SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Private Cloud Storage Management using SMI-S, Windows Server, and System Center

Hector Linares
Principal Program Manager
hectorl@microsoft.com
Microsoft

Rajesh Balwani
Senior Development Engineer
rbalwani@microsoft.com
Microsoft

Agenda

Customer Pain Points

System Center Overview

SAN and NAS Management

Private Cloud Storage Provisioning

Fibre Channel Fabric Management

Applications with FC Storage

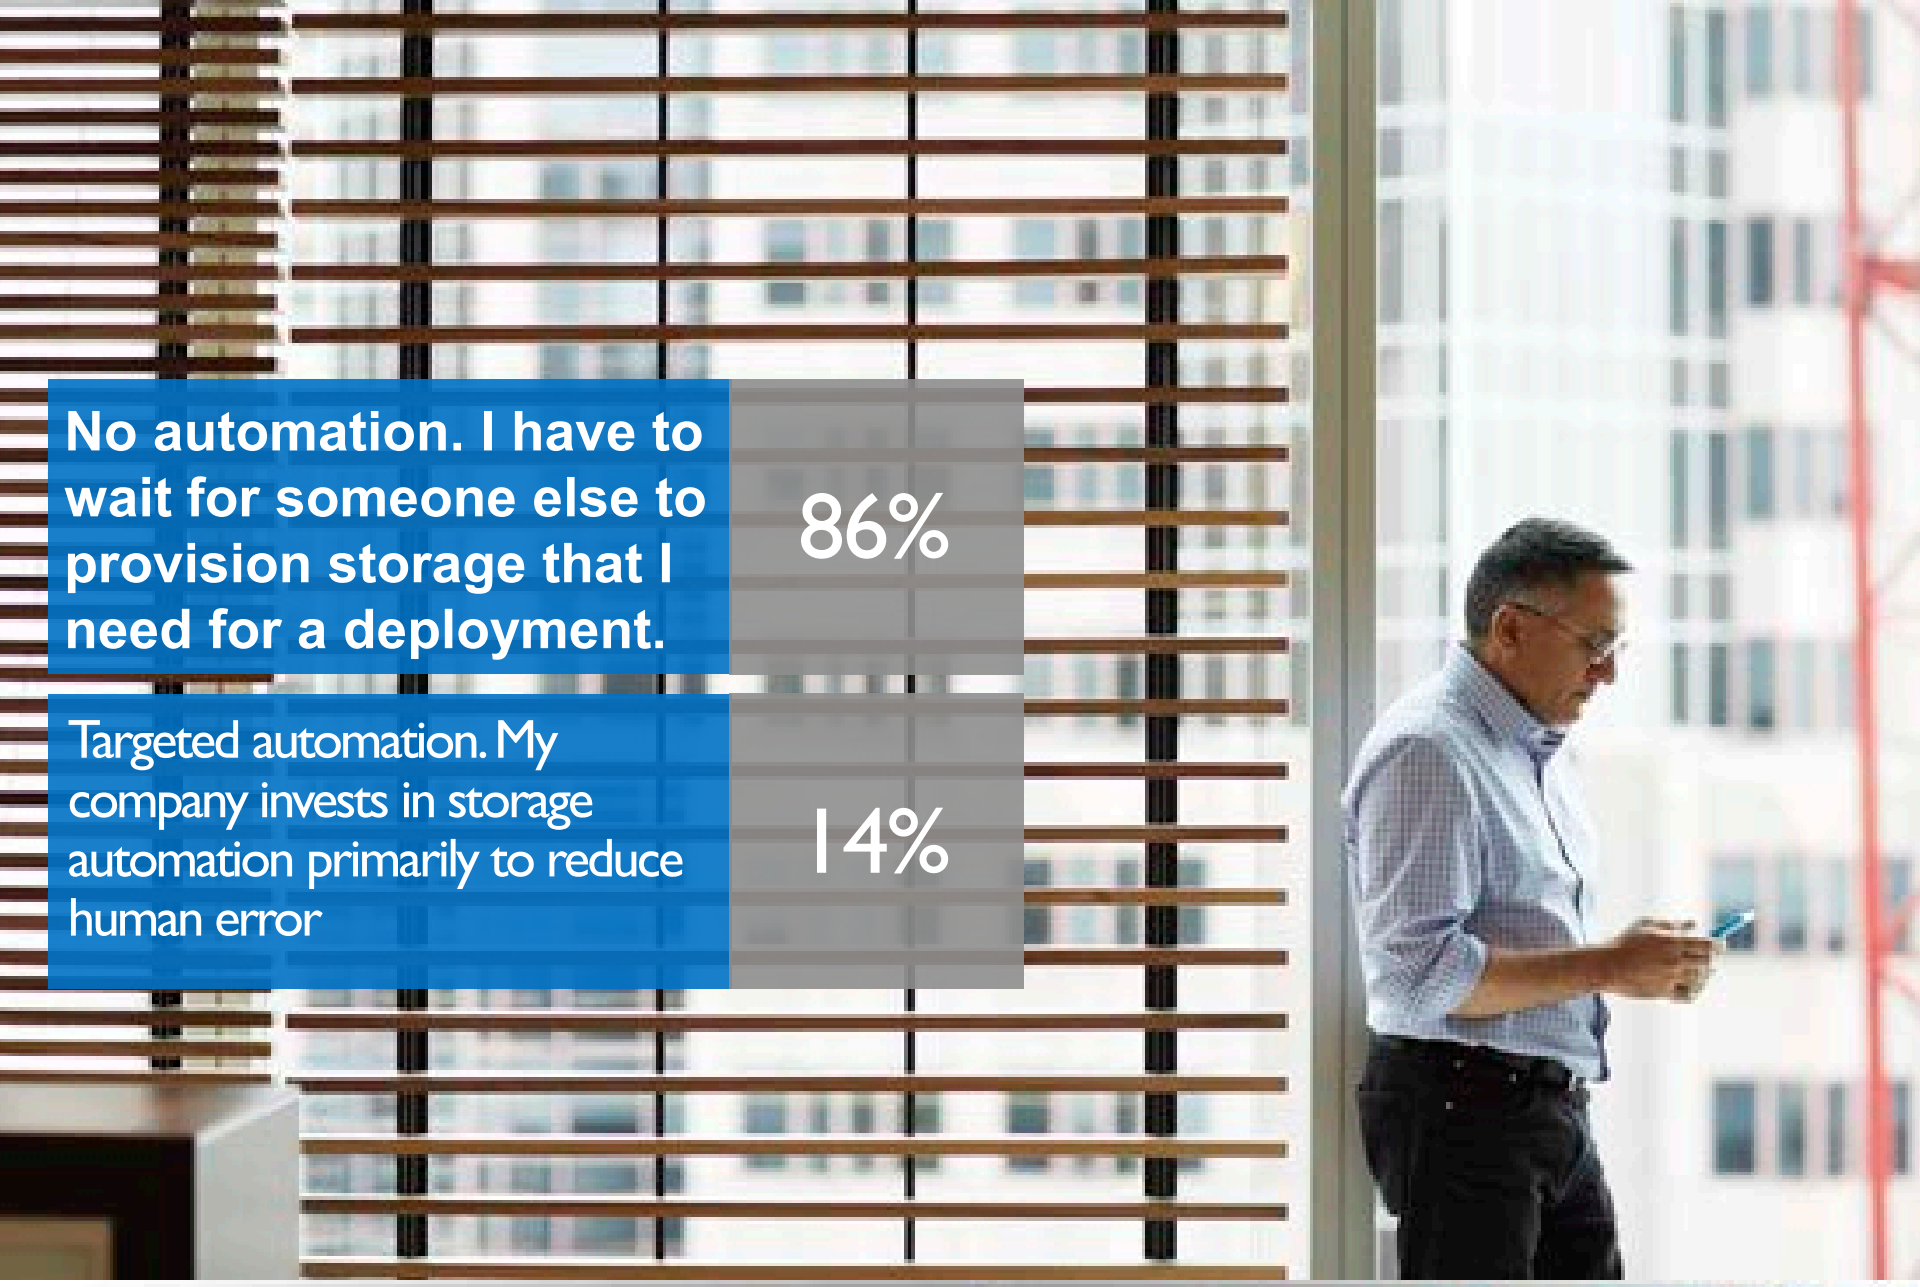

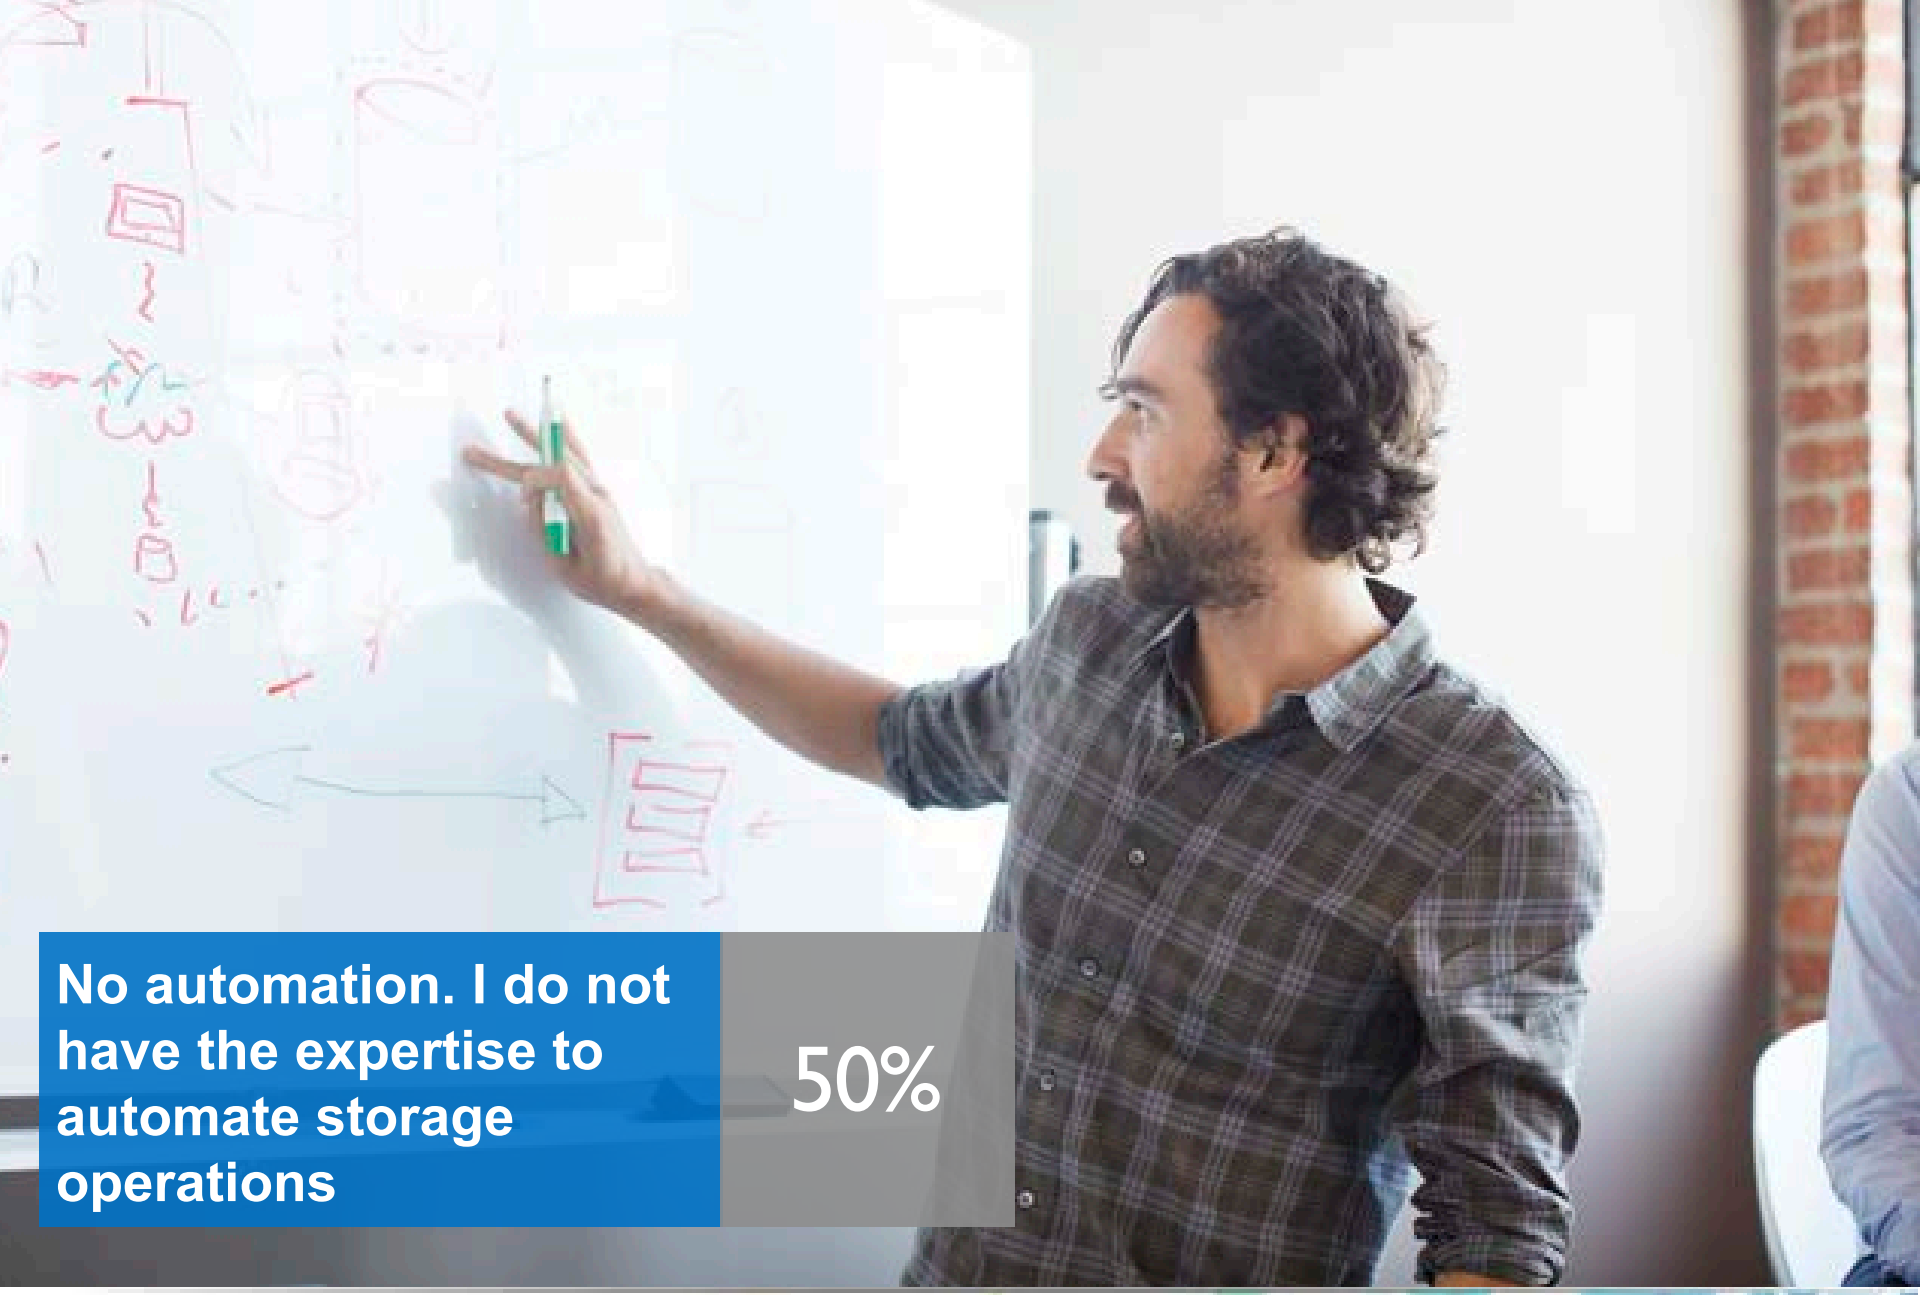

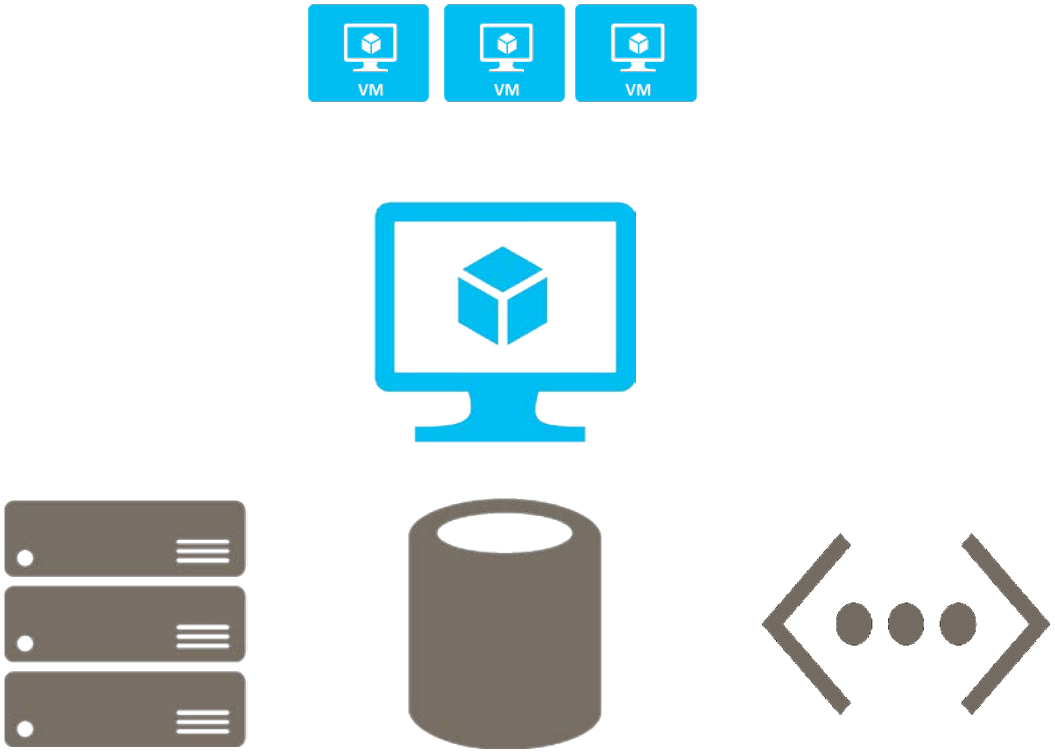

System Center capabilities

Infrastructure provisioning	Infrastructure monitoring	Automation and self-service	Application performance monitoring	IT service management
Enterprise-class multi-tenant infrastructure for hybrid environments	Comprehensive monitoring of physical, virtual, and cloud infrastructure	Application-owner agility while IT retains control	Deep insight into application health	Flexible service delivery

