STO1965

Virtual Volumes Technical Deep Dive

Rawlinson Rivera, VMware, Inc
Suzy Visvanathan, VMware, Inc

vmworld® 2014

Disclaimer

- This presentation may contain product features that are currently under development.
- This overview of new technology represents no commitment from VMware to deliver these features in any generally available product.
- Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind.
- Technical feasibility and market demand will affect final delivery.
- Pricing and packaging for any new technologies or features discussed or presented have not been determined.

Agenda

1	Virtual Volumes: Introduction
2	High level Architecture
3	Data path
4	Managing Storage Capacity
5	Ensuring SLOs through policies
6	VM Objects – Virtual Volumes
7	Data Services
8	Making the transition to Virtual Volumes
9	Learn more

Goals of Virtual Volumes

- Ability to express application (VM/VMDK) granular data services
- Provide easy on-demand Capacity provisioning
- Compliance Monitoring
- Ability to get most out of the storage system

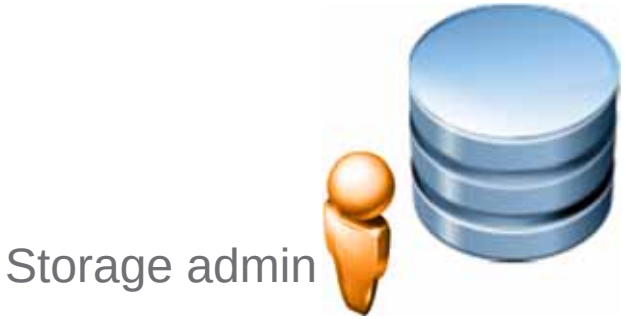

- Easy Capacity management
- Meet VM SLOs
- Access Control and Security

High Level Architecture

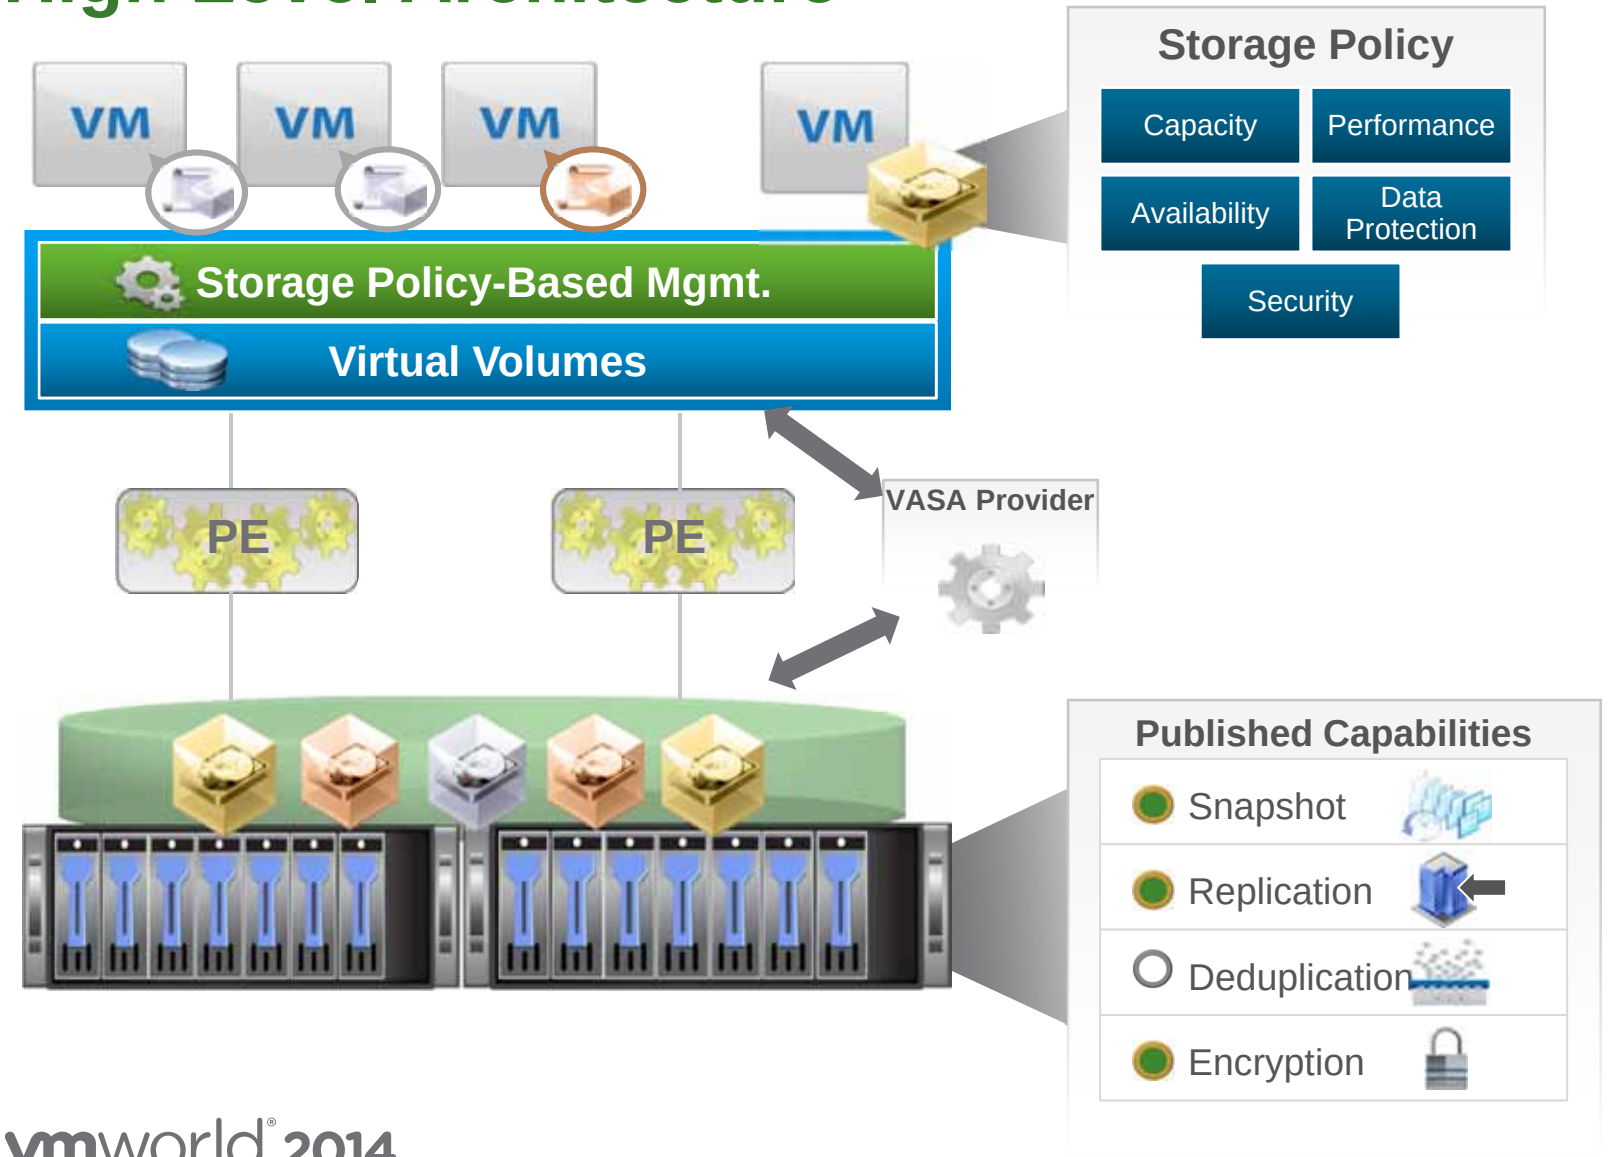

Overview

- No FileSystem
- ESX manages Array through VASA (vSphere APIs for Storage Awareness) APIs.
- Arrays are logically partitioned into containers, called Storage Containers
- VM disks, called Virtual Volumes, stored natively on the Storage Containers.
- IO from ESX to array is addressed through an access point called, Protocol Endpoint (PE)
- Data Services are offloaded to the array
- Managed through storage policy-based management framework

VASA Provider (VP)

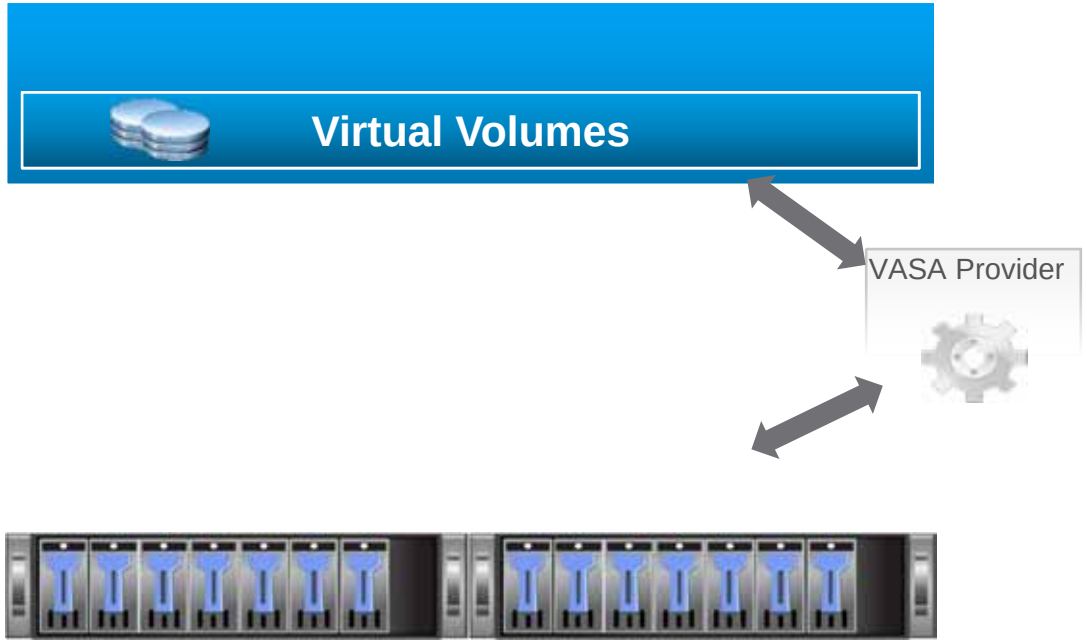

Characteristics

- Software component developed by Storage Array Vendors
- ESX and vCenter Server connect to VASA Provider
- Provides Storage awareness services
- Single VASA Provider can manager multiple arrays
- Supports VASA APIs exported by ESX
- VASA Provider can be implemented within the array's management server or firmware
- Responsible for creating Virtual Volumes

Data Path

Protocol Endpoints (PE)

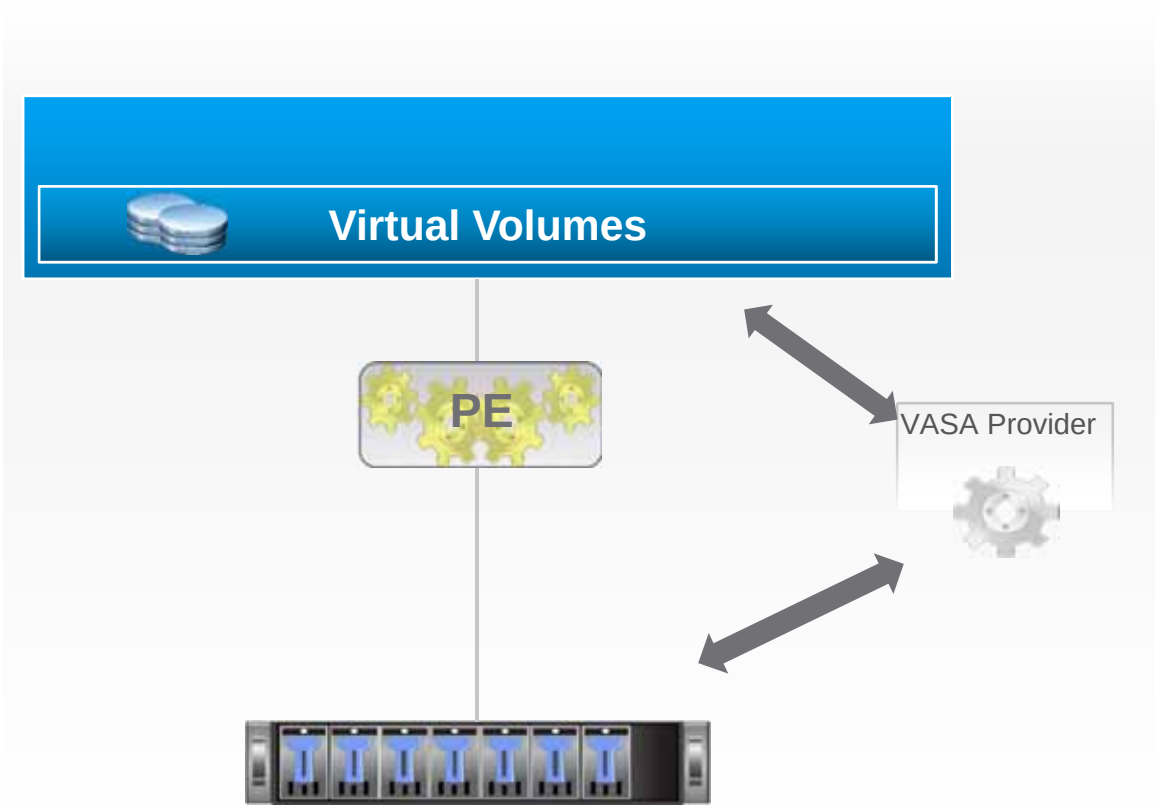

Why Protocol Endpoints?

- Separate the access points from the storage itself
- Can have fewer access points

What are Protocol Endpoints?

