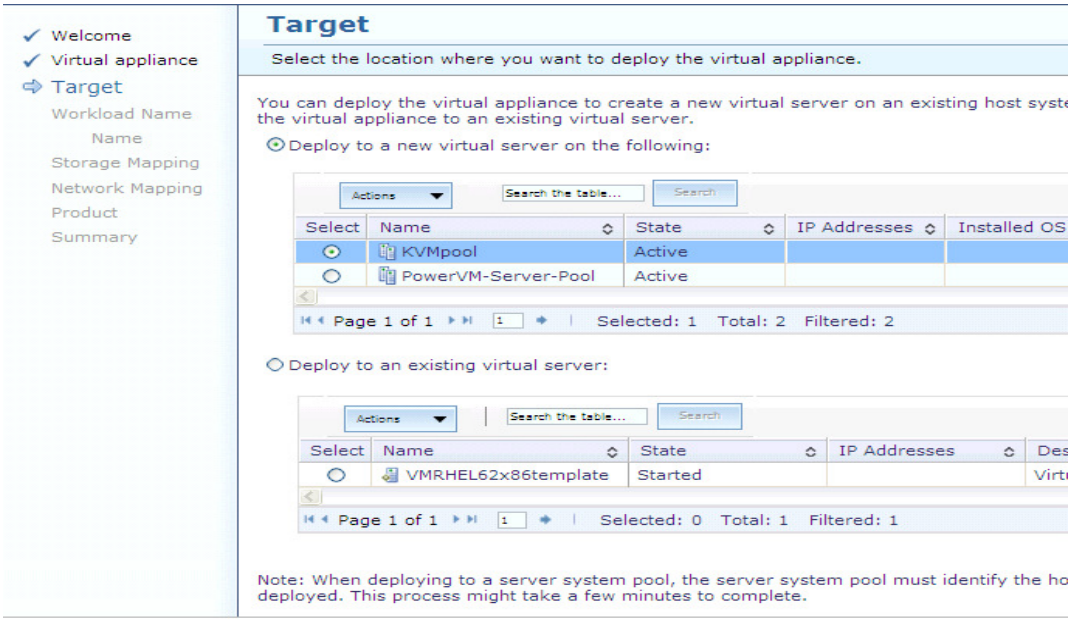

Figure 19: Specify the Deployment Target

As a next step provide the input regarding the storage volume or storage pool for the new virtual appliance.

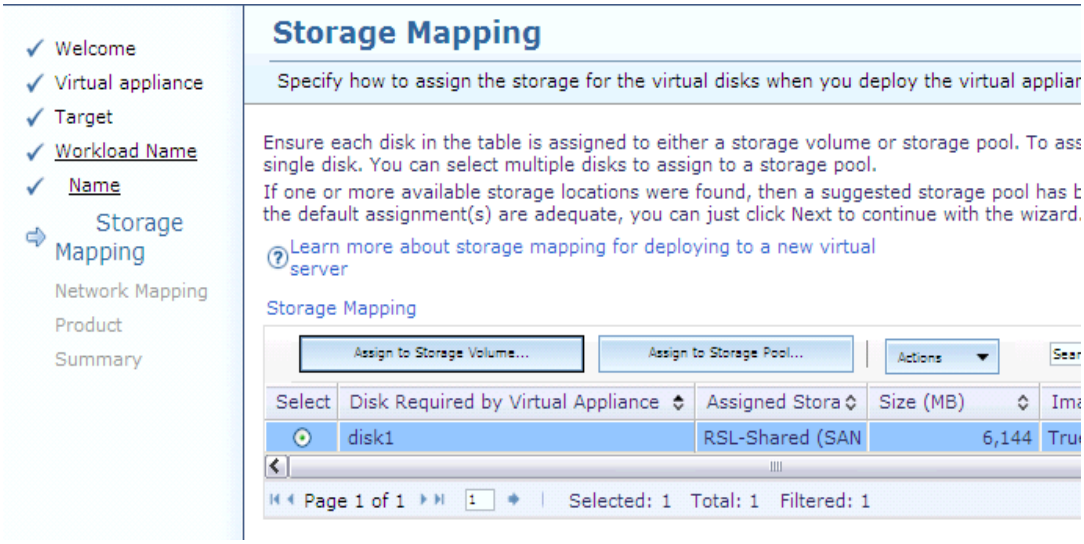

Figure 20: Specify the Storage for Virtual Disk

7.3 Capture a Virtual Machine

You can capture a virtual machine to create an appliance, version it, and store it in an image repository for later use.