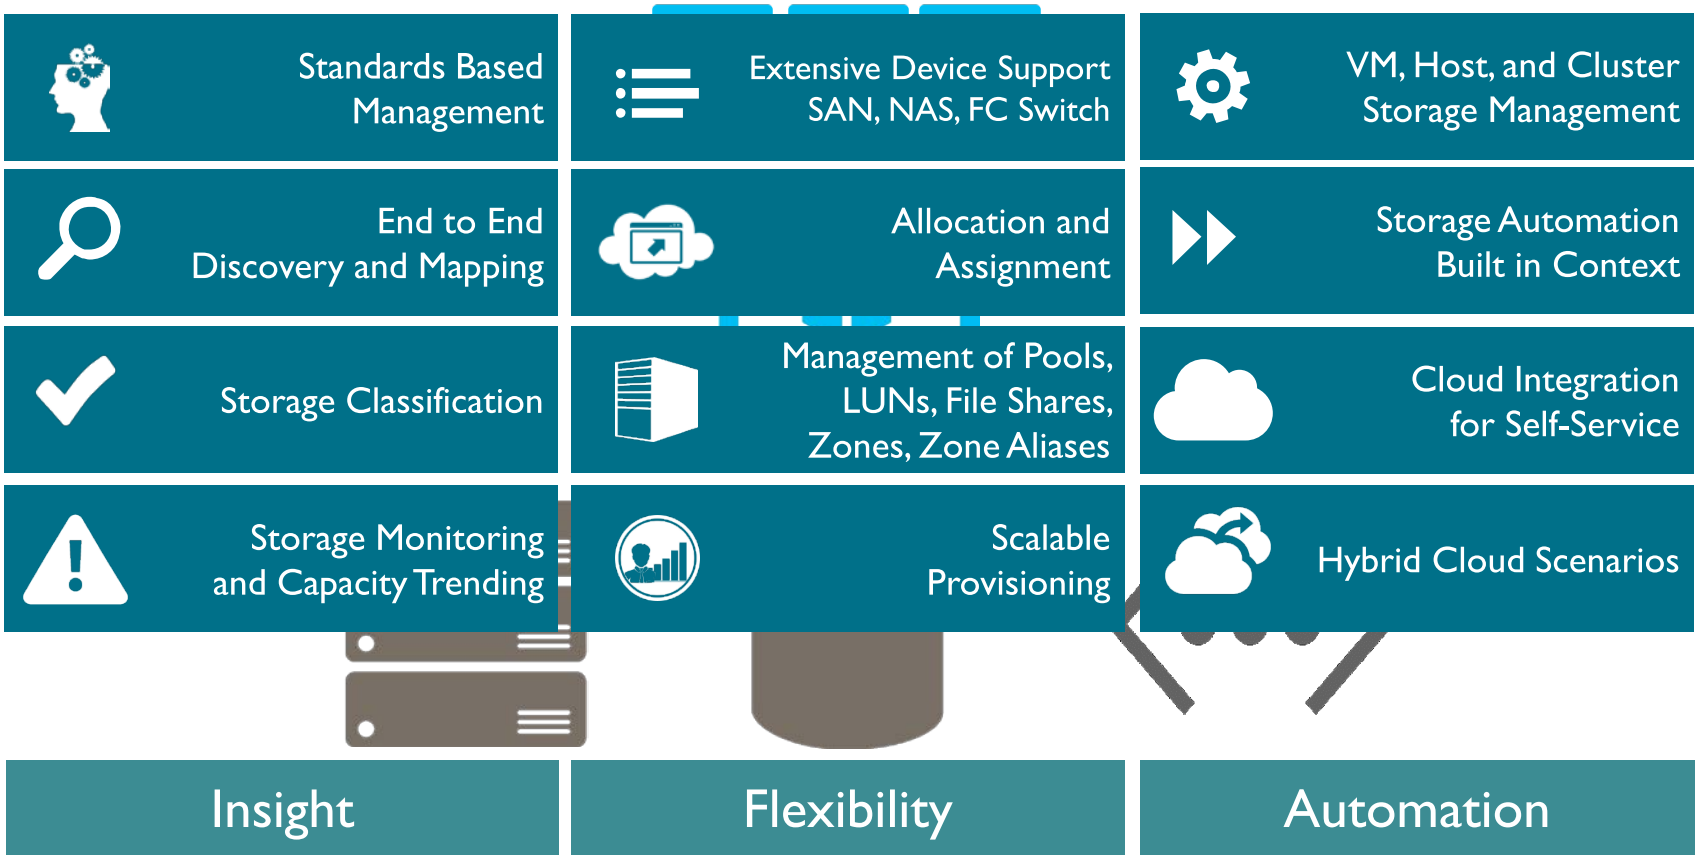

SAN Management

SAN (SMI-S)

Discovery array, pool, storage volume, initiators, targets, masking sets
Create/delete storage pool, storage volume, snapshot clone

Hyper-V Server

Discovery of FC and IQN ports
Connect array to server (zoning, initiator logon)

Virtual Machine

Connect array to server (zoning, initiator logon)
Unmask/mask storage volume to VM

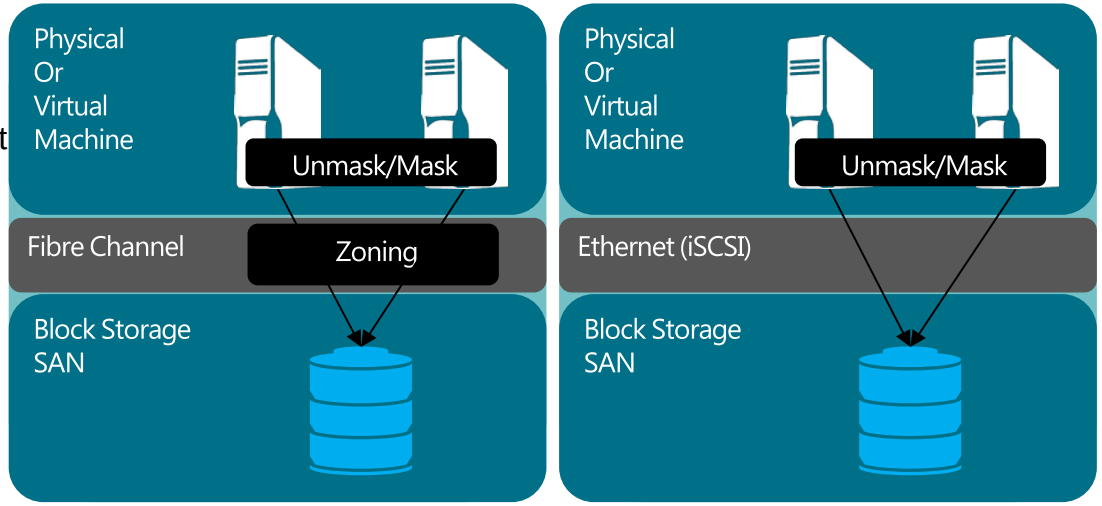

Add Storage Array Provider

```
$raa = Get-SCRunAsAccount -Name “NTAP”

Add-SCStorageProvider -Name "NETAPPSMI" -RunAsAccount $raa -NetworkDeviceName
“http://NETAPPSMI.CONTOSO.LAB" -TCPPort 5988
```

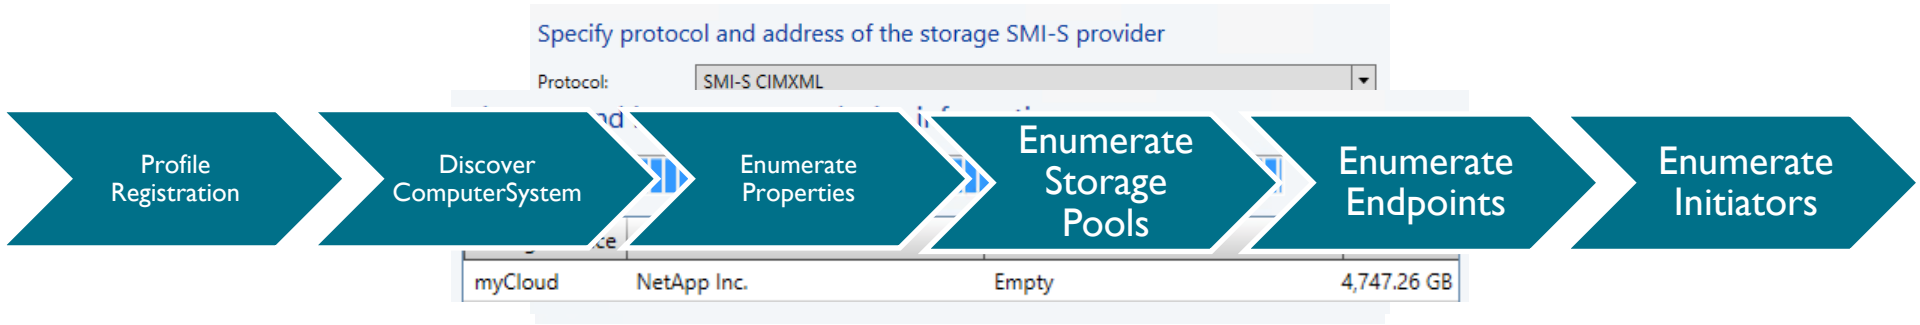

NetworkAddress : http://netappsmi.contoso.lab
TCPPort : 5988
ProviderFlags : StorageArray
RunAsAccount : ntap
StorageArrays : {myCloud}

Registration Requirements (Level 0)

Interop namespace objects

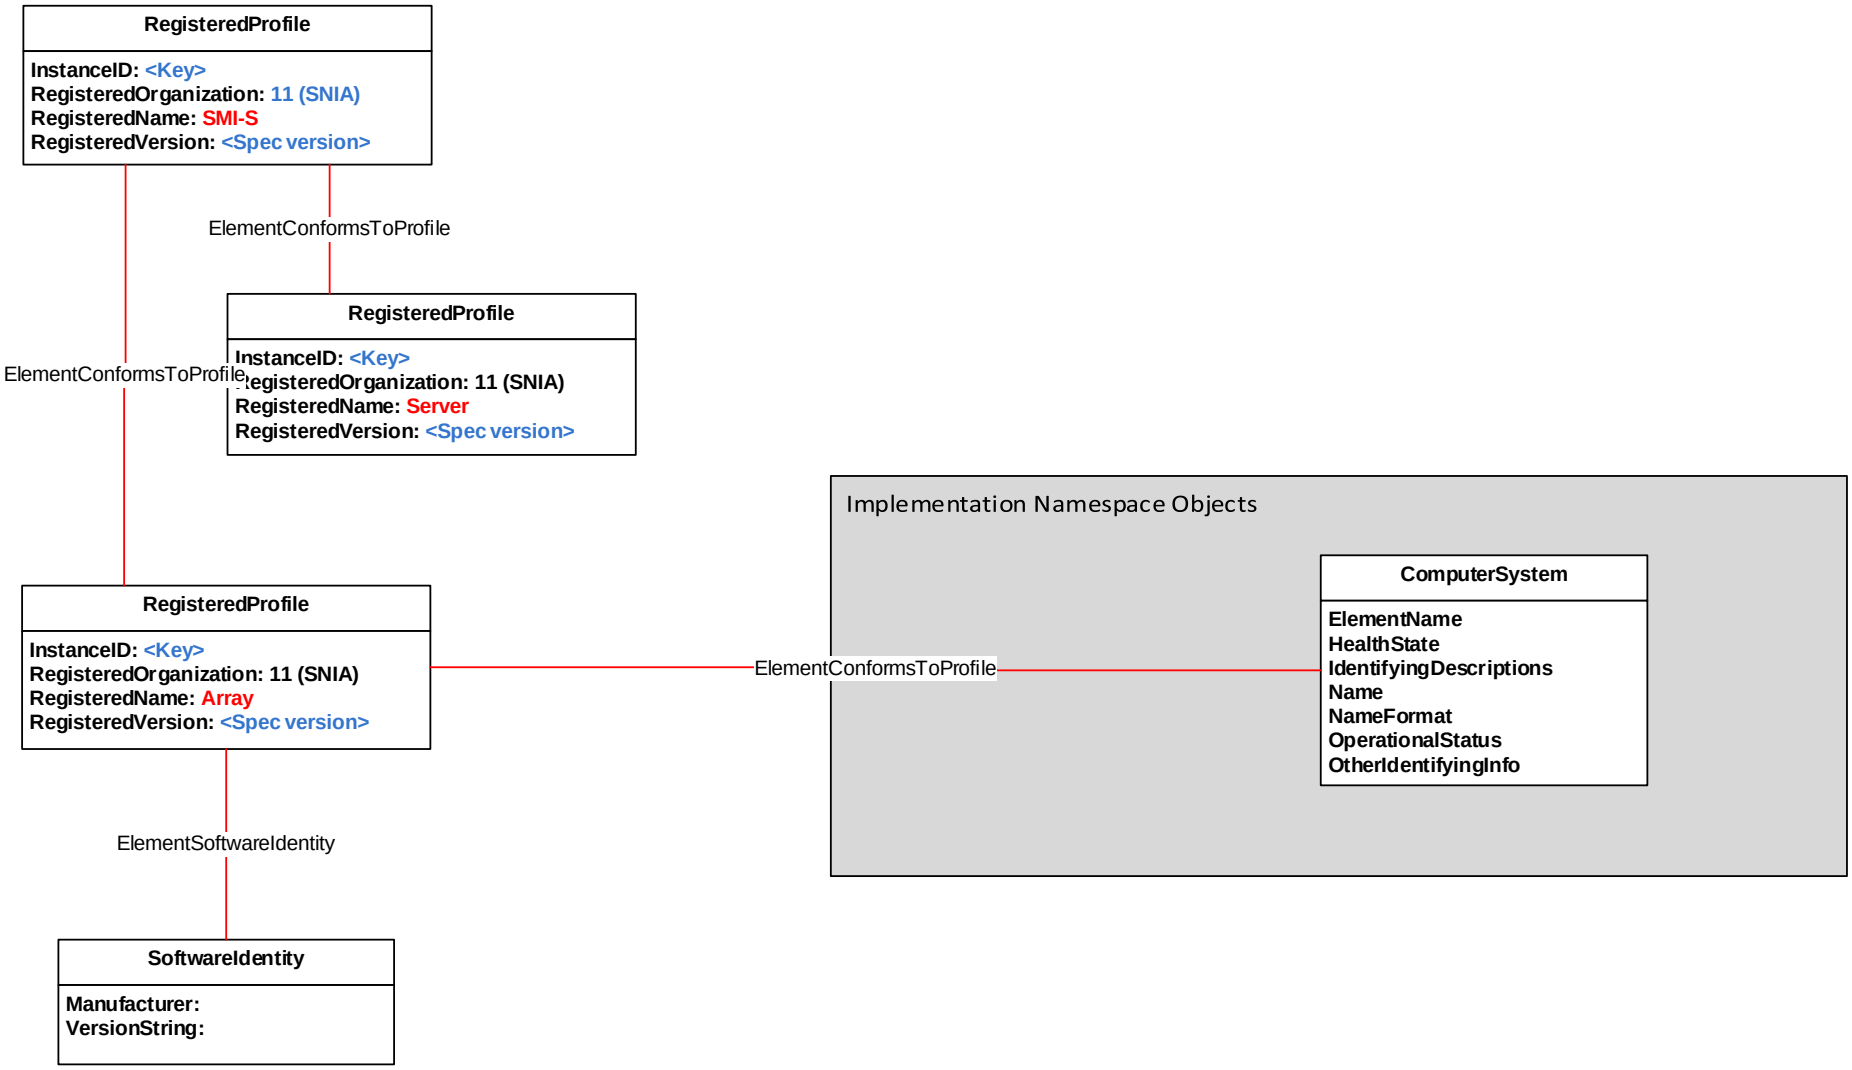

* Make sure **key** properties are filled up appropriately. This diagram shows properties that “Windows Standards-Based Storage Management service” requires.

Capabilities and Services that may be traversed at Level 0

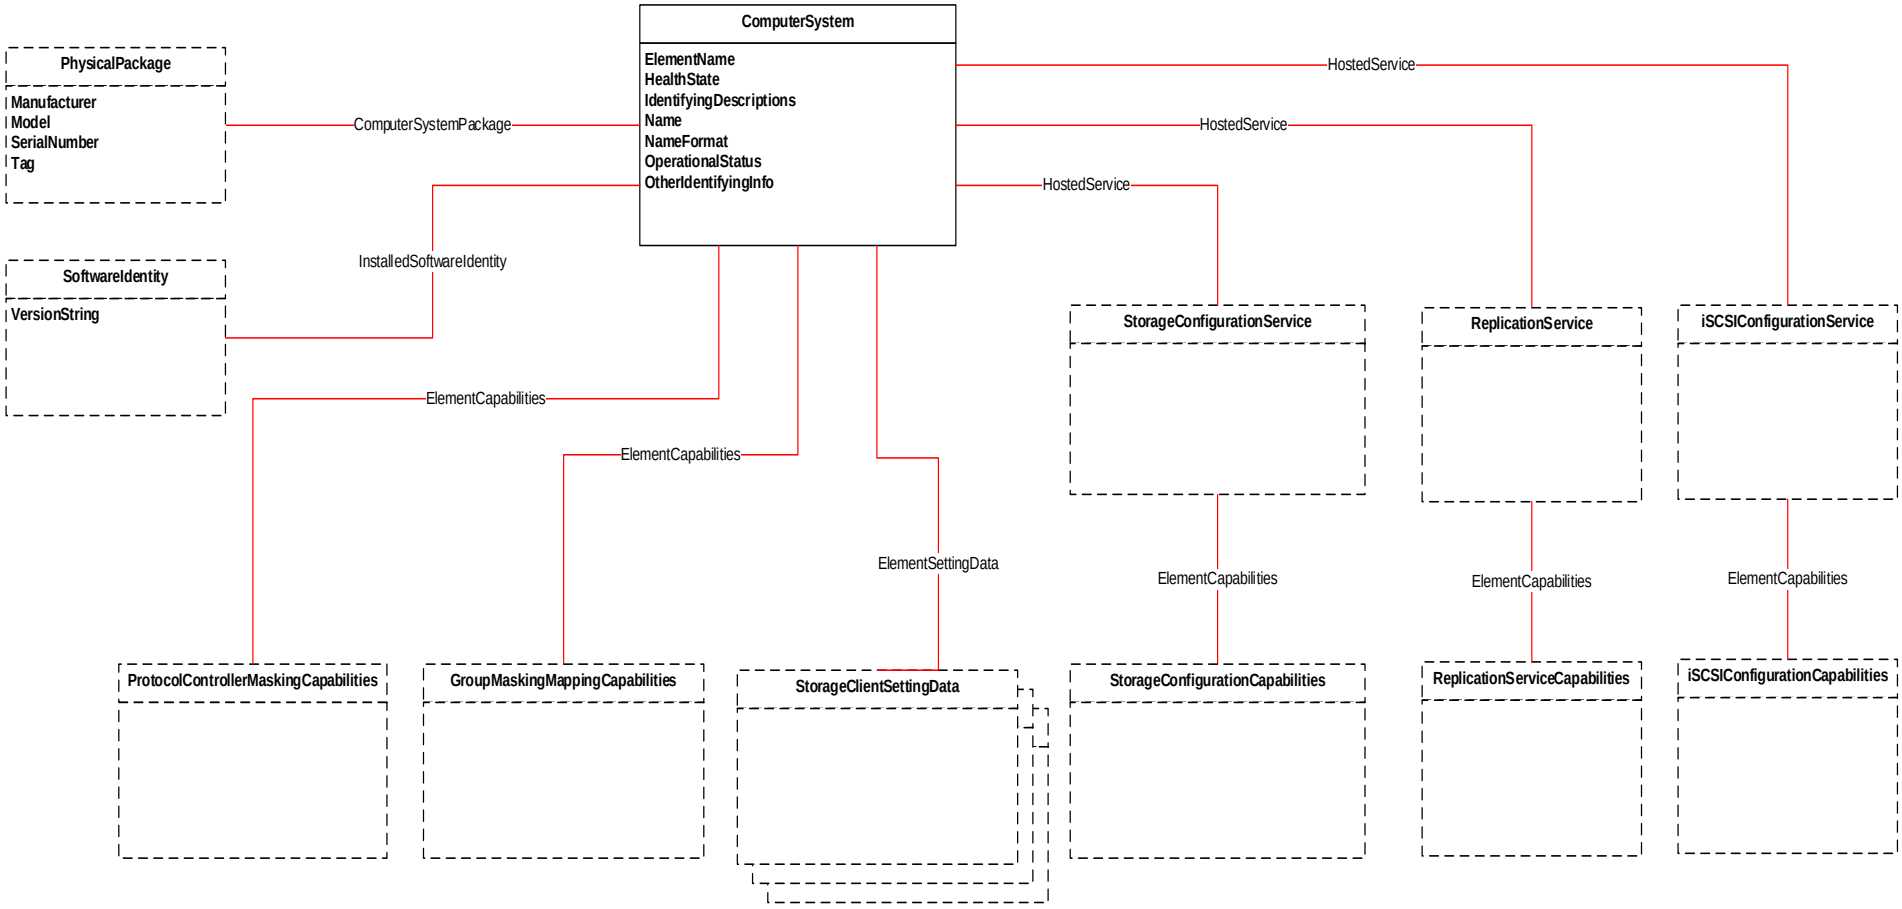

* Make sure **key** properties are filled up appropriately. This diagram shows properties that "Windows Standards-Based Storage Management service" requires.