- Access points that enables communication between ESXi hosts and storage array systems.
- Part of the physical storage fabric
- Created by Storage administrators

Protocol Endpoints (PE)

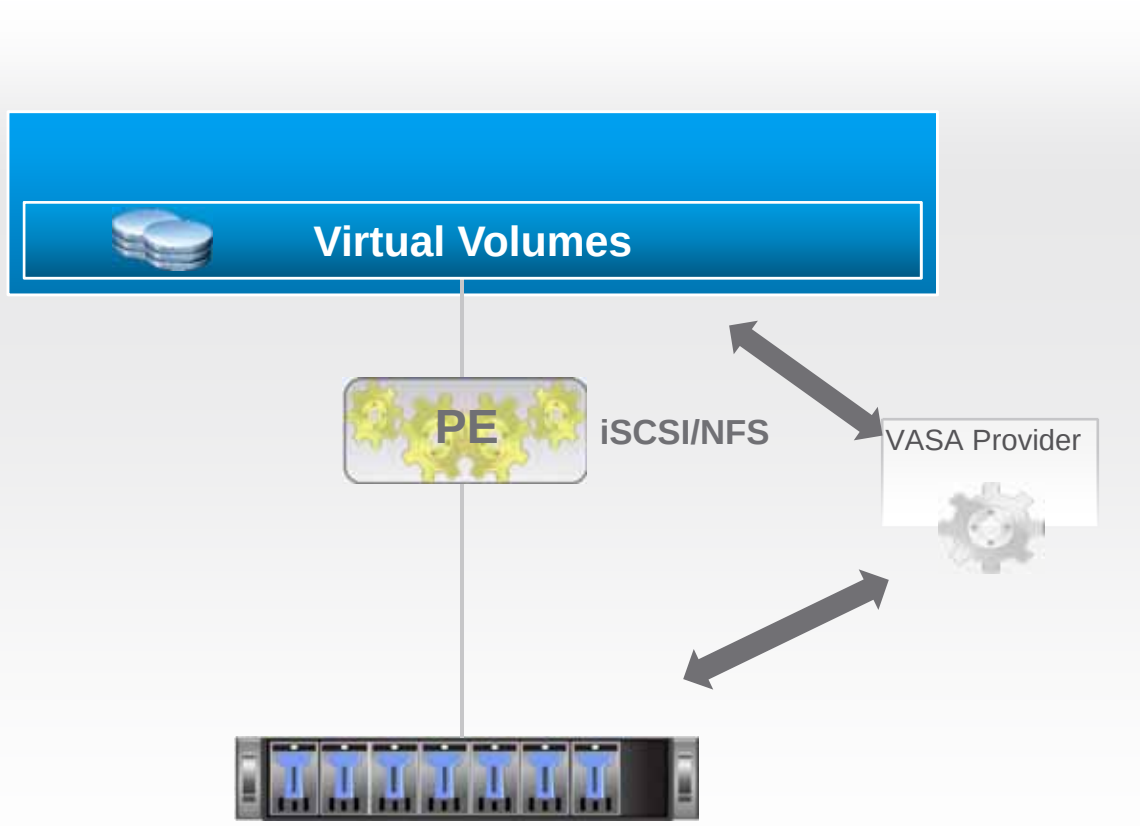

Scope of Protocol Endpoints

- Compatible with all SAN and NAS Protocols:
  - iSCSI
  - NFS v3
  - FC
  - FCoE
- A Protocol Endpoint can support any one of the protocols at a given time
- Existing multi-path policies and NFS topology requirements can be applied to the PE

Protocol Endpoints (PE)

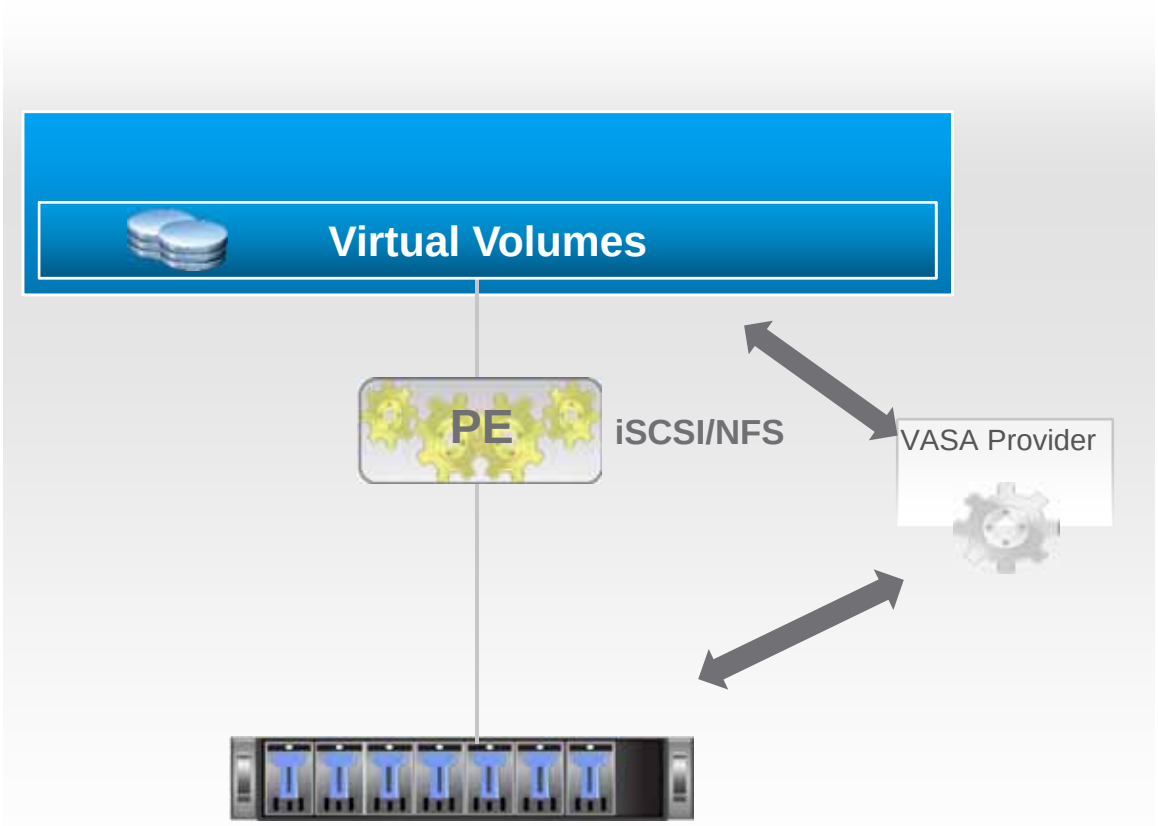

Protocol Endpoint discovery process

- SCSI PEs are discovered during an ESX rescan
- NFS PEs are maintained as IP addresses or file paths
- ESX will identify PE and maintain all discovered PEs in a database.

Managing Storage Capacity

Storage Container (SC)

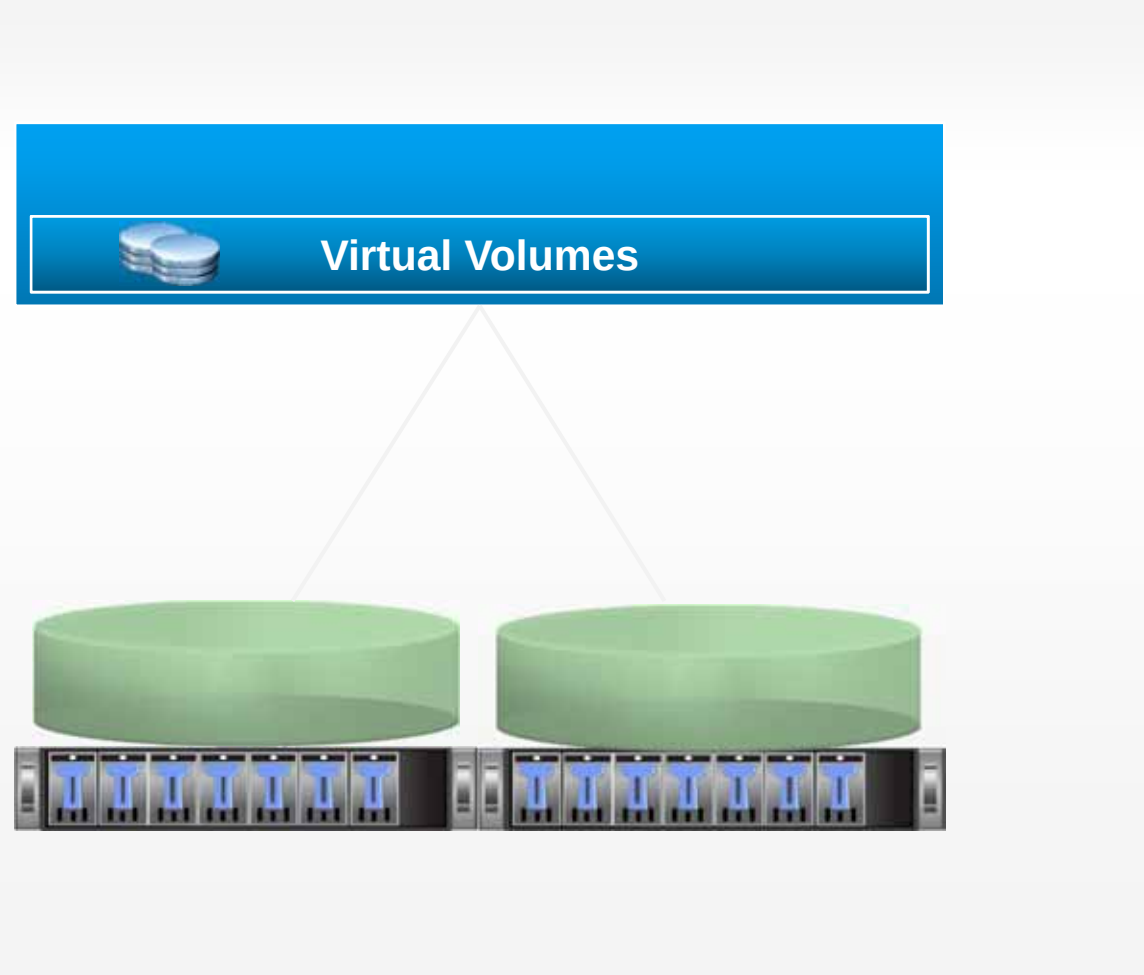

What are Storage Containers?

- Logical storage constructs for grouping of virtual volumes.
- Setup by Storage administrators
- Capacity is based on physical storage capacity
- Logically partition or isolate VMs with diverse storage needs and requirement
- Minimum one storage container per array
- Maximum depends on the array
- A single Storage Container can be simultaneously accessed via multiple Protocol Endpoints

Storage Containers (SC)

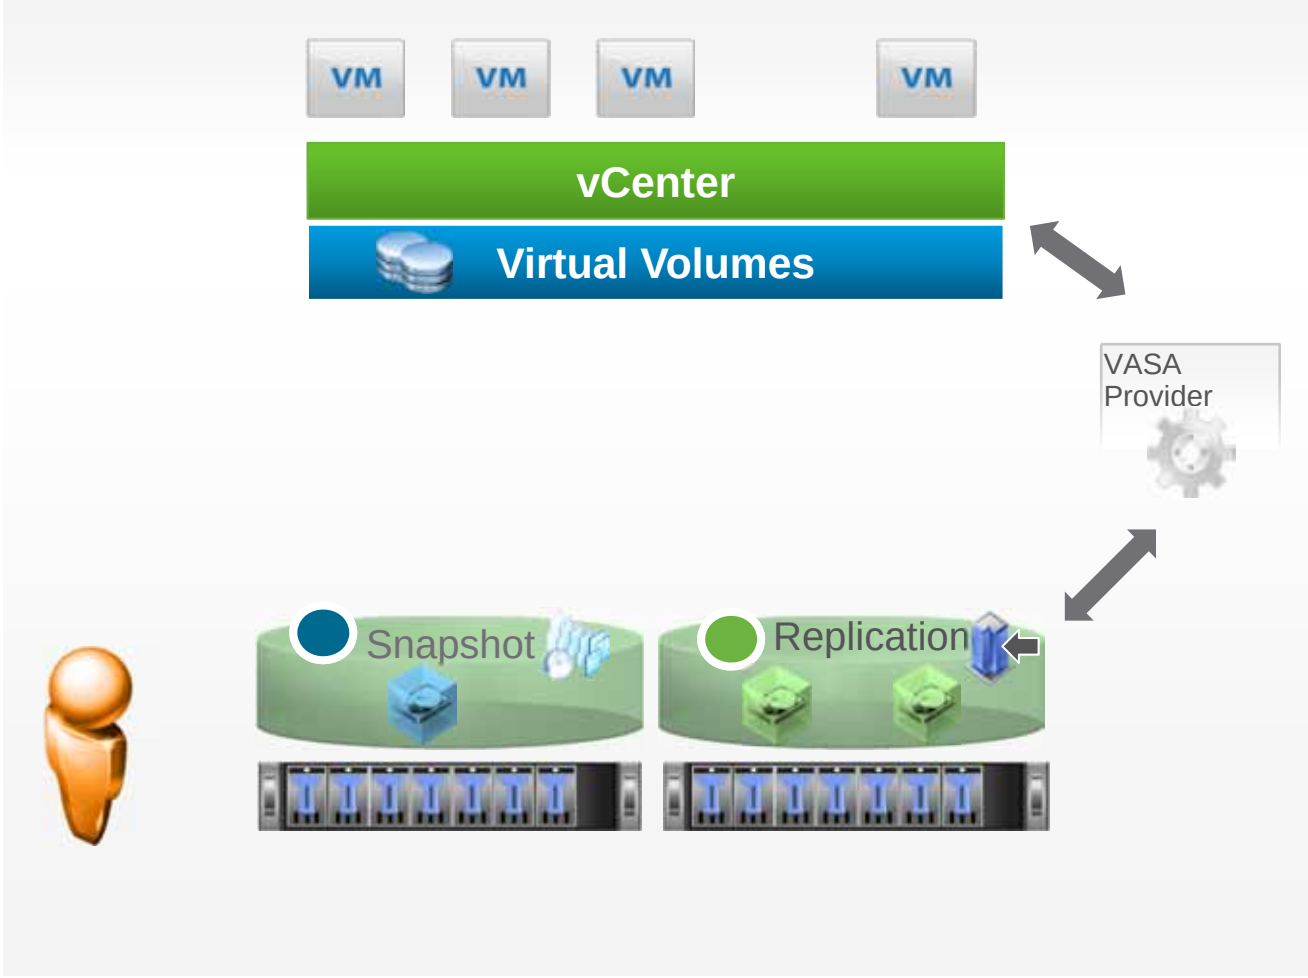

Storage Container Discovery Process

- Storage admin sets up Storage Container with desired capacity
- Desired Capabilities are applied to the Storage Containers
- VASA Provider discovers Storage Container and reports to vCenter
- Any new VMs that are created will subsequently be provisioned in the Storage Container

Differences between Storage Containers and LUNs

Storage Containers

- Size based on array capacity
- Max number of SCs depend only on the array ability
- Size of SC can be extended
- Can distinguish heterogeneous capabilities for different VMs (Virtual Volumes) provisioned in that SC
- Managed by VASA APIs

LUN

- Fixed size
- Fixed size mandates more number of LUNs
- Needs a FileSystem
- Can only apply homogeneous capability on all VMs (VMDKs) provisioned in that LUN.
- Managed by In-band FileSystem commands

Storage Container (SC)

Do I still need to create Datastores?

vSphere Datastore

=

Storage Container

Storage Container (SC)

What do the Admins see?