Capture VM wizard facilitates capturing a virtual machine as shown in the Figure 21 you can select the VM which you want to capture.

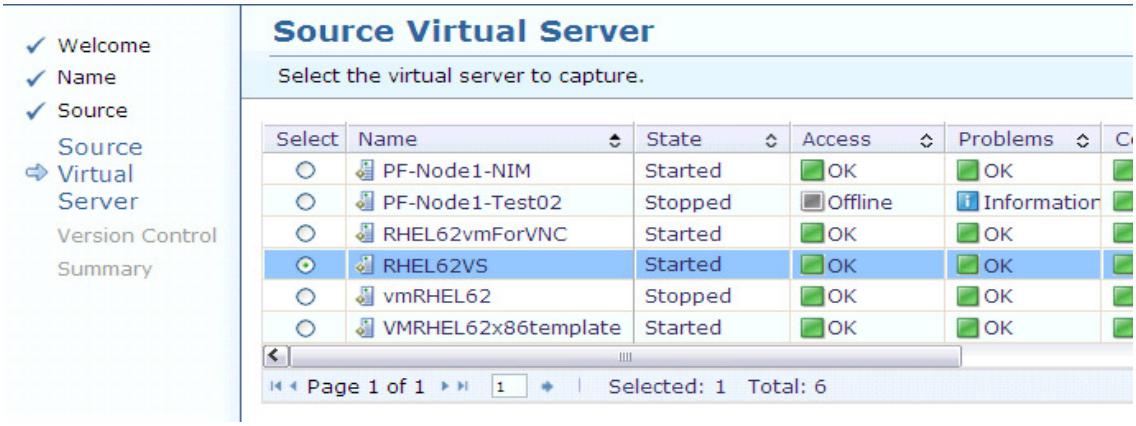

Figure 21: Capture VM Wizard

Capture VM wizard display all the image repositories where the captured image of the VM can be stored. You can specify the repository where you want to store the captured VM.

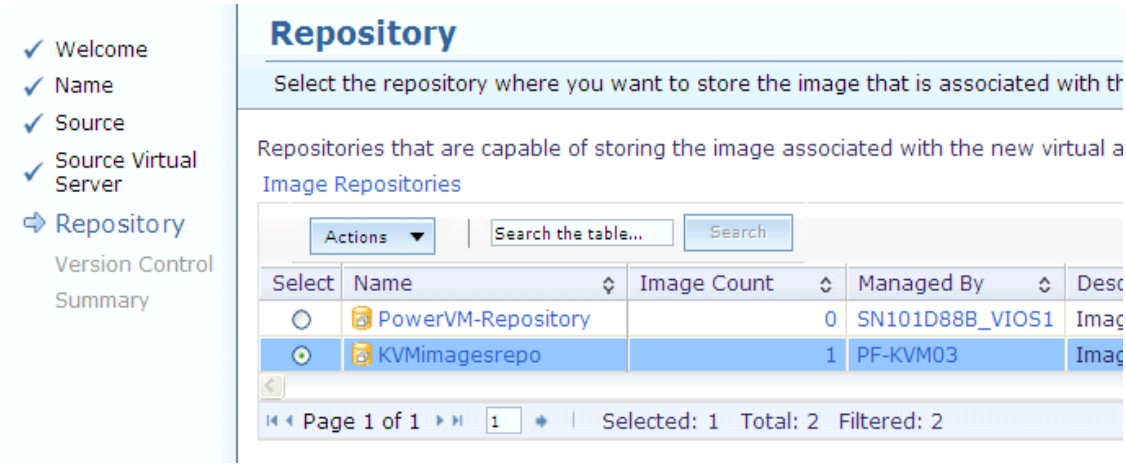

Figure 22: Select the Image Repository for Storing the Captured Virtual Server

If the virtual machine you are capturing has association with a virtual appliance from a previous deployment you can specify version information for the virtual appliance.