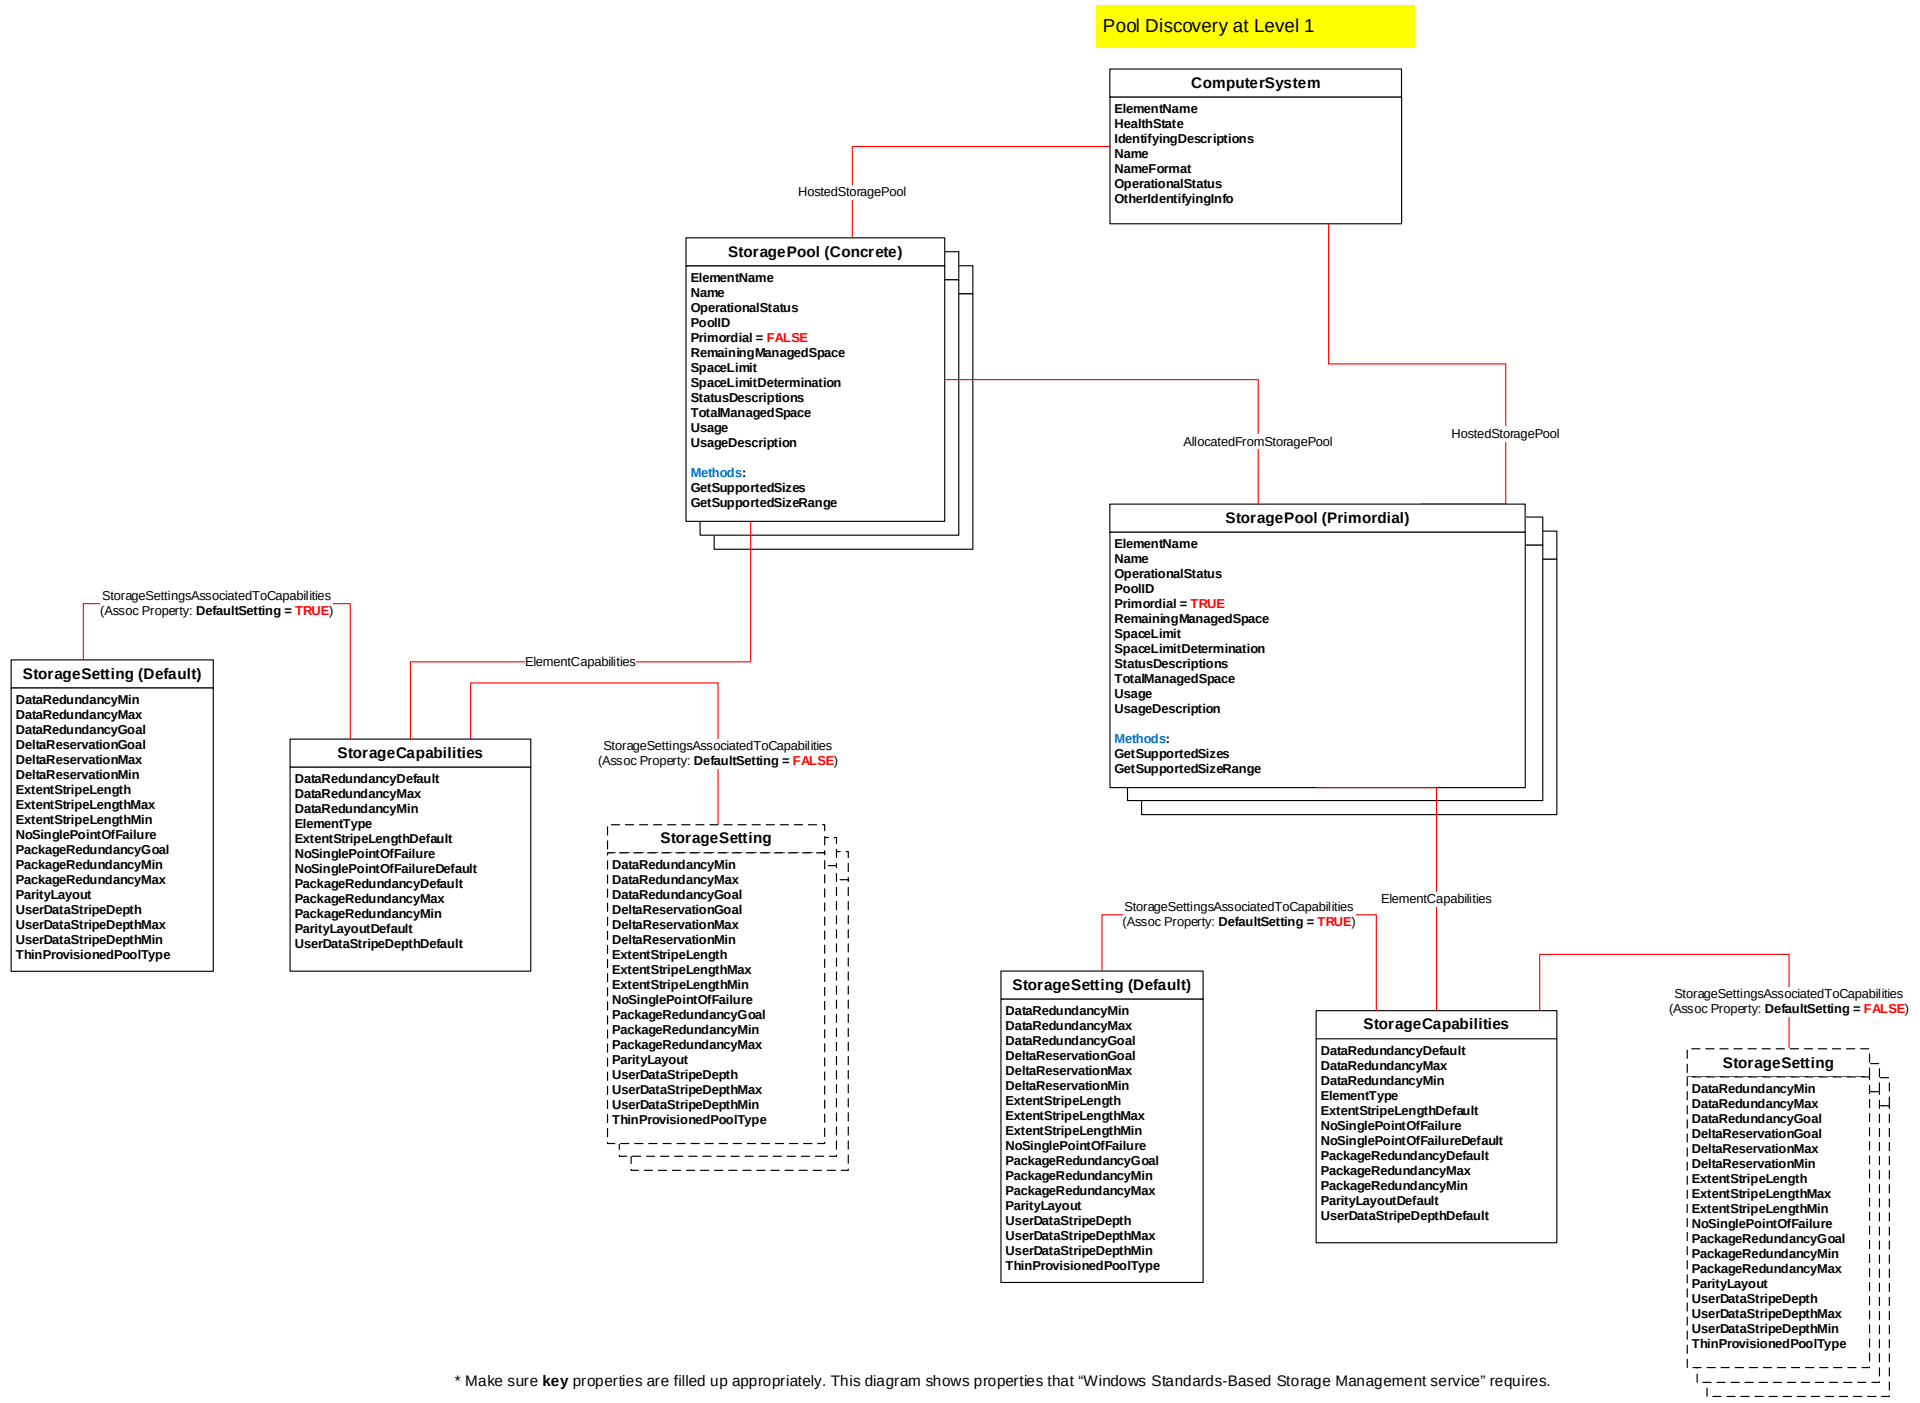

* Make sure **key** properties are filled up appropriately. This diagram shows properties that “Windows Standards-Based Storage Management service” requires.