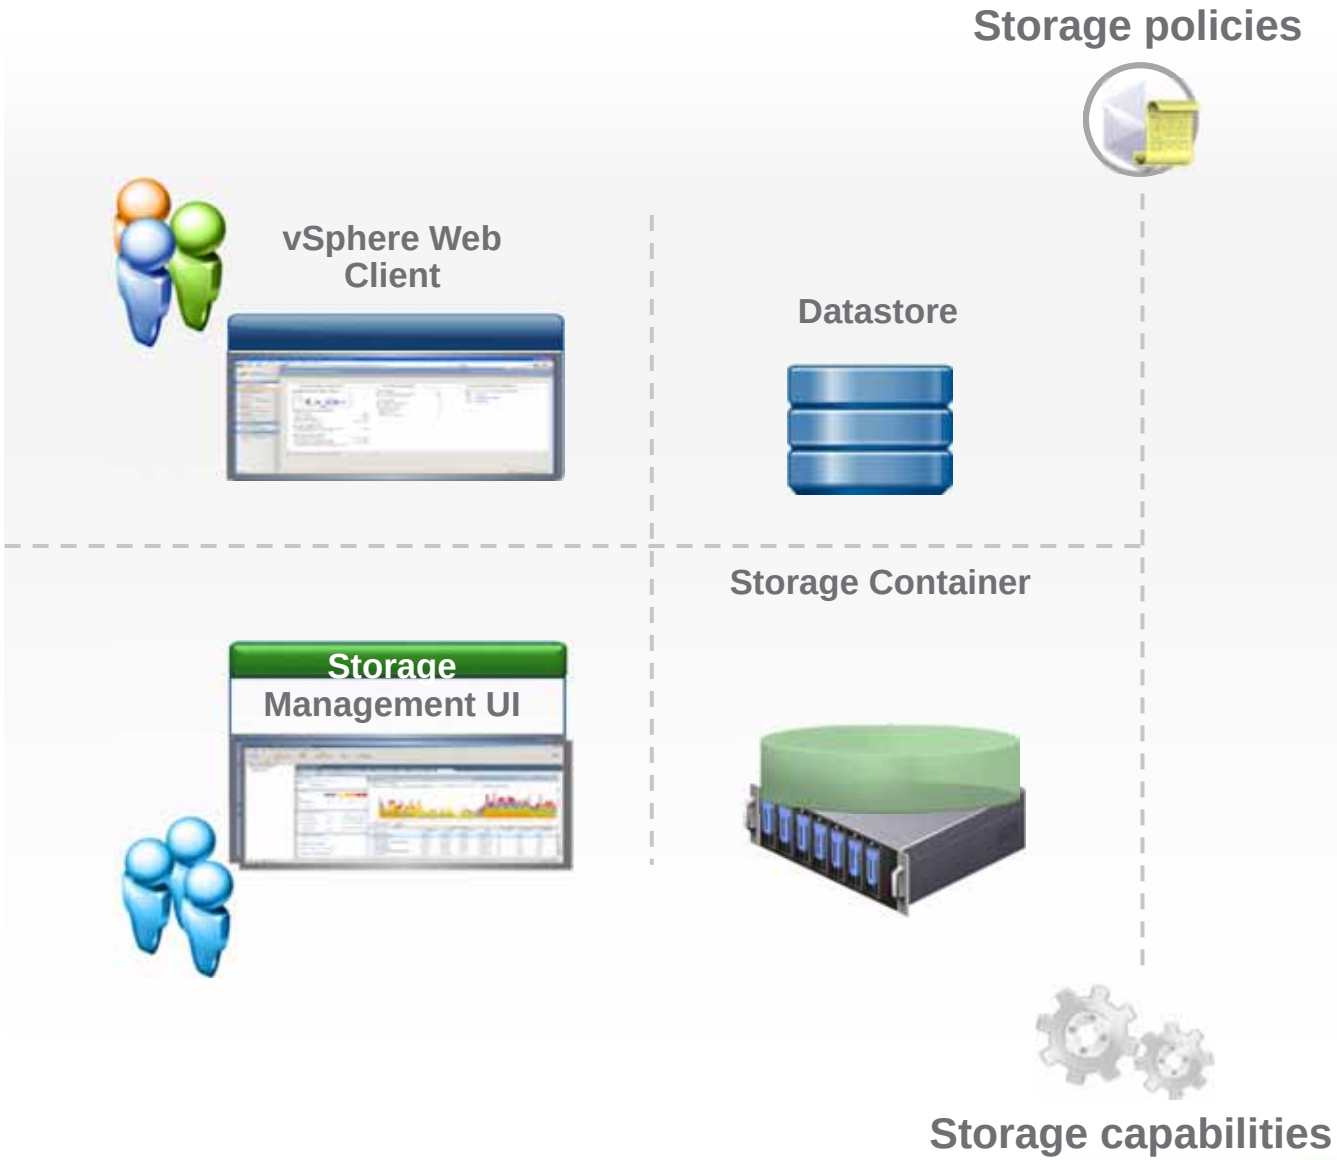

Ensuring SLOs through Policies

Storage Policy Based Management (SPBM) – Array Capabilities

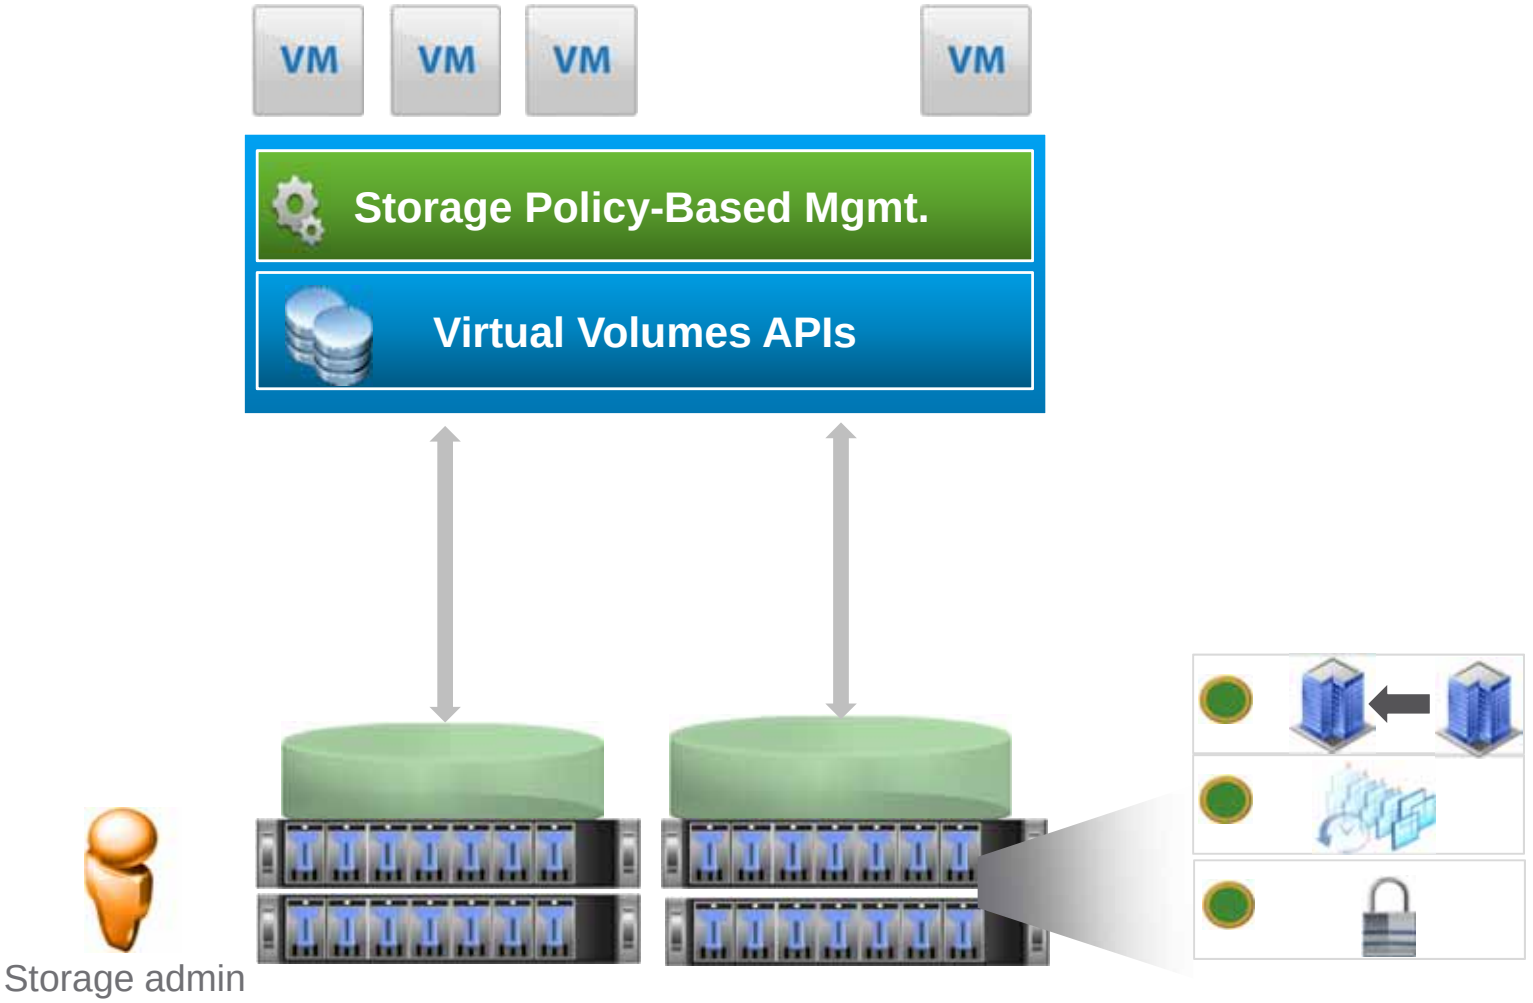

Publish Capabilities

- Array based features and data services
- Defines what an Array can offer
- Advertised to ESX through VASA APIs

Storage Policy Based Management (SPBM) – Array Capabilities

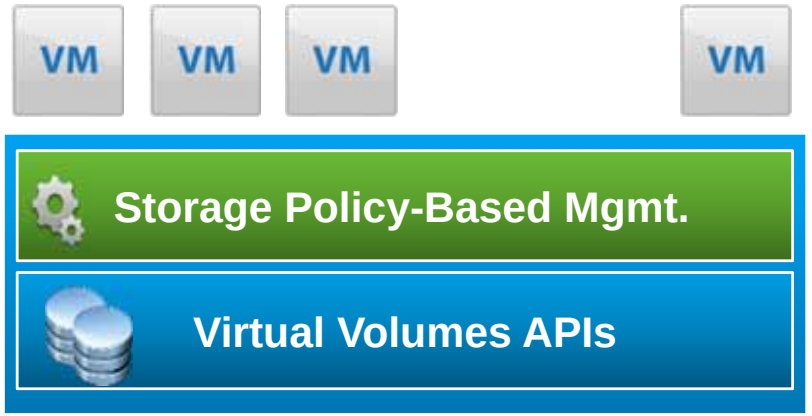

Storage admin

Sample Default Profile for (6090a058-cd89-ffe3-87763007db37] - capabilities	
Disk Types	
Disk Encryption	
Dedupe	
Replication	
Snapshot	