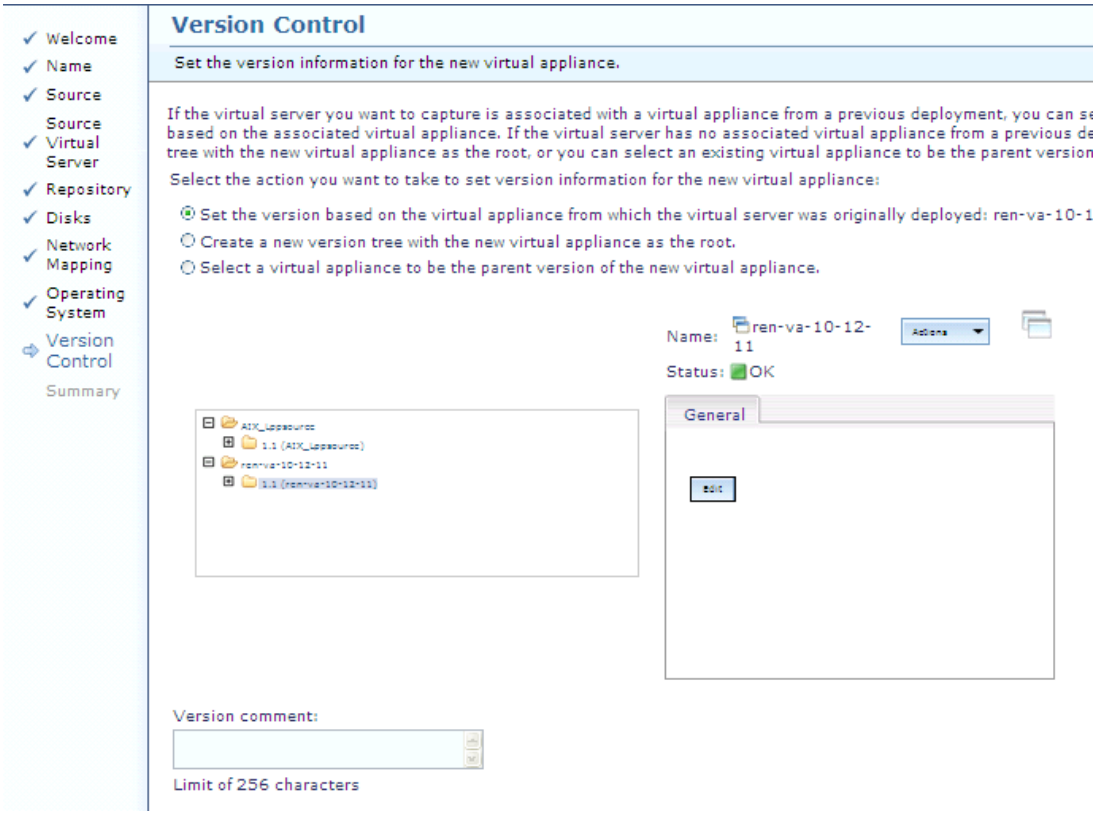

Figure 23: Specify Version Information for the Appliance

8. Working with PureFlex Storwize v7000

This section describes how one can leverage Storwize V7000 capabilities in a PureFlex environment to provide for storage features like thin provisioning, mirroring, higher IOPS, replication, backup, and disaster recovery for KVM.

8.1 Generic, Mirrored, or Thin-Provisioned Volumes

The figure below depicts the Storwize V7000 'New Volume' wizard.