Endpoints Discovery at Level 1

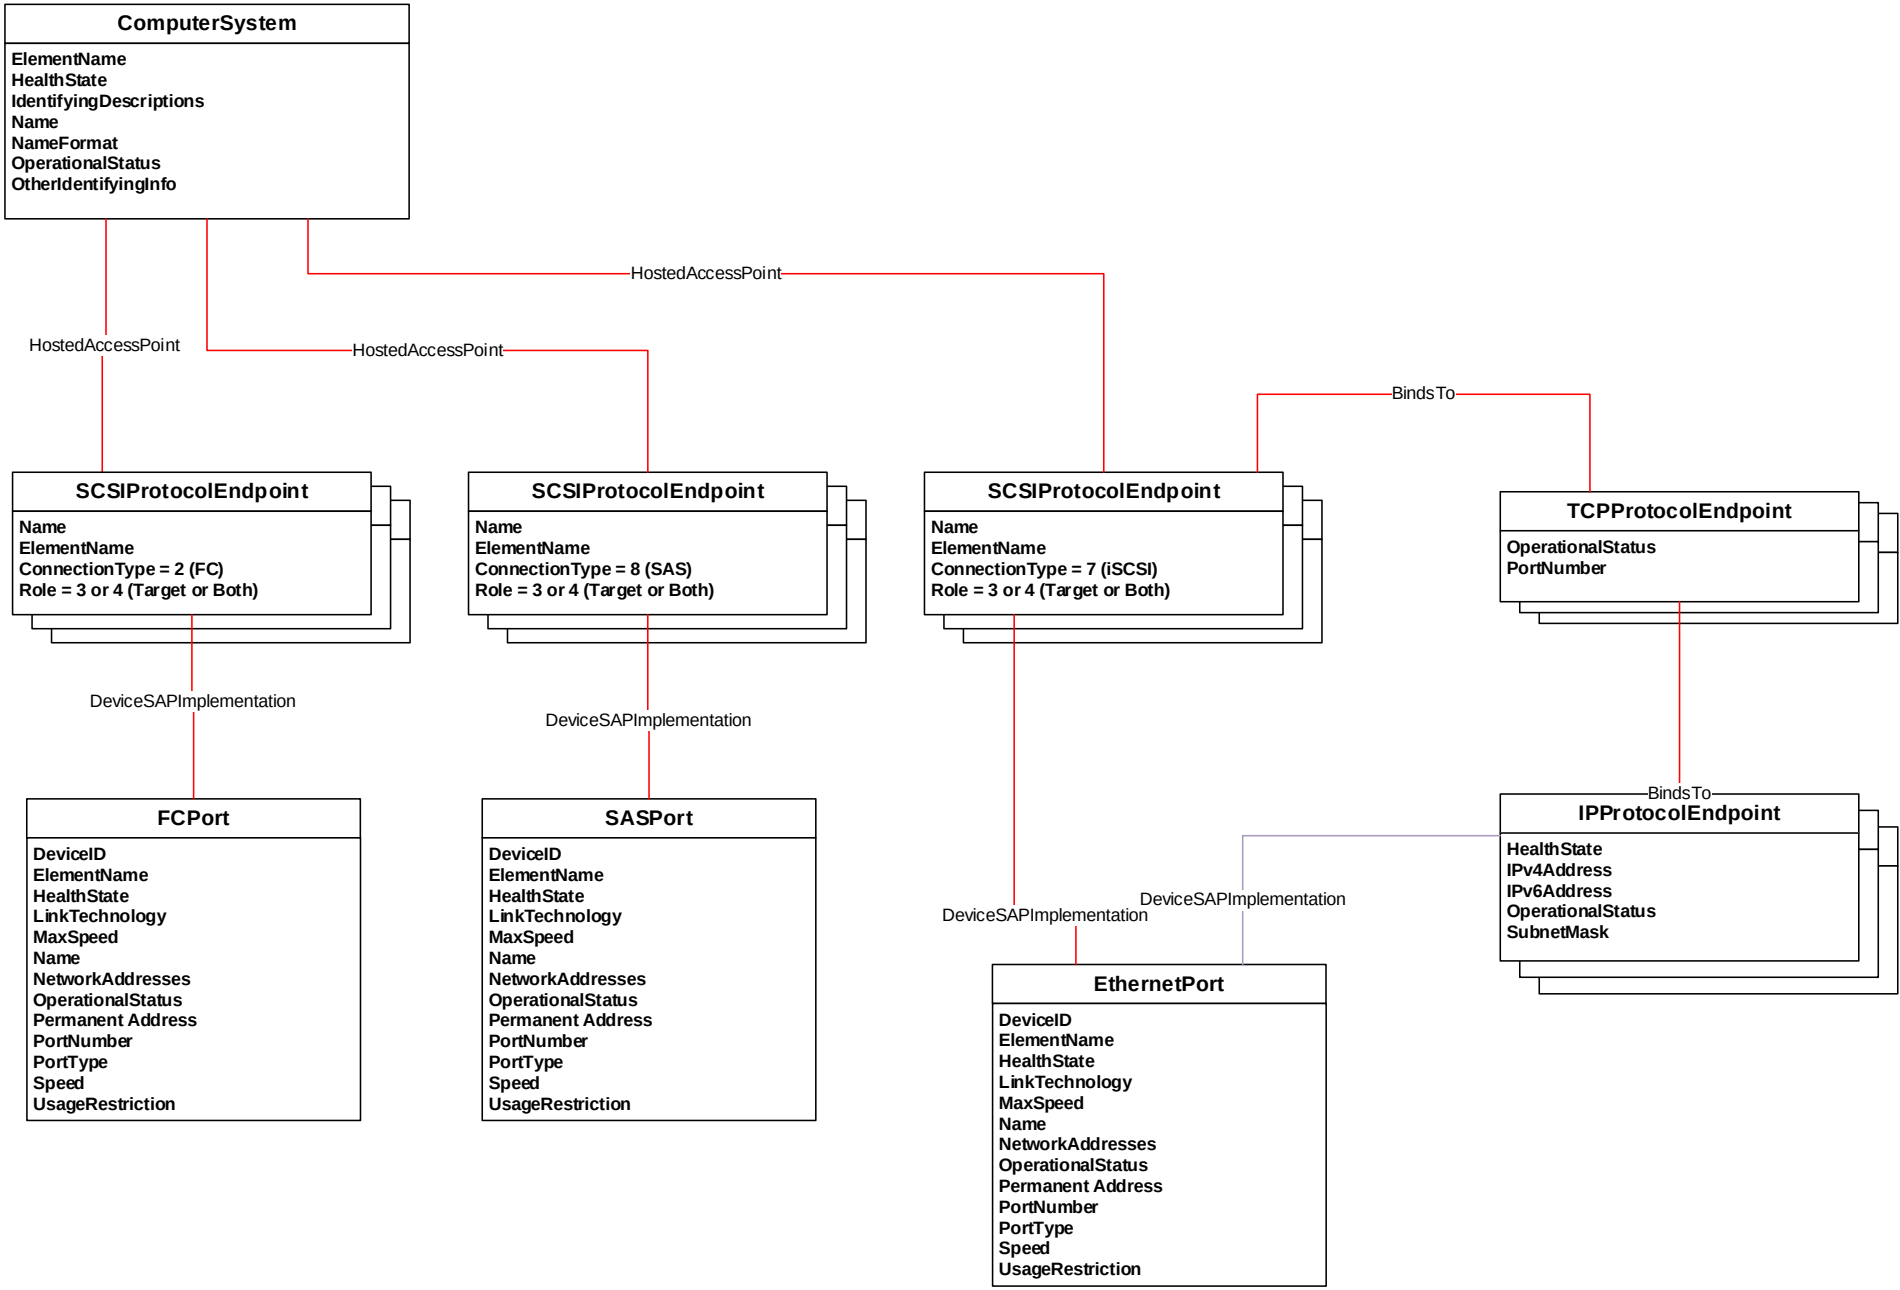

StorageHardwareIDs Discovery at Level 1

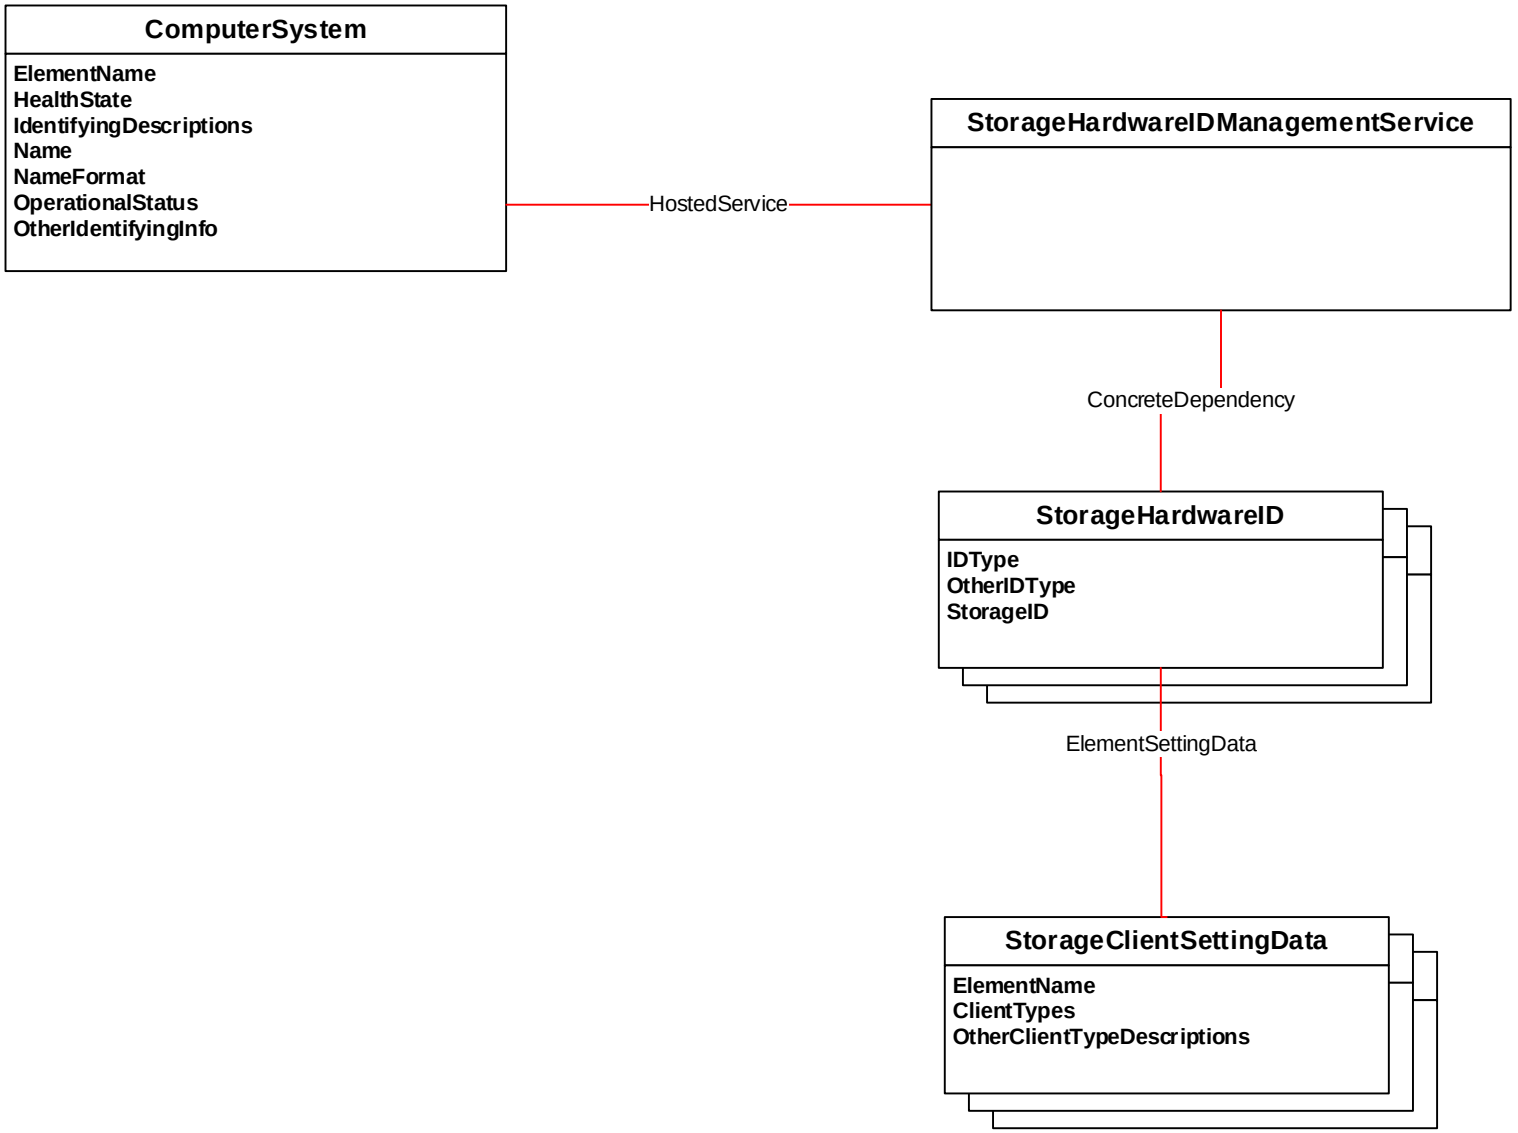

* Make sure **key** properties are filled up appropriately. This diagram shows properties that “Windows Standards-Based Storage Management service” requires.

Manage Storage Pool

```
$pool = Get-SCStoragePool -Name "vol2“
$classification = Get-SCStorageClassification -Name "HI-IOPS“
Set-SCStorageArray $pool.StorageArray -AddStoragePoolToManagement $pool
-StorageClassificationAssociation $classification
```

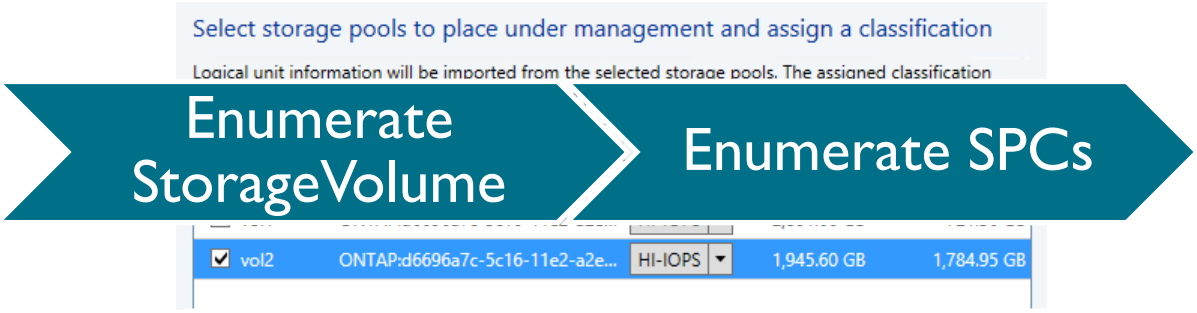

SMDisplayName : vol2
Primordial : False
TotalManagedSpace : 2089072095232
Usage : Unrestricted
ProvisioningTypeDefault : Thin
SupportedProvisioningTypes : {Fixed, Thin}
StorageArray : myCloud
StorageLogicalUnits : {/vol/vol2/LU1-2ed3524-TR1701, /vol/vol2/DISK12,/vol/vol2/DISK03, /vol/vol2/GUFS...}

StorageVolumes Discovery at Level 2

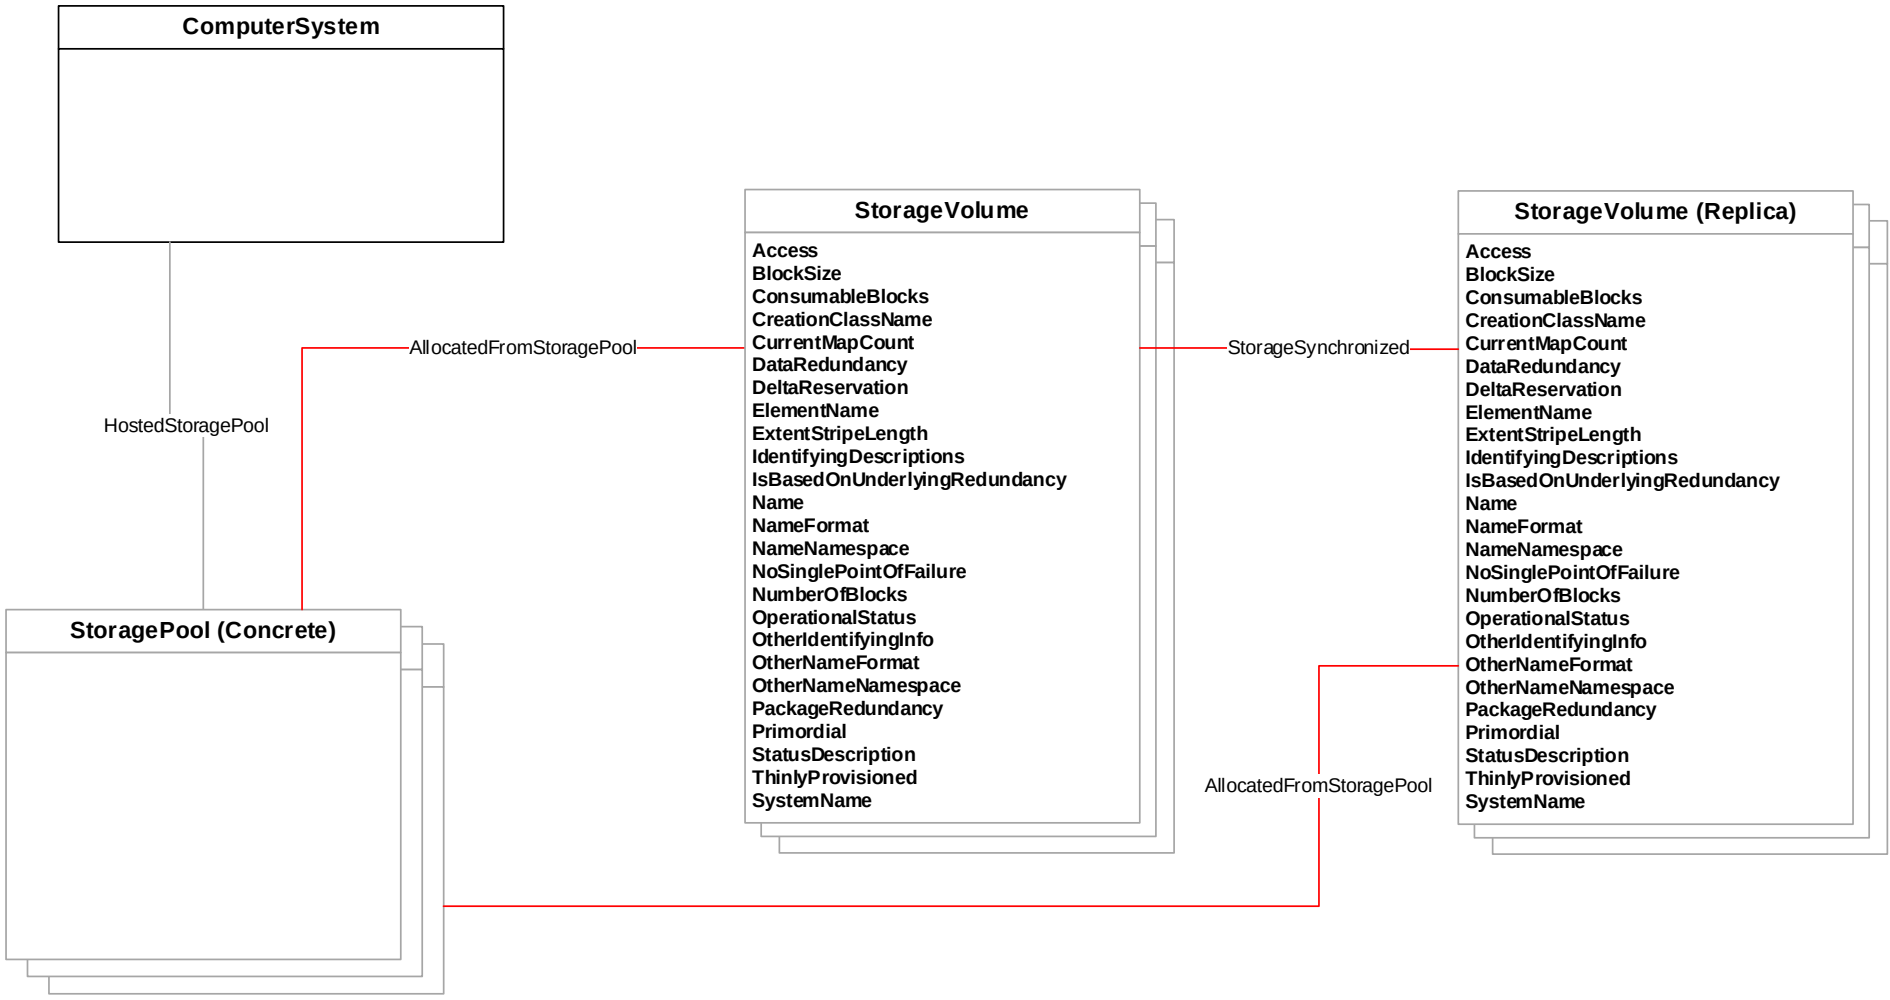

* Make sure **key** properties are filled up appropriately. This diagram shows properties that “Windows Standards-Based Storage Management service” requires.

SCSIProtocolController Discovery at Level 2

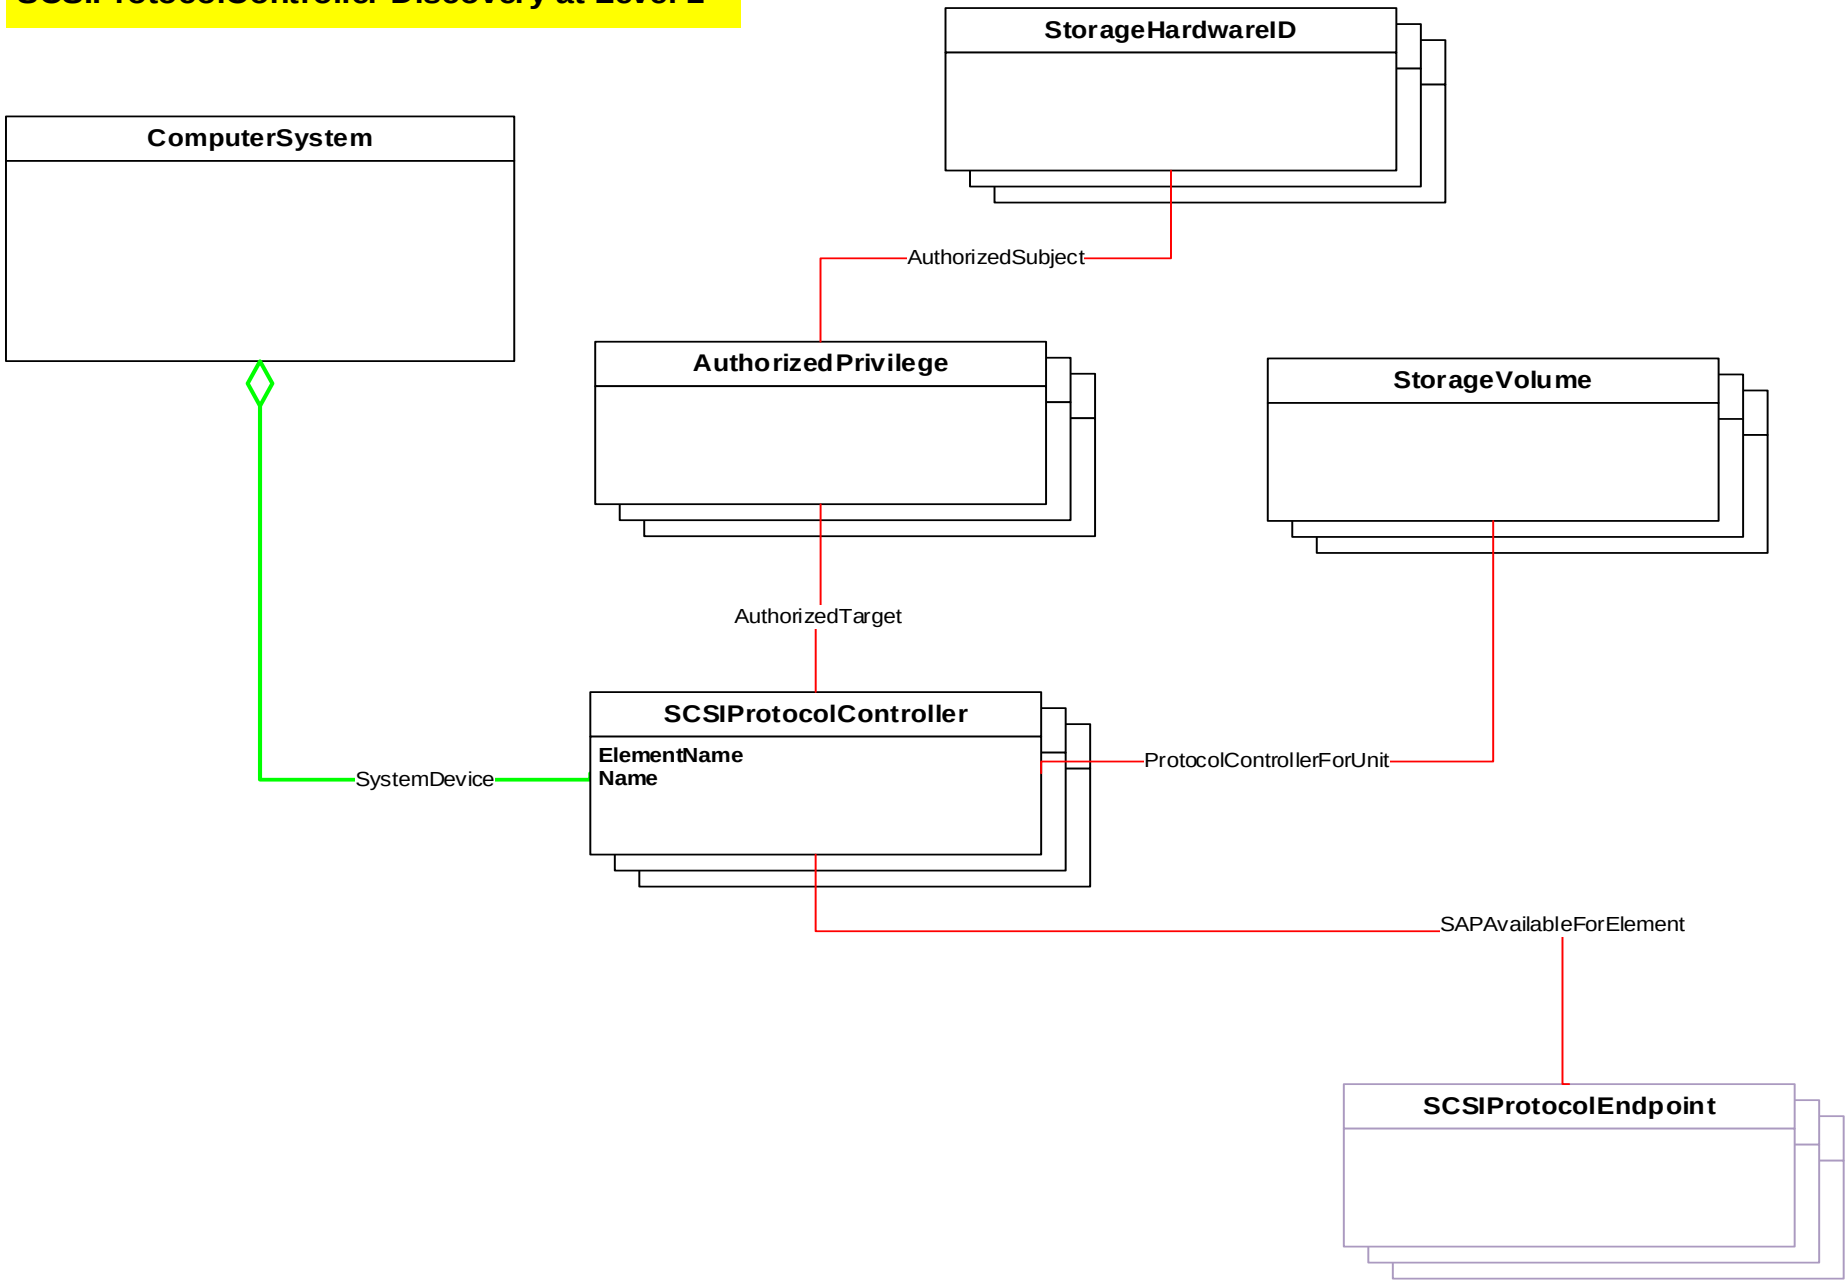

Make sure **key** properties are filled up appropriately. This diagram shows properties that “Windows Standards-Based Storage Management service” requires.

