

# Tintri Deployment and Best Practices Guide

for VMware vCenter™ Site Recovery Manager™

## Revision History

| Version | Date       | Description  | Author                        |
|---------|------------|--|-------------------------------|
| 1.2     | 08/14/2017 | Updated  | Bill Roth                     |
| 1.1     | 05/01/2017 | Updated  | Tintri Technical<br>Marketing |
| 1.0     | 11/01/2014 | Previously known as “Tintri Deployment and Best Practices<br>Guide for VMware vCenter Site Recovery Manager 5.8” | Previous<br>contributors      |

Table 1 - Revision history

## Contents

|  |    |
|--|----|
| Executive Summary .....                    | 4  |
| Intended Audience .....                    | 4  |
| Prerequisites .....                        | 4  |
| Considerations & Limitations.....          | 4  |
| Consolidated List of Practices .....       | 4  |
| Overview .....                             | 5  |
| Tintri Components.....                     | 6  |
| ReplicateVM .....                          | 6  |
| Storage Replication Adapter .....          | 6  |
| Deployment .....                           | 6  |
| Prerequisites .....                        | 7  |
| Initial SRM Configuration.....             | 7  |
| Tintri Storage Replication Adapter.....    | 7  |
| Tintri Hypervisor Settings .....           | 7  |
| Tintri Replication.....                    | 8  |
| ReplicateVM License .....                  | 8  |
| Replication Paths.....                     | 9  |
| Tintri Datastore.....                      | 10 |
| Tintri Service Group.....                  | 13 |
| Array Based Replication.....               | 16 |
| Add Array Manager.....                     | 17 |
| Array Pairs.....                           | 20 |
| Protection Group .....                     | 21 |
| Recovery Plan.....                         | 23 |
| Deployment Options.....                    | 26 |
| Placeholder VMs .....                      | 26 |
| Virtual Machines Not Protected by SRM..... | 28 |
| Stale Folders .....                        | 28 |
| Conclusion.....                            | 29 |
| References .....                           | 29 |
| Tintri.....                                | 30 |
| VMware.....                                | 30 |

## Executive Summary

Tintri VMstore arrays integrate tightly with VMware Site Recovery Manager (SRM) environments by means of the Tintri Storage Replication Adapter (SRA). Tintri SRA functionality is enabled when Tintri ReplicateVM is licensed. ReplicateVM provides highly efficient array based replication between the SRM protected and recovery sites. Replicating only unique deduplicated and compressed data, ReplicateVM offloads the SRM replication process from VMware vSphere and minimizes inter-site replication bandwidth requirements.

Implementation is simple, and is enabled by the intuitive Tintri user interface. All SRM features are supported and execute flawlessly.

This document details a proven deployment methodology in conjunction with known best practices.

## Intended Audience

Focused on the successful deployment of Tintri VMstore arrays into a VMware SRM environment, this document targets key best practices and known challenges. Hypervisor administrators and staff members associated with architecting, deploying, and administering a VMware SRM environment with Tintri VMstore arrays are encouraged to read this document.

## Prerequisites

General knowledge of, and familiarity with the Tintri VMstore product is essential prior to architecting a solution, or integrating Tintri VMstore arrays into a VMware SRM environment. Outside of the deployed technologies, it is critically important to have a comprehensive understanding of disaster recovery requirements within the organization where the solution is being deployed.

## Considerations & Limitations

Product compatibility and support matrices should be referenced to confirm that a given configuration is supported prior to implementation. This includes but is not limited to Tintri and VMware products.

For Tintri support information please visit [support.tintri.com/](http://support.tintri.com/). The Tintri support portal requires login credentials.

Descriptions provided and examples depicted within this document are based on Tintri Operating System version 4.3 and higher in conjunction with VMware SRM version 6.5.

This document does not take the place of Tintri product documentation or VMware product documentation.

The scope of document is constrained to integrating Tintri VMstore arrays into a VMware SRM environment. This document is not a substitute for formal Tintri, VMware, or VMware SRM training.

## Consolidated List of Practices

The table below includes the recommended practices in this document. Click the text on any of the recommendations to jump to the section that corresponds to each recommendation for additional information.

- Add the vCenter server associated with the first VMware SRM site (typically the “Protected Site”) to the Tintri VMstore located at the first site. Add the vCenter server associated with the second VMware SRM site (typically the “Recovery Site”) to the Tintri VMstore located at the second site. Use the “Test all vCenters” button to confirm hypervisor connectivity and login at both sites.
- Assure that a ReplicateVM license has been applied to the Tintri VMstore arrays that will use SRM array based replication.
- Use the “Test Paths” function to confirm that the replication settings and replication path are correct.

- Create a Tintri datastore for each SRM protection group by creating a subfolder within “/tintri” and mounting it as a datastore.
- Make sure that at least one VM resides in a new datastore before configuring that datastore within a Tintri service group.
- VMs created in, or migrated to a datastore that will be configured within a Tintri service group should not have an individual VM snapshot schedule or replication configuration. A Tintri service group will be configured to create snapshots and perform replication between VMstores.
- Do not manually create a destination folder for a Tintri service group on the Tintri VMstore at the recovery site. The destination folder will be automatically created by SRM as required.
- Make sure that the replicated device group associated with a given Tintri service group is able to be discovered and displayed within the array manager hosting the service group before creating a protection group for the replicated device group.

## Overview

VMware SRM is a disaster recovery management solution. The solution automates operations such as disaster recovery testing, failover, reprotect, and failback between two sites.




Figure 1 - SRM component overview

A designated group of VMs (virtual machines) is considered disaster recovery “ready” when properly configured in SRM within a protection group. Simplified, this configuration consists of a protected site and a recovery site. Recent copies of the VMs in the protected site are replicated to the recovery site. VMware provides a replication solution referred to as “vSphere Replication” to facilitate creating recent copies of VMs within the recovery site.

Array based replication augments SRM by leveraging array vendor native replication technology by means of an SRA. The SRA enables SRM to control array functions necessary for recovery plan testing, failover, reprotect, and failback operations.

Array vendor native replication technology is generally considered to be more efficient than vSphere replication because it bypasses VMware vSphere. In effect, the replication data path required for maintaining recent copies of VMs on the protected site is direct, between two arrays, and does not place additional resource load on hypervisor host nodes. Additionally, when more functions are handled by the hypervisor, more complexity is introduced to hypervisor planning, maintenance, upgrades, and scaling. Array vendor native replication may introduce additional efficiencies dependent on the array vendor’s replication implementation.

## Tintri Components

In addition to VMware SRM deployment requirements, integration of Tintri VMstore arrays includes a small number of components. Specifically, those components consist of:

- Two or more Tintri VMstore arrays
- Tintri ReplicateVM licenses for each Tintri VMstore array
- Tintri SRA installation on VMware SRM Server hosts

### ReplicateVM

Tintri asynchronous replication is trademarked as “ReplicateVM”. Replication occurs at the VM level, and is based on individual VM snapshots. ReplicateVM is a licensed feature that requires enablement through the use of a license key.




Figure 2 - ReplicateVM conceptual diagram

Replication efficiency is important from a number of perspectives. One perspective deals with the fact that a replicated snapshot becomes a recovery point for an individual VM on a destination VMstore. It represents a disaster recovery copy of the VM, and the point at which the replication operation completes becomes the point at which the VM enters a “disaster recovery ready” state. Another perspective involves network bandwidth usage. Tintri replication sends only changed blocks between snapshots after deduplication and compression to dramatically reduce the amount of data sent over a WAN connection by up to 95 percent.

### Storage Replication Adapter

Tintri SRA functionality is enabled when a ReplicateVM license is installed. The Tintri SRA integrates with VMware SRM and ReplicateVM to enable SRM workflows that automate the recovery plan test, failover, reprotect, and failback operations.




Figure 3 - Tintri SRA installed on an SRM Server host

The Tintri SRA is transparent, in that once it is installed on the SRM Server hosts, no further action or direct administration is required.