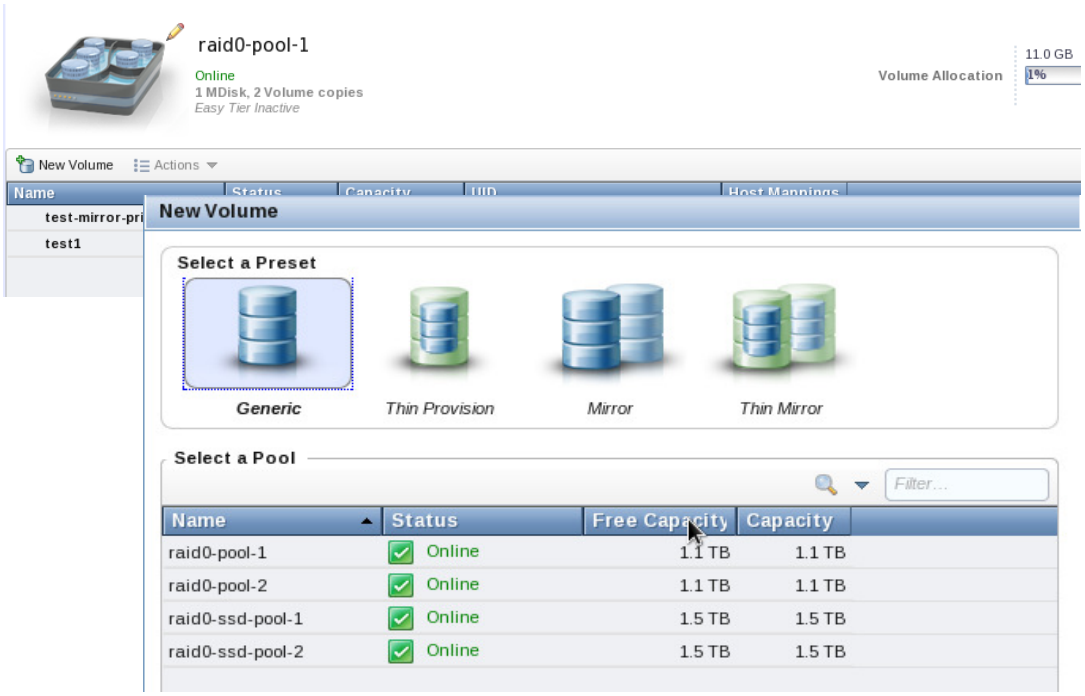

Figure 24: New Volume Wizard

Using Storwize V7000 GUI you can create a new volume by selecting New Volume option. You can select one of the following type of the new volume as shown in Figure 24

- **Generic**: A fully provisioned volume according to the RAID type of the storage pool
- **Thin provision**: A space efficient volume that grows based on usage
- **Mirror**: A synchronized copy is kept to prevent any data loss arising if the primary copy is lost
- **Thin Mirror**: Two synchronized copies, both thin-provisioned

After successfully creating the volume with required characteristics, add the same as a virtual storage for a KVM VM.

You can leverage mirror volumes for VM backup and recovery also.

Additionally, you can leverage Metro Mirror or Global Mirror storage replication features of Storwize V7000 for disaster recovery solutions.

8.2 Leverage Storage Tiering for providing better IOPS in a KVM VM

By creating a storage pool (managed disk group) with both generic SSD MDisks and generic HDD MDisks, Easy Tier is automatically turned on for pools with both SSD MDisks and HDD MDisks.

Create a KVM virtual machine and assign a volume created on an EasyTier enabled pool to get higher IOPS for the specific virtual machine.

Further, you can assign the Easy Tier enabled pool to a KVM server system pool.

Any workloads/VMs requiring higher IOPS can be deployed on this system pool.

While creating a new volume a list of configured storage pools is displayed you can select a storage pool for which easy tier is enabled which will give performance benefits for the workload deployed.