List Storage Volumes

```
(Get-SCStorageLogicalUnit | where {$_.StoragePool -eq $pool})[0]
```

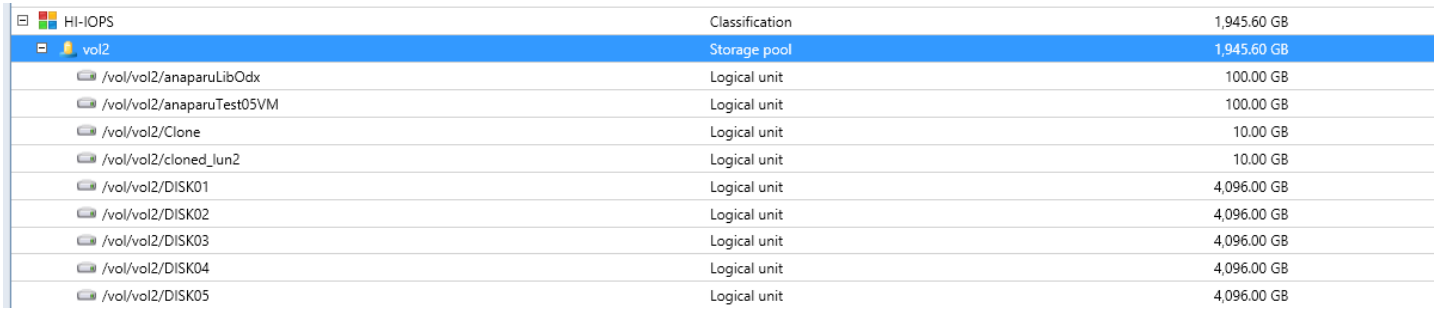

HI-IOPS	Classification	1,945.60 GB
vol2	Storage pool	1,945.60 GB
/vol/vol2/anaparuLibOdx	Logical unit	100.00 GB
/vol/vol2/anaparuTest05VM	Logical unit	100.00 GB
/vol/vol2/Clone	Logical unit	10.00 GB
/vol/vol2/cloned_lun2	Logical unit	10.00 GB
/vol/vol2/DISK01	Logical unit	4,096.00 GB
/vol/vol2/DISK02	Logical unit	4,096.00 GB
/vol/vol2/DISK03	Logical unit	4,096.00 GB
/vol/vol2/DISK04	Logical unit	4,096.00 GB
/vol/vol2/DISK05	Logical unit	4,096.00 GB

```
SMName                          : b2c70854-874c-11e2-a2e0-123478563412
SMLunId                         : 600a09802d655533655d43423570646d
SMLunIdFormat                   : 9
SMLunIdFormatDescription        : NAA
SMLunIdNamespace                : 2
SMLunIdNamespaceDescription     : VPD83Type3
TotalCapacity                   : 107374182400
StoragePool                     : vol2
Classification                  : HI-IOPS
StorageGroups                   : {vmmlab1823n3}
```

SAN and NAS Management

Storage Management API (SMAPI)