## Deployment

The following subsections detail the procedural steps undertaken in a typical deployment of Tintri VMstore arrays into an existing SRM environment.

Naming conventions used in the example deployment have been designed to highlight the relationship of components within a given SRM protection group. Folder name, datastore name, Tintri service group name, protection group name, and recovery plan names are all suffixed with a numeric value to demonstrate the

relationship between each of these components. Users are encouraged to implement a naming convention that aligns with their own organizational requirements.

## Prerequisites

Prior to integrating the Tintri VMstore product into an SRM environment, two VMware SRM sites should be deployed and configured. Each site should include a vCenter Server, vSphere host, and an SRM Server. In addition, a minimum of one Tintri VMstore array should be deployed and configured for each of the two VMware SRM sites.

For VMware and VMware SRM deployment, configuration, and licensing information please visit [www.vmware.com](http://www.vmware.com).

For Tintri VMstore deployment and configuration information, please visit the Tintri support portal at [support.tintri.com](http://support.tintri.com).

## Initial SRM Configuration

In addition to the prerequisites already stated, SRM should be configured to the point where array managers and arrays pairs can be configured. This includes:

- Pairing sites
- Configuring inventory mappings
  - Creating resource mappings
  - Creating folder mappings
  - Creating network mappings
- Configuring placeholder datastores
  - Optional configuration of a placeholder datastore on a Tintri VMstore array is presented later in this document

## Tintri Storage Replication Adapter

The Tintri SRA can be downloaded from the Tintri support portal at [support.tintri.com](http://support.tintri.com) or alternatively from the “Drivers & Tools” section of the VMware SRM download portal. Download and install the SRA on both SRM Server hosts.

## Tintri Hypervisor Settings

In this subsection, each Tintri VMstore array will be configured to serve one of the two VMware SRM sites. From the Tintri VMstore located at the first VMware SRM site, launch the graphical user interface and click the “Settings” button.




Figure 4 - VMstore Settings

When the “Adjust your VMstore’s settings” window appears, click “Hypervisor managers”. By default, the “vCenter” hypervisor type should be selected. Check to see if the vCenter server associated with the first VMware SRM site has already been added to the VMstore configuration. If the vCenter has already been added no further action is required within this dialog window.

If the appropriate vCenter host has not yet been added to the VMstore configuration, click the “Add vCenter” button and input the required parameters to continue.

Figure 5 - Hypervisor managers – “Add vCenter”

Testing vCenter connectivity and login is accomplished by clicking the “Test all vCenters” button.

Figure 6 - Hypervisor managers - "Test all vCenters"

**Note:** The same action(s) should also be performed from the Tintri VMstore located at the second VMware SRM site. On the Tintri VMstore located at the second site confirm that the vCenter server associated with the second site has been added to the VMstore configuration. If the appropriate vCenter host has not yet been added to the VMstore configuration, click the “Add vCenter” button and input the required parameters to continue.

- Add the vCenter server associated with the first VMware SRM site (typically the “Protected Site”) to the Tintri VMstore located at the first site. Add the vCenter server associated with the second VMware SRM site (typically the “Recovery Site”) to the Tintri VMstore located at the second site. Use the “Test all vCenters” button to confirm hypervisor connectivity and login at both sites.

## Tintri Replication

ReplicateVM needs to be licensed, and have replication paths configured on each VMstore array.

### ReplicateVM License

The “Adjust your VMstore’s settings” dialog window should still be open from the prior configuration step. If the dialog window has been closed, click the “Settings” button to launch it. Click the “more” button in the menu, and then select the “Licenses” menu item. The “Manage feature licenses” window will appear.

Figure 7 - Manage feature licenses

Type or paste a valid ReplicateVM license into the “Add new license” field. Valid licenses will then be automatically added to the VMstore. Within the “Installed licenses” area of the window, assure that ReplicateVM is shown with the license key that was entered.



**Note:** The same action(s) should also be performed on the Tintri VMstore located at the second VMware SRM site.

- 
- *Assure that a ReplicateVM license has been applied to the Tintri VMstore arrays that will use SRM array based replication.*
- 

### **Replication Paths**

The “Adjust your VMstore’s settings” dialog window should still be open from the prior configuration step. If the dialog window has been closed, click the “Settings” button to launch it. Click the “more” button in the menu, and then select the “Replication” menu item. The “Change how VMs replicate” window will appear.

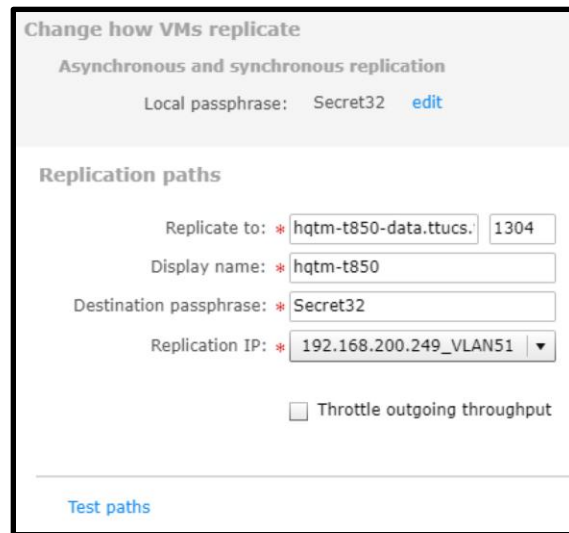


Figure 8 - Change how VMs replicate

The Tintri VMstore array at the first VMware SRM site needs to have a replication path configured to the VMstore array at the second VMware SRM site. If the replication path has not yet been configured, click the plus (+) icon to add the replication path. Replication paths are defined on a source VMstore by specifying parameters associated with a given replication destination:

- “Replicate to” is the Fully Qualified Domain Name (FQDN) or IP address of the destination Tintri VMstore array, and the listening port number
- “Display name” is a user supplied value that identifies the replication target, usually a hostname
- “Destination passphrase” is the passphrase of the destination Tintri VMstore array
- “Replication IP” is the outgoing replication IP address of the local Tintri VMstore

The “Test Paths” function can be employed to test the replication settings and replication path.

**Note:** The same action(s) should also be performed on the Tintri VMstore located at the second VMware SRM site, ensuring that there is a replication path back to the Tintri VMstore located at the first VMware SRM site.

- 
- *Use the “Test Paths” function to confirm that the replication settings and replication path are correct.*
- 

For additional information about Tintri replication, see the Tintri white paper titled, *Data Protection Overview and Best Practices*.

## Tintri Datastore

Before continuing, the Tintri VMstore at the VMware SRM site designated to be the “Protected Site” (typically the first site) should be mounted as a datastore. For additional detail, see the *Tintri VMstore All Flash/Hybrid System Administration Manual*, section titled, “Adding datastores using VMware vCenter”.

A subfolder within the default Tintri datastore mount point “/tintri” should be created for each protection group that will be configured within SRM. For instance, users seeking two protection groups should create two subfolders within the default “/tintri” datastore mount point.

**Note:** The default “/tintri” NFS export or folder cannot be used in the creation of a protection group within SRM. This restriction is enforced, as a datastore mounted on “/tintri” cannot be selected when creating a Tintri service group. Creation of a Tintri service group for use with SRM is detailed in a subsequent section of this document.

Subfolder creation is a simple and straightforward process. From the vSphere web client storage view select the datastore mounted on “/tintri”. Selecting the “Configure” tab and clicking, “Device Backing” will display the “Server” and “Folder” parameters. The server should correlate to the correct Tintri VMstore array by means of the “Datastore data IP” or FQDN of the “Datastore data IP”. The “Folder” name should be “/tintri”.