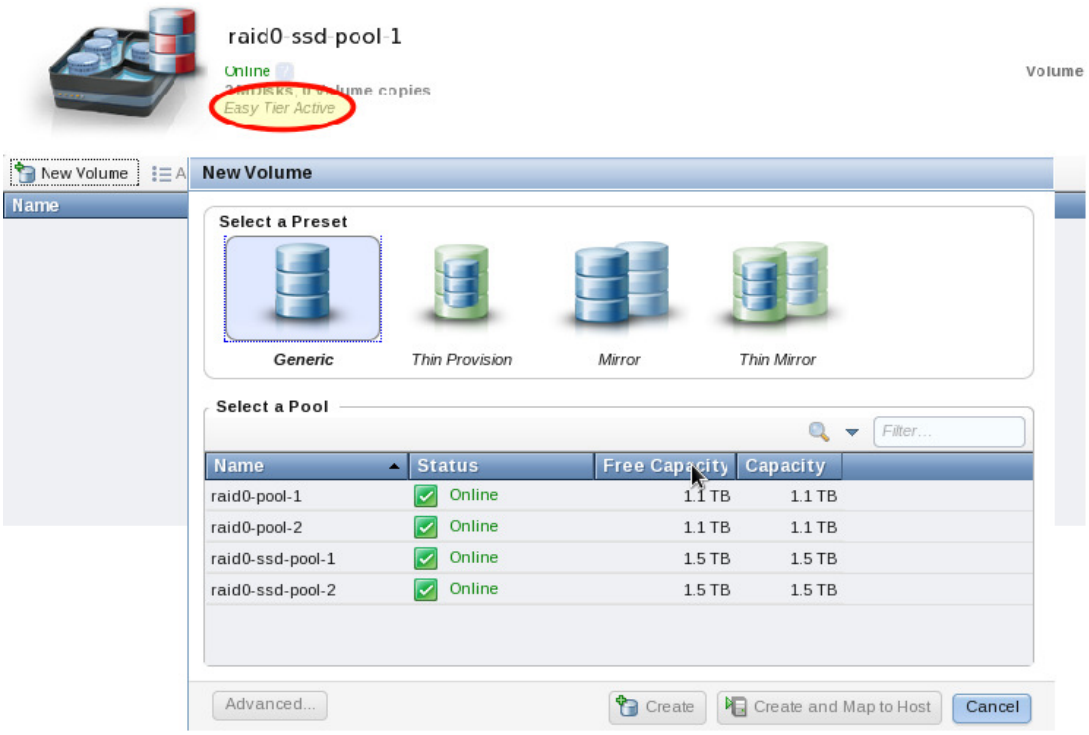

Figure 25: Create New Volume in an Easy Tier Enabled Pool

Using Flex System Manager you can assign a specific Storwize V7000 volume to a virtual machine to provide storage features benefits to the specific workload.

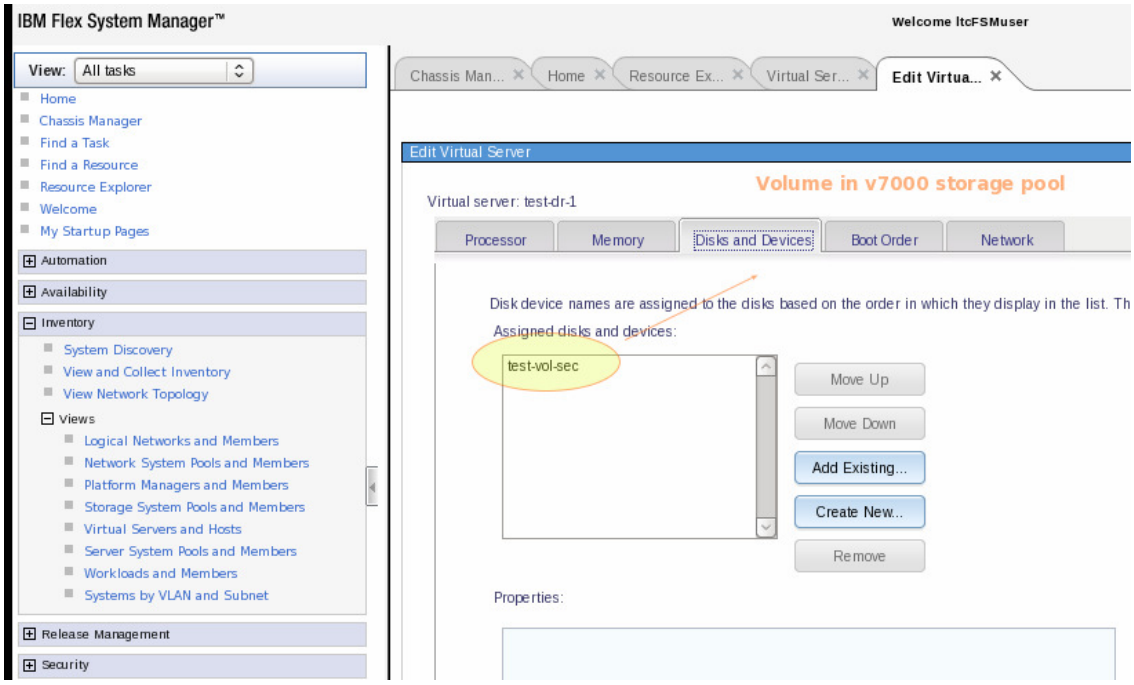

Figure 26: Assign a Specific Storwize V7000 Storage Volume to a Virtual Machine

9. Summary

The intent of this paper was to showcase the capabilities of IBM PureFlex when used in the larger context of automating the datacenter infrastructure. Many customers today are looking very hard on Linux based virtualization because of the high level of security, performance and availability this platform brings – while still keeping costs under control. By being able to get both the hypervisor and the hardware management platform together, and automating it, customers can derive a very high level of automation and flexibility in today's increasingly complex virtualized environment.

IBM home page: http://www.ibm.com

IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
November 2011