Storage Policy Based Management (SPBM) – VM Policies

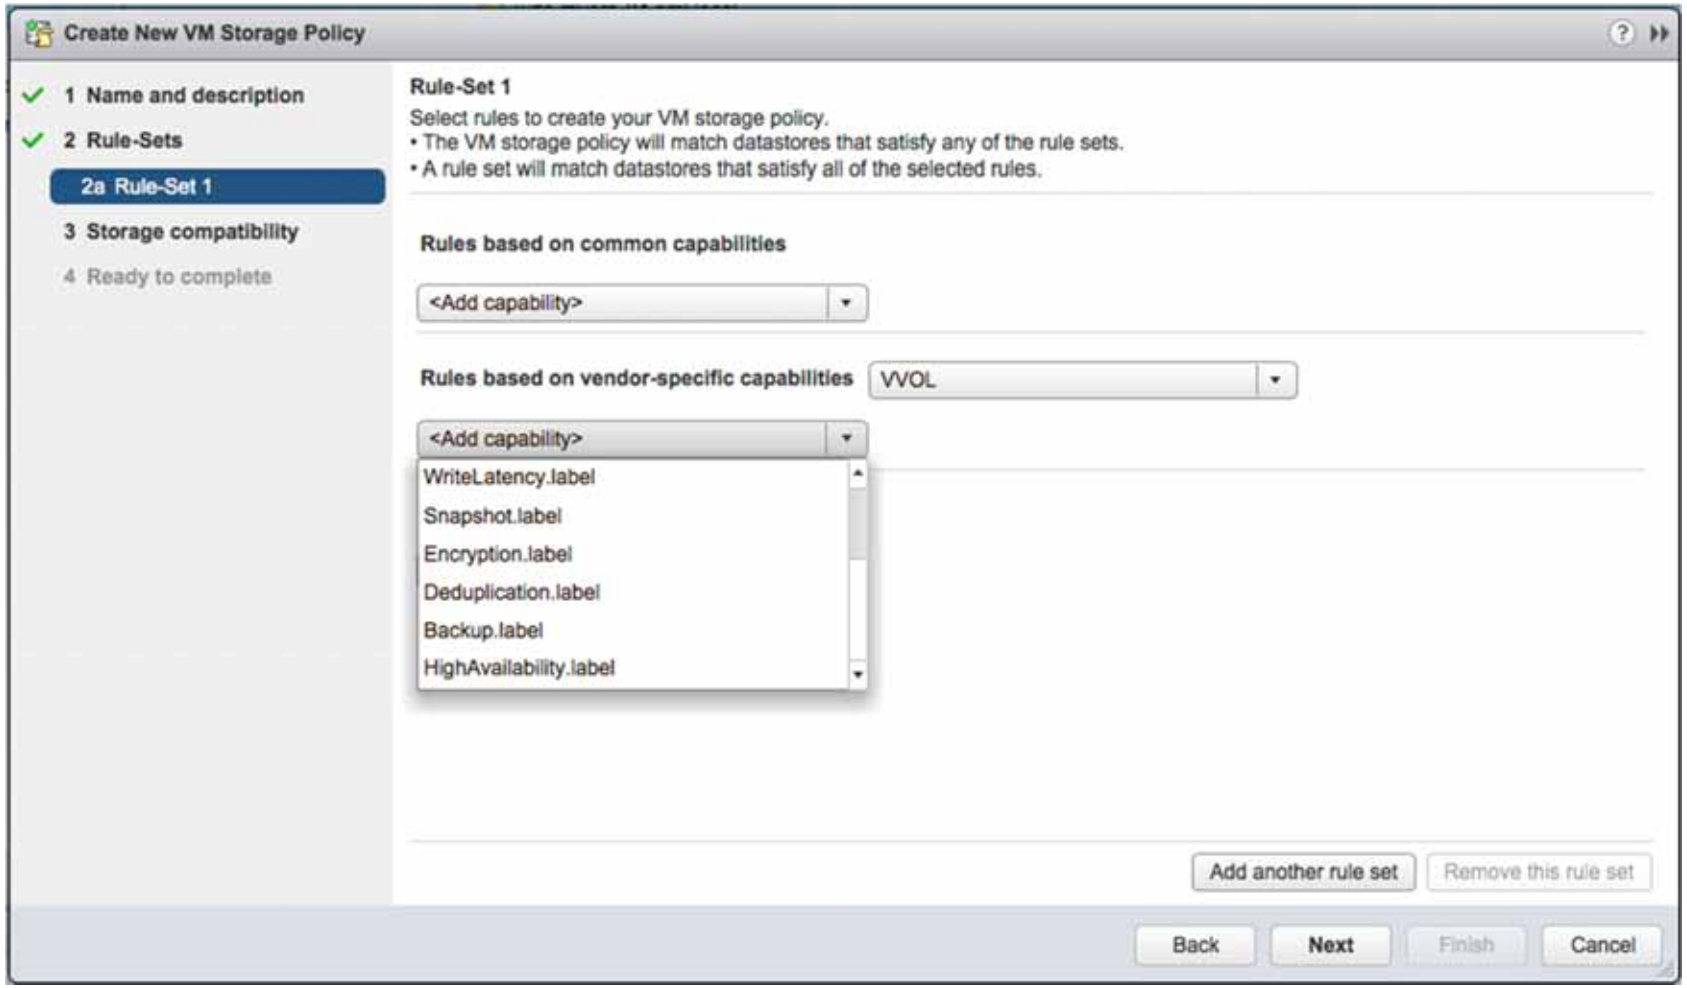

Storage Policy Based Management (SPBM)

VM Objects - Virtual Volumes

Virtual Volumes – An Introduction

- **Virtual Volumes**
  - Virtual machine objects stored natively on the array storage containers.
  - No Filesystem on-disk formatting required
- There are **five** different types of recognized Virtual Volumes:
  - **Config-VVol** – Metadata
  - **Data-VVol** – VMDKs
  - **Mem-VVol** – Snapshots
  - **Swap-VVol** – Swap files
  - **Other-VVol** – Vendor solution specific

vSphere Web Client View

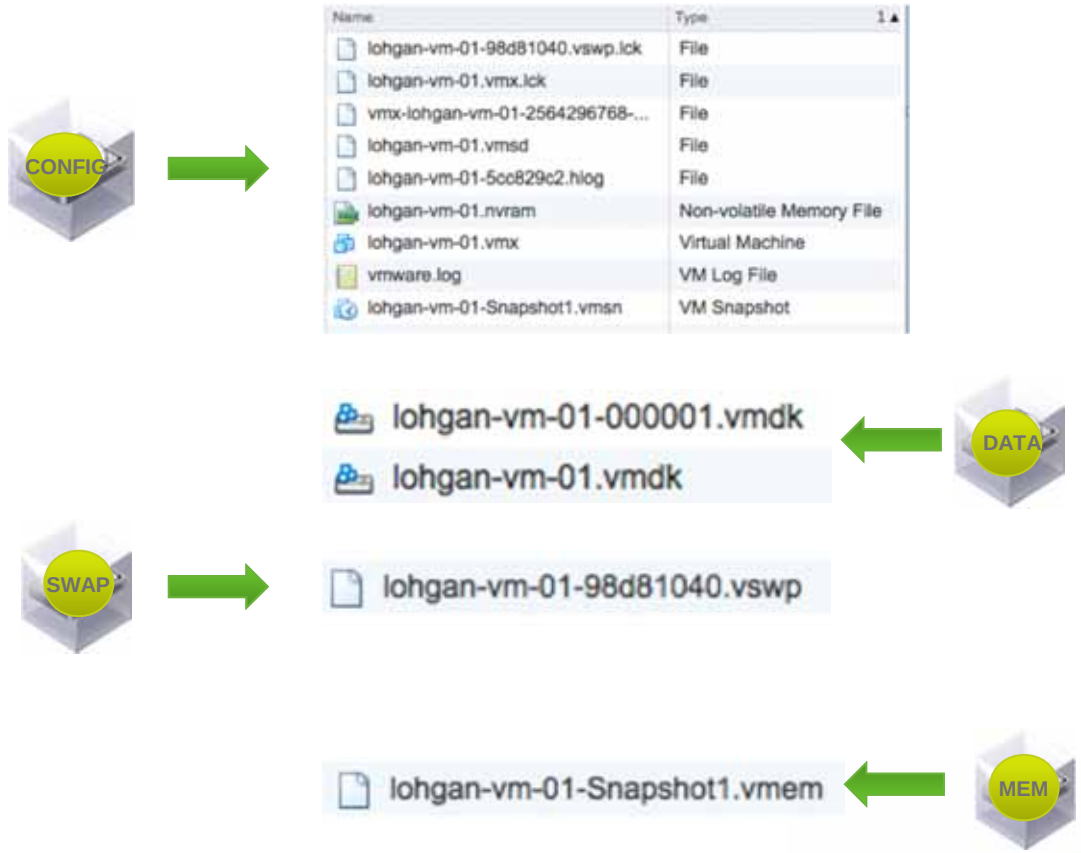

Virtual Volumes

What do the admins need to get familiar with?