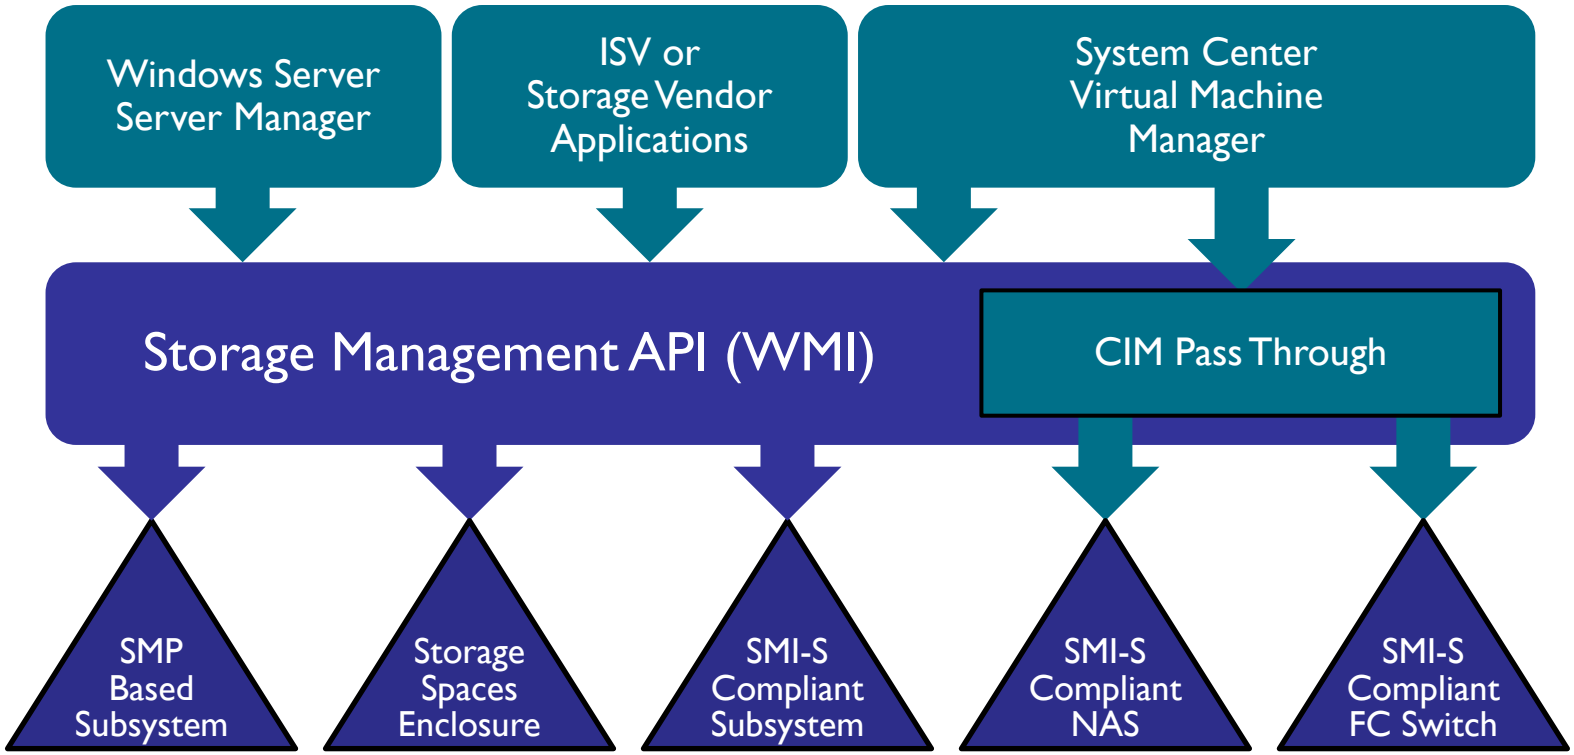

Architecture

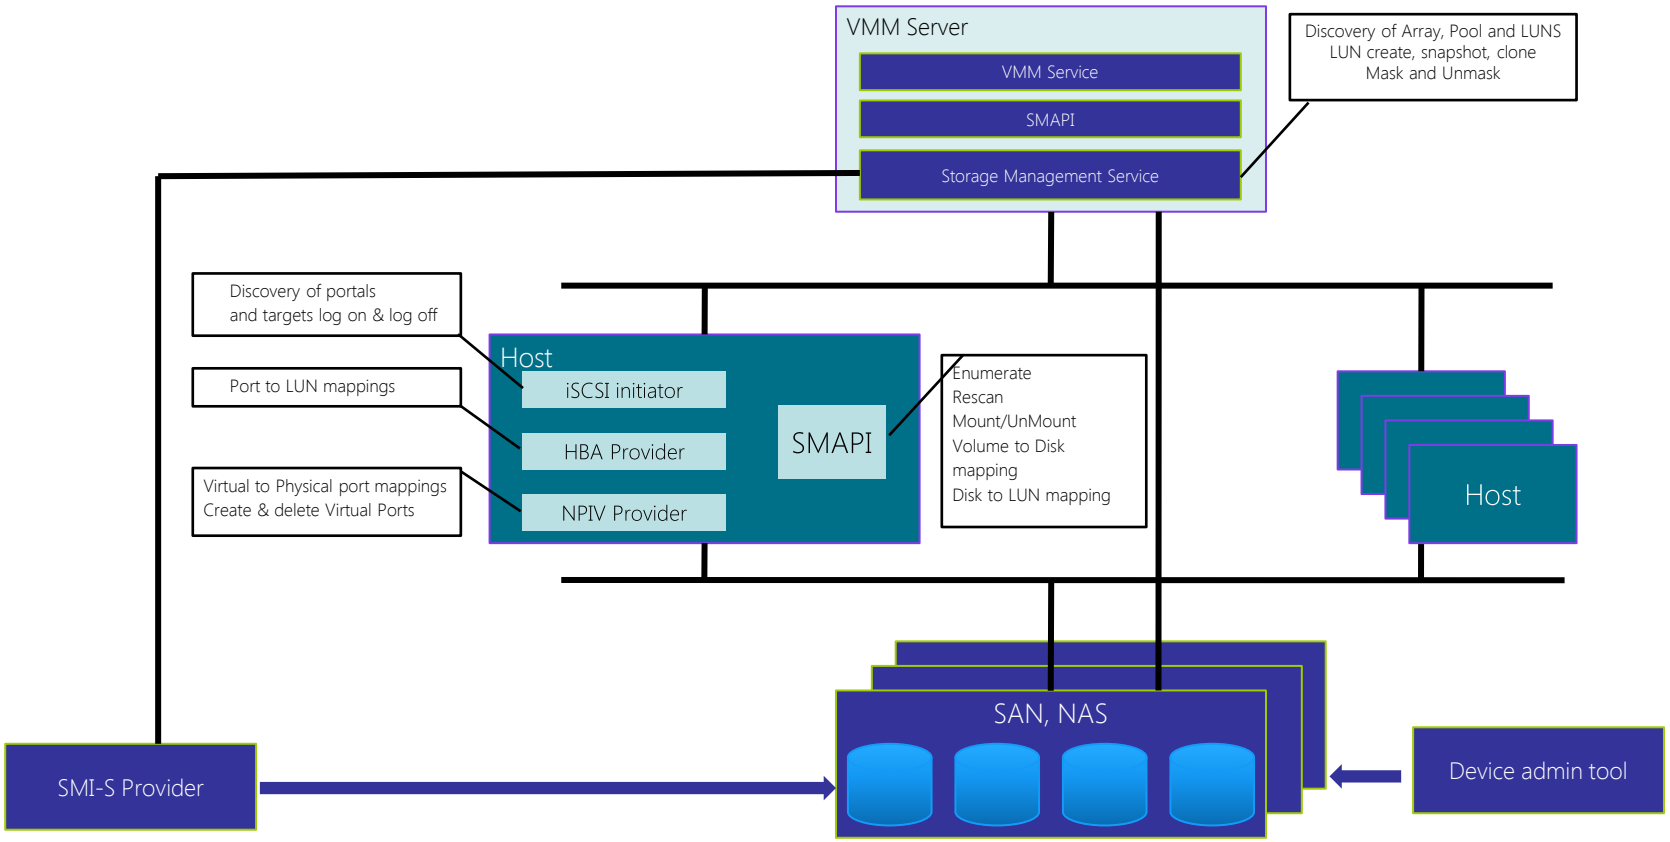

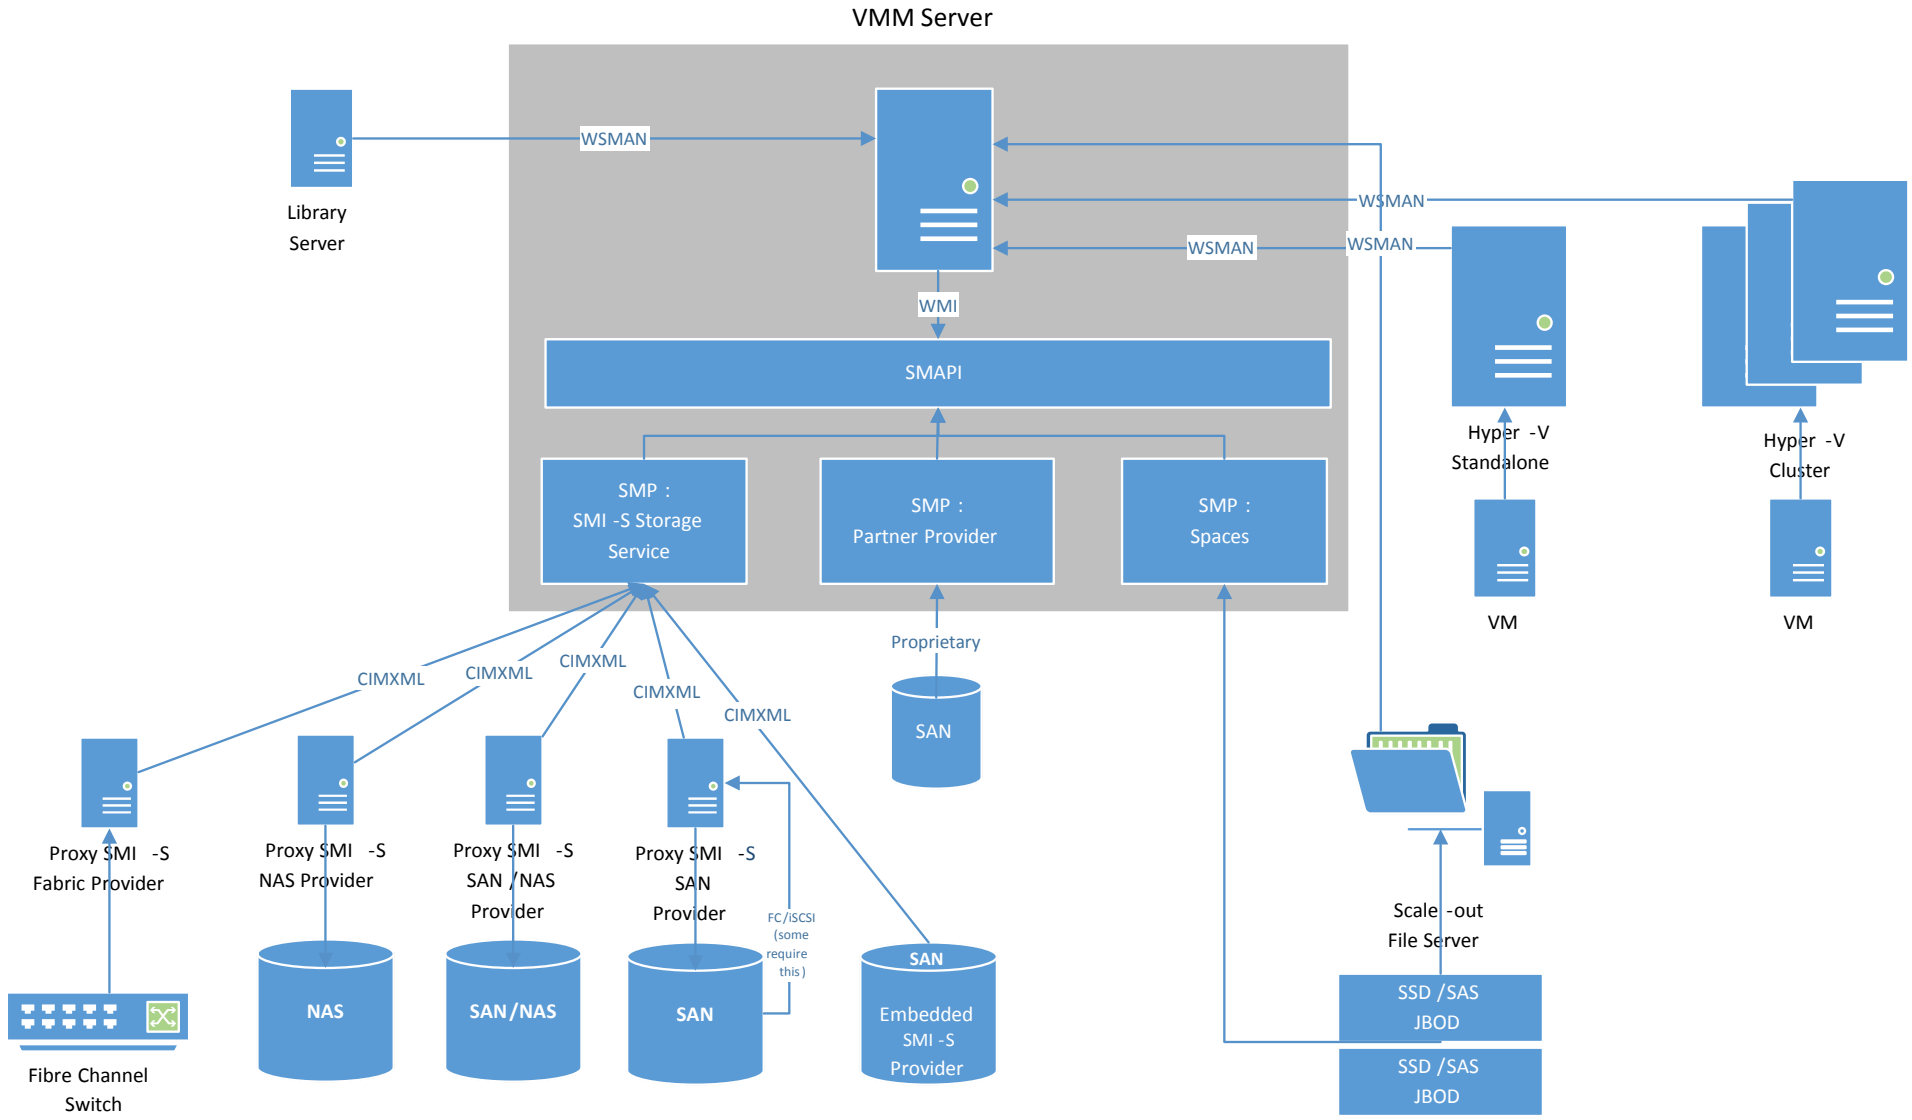

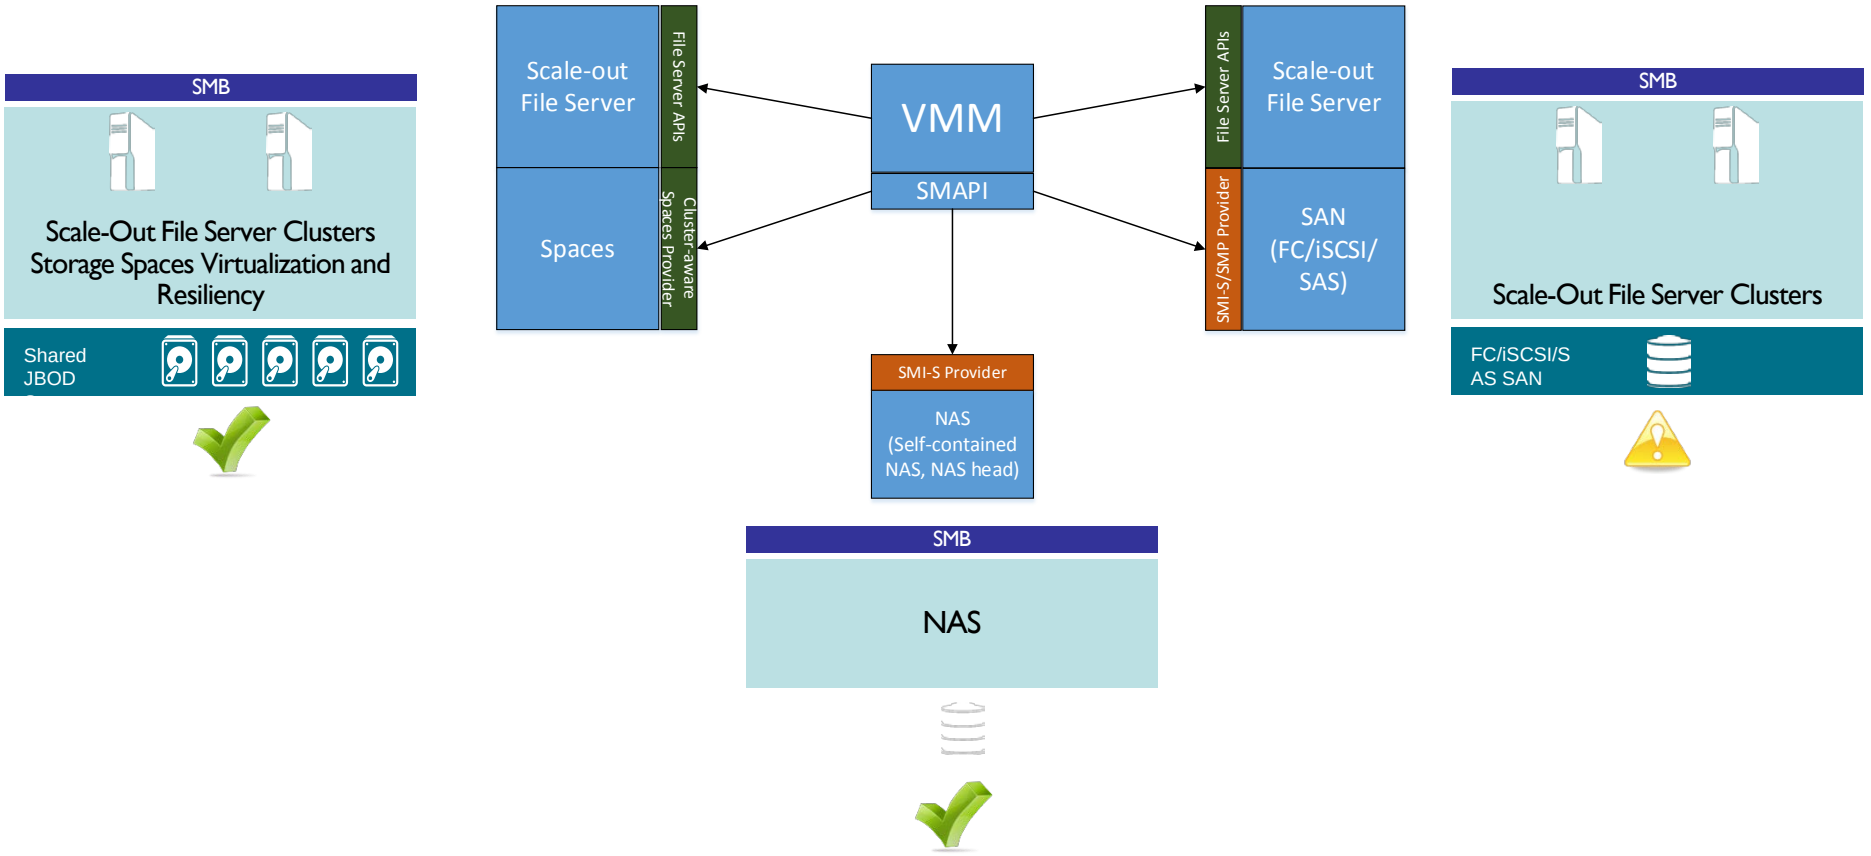

Private Cloud Storage Provisioning

Storage Classification

Assign Classification to Storage

Create multiple classifications

Classify discovered storage pool, disk, and file share

Storage classification inheritance

Streamline VM Deployment

Express storage intent in templates

Placement is classification aware

Differentiate Capacity in Multi-Tenant Environment

Create clouds with specific classification

Self-service users restricted to allocated storage based on classification

Assign Classification to Storage

Storage Pool Classification

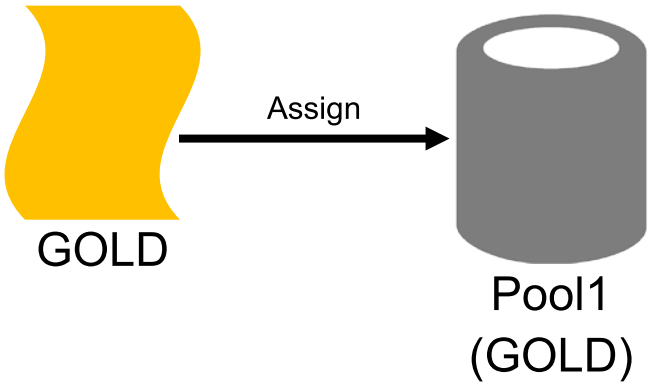

Assign Classification to Storage

Storage Pool Discovery

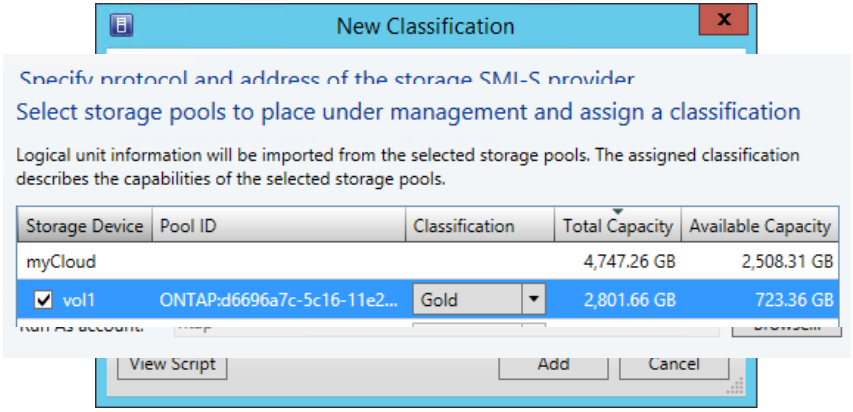

Assign Classification to Storage

Inherited Classification

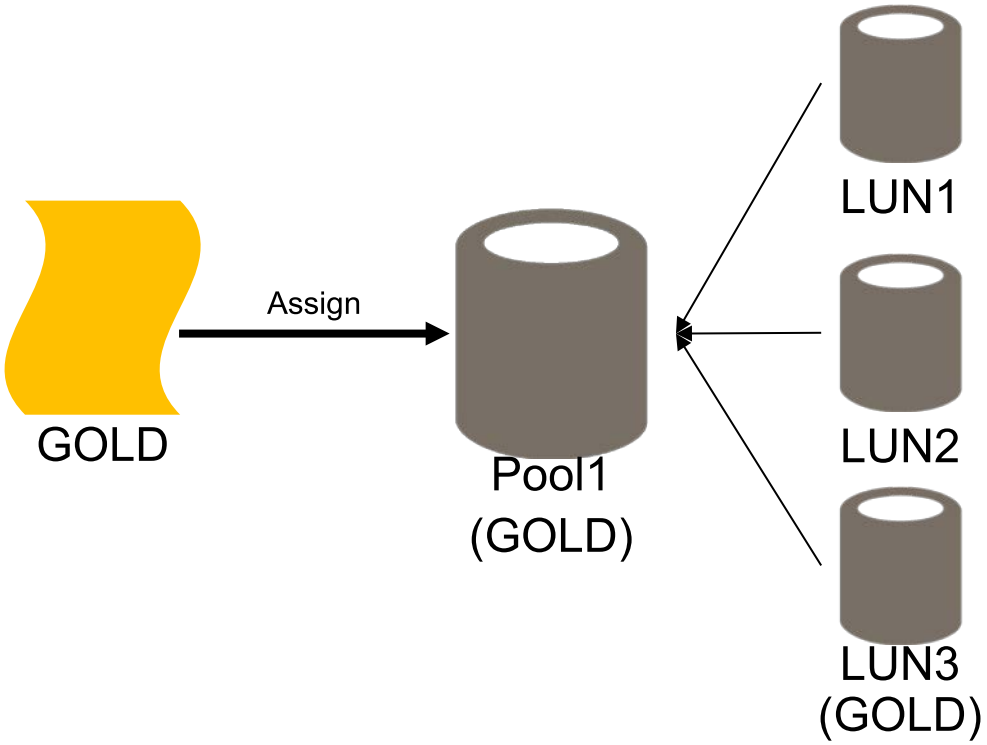

Assign Classification to Storage

Storage Disk Classification

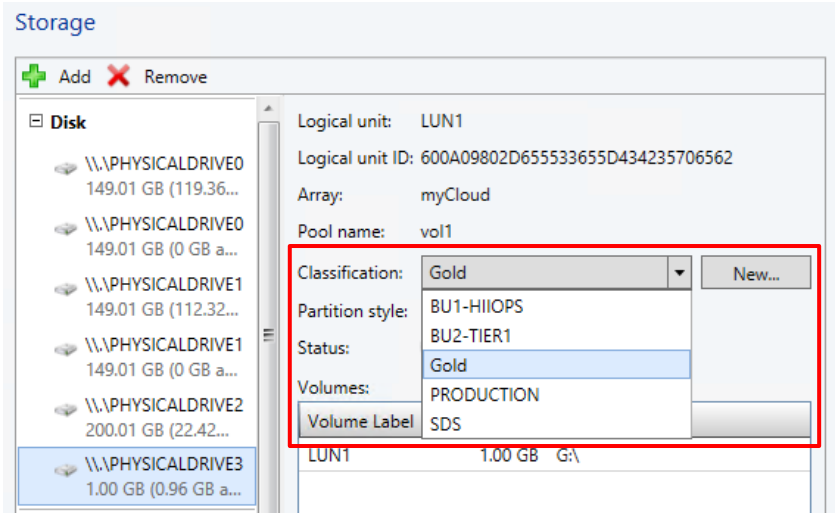

Allocate Storage to Cloud

Cloud Scope	Template Scope	Instance Scope
No classification assigned to cloud	All classifications available on host group	Placement selected storage based on available capacity
Gold classification assigned to cloud	Limited to Gold classification	Placement selected Gold storage with available capacity
Gold and Silver classification assigned to cloud	Gold and Silver classification available	Placement selected storage based on selection in template and available capacity

Assign Classification to Storage

File Share Classification

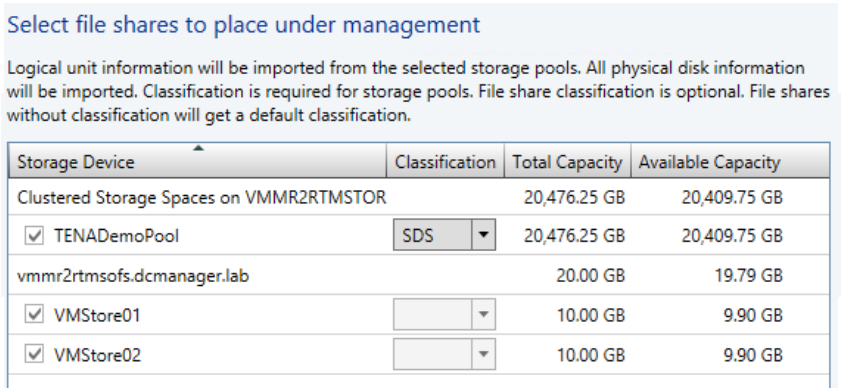

Assign Classification to Storage

Inherited Classification

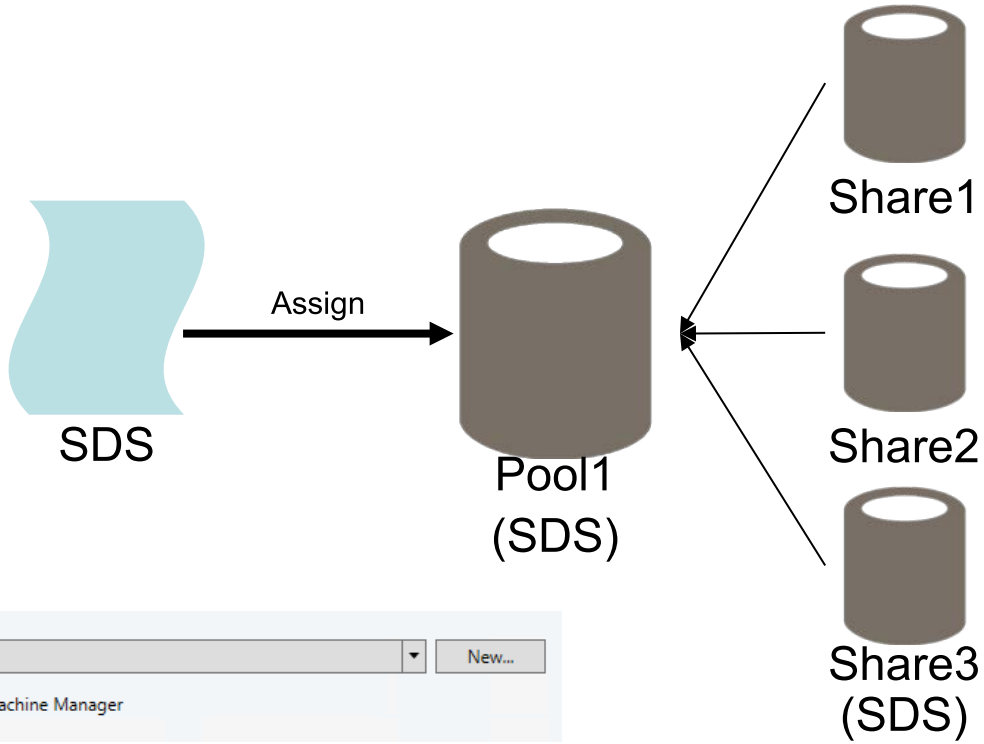

Fibre Channel Fabric Management

Topologies

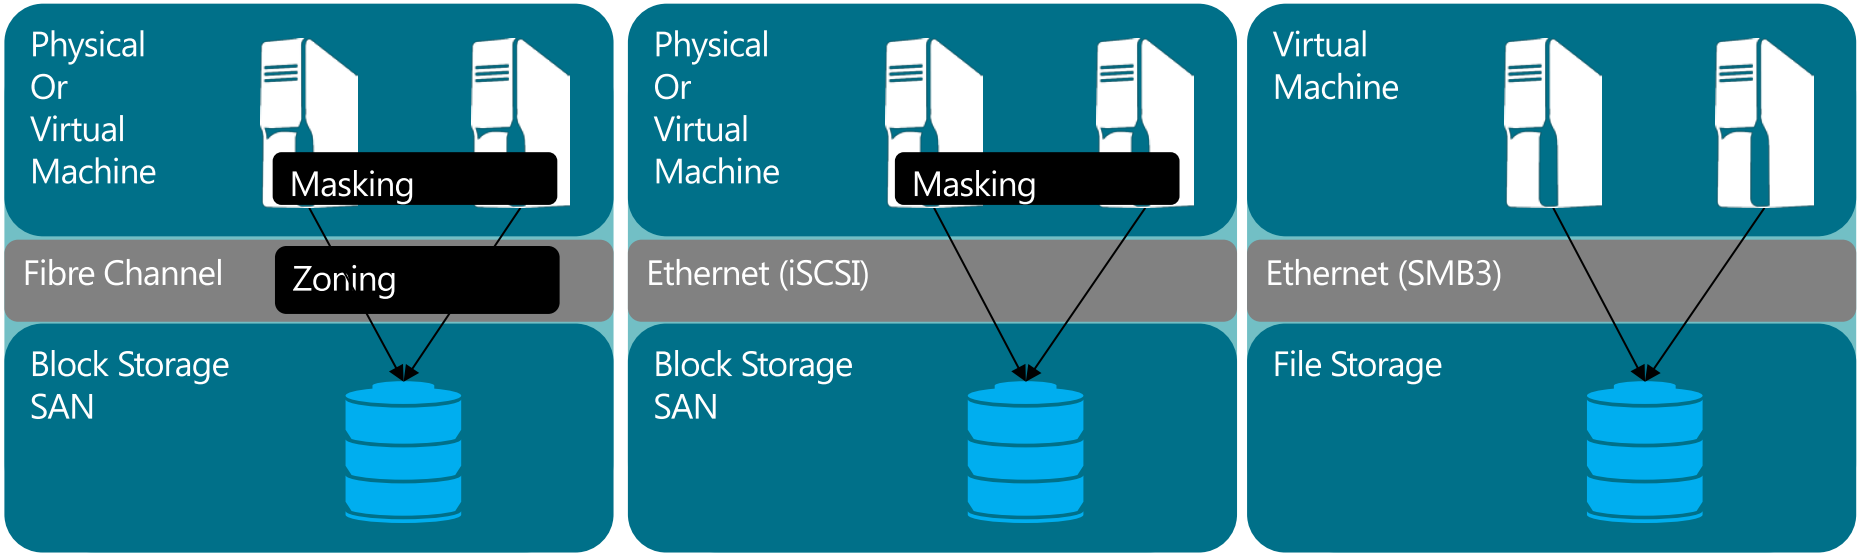

Storage Management API (SMAPI)

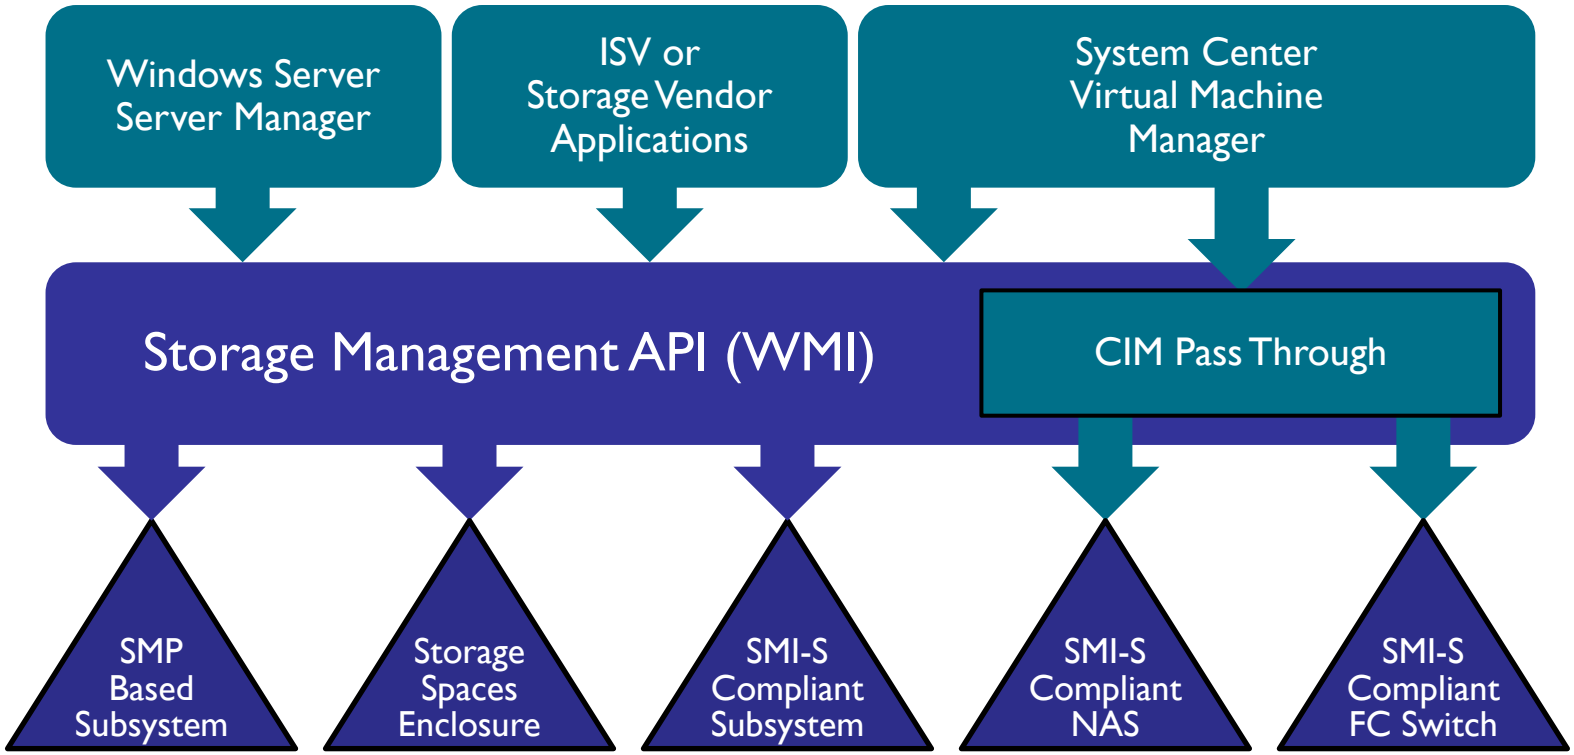

Manage FC Fabric

```
$provider = Get-SCStorageProvider -Name "DCMRRSMISCISCO"
$fabric = Get-SCStorageFabric -Name "2002547FEE79EE81"
$classification = New-SCStorageFabricClassification -Name "PRODUCTION"
Set-SCStorageFabric -EnableManagement -StorageFabric $fabric -StorageFabricClassification
$classification
```

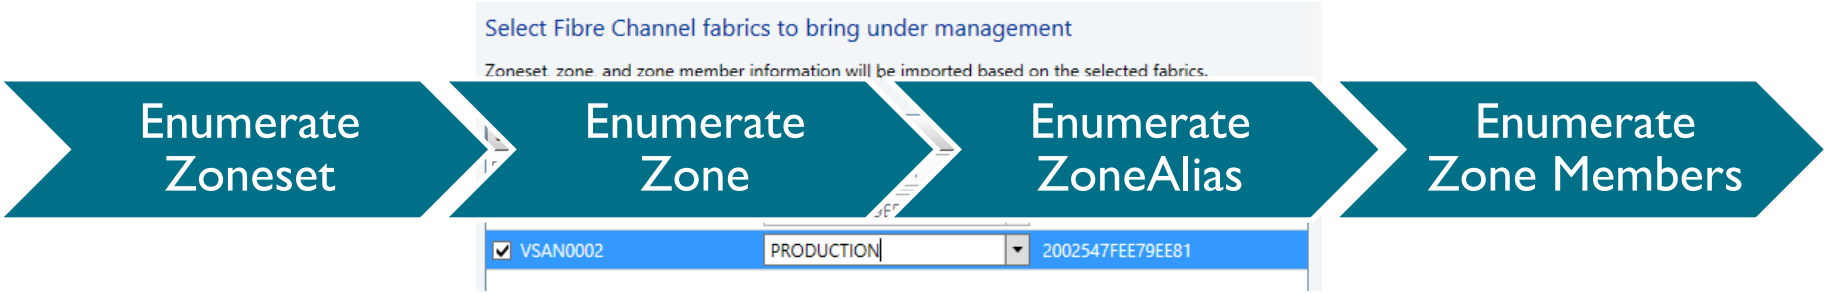

ObjectId : CISCO_Vsan.CreationClassName="CISCO_Vsan",Name="2002547FEE79EE81"
ElementName : VSAN0002
Classification : PRODUCTION
StorageZoneSets : {Vsan2Active, Vsan2Active}
StorageZoneAliases : {t_alias1, t_alias3}
StorageZoneMemberships :
{CISCO_ZoneMemberSettingData.InstanceID="2\2002547FEE79EE81\vm_0_zone\2\500507680220214F\true" …}
StorageZones : {dcmrr25r19n11, sfcvm1_zone, dcmrr25r19n11, sfcvm1_zone…}

Storage Fabric Classification

FC Fabric Classification

Identify fabric using friendly name

Classification aware Placement

Select Fibre Channel fabrics to bring under management

Zoneset, zone, and zone member information will be imported based on the selected fabrics.

