

STO2063BU - Architecting Site Recovery Manager to Meet Your Recovery Goals

GS Khalsa (@gurusimran), VMware

Terminology

RPO - Last viable restore point

RTO - How long it will take before all functionality is recovered

You should break these down to an application, or a service tier level

*Protection Groups and Recovery Plans

What is a Protection Group?

Group of VMS that will be recovered together

- Application
- Department
- System Type
- ?

Different depending on replication type

A VM can only belong to one Protection Group

How do protection groups fit into Recovery Plans?

[1431105945_thumb.png] image via <https://blogs.vmware.com/vsphere/2015/05/srm-protection-group-design.html>

vSphere Replication Protection Groups

- group VMs as desired into Protection Groups
- what storage they are located on doesn't matter

Array Based Protection Groups

If you want your protection groups to align to your applications - you'll need to shuffle storage around

Policy Driven Protection

- new style Protection Group leveraging storage profiles
- high level of automation compared to traditional protection groups
- policy based approach reduces OpEx
- simpler integration of VM provisioning, migration, and decommissioning

Introduced in SRM 6.1

How Should You Organise Your Protection Groups?

More Protection Groups

- Higher RTO
- Easier testing
- Only what is needed
- More granular and complex

Fewer Protection Groups

- Lower RTO
- Less granular, complex and flexible

This varies by customer and will be dictated by the appropriate combination of

complexity and flexibility

*Topologies

SRM Supports Multiple DR Topologies

Active-Passive Failover

- dedicated resources for recovery

Active-Active Failover

- run low priority apps on recovery infrastructure

Bi-directional Failover

- Production applications at both sites

- Each site acts as the recovery site for the other

Multi-site

- Many-to-one failover

- Useful for Remote Office / Branch Office

Enhanced Topology Support

[SRM-Shared-Protection-Recovery.png]

Enhanced Topology Support 2

[SRM-Shared-Recovery-Central-VC.png]

Enhanced Topology Support 3

[SRM-3-site-config.png]

Images via <https://blogs.vmware.com/virtualblocks/2016/07/28/srm-multisite/>

10 SRM pairs per vCenter

A VM can only be protected once

SRM & Stretched Storage

[SRM_stretched_storage.png] - image via <https://blogs.vmware.com/virtualblocks/2015/09/01/srm-6-1-whats-new/>

Supported as of SRM 6.1

SRM and vSAN Stretched Cluster

[vsan-vr-srm.png] - image via <https://blogs.vmware.com/virtualblocks/2015/08/31/whats-new-vmware-virtual-san-6-1/>

Failover to the third site (not the 2 sites comprising the cluster)

Enhanced Linked Mode

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.install.doc/GUID-91EF7282-C45A-4E48-ADB0-5A4230A91FF2.html>

Makes it easier to manage your environment. Introduced in vSphere 6.0.

*Impacts to RTO

Decision Time

How long does it take to decide to failover?

IP Customisation

Workflow without customisation

- Power on VM and wait for VMtools heartbeats

Workflow with IP customisation

- Power on VM with network disconnected
- Customise IP utilising VMtools
- Power off VM
- Power on VM and wait for VMtools heartbeats

Alternatives

- Stretched Layer 2
- Move VLAN / Subnet

It's going to take some time to do when you failover a guest

Priorities and Dependencies

Priorities Only

Organisation for lower RTO

- fewer/larger NFS datastore / LUNs
- fewer protection groups
- don't replicate VM swap files
- fewer recovery plans

VM Configuration

- VMware Tools installed in all VMs
- suspend VMS on Recovery vs PowerOff VMs
- array-based replication vs vSphere Replication

Recovery Site Sizing

- vCenter sizing - it works harder than you think
- Number of hosts - more is better
- Enable DRS - why wouldn't you?
- Different recovery plans target different clusters

*Recommendations

Be Clear with the business

What is/are their

- RPOs?
- RTOs?
- Cost of downtime?
- Application priorities?
- Units of failover?
- Externalities?

Do you have Executive buy-in?

Risk with Infrequent DR Plan Testing

- Parallel and cutover tests provide the best verification, but are very resource intensive and time consuming
- cutover tests are disruptive, may take days to complete and leaves the business at risk

Frequent DR Testing Reduces Risk

- increased confidence that the plan will work
- recovery can be tested at anytime without impact to production

Test Network

Use VLAN or isolated network for test environment

- default "auto" setting does not allow VM communication between hosts

Different PortGroup can be specified in SRM for test vs actual run

- specified in Network Mapping and / or Recovery Plan

Test Network - Multiple Options

Two Options

- Disconnect NSX Uplink (this can be easily scripted)

- Use NSX to create duplicate "Test" networks

RTO = dollars

*Demos

<https://storagehub.vmware.com/#!/site-recovery-manager-3/srm-6-5-technical-overview>