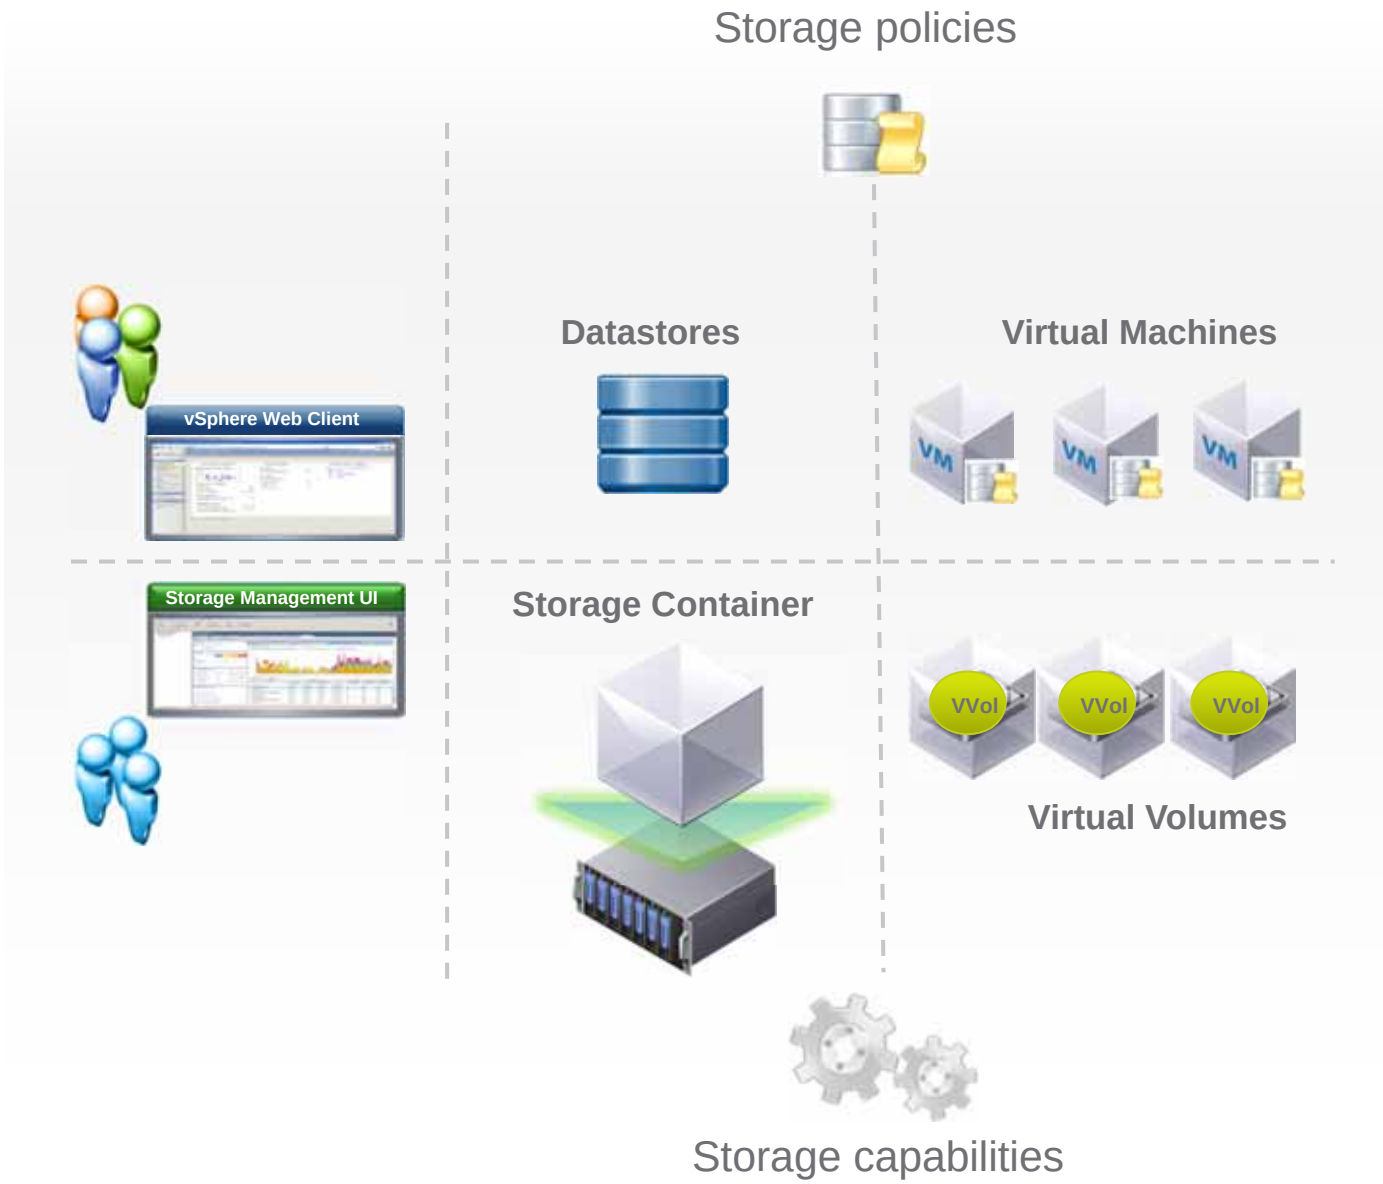

Virtual Volumes – Bind requests

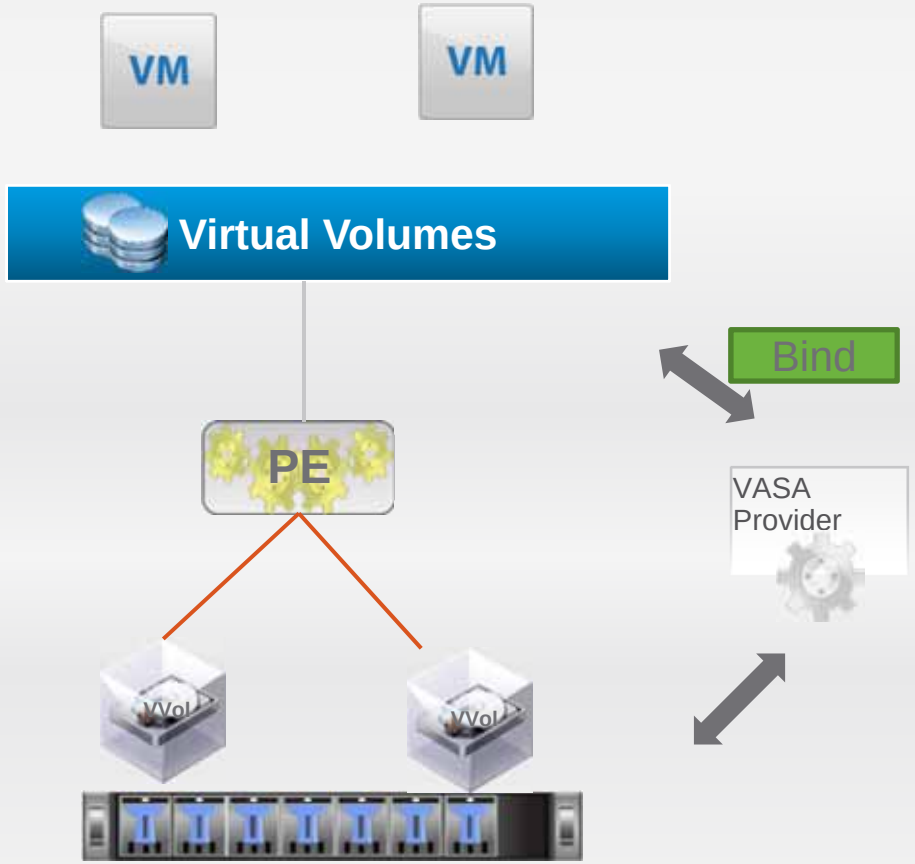

- IO path is established through a VASA Bind request
- VASA Provider does two things upon receiving a Bind request
  - Returns a PE ID to which the VVol is bound
  - A unique secondary ID to be used for IO between the bound VVol and PE
- For SCSI, the secondary ID is the secondary LUN ID
- For NAS, the secondary ID is the file path
- M:M Relationship between VVol and PE

Data Services

Snapshots

- Offloaded to Array
- Copy on write image of a Virtual Volume
- Two type of snapshots supported:
  - Managed Snapshot – Managed by ESX.
    - A maximum of 32 snapshot are supported for fast clones
  - Unmanaged Snapshot – Managed by the storage array.

Managed Snapshot - vSphere

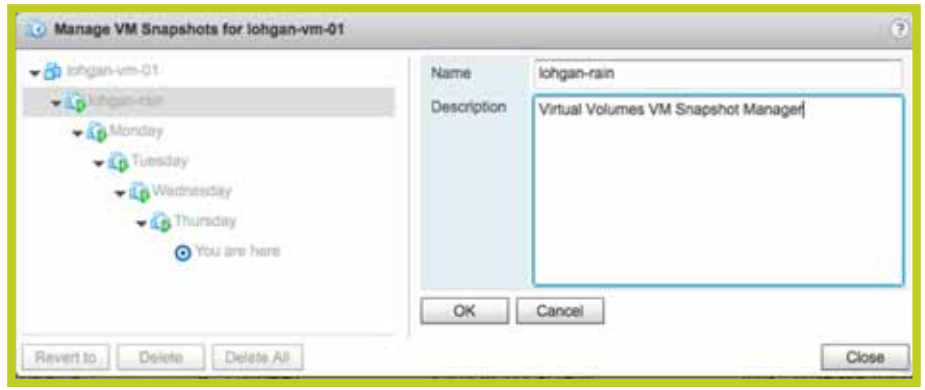

Unmanaged Snapshot - Array

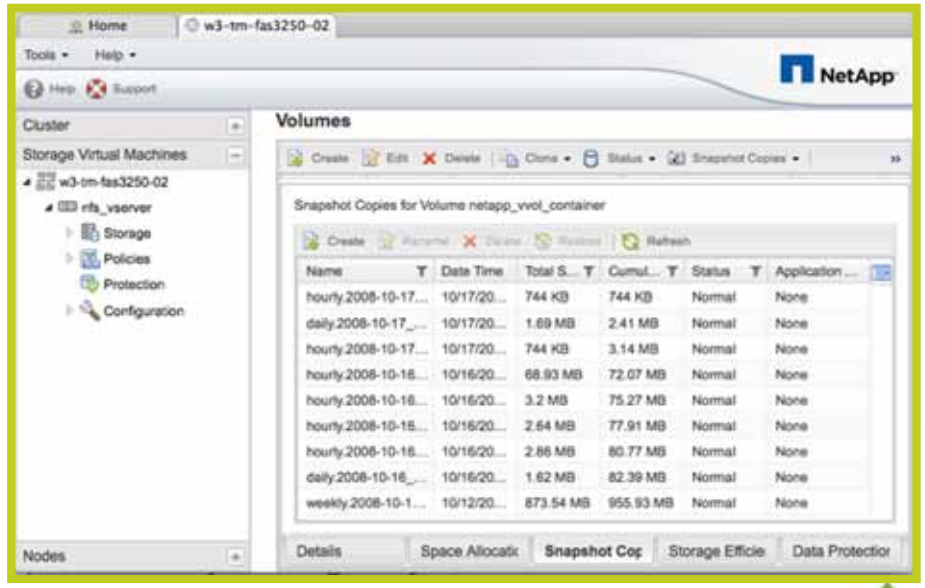

Piecing It All Together

Virtual Volumes – The New De-facto Storage Paradigm

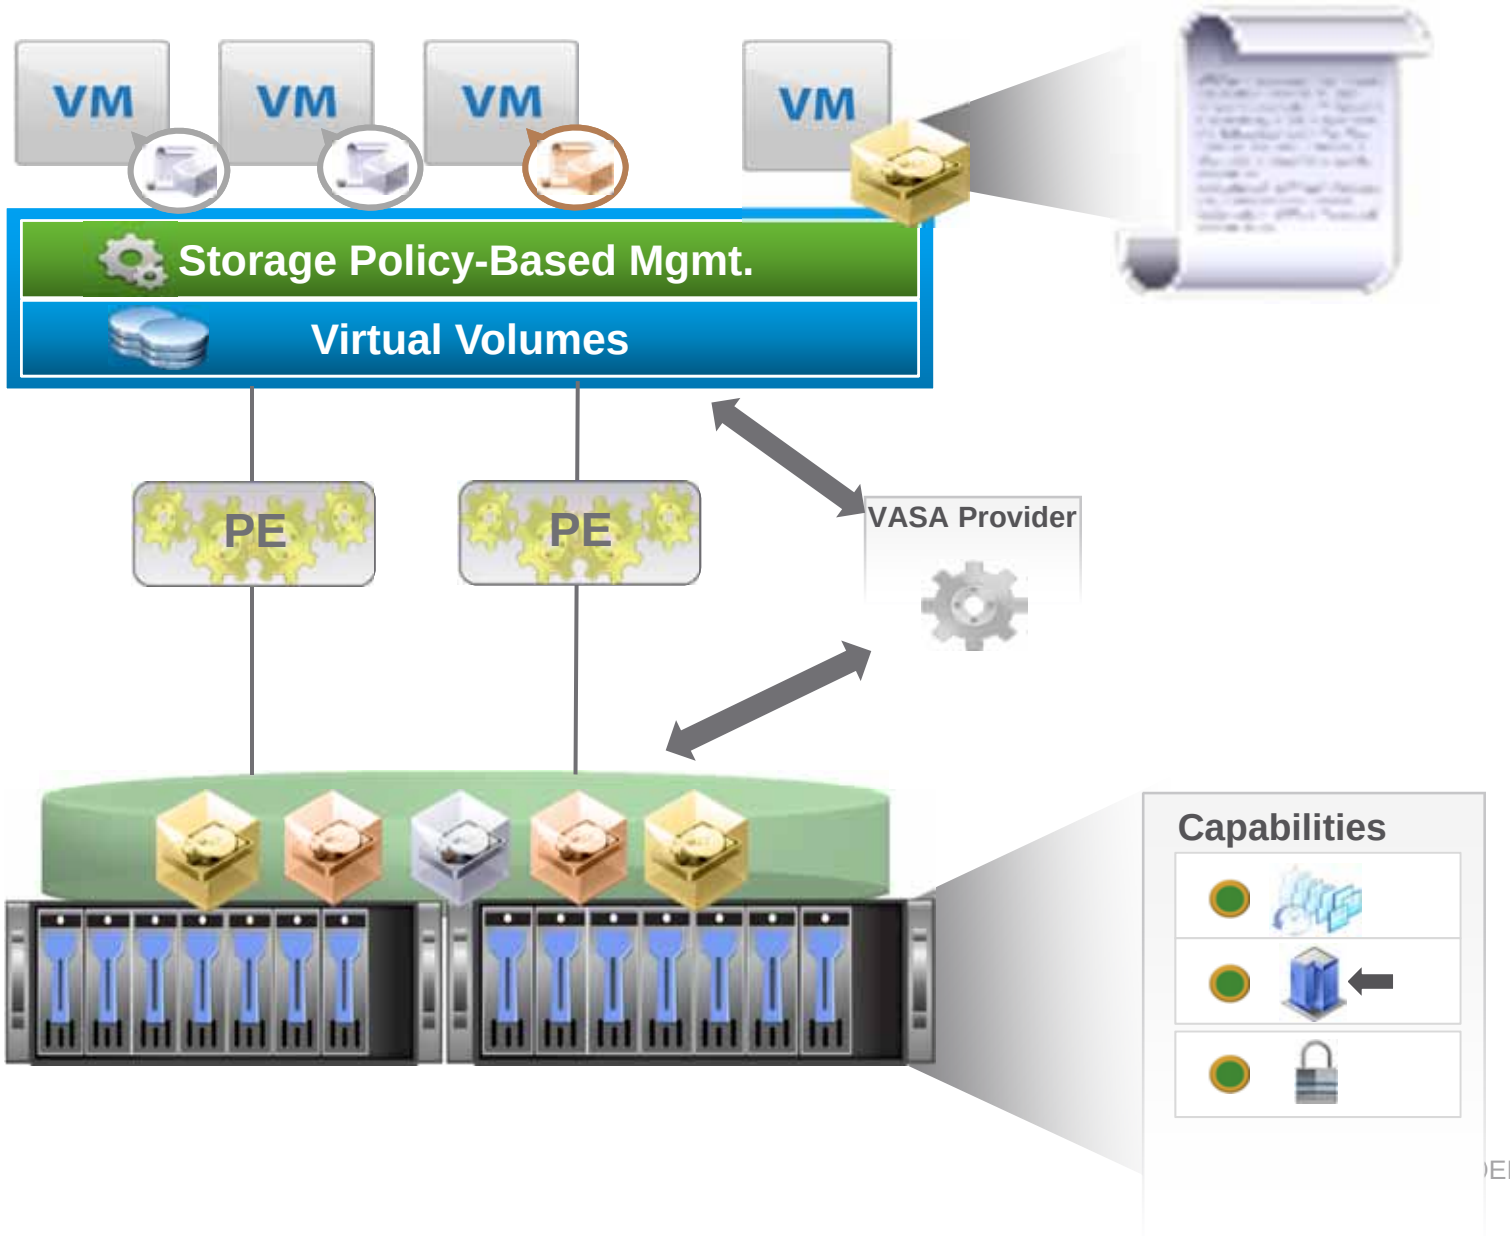

Making the Transition to Virtual Volumes

Support by Broad Ecosystem Makes Transition Smooth

And Many More…

More than 20 VVOL Partners

Partners Announcing GA

Virtual Volumes in Beta

Learn More…..