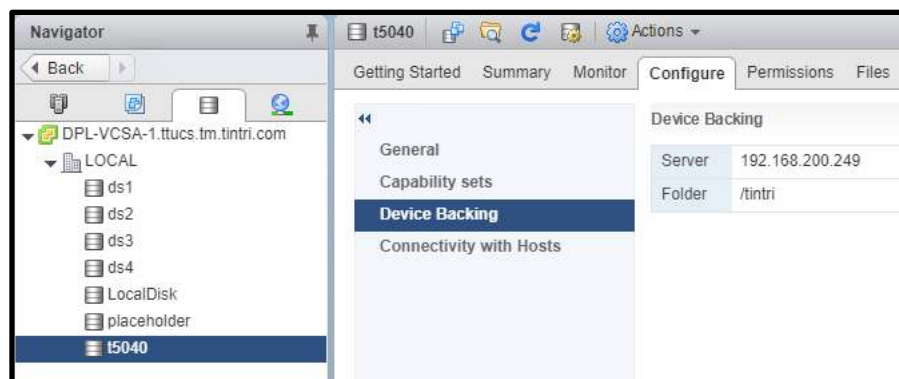


Figure 9 – Datastore – Configure - Device Backing

When the correct device backing has been confirmed select the “Files” tab.

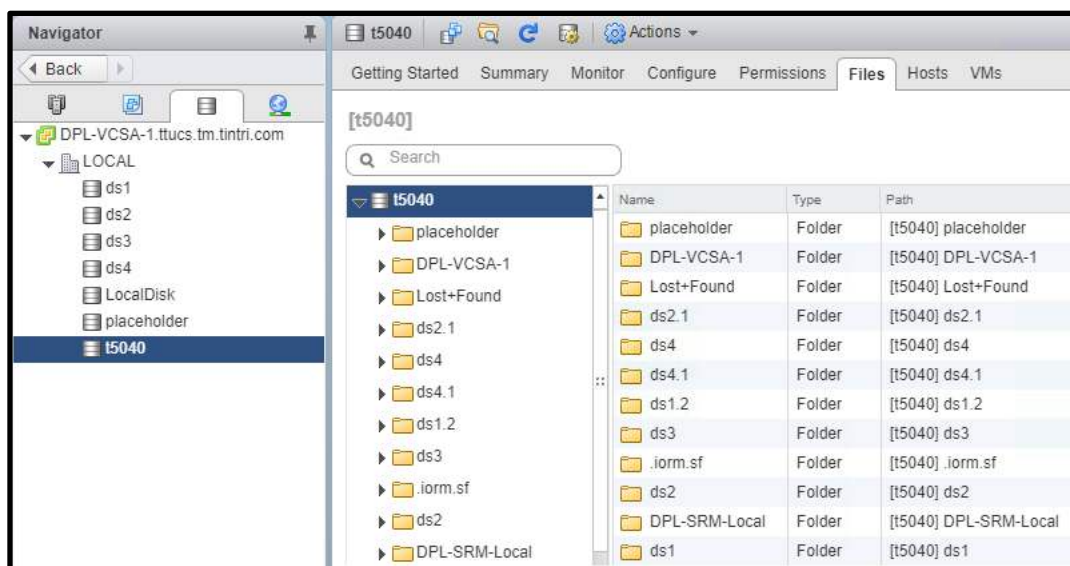


Figure 10 - Datastore - Files

To create a subfolder for use within an SRM protection group click the “Create a new folder” icon.




Figure 11 - New folder creation and naming

Input a name for the new folder. The folder name should comply with organizational naming requirements with the understanding that the folder will be associated with a datastore name. Click the “Create” button to create the folder.

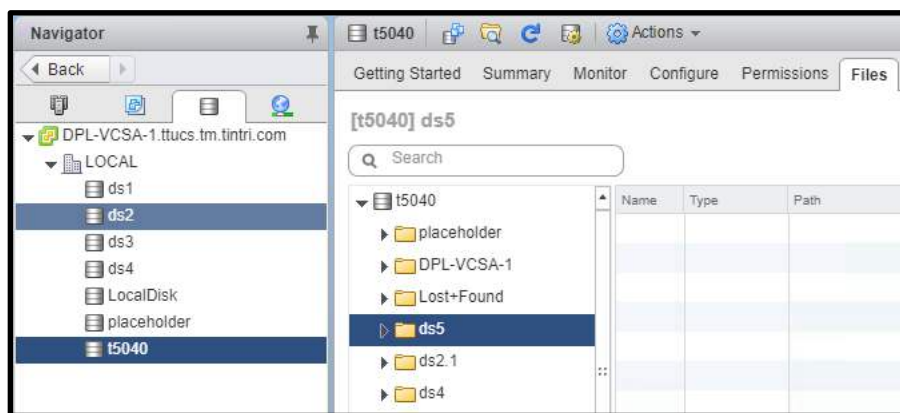


Figure 12 - New folder confirmation

After the new folder has been created it should be configured as a datastore. From the “Hosts and Clusters” view, click the “Create a new datastore” icon.



Figure 13 – Create a new datastore

The “New Datastore” workflow prompts the user for parameters.

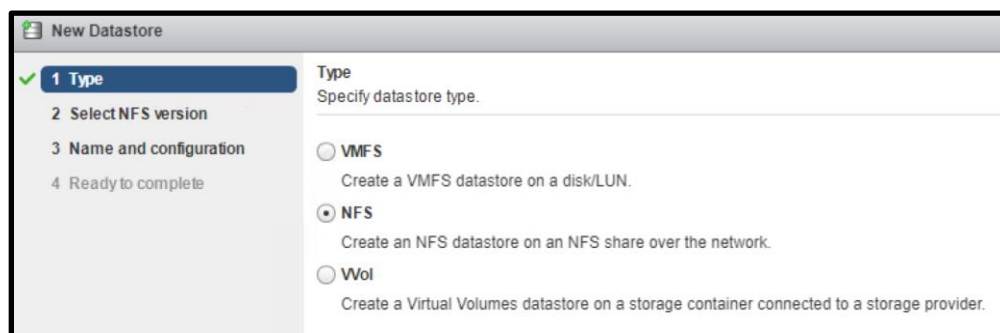


Figure 14 - New Datastore - Type

Click the “NFS” radio button and then click the “Next” button to continue.

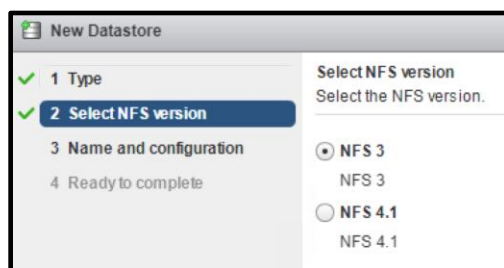


Figure 15 - New Datastore - Select NFS version

Click the “NFS 3” radio button and then click the “Next” button.

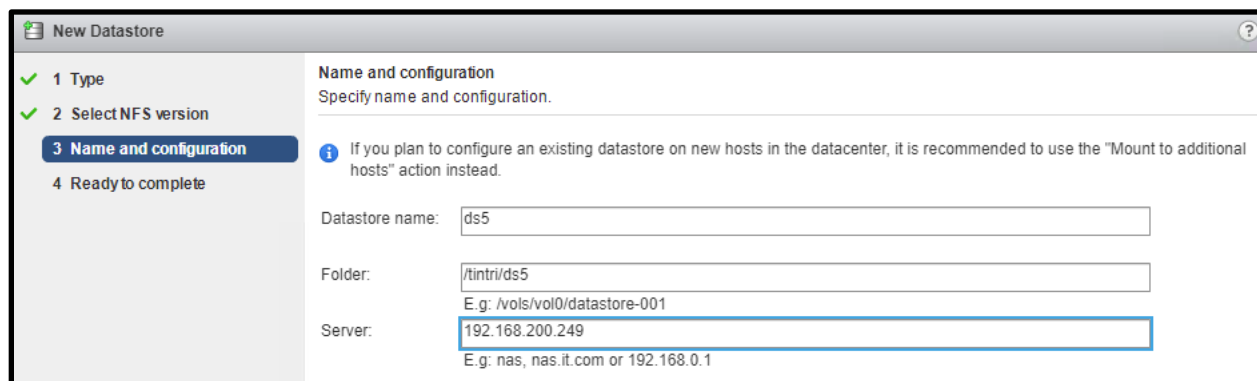


Figure 16 – New Datastore - Name and configuration

Input a “Datastore name”. The datastore name should comply with organizational naming requirements with the understanding that the datastore will be associated with a Tintri service group. Input the “Folder” path. The folder path will be “/tintri/new\_folder\_name”, where “new\_folder\_name” equals the name of the new folder that was created earlier.