Storage Device	Classification	Fabric WWNN
Dummy	2003547FEE79EE81	2003547FEE79EE81
VSAN0001	2001547FEE79EE81	2001547FEE79EE81
VSAN0002	PRODUCTION	2002547FEE79EE81

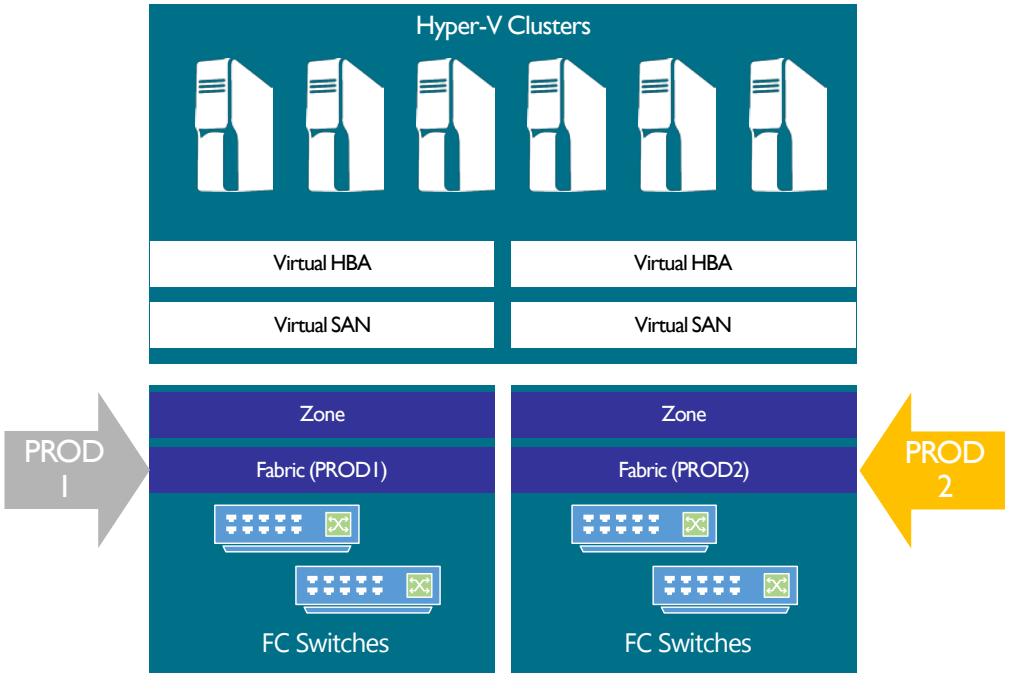

Create Zone and Add Members (VMM)

```
#add list of zone members
$zoneset_newZone = Get-
SCStorageZoneSet -Name "Active"

$newZone = New-SCStorageZone -Name
"MyZone" -StorageZoneSet
$zoneset_newZone -
AddZoneMembership
@("21240002AC000C63",
"20230002AC000C63",
"10000000C9C17FCA",
"10000000C9C17FCB")
```

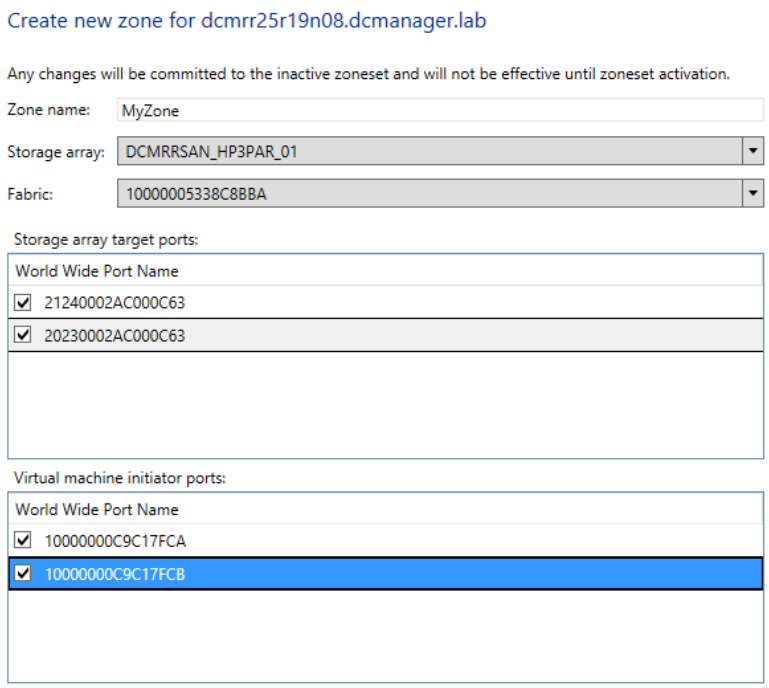

Activate Zoneset (VMM)

```
#Activate zoneset
If($zoneset_newZone.Active -eq $false)
{
        Set-SCStorageZoneSet -StorageZoneSet $zoneset_newZone -Enable
}
```

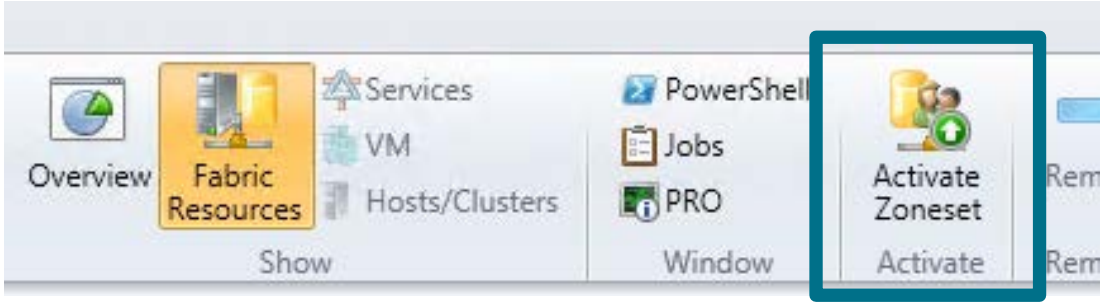

Applications with FC Storage

Problem Statement

Multi-tier application that requires a SQL database deployed on a failover cluster with shared FC storage from a SAN

Application Environment Details

Fibre Channel Switches

Dual redundant fabrics between hosts and storage

Hyper-V Host

Hosts with at least 2 FC ports

Connect each FC port to a different fabric

Create a virtual SAN per fabric

SQL Tier Service deployment

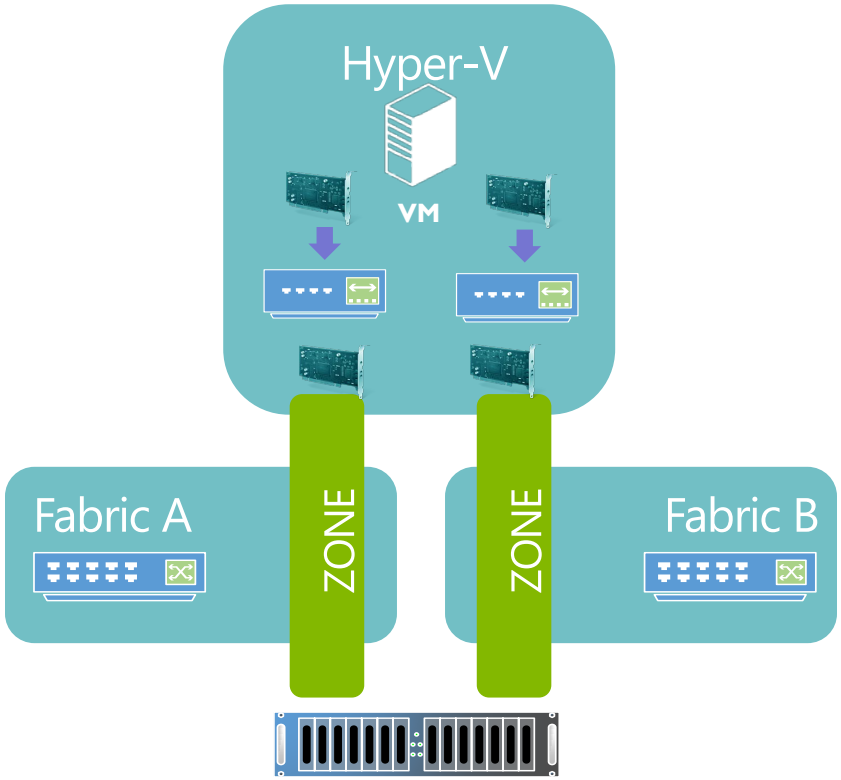

Partnerships

Storage Manageability Certification

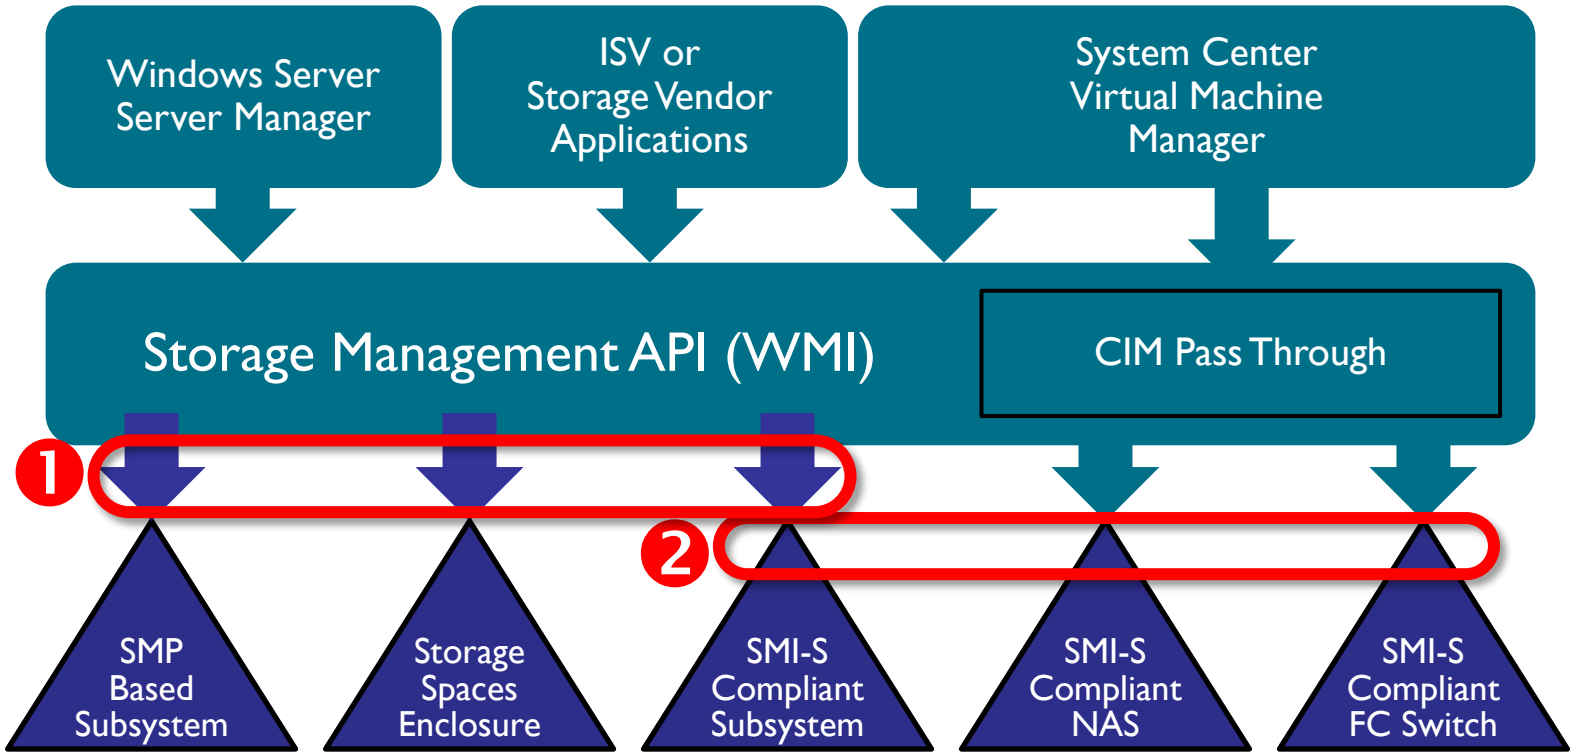

VMM Based Scenario Tests

Native SMAPI Functionality Tests

SNIA SMI CTP

Thank you!

Hector Linares
Principal Program Manager
hectorl@microsoft.com
Microsoft

Rajesh Balwani
Senior Development Engineer
rbalwani@microsoft.com
Microsoft

MASKING and MAPPING IN-DEPTH

MASKING IN-DEPTH

- SCSI Protocol Controller
  - a.k.a. Storage Groups
  - a.k.a. Masking Sets
  - a.k.a. Host Groups
  - a.k.a. View
- Ports Per View
  - Array target ports
- One Hardware ID Per View
  - Server initiator ports

SCSI Protocol Controller

- SCSI protocol controller (SPC) is a grouping mechanism on the array using for exposing a storage volume (LUN) to a Windows server
- SPC contains target ports (array), storage volumes (LUN), initiator ports (Windows host)

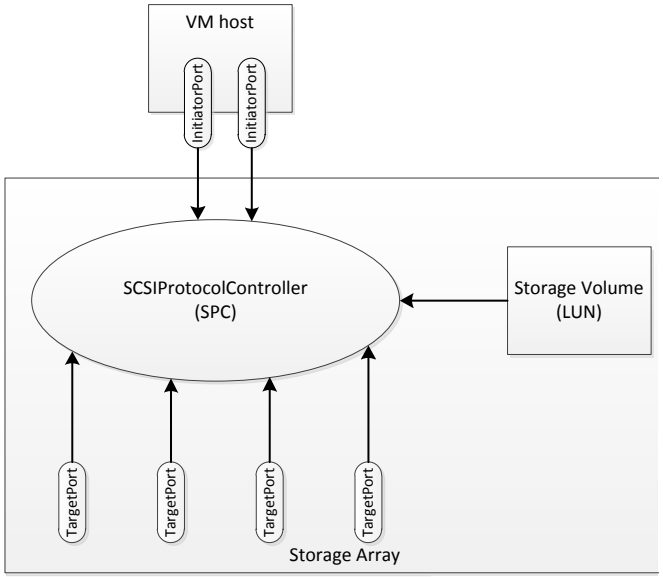

Ports per SPC

- Controls how many array target ports an SPC can have
  - All Ports Share Same View – All targets ports in SPC
  - Multiple Ports Per View – one, some or all target ports in SPC
  - One Port Per view – one target port in SPC

All Ports Share Same SPC

- SPC includes all of the target ports on the array

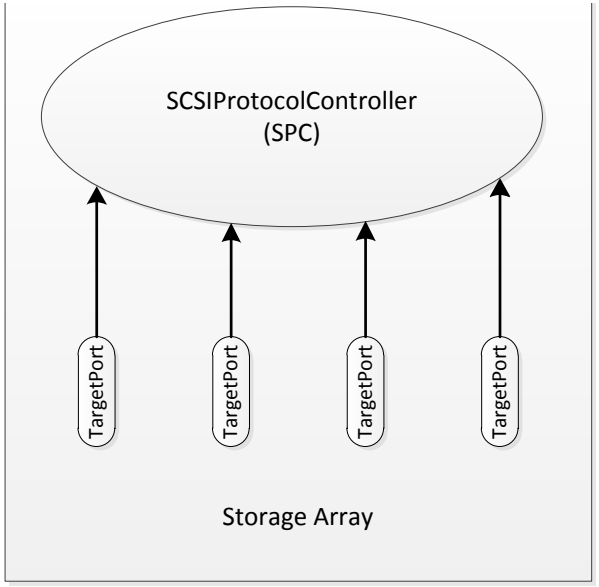

Multiple Ports per SPC

- SPC includes all, some, or one of the target ports on the array

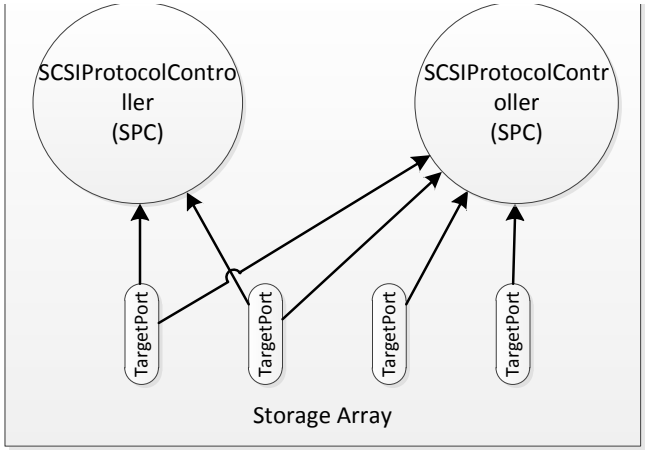

One Port per SPC

- SPC includes only one target port on the array

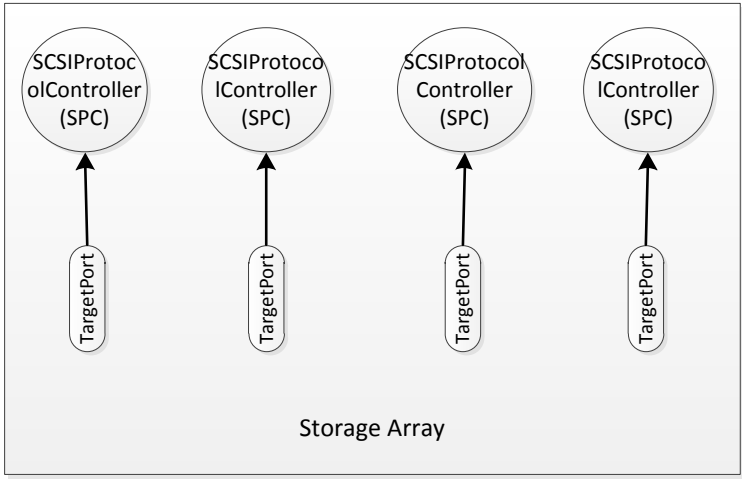

Ports per SPC

Setting	Implementation Notes	Examples
All Ports Share Same View	**Simplicity**	NetApp FAS HP EVA EMC Clariion Dell Compellent
Multiple Ports Per View	**Flexibility** (can mimic both One Port per View and All Ports Share Same View)	EMC VMAX HP 3PAR
One Port Per View	**Traditional**	Hitachi VSP HP P2000 Hitachi AMS