VVols and SPBM in Nimble – Architecture Overview

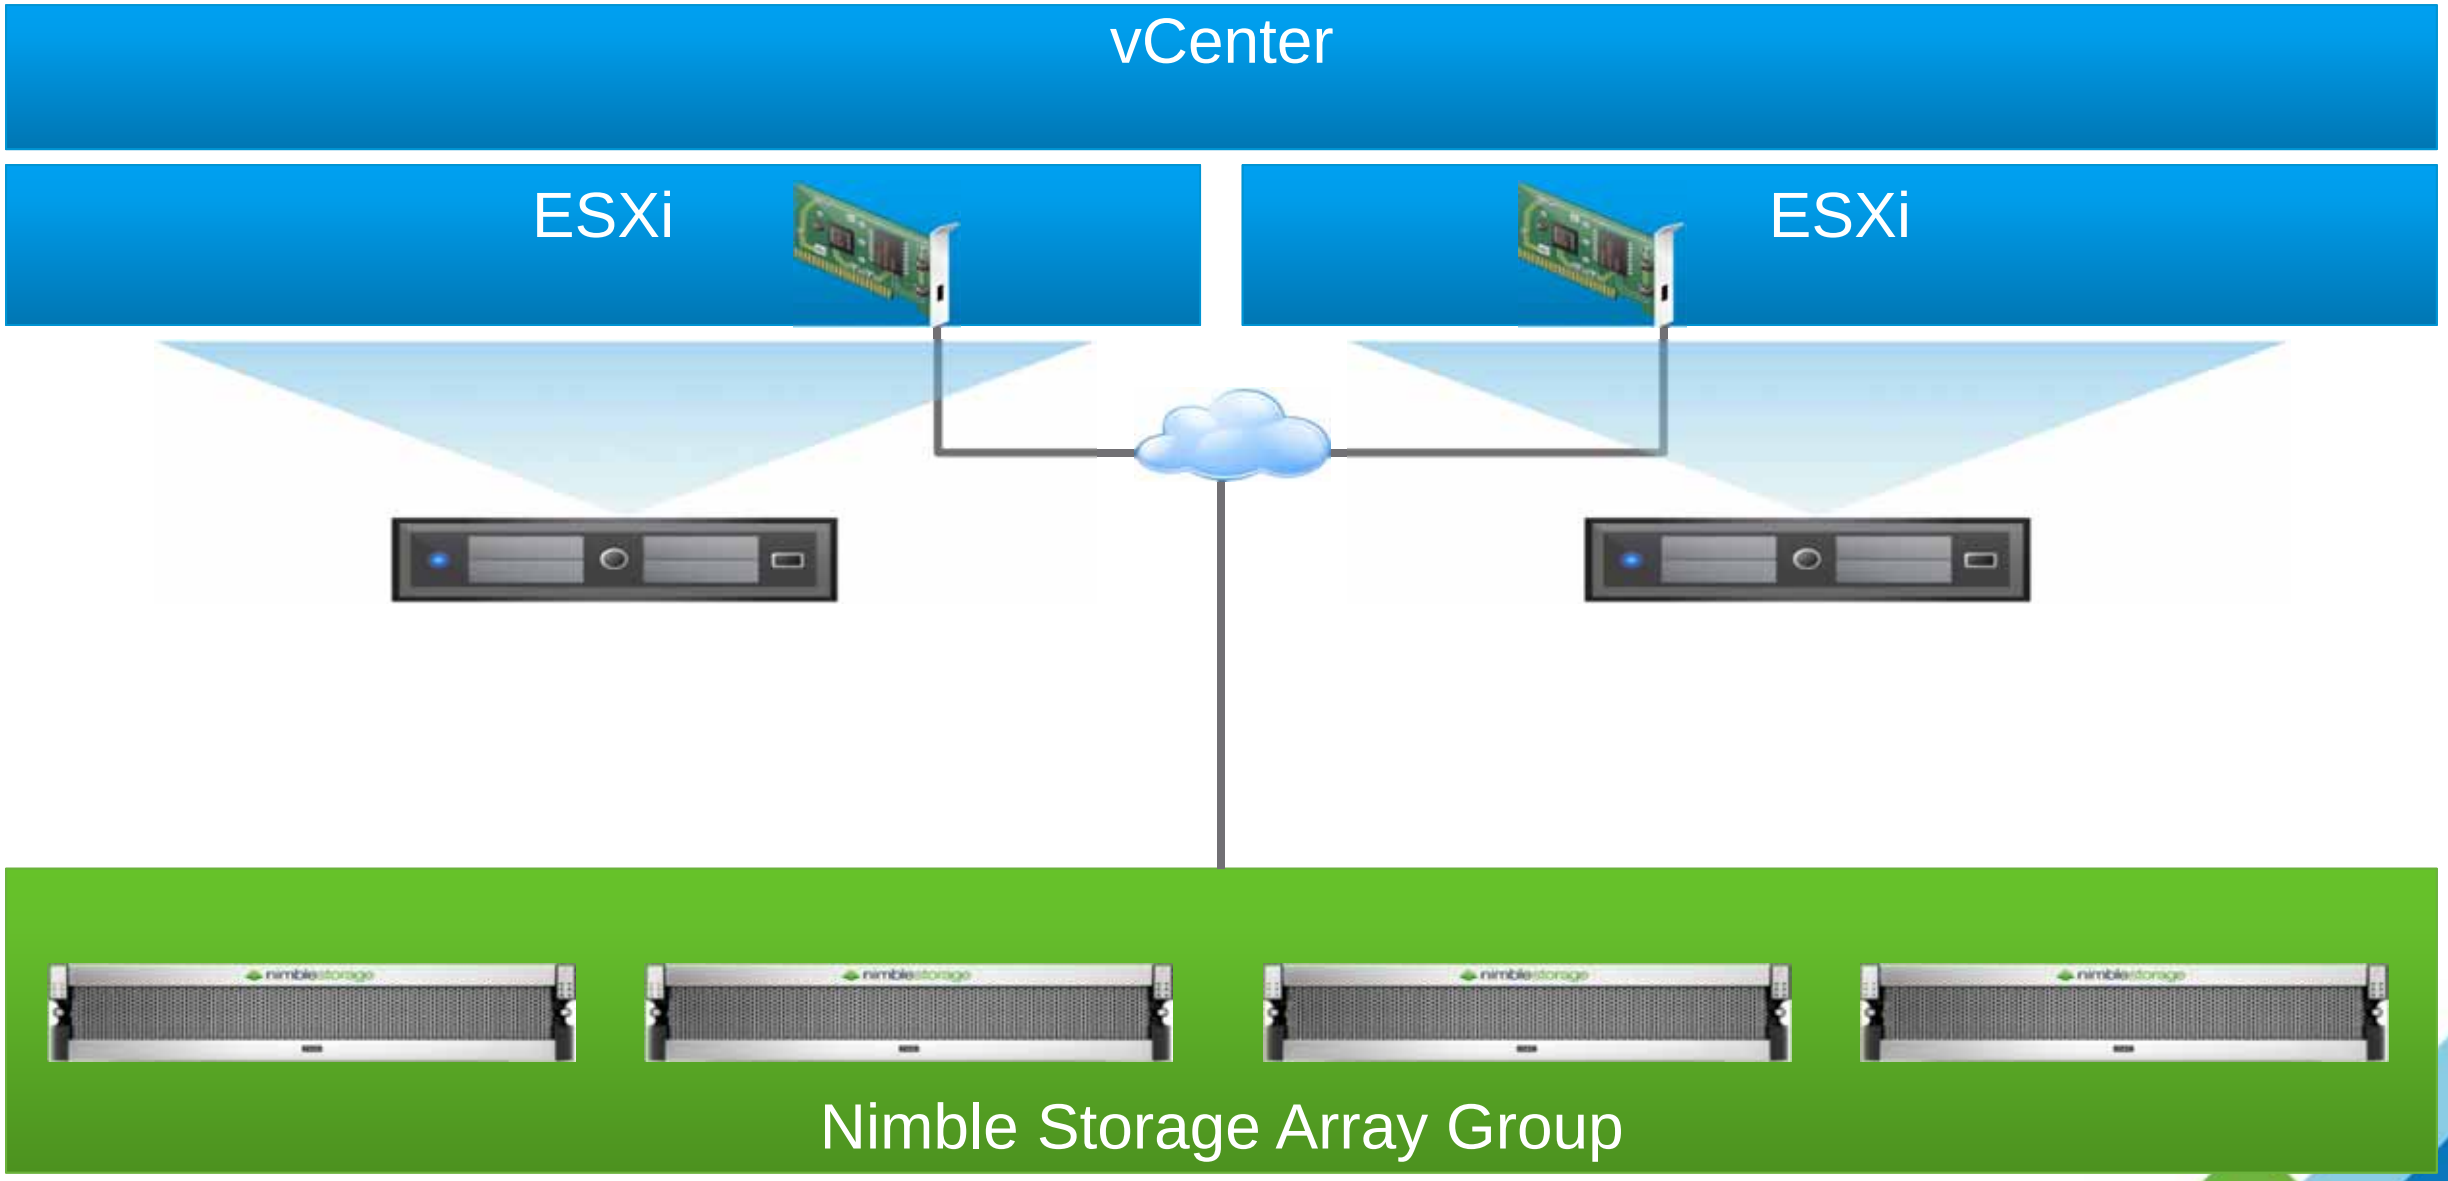

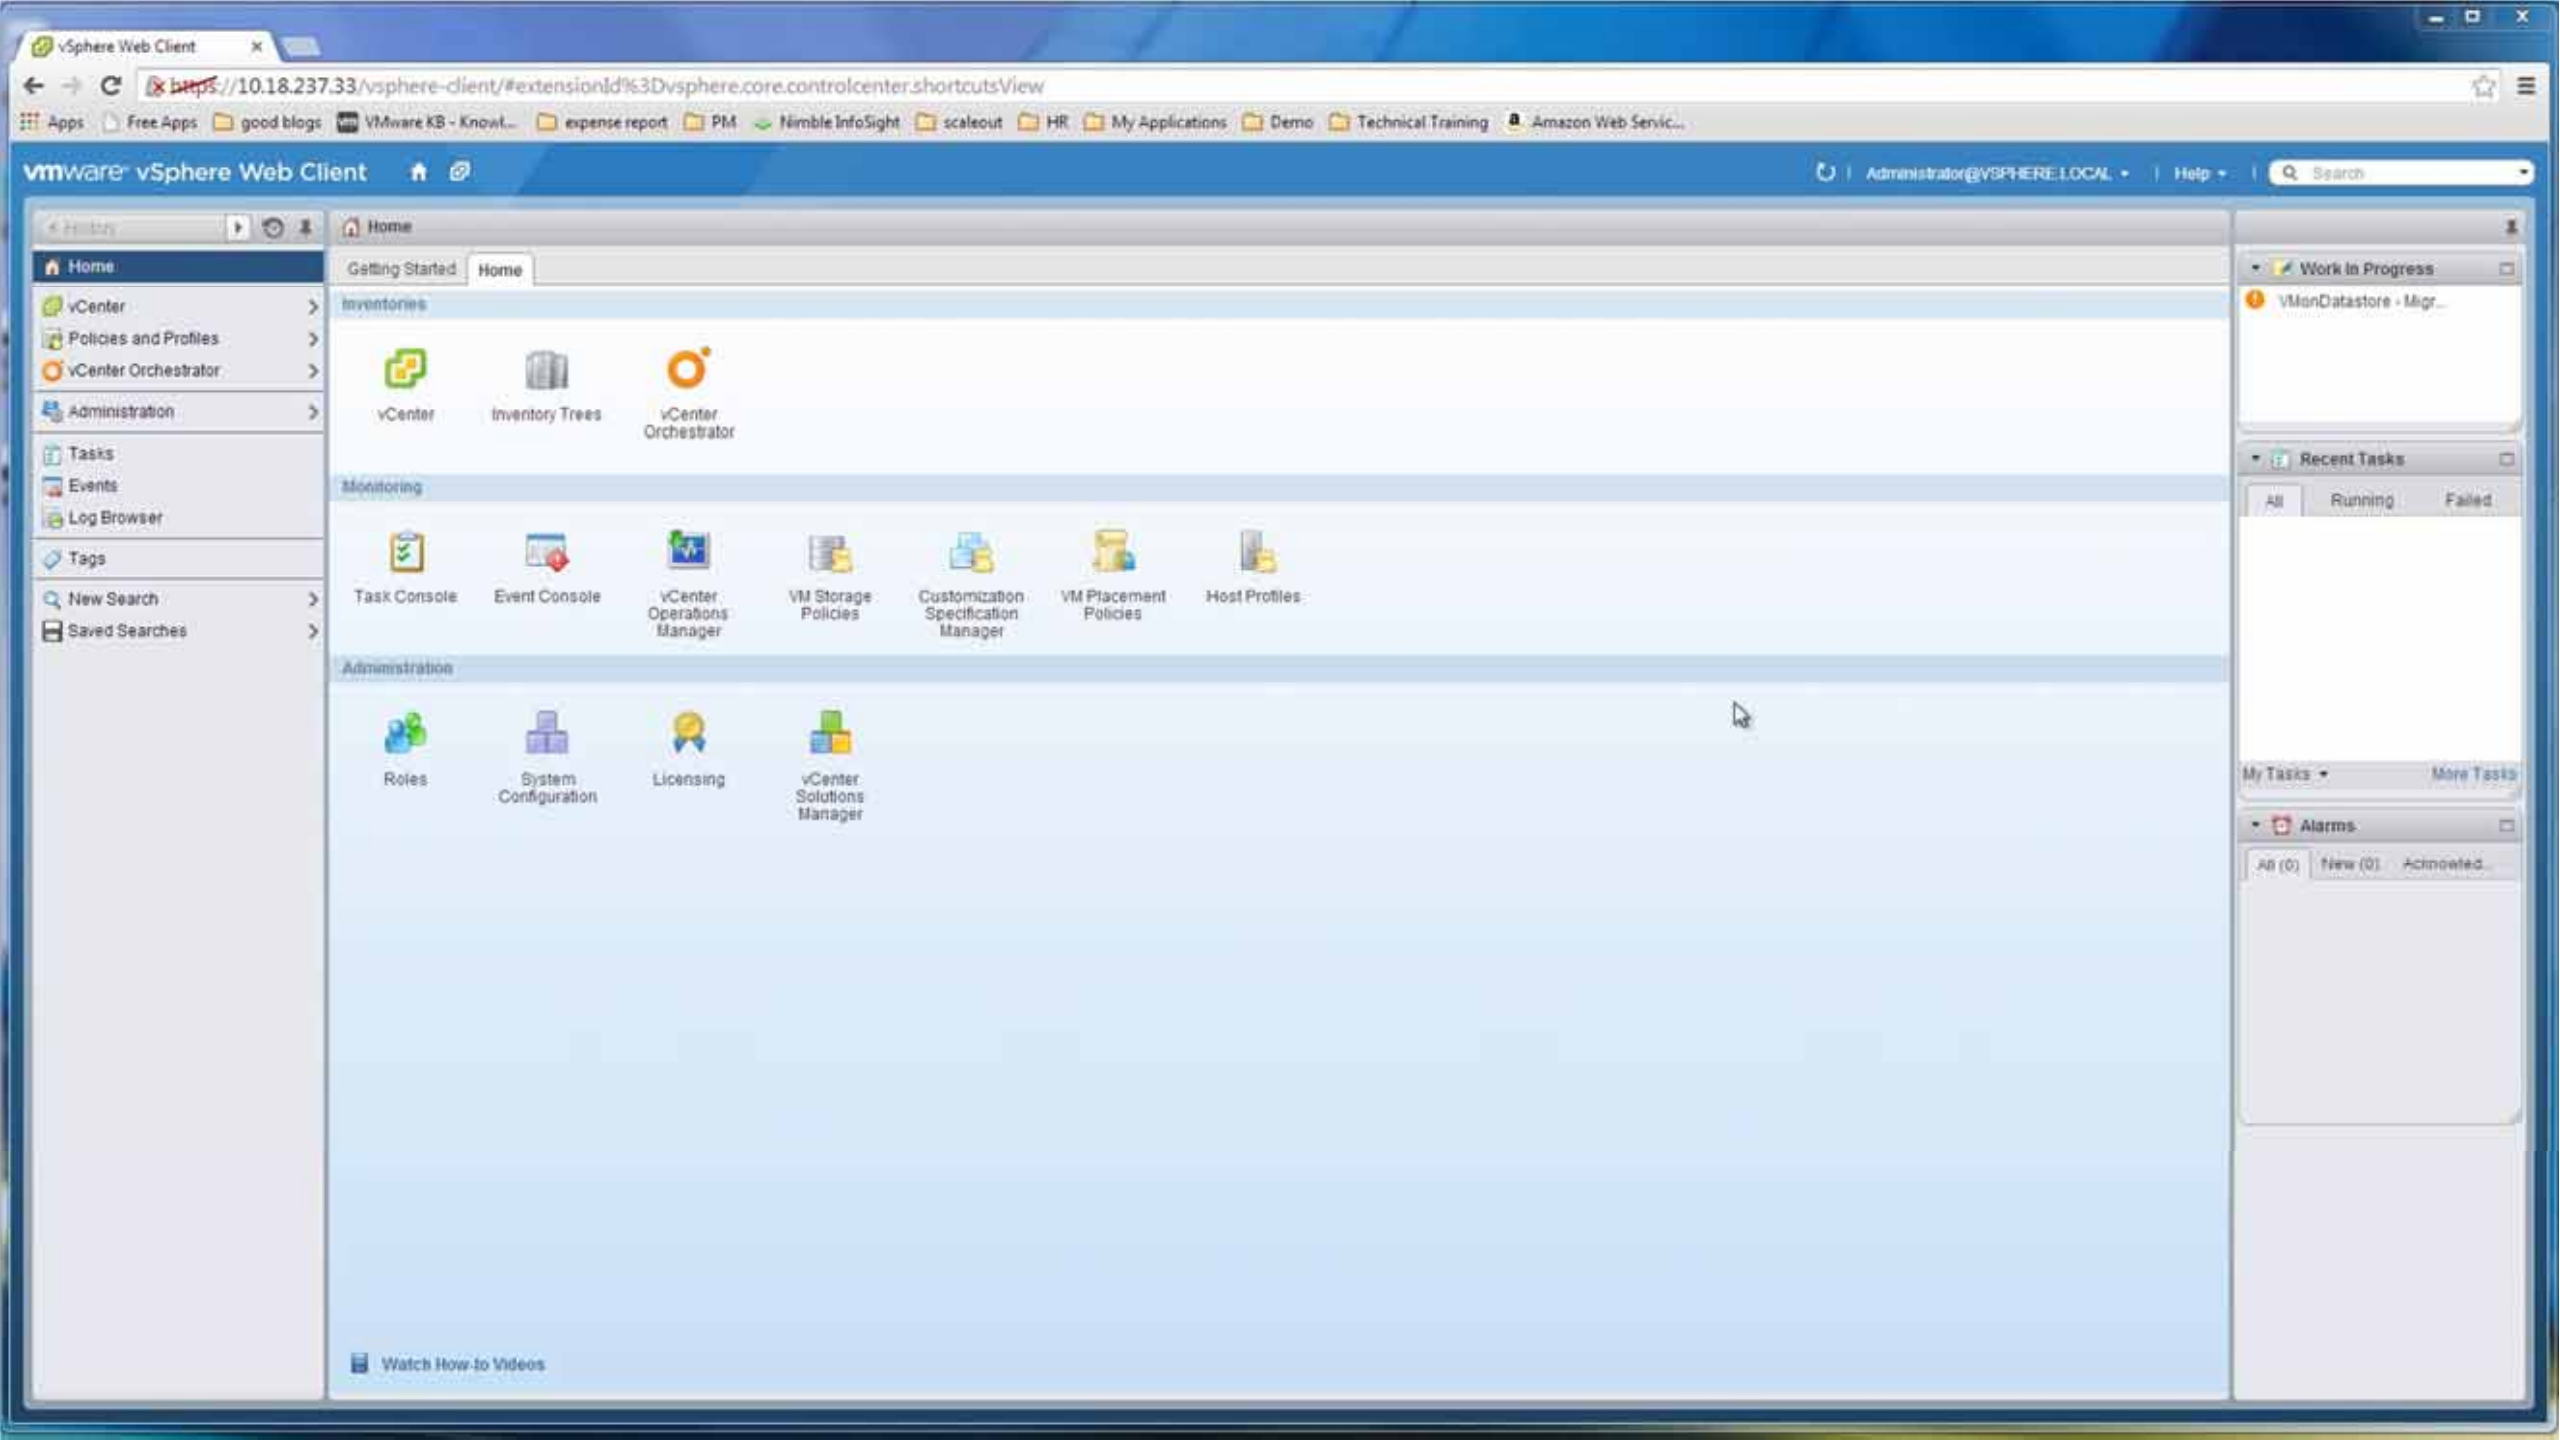
vSphere Web Client
https://10.18.237.33/vsphere-client/#extensionId%3Dvsphere.core.controlcenter.shortcutsView
Apps
Free Apps
good blogs
VMware KB - Knowl...
expense report
PM
Nimble InfoSight
scaleout
HR
My Applications
Demo
Technical Training
Amazon Web Servic...
vmware vSphere Web Client
Administrator@VSPHERE.LOCAL
Help
Search
Home
Home
vCenter
Policies and Profiles
vCenter Orchestrator
Administration
Tasks
Events
Log Browser
Tags
New Search
Saved Searches
Home
Getting Started
Home
Inventories
vCenter
Inventory Trees
vCenter Orchestrator
Monitoring
Task Console
Event Console
vCenter Operations Manager
VM Storage Policies
Customization Specification Manager
VM Placement Policies
Host Profiles
Administration
Roles
System Configuration
Licensing
vCenter Solutions Manager
Watch How-to Videos
Work In Progress
VMonDatastore - Migr...
Recent Tasks
All
Running
Failed
My Tasks
More Tasks
Alarms
All (0)
New (0)
Acknowled...