Also input the correct parameters for the “Server”. The server name can be the FQDN of the Tintri VMstore data IP, or it can be specified as the VMstore data IP address.

**Note:** Domain Name Server (DNS) issues or challenges may prevent the successful use of FQDN values in conjunction with SRM. Using the Tintri VMstore data IP address may mitigate any DNS related issues or challenges.

Click the “Next” button to continue.

Complete the “New Datastore” workflow by clicking the “Finish” button within the “Ready to complete” window.

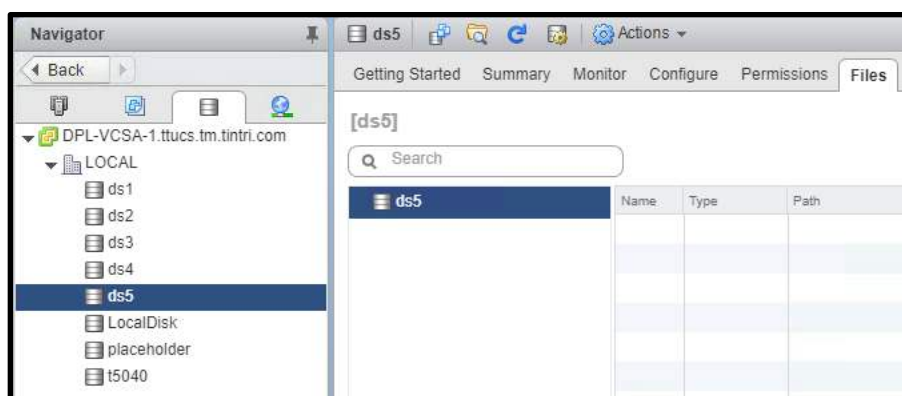


Figure 17 – Storage

The new datastore will be visible within the “Storage” view.

At this point new VMs can be created in the new datastore, or existing VMs to be protected by SRM can be migrated to the new datastore. At a minimum, at least one VM should reside in the new datastore before continuing with the deployment process.

**Note:** For clarification, VMs created within or migrated to a datastore designated for use within an SRM protection group should not be configured with a Tintri protection policy. These VMs should not be configured with a snapshot schedule. These VMs should not be configured to replicate to another Tintri VMstore array. When a Tintri service group is created, the service group will perform all required tasks associated with snapshot creation and replication to the SRM recovery site. Users should not create Tintri VM protection policies, snapshot schedules, or replication destinations for individual VMs that will be protected within an SRM protection group.

- 
- *Create a Tintri datastore for each SRM protection group by creating a subfolder within “/tintri” and mounting it as a datastore.*
  - *Make sure that at least one VM resides in a new datastore before configuring that datastore within a Tintri service group.*
  - *VMs created in, or migrated to a datastore that will be configured within a Tintri service group should not have an individual VM snapshot schedule or replication configuration. A Tintri service group will be configured to create snapshots and perform replication between VMstores.*
- 

### **Tintri Service Group**

A Tintri service group consists of a source datastore, replication destination, and destination folder name. The source datastore is a datastore on the local Tintri VMstore array used with SRM at the protected site. The replication destination is the display name of a configured replication path. This should be the name of a Tintri VMstore array used with SRM at the recovery site. The destination folder is the name of a folder that will be created automatically on the Tintri VMstore array at the recovery site. The destination folder will be created by SRM within the “/tintri” NFS export. When recovery plan test, or failover operations are performed, this folder will be mounted as a datastore at the recovery site.

**Note:** Configurations temporarily deployed to seed a replication destination for use in conjunction with SRM should conduct the seeding process within a Tintri service group. Replica snapshots not contained within a service group cannot be relocated into a service group at the recovery site.

Create one Tintri service group for each SRM protection group. Perform this operation on the Tintri VMstore at the protected site. The Tintri service group should be created prior to the creation of an SRM protection group.



Figure 18 - Create group

Create a Tintri service group from within the Tintri VMstore array graphical user interface by clicking:

1. The “Search VM” button
2. The “Service groups” button
3. The “Create group” button

Figure 19 - Create new SRM group

When the “Create new SRM group” window appears, input the desired name of the SRM group. The “All VMs on datastore” field is populated by selecting a datastore from the pull-down menu. The datastore selected should be the name of a datastore created for use by SRM. The “Replicate to” field is populated by selecting a replication destination from the pull-down menu. The replication destination selected should be the display name of a Tintri VMstore array used with SRM at the recovery site. In the “Destination folder” field, input the desired name of a folder that will be created automatically on the Tintri VMstore array at the recovery site.

**Note:** For clarification, the “Destination folder” will be created automatically by SRM and will appear as a folder within the “/tintri” NFS export at the recovery site. The folder is not created immediately in conjunction with Tintri service group creation. When a recovery plan test operation is performed, this folder will be created and mounted as a datastore. When a recovery plan cleanup operation is performed, the datastore will be dismounted and the folder will be deleted automatically by SRM. Similarly, when a recovery plan failover operation is executed, this folder will be created and mounted as a datastore. The user should not create this folder, or mount the folder as a datastore at the recovery site. SRM automates these tasks.

- 
- *Do not manually create a destination folder for a Tintri service group on the Tintri VMstore at the recovery site. The destination folder will be automatically created by SRM as required.*
-

RPO: ☐ Every minute ☐ 1 hour ☐ 1 day ☒ custom

Custom schedule:

☒ Hourly  on  Keep   hours

☐ Daily

Alert RPO threshold:

Figure 20 – SRM group snapshot and replication schedule

The lower portion of the “Create new SRM group” window contains four radio buttons that facilitate selection of Recovery Point Objective (RPO). This parameter set dictates the frequency at which snapshots of VMs contained within the Tintri service group will be taken. Snapshots are replicated to the destination Tintri VMstore at the recovery site based on this schedule. The RPO parameters selected should align with organizational RPO requirements.

Selecting the “custom” radio button enables the ability to create a customized schedule. The example provided includes a schedule that will create snapshots at 5, 20, 35, and 50 minutes past the hour, every hour. The example schedule executes everyday, and retains local and remote snapshots for a period of one hour. Snapshot consistency has been set to “Crash-consistent”.

**Note:** Tintri service group RPO parameters can be altered after the service group has been created. Right-clicking the Tintri service group name and selecting “Update” from the pop-up menu enables viewing or changing snapshot scheduling parameters. Changes can only be made from the Tintri VMstore at the protected site.

For additional information about Tintri snapshots and replication, see the Tintri white paper titled, *Data Protection Overview and Best Practices*.