One Hardware ID per SPC

- One Hardware ID per SPC == True

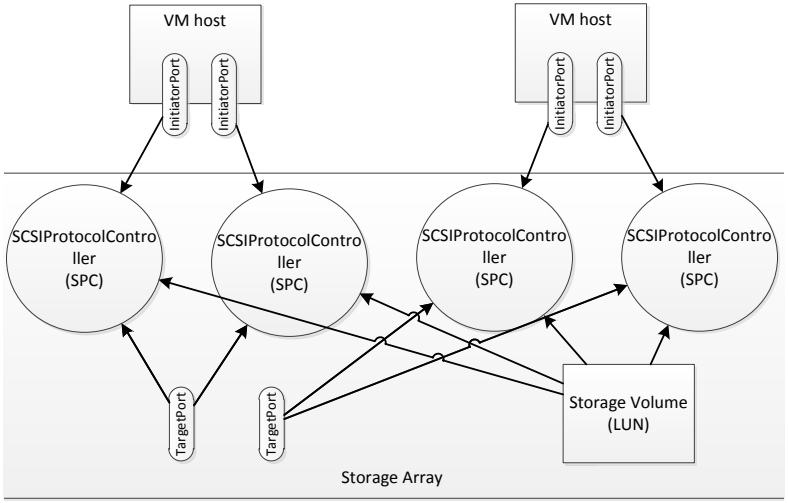

- One Hardware ID per SPC == False

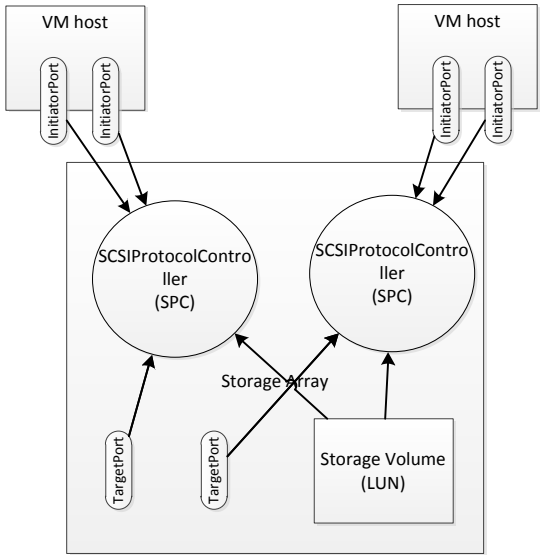

One Hardware ID per SPC

Setting	Implementation Notes	Examples
TRUE	Traditional	Hitachi AMS HP P2000 Dell Compellent
FALSE	Flexibility (allows for one SPC per cluster)	EMC VMAX, Clariion IBM XIV, SVC, V7000 HP EVA, 3Par NetApp FAS Hitachi VSP

SPC Creation for Clusters

- Create SPCs per Cluster == True
  - Creates storage groups that contains all initiator ports for all nodes in the cluster

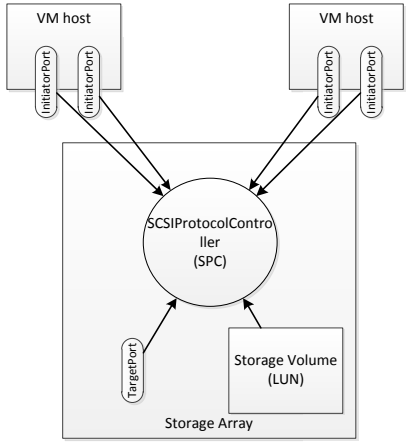

- Create SPCs per Cluster == False
  - Creates a storage group that contain all initiator ports for each node in the cluster

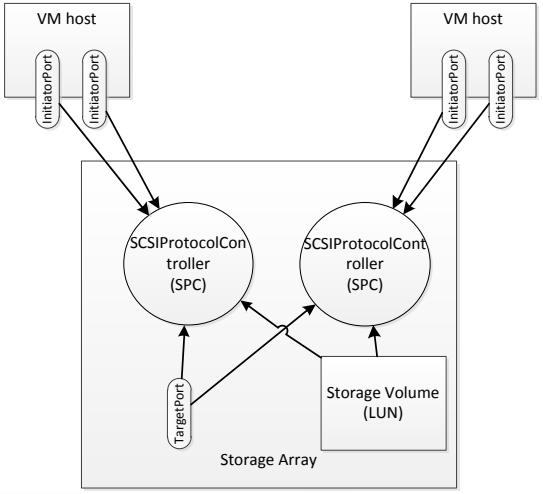

SPC Creation for Clusters

Setting	Implementation Notes	Examples
TRUE	On some arrays, unmasking operations get serialized so the time required to unmask a LUN to multiple nodes in the cluster increases (minutes) No flexibility to have a separate LUN for boot from SAN	EMC VMAX Hitachi VSP
FALSE	Offers the most flexibility if you want to expose LUNs to a subset of nodes in the cluster (e.g. to enable boot from SAN).	NetApp FAS HP 3AR Dell Compellent

	All Ports Share Same SPC	Multiple Ports Per SPC	One Port Per SPC
One Hardware ID Per SPC == FALSE	Create SPCs per Cluster == True or False	Create SPCs per Cluster == True or False	Create SPCs per Cluster == True or False
One Hardware ID Per SPC == TRUE	Create SPCs per Cluster – Not Applicable	Create SPCs per Cluster - Not Applicable	Create SPCs per Cluster - Not Applicable

- **All/Multi** Ports per SPC
- One Hardware ID per SPC == **False**
- Create SPCs per Cluster == **True**

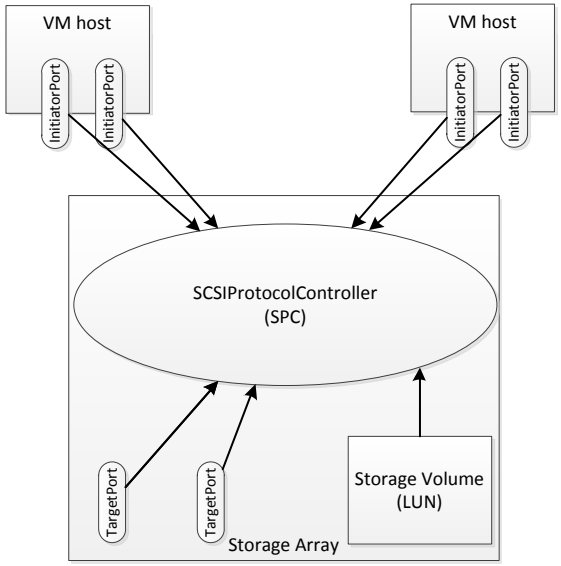

- **All/Multi** Ports per SPC
- One Hardware ID per SPC == **False**
- Create SPCs per Cluster == **False**

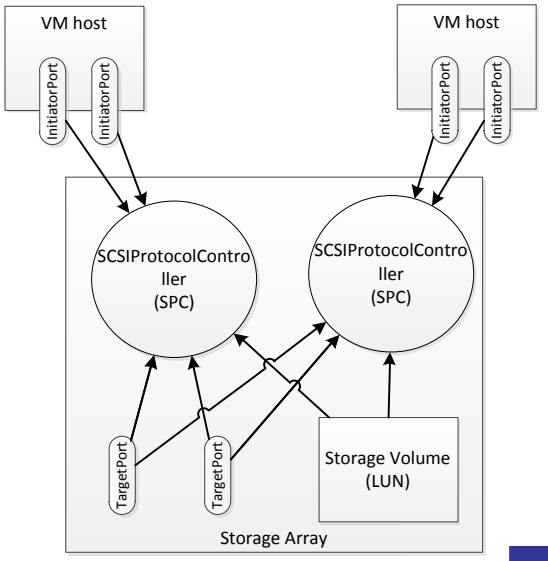

- **One** Port per SPC
- One Hardware ID per SPC == **False**
- Create SPCs per Cluster == **True**

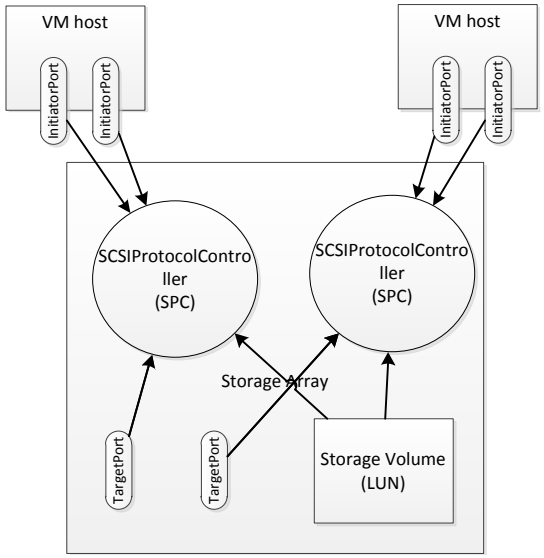

- **One** Port per SPC
- One Hardware ID per SPC == **False**
- Create SPCs per Cluster == **False**

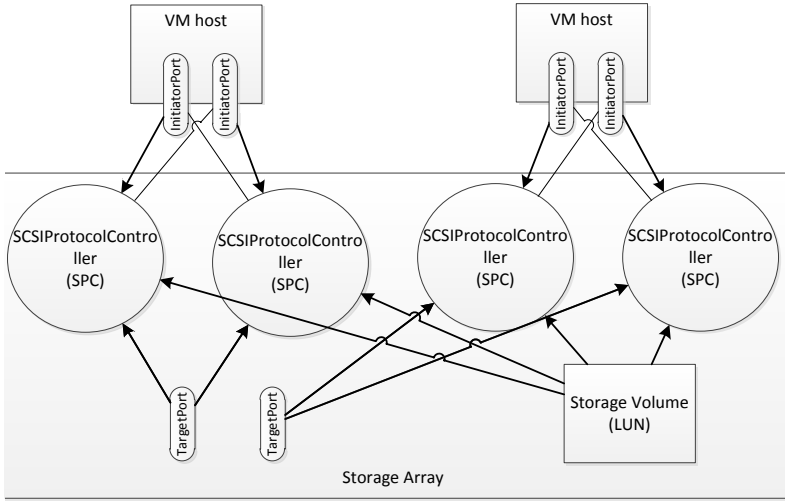

- **All/Multi** Ports per SPC
- One Hardware ID per SPC == **True**
- Create SPCs per Cluster – **N/A**

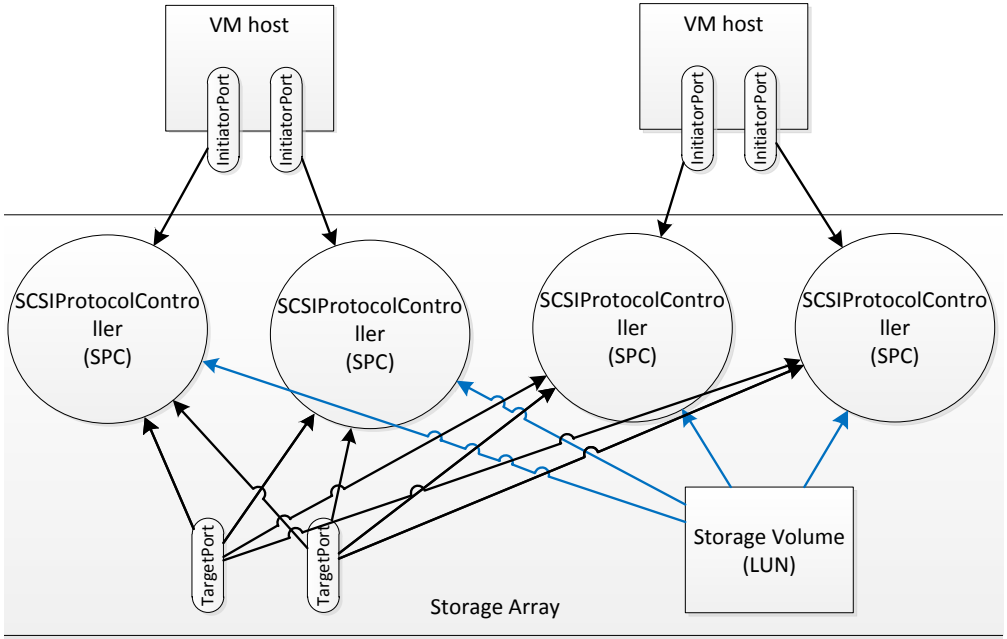

- **One** Ports per SPC
- One Hardware ID per SPC == **True**
- Create SPCs per Cluster – **N/A**

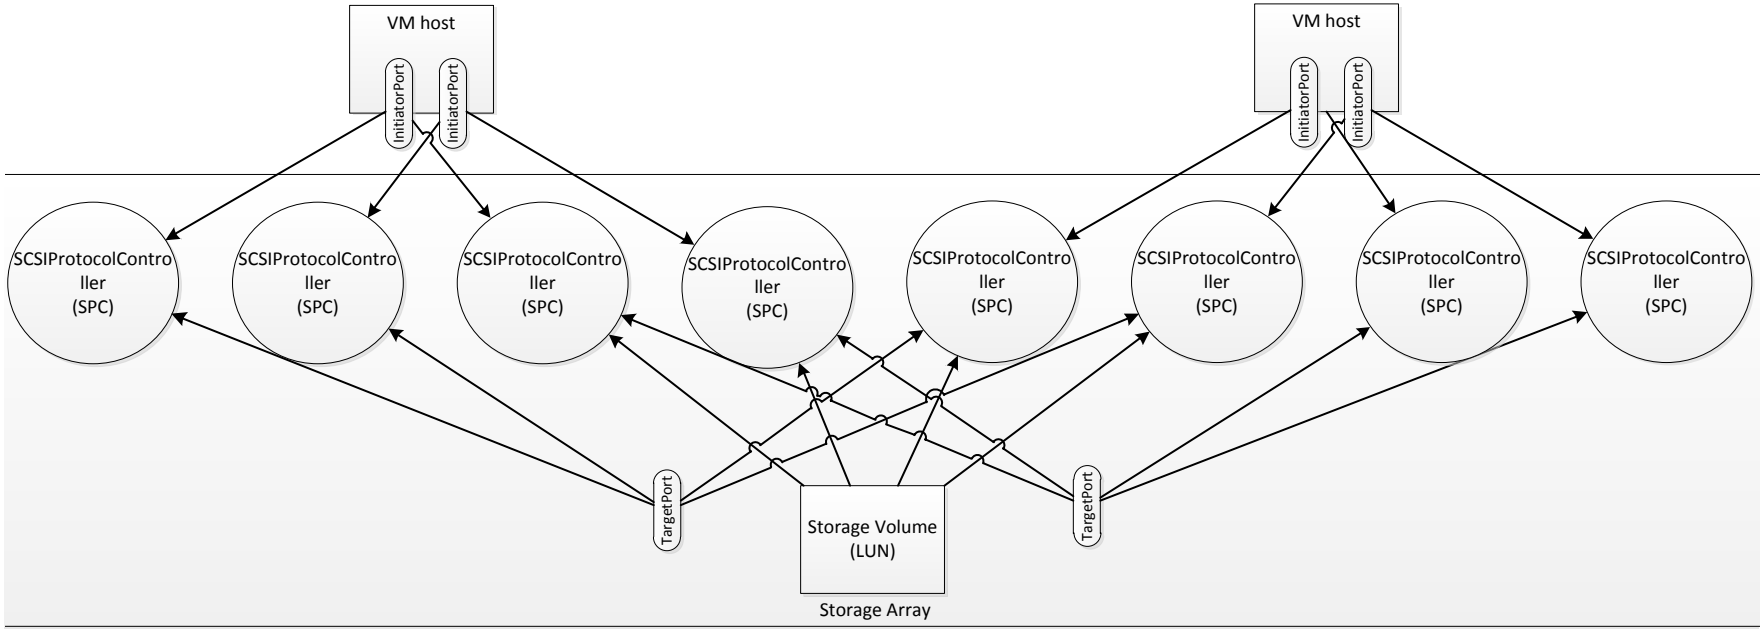

ISCSI Target Models

Masking/Mapping depends on Target model

Types of target models

- Static

  Targets are static as fiber channel targets usually one per Ethernet port or pre setup. #of static iscsi targets won't change for life time of array unless there configuration change.

- Dynamic

  Targets are dynamic. Two type of dynamic target systems.

  - Automatic

    Targets are created automatically one per volume upon volume creation.

  - Manual

    Targets can be created by user based as they need.

ISCSI Target Model Discovery

Using ISCSI Configuration Service Capabilities

Static target model

Absence of iscsi configuration capabilities instance or following values would be discovered as static target model.

iscsiConfigurationCapabilitiesInstance->ISCSIProtocolEndpointCreationSupported = false

iscsiConfigurationCapabilitiesInstance->ISCSINodeCreationSupported = false

Automatic target model

iscsiConfigurationCapabilitiesInstance->ISCSIProtocolEndpointCreationSupported = true

iscsiConfigurationCapabilitiesInstance->ISCSINodeCreationSupported = false

Manual target model

iscsiConfigurationCapabilitiesInstance->ISCSIProtocolEndpointCreationSupported = false

iscsiConfigurationCapabilitiesInstance->ISCSINodeCreationSupported =

ISCSI Masking

Static target model

Mask/unmask same as Fiber channel target ports described earlier i.e. depends on PortsPerView, OneHardwareIDPerView and MaxMapCount properties. VMM expects iscsi login as precheck for provisioning which can be automated as part of array onboarding per host.

Automatic target model

ISCSI targets are 1:1 to volume hence one SPC per volume i.e. volume is key for SPC. Masking is nothing but adding and removing initiators from SPC.

SPC is created by VMM upon first time volume is unmasked, provider needs to establish its association with specific target automatically. VMM will issue ISCSI login/logoff per target upon provisioning.

Manual target model

Provides flexibility to create targets as needed. SMAPI doesn't expose target creation api's. Solution is for provider to create target automatically per SPC creation from VMM. Here initiators become key for SPC. Once SPC created masking is nothing but adding removing volumes from SPC. VMM issues iscsi login/logoff upon provisioning.

Notes for ISCSI Masking

Masking capabilities not per protocol

This could be tricky for arrays which implement different target models for iscsi vs fiber channel. For example FC is static target system where as ISCSI is dynamic target system.

Some workaround available. To mitigate VMM relies on ISCSI configuration capabilities for ISCSI before looking into Masking capabilities. However some combinations may not work

Automatic target model not scalable.

Since there is 1:1 SPC per volume system could end up with too many SPCs potentially max number of volumes. This causes scalability and performance issues not only from provider but also from host initiator.

Group Masking

Basically group masking allows SPC to be associated with initiator group, volume group and port group. These groups can be nested.

Support for group masking with following limitations

Only single level groups are supported i.e. nested groups not supported.

New initiator group and volume group are automatically created per spc creation and destroyed as spc goes away.

Target port group are not created automatically. This needs to be created out of band. This is done to facilitate reuse of target port groups.

Recommendation

Use masking/mapping profile as all scenarios can be addressed with this.