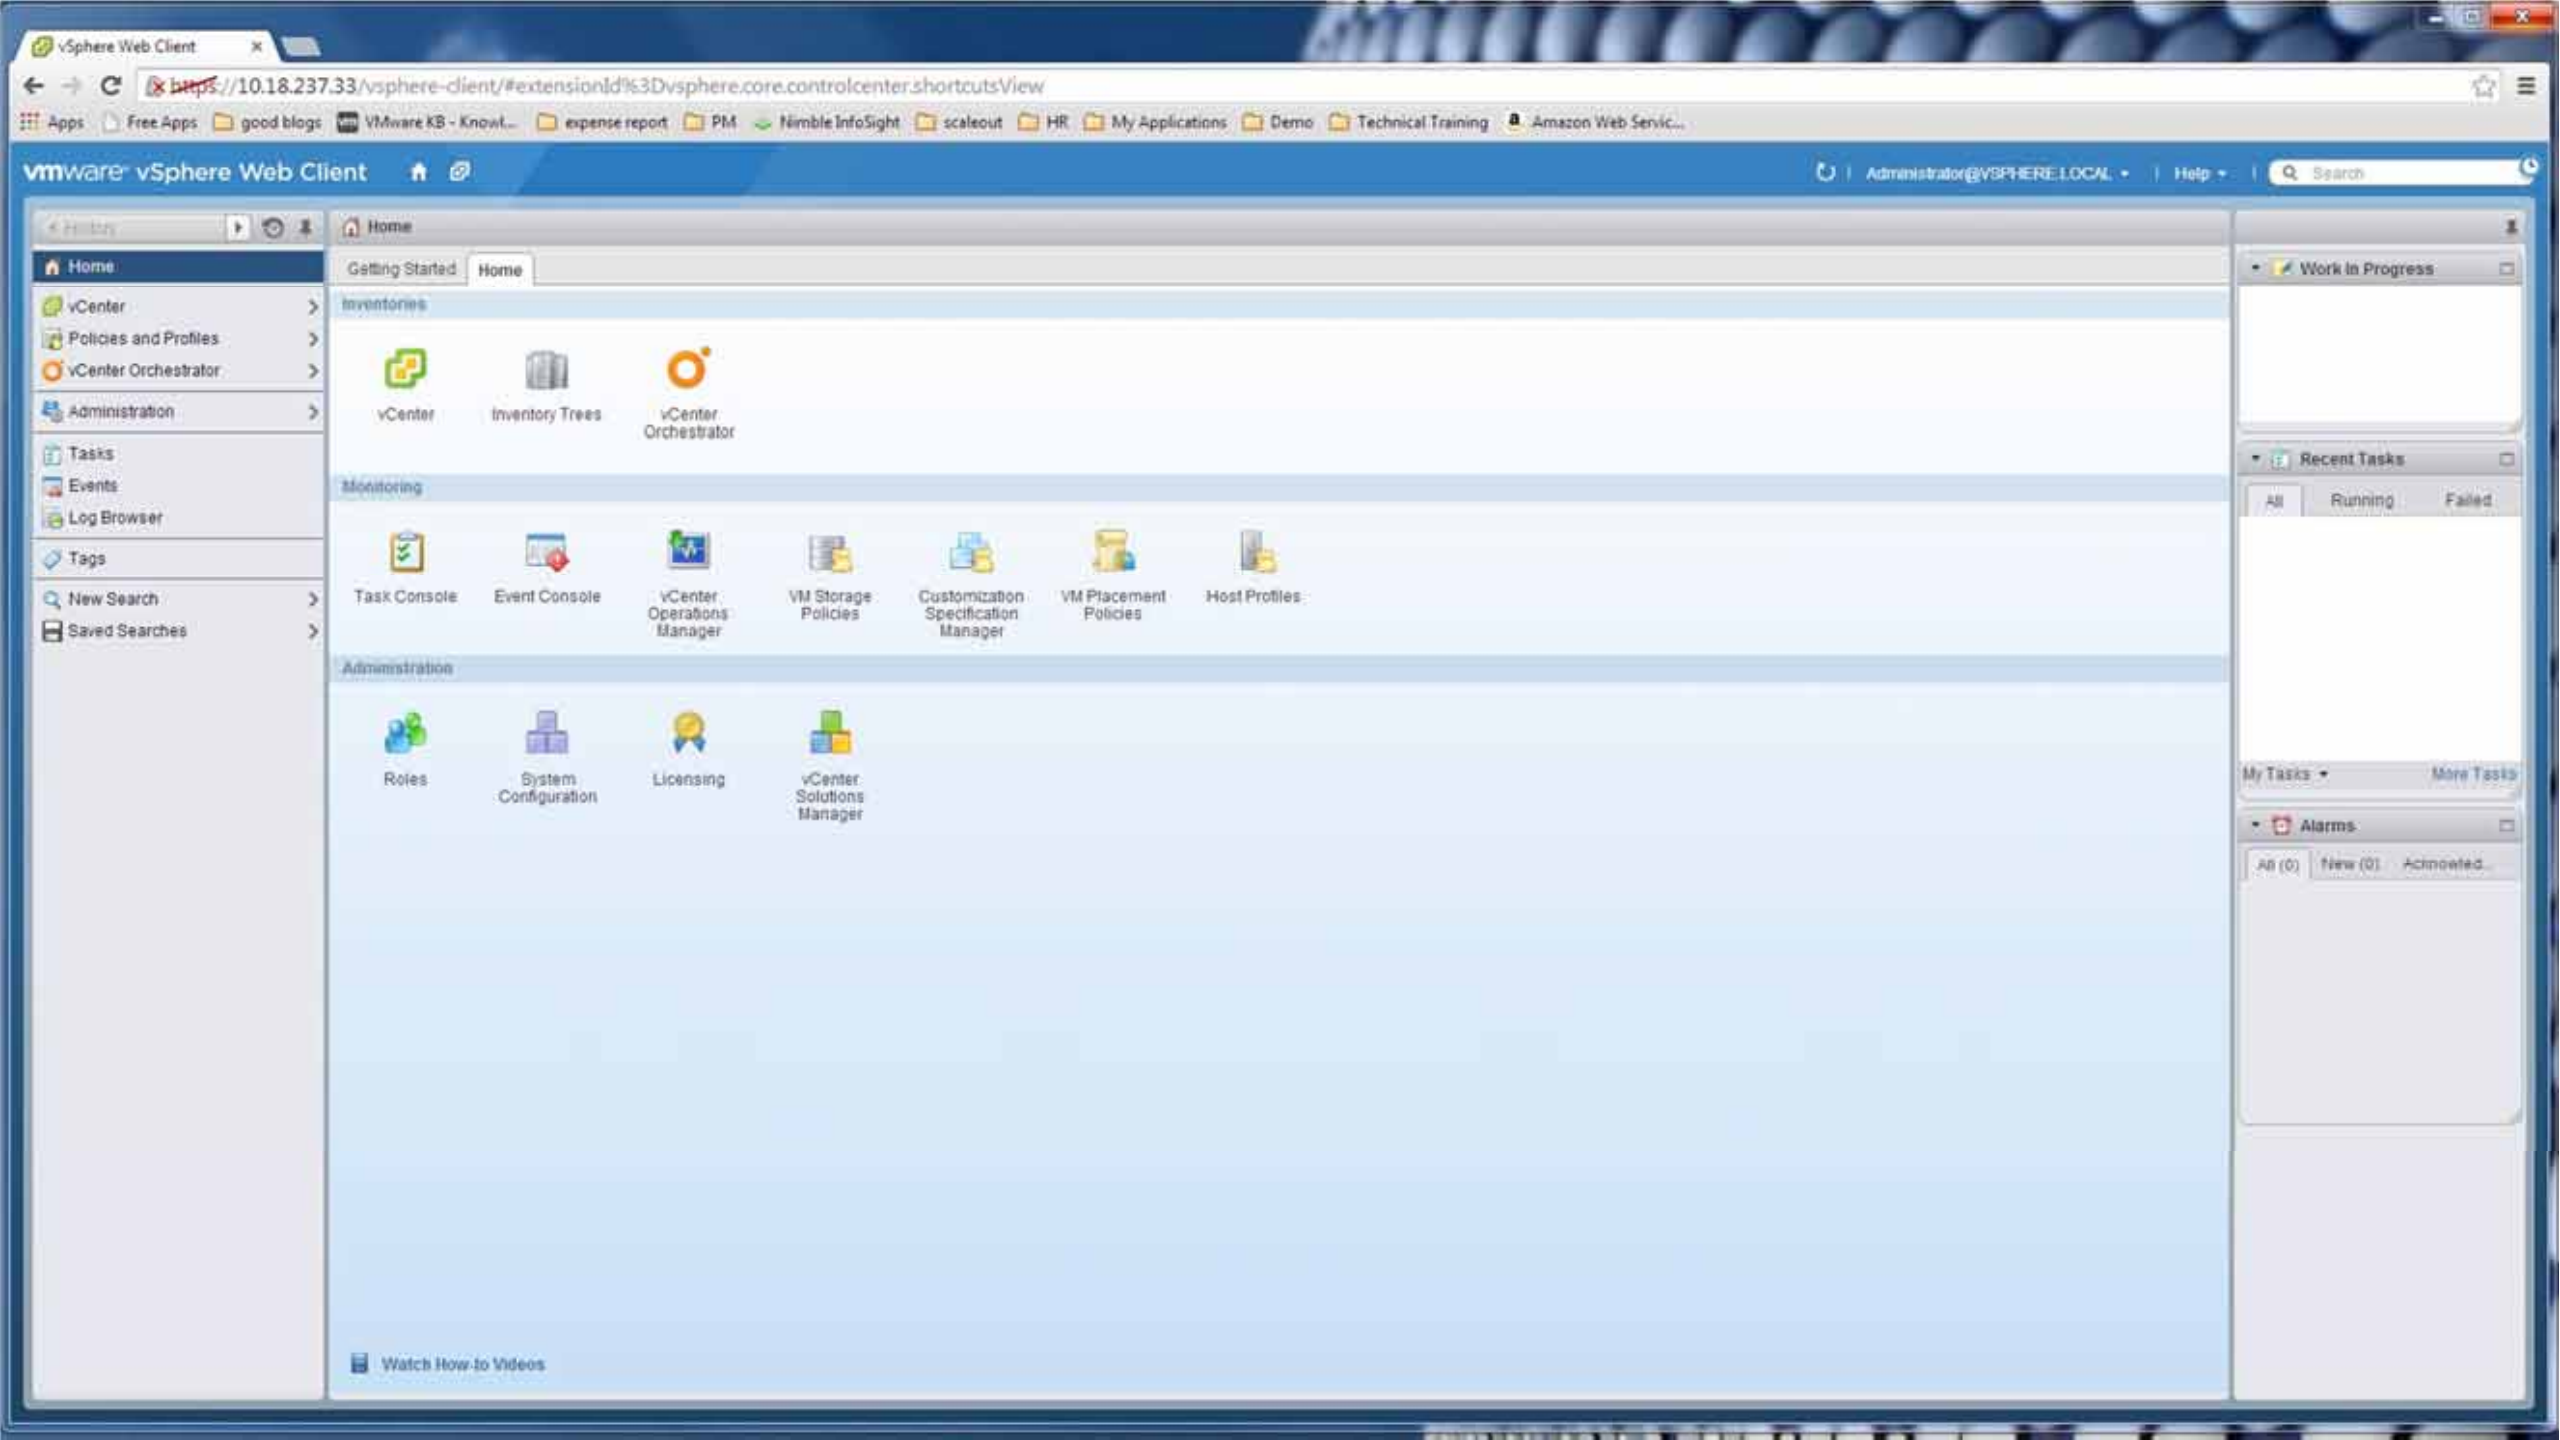
vSphere Web Client
https://10.18.237.33/vsphere-client/#extensionId%3Dvsphere.core.controlcenter.shortcutsView
Apps
Free Apps
good blogs
VMware KB - Knowl...
expense report
PM
Nimble InfoSight
scaleout
HR
My Applications
Demo
Technical Training
Amazon Web Servic...
vmware vSphere Web Client
Administrator@VSPHERE.LOCAL
Help
Search
Home
Home
vCenter
Policies and Profiles
vCenter Orchestrator
Administration
Tasks
Events
Log Browser
Tags
New Search
Saved Searches
Home
Getting Started
Home
Inventories
vCenter
Inventory Trees
vCenter Orchestrator
Monitoring
Task Console
Event Console
vCenter Operations Manager
VM Storage Policies
Customization Specification Manager
VM Placement Policies
Host Profiles
Administration
Roles
System Configuration
Licensing
vCenter Solutions Manager
Watch How-to Videos
Work In Progress
Recent Tasks
All
Running
Failed
My Tasks
More Tasks
Alarms
All (0)
New (0)
Acknowled...