Click the “Create” button to finish creating the Tintri service group.

| Name      | VMs from | Replicated to | State           | VMs | Schedules   | RPO                |
|-----------|----------|---------------|-----------------|-----|-------------|--------------------|
| sg1 (SRM) | ds1.2    | ds1           | FailoverPrimary | 1   | Keep 1 hour | Replica up to date |
| sg2 (SRM) | ds2.1    | ds2           | FailoverPrimary | 1   | Keep 1 hour | Replica up to date |
| sg3 (SRM) | ds3.1    | ds3           | FailoverPrimary | 1   | Keep 1 hour | Replica up to date |
| sg4 (SRM) | ds4.1    | ds4           | FailoverPrimary | 1   | Keep 1 hour | Replica up to date |
| sg5 (SRM) | ds5      | ds5           | FailoverPrimary | 1   | Keep 1 hour | Scheduled          |

Figure 21 – Tintri service groups

The newly created Tintri service group will appear on the VMstore array at the protected site. Note that the Tintri service group name will appear in a dark font when viewed on the protected site’s VMstore array.





| Tintri   |                      |                  |               |                         |     |
|--|----------------------|------------------|---------------|-------------------------|-----|
| Dashboard  |                      | Virtual Machines |               |                         |     |
| Service groups   VMs   Virtual disks   Snapshots |                      |                  |               |                         |     |
| Name   |                      | VMs from         | Replicated to | State                   | VMs |
| sg1 (SRM)  | from 192.168.200.249 | ds1.2            | ds1           | FailoverNormalSecondary | 1   |
| sg3 (SRM)  | from 192.168.200.249 | ds3.1            | ds3           | FailoverNormalSecondary | 1   |
| sg4 (SRM)  | from 192.168.200.249 | ds4.1            | ds4           | FailoverNormalSecondary | 1   |
| sg2 (SRM)  | from 192.168.200.249 | ds2.1            | ds2           | FailoverNormalSecondary | 1   |
| sg5 (SRM)  | from 192.168.200.249 | ds5              | ds5           | FailoverNormalSecondary | 1   |

Figure 22 – Tintri service groups

Note that the Tintri service group will also appear on the VMstore array at the recovery site after replication occurs. Also note that the Tintri service group name will appear in a light gray font when viewed on the recovery site VMstore array.

## Array Based Replication

Prior to configuring array based replication, the following tasks should have already been completed:

- The Tintri SRA has been installed on both SRM Server hosts
- ReplicateVM licenses have been installed on the Tintri VMstore arrays that will be used with SRM array based replication
- Replication paths have been configured on the Tintri VMstore arrays that will be used with SRM array based replication
- The Tintri VMstore at the VMware SRM “Protected Site” has been mounted as a datastore
- At least one subfolder within the “/tintri” NFS export has been created and mounted as a datastore at the “Protected Site”
- At least one VM has been created in or migrated onto each datastore that will be used with SRM
  - The VMs created in or migrated onto each datastore that will be used with SRM should not have individual VM protection policies to create or replicate snapshots
- At least one Tintri service group has been created on the Tintri VMstore at the “Protected Site”
- Tintri service group replication has successfully occurred between the “Protected Site” VMstore and the “Recovery Site” VMstore




Figure 23 - Site Recovery

From the VMware vSphere Web Client “Home” view, click the “Site Recovery” icon.



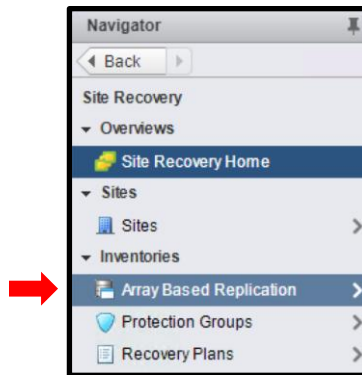


Figure 24 - Array Based Replication

Within the “Navigator” pane click “Array Based Replication”.

### **Add Array Manager**

An array manager is a logical entity that enables SRM to control a specific array by means of the appropriate SRA. For each Tintri VMstore array at either the protected site or recovery site, an array manager should be configured.




Figure 25 - Add Array Manager

From within the “Array Based Replication” pane click “Add Array Manager”. This action will launch the “Add Array Manager” workflow.

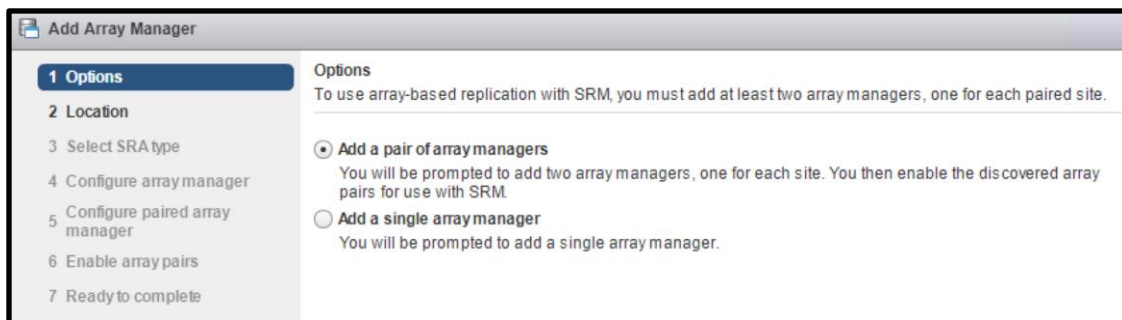


Figure 26 - Add Array Manager – Options

By default, the “Add a pair of array managers” radio button will be selected. This is the correct setting when initially configuring two Tintri VMstore arrays with SRM. Click the “Next” button to continue.

| Sites   | SRM ID          |
|---|-----------------|
| dpl-vcsa-1.ttucs.tm.tintri.com - dpl-vcsa-2.ttucs.tm.tintri.com | com.vmware.vcDr |
|   |                 |
|   |                 |
|   |                 |

Figure 27 - Add Array Manager – Location

By default, the sites that were previously paired within SRM will be displayed and will be selected. Click the “Next” button to continue.

| SRA Type   | Status |
|------------|--------|
| Tintri SRA | OK     |
|            |        |
|            |        |
|            |        |

Figure 28 - Add Array Manager - Select SRA type

The “Tintri SRA” should be selected. Click the “Next” button to continue.

Specify parameters for site 'dpl-vcsa-1.ttucs.tm.tintri.com'

Display Name: dpl-vcsa-1-t5040

**VMstore**

VMstore connection parameters

Hostname or IP address: hqtin-t5040.ttucs.tm.tintri.com  
Enter the hostname or the IP address of the VMstore

User name: admin  
Enter user name for VMstore

Password: \*\*\*\*\*  
Enter password for VMstore

Figure 29 - Add Array Manager - Configure array manager

The first parameter to be input is the “Display Name” field. A descriptive “Display Name” assists in identifying the VMware SRM site and array manager it is associated with. In the example provided, a “Display Name” has been provided that identifies the VMware SRM site name (dpl-vcsa-1) and the array name (t5040).

In the “VMstore” section of the window input the FQDN or IP address of the Tintri VMstore array management port.

**Note:** DNS issues or challenges may prevent the successful use of FQDN values in conjunction with SRM. Using the Tintri management IP address may mitigate any DNS related issues or challenges.

Also supply the “User name” and “Password” values.

Click the “Next” button to continue.

**Add Array Manager**

- 1 Options
- 2 Location
- 3 Select SRA type
- 4 Configure array manager
- 5 Configure paired array manager**
- 6 Enable array pairs
- 7 Ready to complete

**Configure paired array manager**  
Enter the name and connection parameters for the paired array manager.

Specify parameters for site 'dpl-vcsa-2.ttucs.tm.tintri.com'

Display Name:

**VMstore**

VMstore connection parameters

Hostname or IP address:   
Enter the hostname or the IP address of the VMstore

User name:   
Enter user name for VMstore

Password:   
Enter password for VMstore

Figure 30 - Add Array Manager - Configure paired array manager

The parameter requirements for the Tintri VMstore array at the second VMware SRM site are similar to what was entered for the first array. Care should be taken to use a descriptive “Display Name” that assists in identifying the VMware SRM site and array. In the example provided, a “Display Name” has been provided that identifies the VMware SRM site name (dpl-vcsa-2) and the array name (t850).

Click the “Next” button to continue.

**Add Array Manager**

- 1 Options
- 2 Location
- 3 Select SRA type
- 4 Configure array manager
- 5 Configure paired array manager
- 6 Enable array pairs**
- 7 Ready to complete

**Enable array pairs**  
Select array pair that would be enabled on finish.

| Array Pair  | Status              | Local Array Manager |
|---|---------------------|---------------------|
| <input checked="" type="checkbox"/> HQTM-T5040.ttucs.tm.tintri.com... | Ready to be enabled | dpl-vcsa-1-t5040    |

Figure 31 - Add Array Manager - Enable array pairs

In the example provided, the local array manager labeled “dpl-vcsa-1-t5040” has been paired with one remote array manager.

Click the “Next” button to continue.

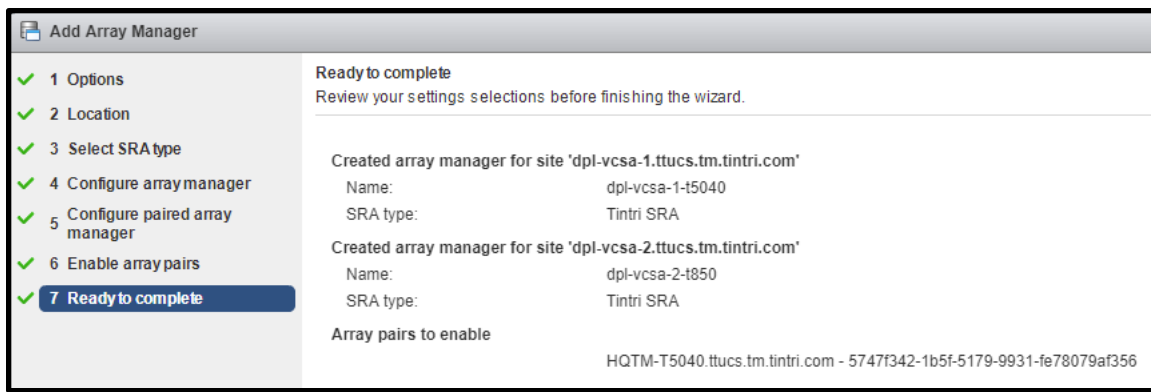


Figure 32 - Add Array Manager - Ready to complete

Click the “Finish” button to continue.




Figure 33 - Array Based Replication – Objects

At this point of the configuration process, there should be a minimum of two (2) configured array managers and each should report a status of “OK” with the appropriate Tintri SRA version. In the example provided:

- One array manager is associated with site “dpl-vcsa-1”
  - “dpl-vcsa-1-t5040” is the array manager for site “dpl-vcsa-1” and array “t5040”
- One array manager is associated with site “dpl-vcsa-2”
  - “dpl-vcsa-2-t850” is the array manager for site “dpl-vcsa-2” and array “t850”

### Array Pairs

A given array pair should have at least one associated local device and remote device. Clicking an array manager in the “Array Based Replication – Objects” pane will open a new pane detailing array pairing and replicated devices.

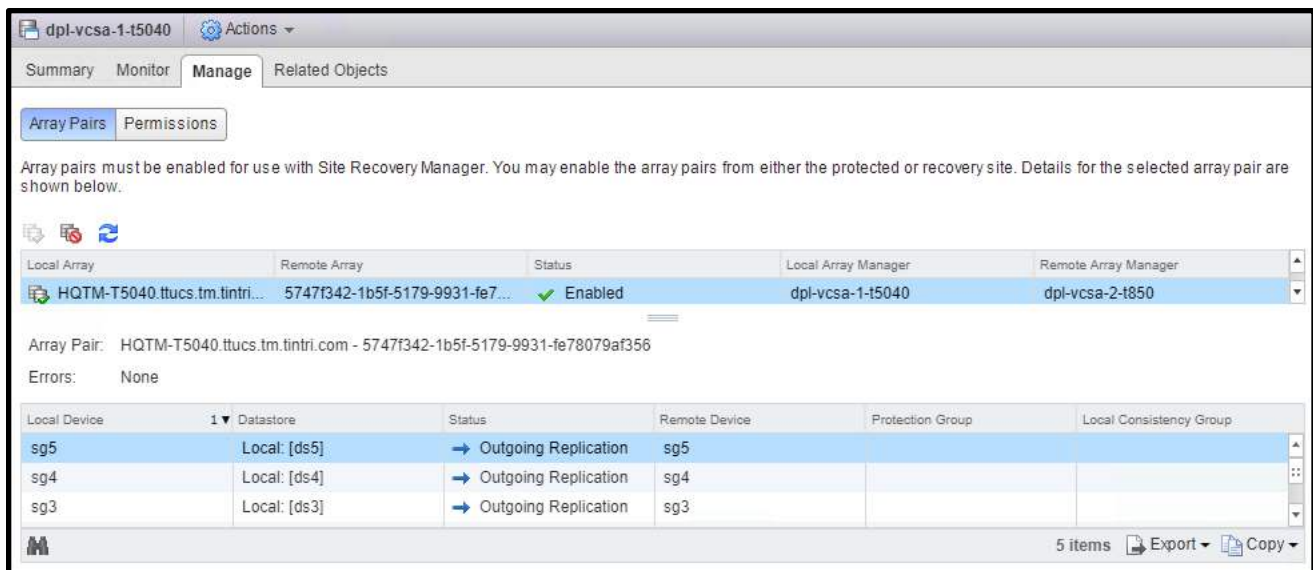


Figure 34 - Array Manager - Manage Array Pairs

At this stage of the deployment process there should be at least one array pair. Dependent on the array manager naming convention employed, it may be easy to identify the site name and array name for both the local and remote array managers.

Also present at this stage of the deployment process, there should be at least one associated replicated device group. A replicated device group is the VMware SRM abstraction of a Tintri service group. A replicated device group will include “Local Device”, “Datastore”, “Status”, and “Remote Device” information. Note that the “Local Device” name should equate to the name of a Tintri service group. The “Datastore” name should equate the name of a Tintri datastore created for use by SRM within the “/tintri” NFS export and configured within a Tintri service group.

In the example provided, there is one array pair. The array pair has five associated replicated device groups, although only three of the five replicated device groups is visible in the display without scrolling downward.



Figure 35 - Discover Replicated Devices

It may be necessary to manually invoke discovery of replicated devices to view recently created replicated device groups (Tintri service groups). With the desired array pair highlighted, click the “Discover Replicated Devices” icon to invoke the discovery process.

- *Make sure that the replicated device group associated with a given Tintri service group is able to be discovered and displayed within the array manager hosting the service group before creating a protection group for the replicated device group.*

## Protection Group

Protection groups are objects related to a particular array manager pair and at least one replicated SRM device group.




Figure 36 - Create Protection Group

Within the “Array Manager” pane select the “Related Objects” tab and click “Create Protection Group”. This will launch the “Create Protection Group” workflow.

Figure 37 - Create Protection Group - Name and location

Provide a name and description for the protection group. Click the “Next” button to continue.

Figure 38 - Create Protection Group - Protection group type

The appropriate “Direction of protection” radio button should be selected to reflect the correct VMware SRM protected and recovery sites. For example, if the VMs in the replicated device group reside on the site named “dpl-vcsa-1” and this device group is being replicated to site “dpl-vcsa-2”, then the radio button that reflects this replication direction should be selected.

The “Protection group type” radio button “Datastore groups (array-based replication)” should be selected.

The array pair should be selected based on the array manager name and the desired array pair.

**Note:** It is not possible to select more than a single array pair for a given protection group.

Click the “Next” button to continue.

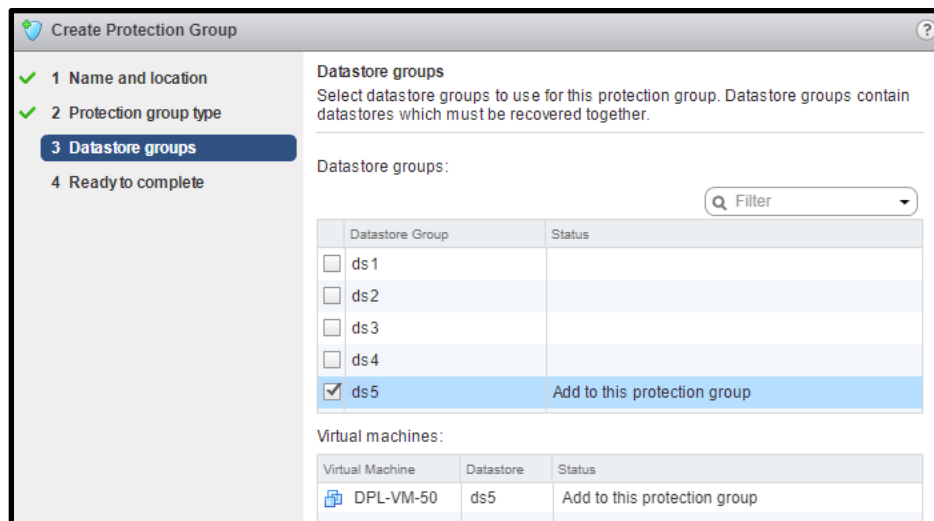


Figure 39 - Create Protection Group - Datastore groups

Click the appropriate checkbox or checkboxes to add one or more datastore groups to the protection group. The VMs contained within the selected datastore group are displayed in the “Virtual machines” section of the window.