Learn More…….

Breakout Sessions

STO1963 – Virtual Volumes Business Overview

STO 3163 – Virtual Volumes and Scalable Data Protection in a Software Defined Enterprise

STO3162 – Satisfy requirements of your application at the granularity of VVols (Nimble)

STO3246 – Scalable Virtual Volumes Storage Management with IBM XIV storage

STO3161 – What can Virtual Volumes do for you? (EMC)

STO2142 – Hypervisor & Storage QOS; Two Great Tastes that Taste Great Together (SolidFire)

STO3247 – VVol Technical Preview with DELL storage

STO2554 – How Virtual Volumes will provide Shared Storage with X-ray vision (HP)

STO2752 – Deploying VVol with Hitachi Data Systems

Misc

HOL-SPL-1429
VVol Tech Preview

Focused Sessions

Engage with VMware

Engage with Vendors

Attend NDA sessions

Demos at Booth

VMware
HP
Netapp
Dell
EMC
IBM
HDS
Nimble
Tintri
SolidFire
Atlantis Computing
SANBlaze

★ Participate in Virtual Volumes Beta

Q&A


Thank You
vmworld 2014


Fill out a survey
Every completed survey is entered into a drawing for a $25 VMware company store gift certificate
vmworld 2014

STO1965

Virtual Volumes Technical Deep Dive

Rawlinson Rivera, VMware, Inc
Suzy Visvanathan, VMware, Inc

vmworld® 2014

BACKUP

Snapshots

- Snapshots are a point in time copy on write image of a Virtual Volume with a different ID from the original
- Virtual Volumes snapshots are useful in the contexts of creating:
  - a quiesced copy for backup or archival purposes, creating a test and rollback environment for applications, instantly provisioning application images, and so on
- Two type of snapshots supported:
  - Managed Snapshot – Managed by ESX
    - A maximum of 32 snapshot are supported for fast clones
  - Unmanaged Snapshot – Manage by the storage array
    - Maximum snapshot dictated by the storage array

Managed Snapshot - vSphere

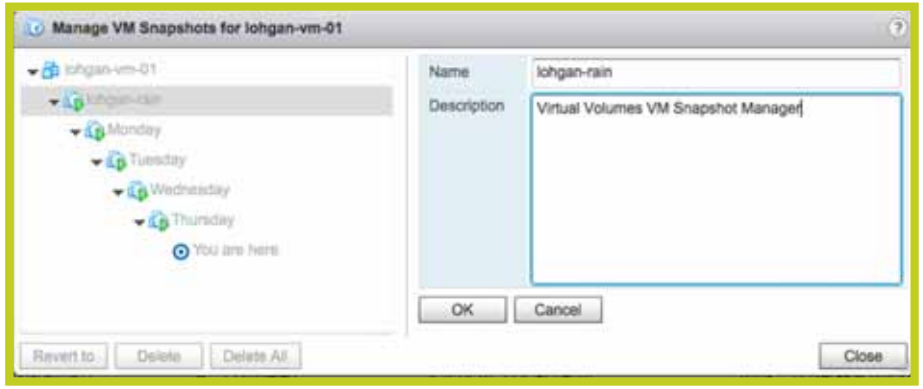

Unmanaged Snapshot - Array

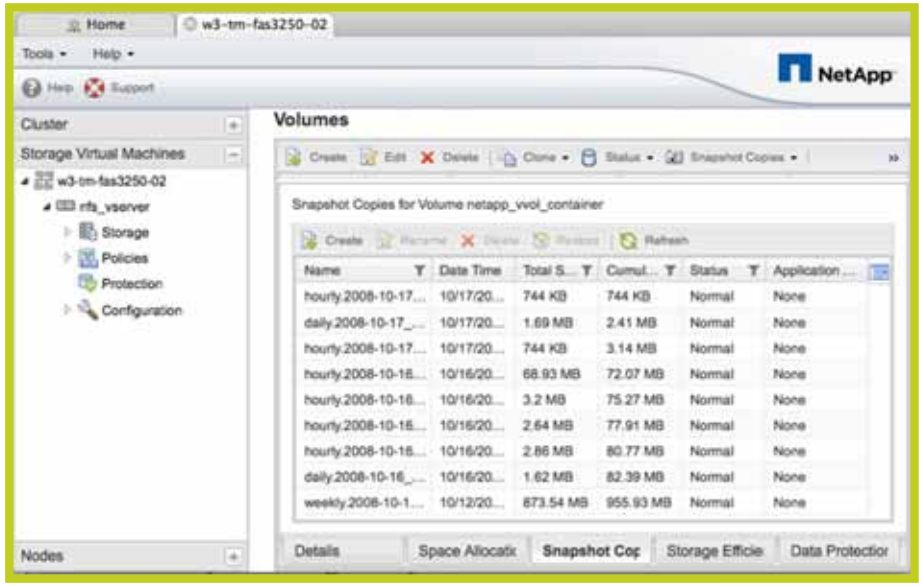

Fast Clone

- **Fast-clones** are an out-of-band space efficient cloning operations, performed exclusively on VMDKs
  - performed on the same storage container using the storage profile of the original virtual volume, or the profile specified
- ESXi hosts guarantee not to issue any I/O to the source or destination virtual volumes during fast clone operations
- **Fast-clones** are similar to the snapshot operation with a few differences:
  - There are no revert operations
  - Space-efficient

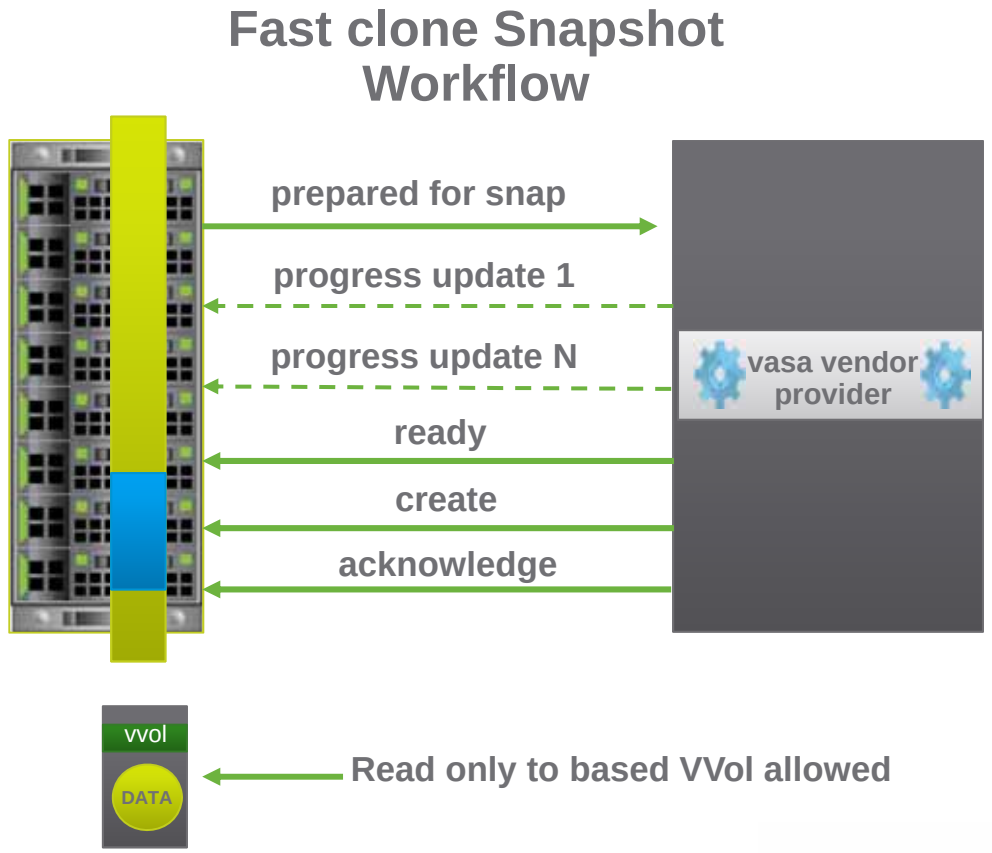

- General
- **Capability sets**
- Default Profiles
- Connectivity with Hosts
- Protocol Endpoints

Name	Description
Sample_Default_Profile for (6090a058-cd..	Default capability profile for storage container

Sample Default Profile for (6090a058-cd89-ffe3-87763007db37] - capabilities

Disk Types	
Disk Encryption	
Dedupe	
Replication	

VAAI vs VVol

1 clone VM from VVol container to different VVol container

vSphere

2 offload to array

cloning operation

3

SC-A SC-B

vendor native clone utilized with VASA primitives

Fully VAAI & VASA APIs Compatible Array

always attempt default operation with VASA API primitives

1 default

fail back 2

If default operations fails VAAI API primitives are used

1 clone VM from VVol container to different VVol container

vSphere

3

2 offload to array

vmkernel data mover uses VAAI primitives for cloning operation

SC-A SC-B

Fully VAAI & VASA APIs Compatible Array