Click the “Next” button to continue.

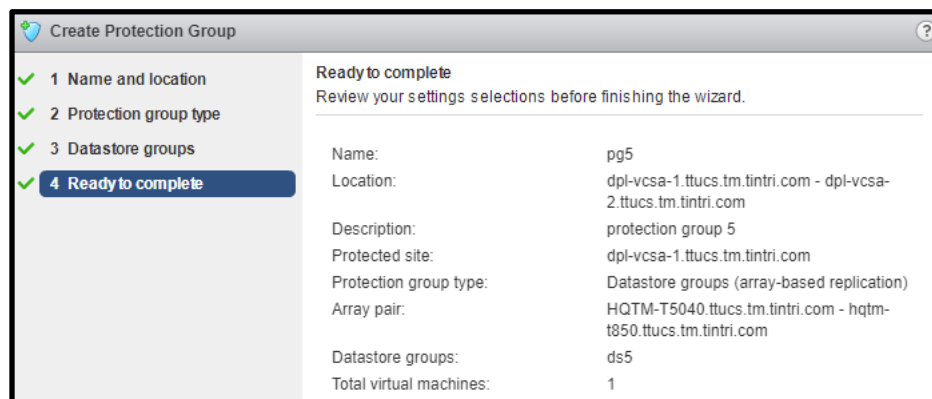


Figure 40 - Create Protection Group - Ready to complete

Click the “Finish” button to continue.




Figure 41 – Protection Groups

At this point the protection group has been created.

### **Recovery Plan**

A recovery plan includes information about the paired sites, which site will be the recovery site, and the protection group (or groups) that will be recovered. A recovery plan also includes information about test networks that will be used during testing of the plan.

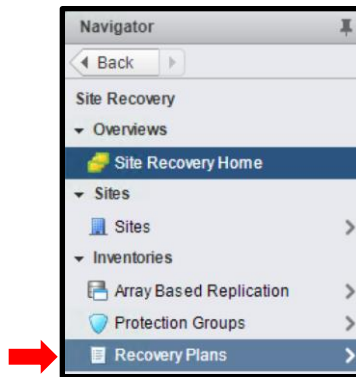


Figure 42 – Recovery Plans

Within the “Navigator” pane click “Recovery Plans”.




Figure 43 - Create Recovery Plan

From within the “Recovery Plans” pane, click “Create Recovery Plan”. This will launch the “Create Recovery Plan” workflow.

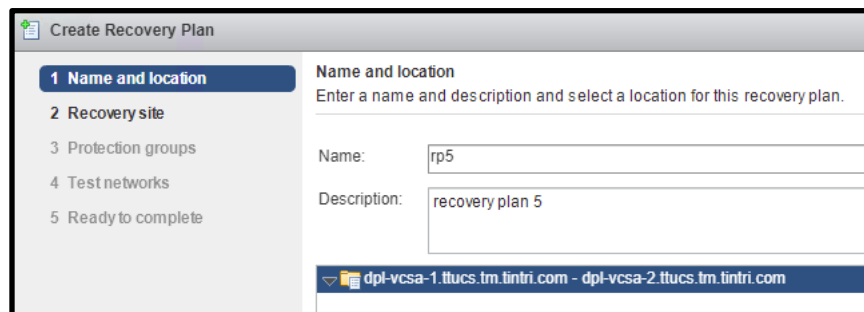


Figure 44 - Create Recovery Plan - Name and location

Enter a name and description for the recovery plan. Click the “Next” button to continue.




Figure 45 - Create Recovery Plan - Recovery site

Select the radio button corresponding to the recovery site. The recovery site is the site that the replicated device group is replicated to. Click the “Next” button to continue.





Figure 46 - Create Recovery Plan - Protection groups

The “Protection groups” window will appear. The “Group type” pull-down menu can be used to select “VM protection groups”. Also select the protection group (or groups) to be used with the recovery plan. Click the “Next” button to continue.

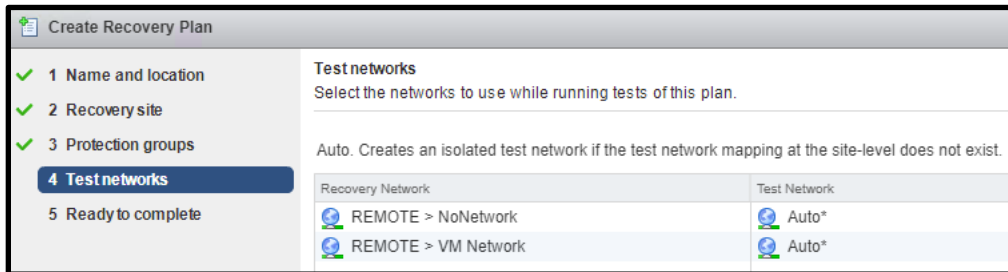


Figure 47 - Create Recovery Plan - Test networks

Accept the default values or click the “Test Network” column and row to select an alternate network from the pull-down menu that will appear. Click the “Next” button to continue.

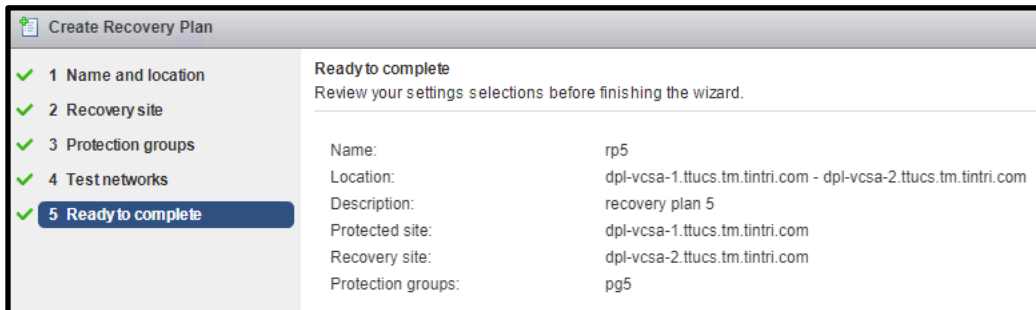


Figure 48 - Create Recovery Plan - Ready to complete

Click the “Finish” button to continue.




Figure 49 - Recovery Plans

The newly created recovery plan will appear, and is ready for use.

# Deployment Options

## Placeholder VMs

SRM automatically creates placeholder VMs in a non-replicated datastore designated for use as a placeholder datastore. In order for recovery plan failover with planned migration and reprotect operations to execute successfully, a placeholder datastore must exist at both sites. A placeholder datastore can be configured to reside on a Tintri VMstore array.

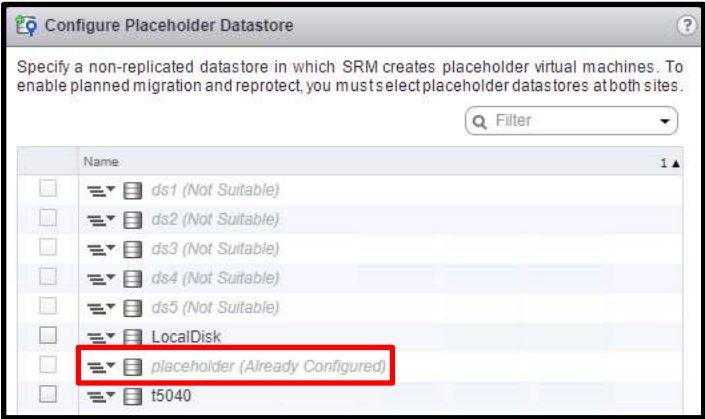


Figure 50 - Configure Placeholder Datastore

In the example provided a placeholder datastore has been configured on the datastore named “placeholder”. This a non-replicated datastore residing on a Tintri VMstore array. Note that replicated datastores used with SRM array based replication have “(Not Suitable)” status. Also note in the example provided that the “t5040” or “LocalDisk” datastores are available to be selected for use as a placeholder datastore if desired.

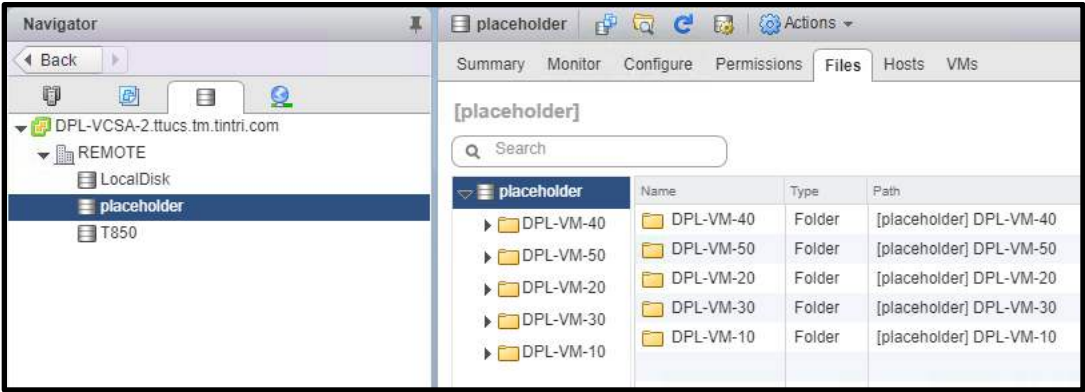



Figure 51 - Placeholder datastore with placeholder VMs on a Tintri VMstore array




| Tintri                   |                  |                    |
|--------------------------|------------------|--------------------|
| Dashboard                | Virtual Machines | VMstore            |
| Service groups           | VMs              | Virtual disks      |
| Snapshots                |                  |                    |
| VM                       | Replicate        | Replication state  |
| DPL-SRM-Local            |                  | Not configured     |
| DPL-VCSA-1               |                  | Not configured     |
| <i>DPL-VM-10</i> 12:50PM | hqtm-t850        | Replica up to date |
| <i>DPL-VM-10</i>         |                  | Not configured     |
| <i>DPL-VM-20</i>         |                  | Not configured     |
| <i>DPL-VM-20</i> 12:50PM | hqtm-t850        | Replica up to date |
| <i>DPL-VM-30</i> 12:50PM | hqtm-t850        | Replica up to date |
| <i>DPL-VM-30</i>         |                  | Not configured     |
| <i>DPL-VM-40</i>         |                  | Not configured     |
| <i>DPL-VM-40</i> 12:51PM | hqtm-t850        | Replica up to date |
| <i>DPL-VM-50</i> 12:50PM | hqtm-t850        | Replica up to date |
| <i>DPL-VM-50</i>         |                  | Not configured     |

Figure 52 - Placeholder VMs after a failover & prior to reprotect

In the example provided the VMs appearing with a red colored font are placeholder VMs in a state immediately following a successful failover. Hovering over the VM generates a pop-up informational dialog indicating that the VM is powered off and missing disks. This represents the expected behavior of placeholder VMs prior to a recovery plan reprotect operation being performed.

Also note that within the example provided there are VMs that appear with an italicized gray colored font. These are synthetic VMs that represent snapshots created as the result of Tintri service group RPO configuration settings.



| Tintri                          |                   |                 |
|---------------------------------|-------------------|-----------------|
| Dashboard                       | Virtual Machines  | VMstore         |
| Service groups                  | VMs               | Virtual disks   |
| Snapshots                       |                   |                 |
| VM                              | Replication state | Replicated from |
| DPL-SRM-Remote                  | Not configured    |                 |
| DPL-VCSA-2                      | Not configured    |                 |
| <i>DPL-VM-10</i> 12:35PM        | Incoming          | 192.168.200.249 |
| <i>DPL-VM-10 placeholder VM</i> | Not configured    |                 |
| <i>DPL-VM-20</i> 12:38PM        | Incoming          | 192.168.200.249 |
| <i>DPL-VM-20 placeholder VM</i> | Not configured    |                 |
| <i>DPL-VM-30 placeholder VM</i> | Not configured    |                 |
| <i>DPL-VM-30</i> 12:38PM        | Incoming          | 192.168.200.249 |
| <i>DPL-VM-40 placeholder VM</i> | Not configured    |                 |
| <i>DPL-VM-40</i> 12:38PM        | Incoming          | 192.168.200.249 |
| <i>DPL-VM-50 placeholder VM</i> | Not configured    |                 |
| <i>DPL-VM-50</i> 12:38PM        | Incoming          | 192.168.200.249 |

Figure 53 - Placeholder & synthetic VMs on a Tintri VMstore array

In the example provided, a recovery plan reprotect operation has been performed. Placeholder VMs are clearly identified with a “placeholder VM” designation. This represents the expected behavior of placeholder VMs after a recovery plan reprotect operation has been performed.

**Note:** When a placeholder datastore is configured on a Tintri VMstore array, users should avoid performing any Tintri snapshot, replication, recovery, or cloning operations on placeholder VMs. SRM manages placeholder VMs automatically, and there is no administrative action required on the part of the user.

**Note:** Placeholder VMs may or may not reside in a given placeholder datastore dependent on the status of the protection group. The protected site for a given protection group will not contain any placeholder VMs. The recovery site for a given protection group will contain placeholder VMs. SRM manages placeholder VMs automatically, and there is no administrative action required on the part of the user.

**Note:** This document purposefully provides no guidance on the subject of moving or migrating existing placeholder VMs to a new placeholder VM datastore. Users are advised to consult VMware SRM product documentation prior to undertaking any movement or migration of placeholder VMs.

### **Virtual Machines Not Protected by SRM**

VMs that are not included in an SRM protection group can also reside on a Tintri VMstore array.

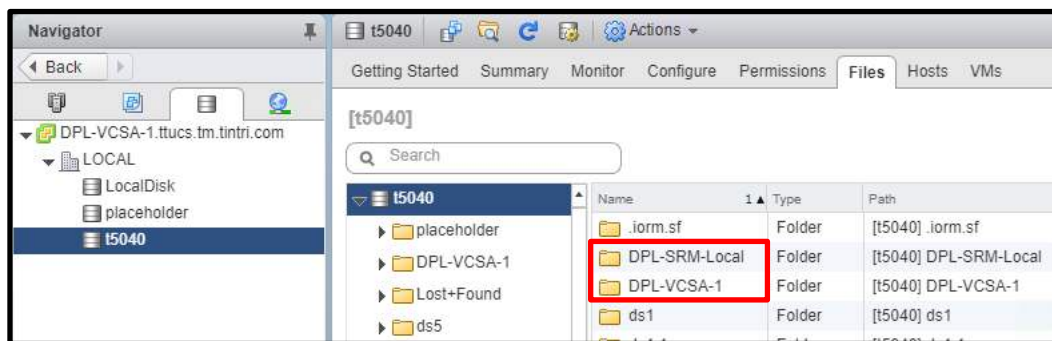


Figure 54 - VMs not protected by SRM on a Tintri VMstore array

In the example provided, the SRM Server host “DPL-SRM-Local” and the vCenter Server Appliance “DPL-VCSA-1” reside on a Tintri VMstore array and are not members of an SRM protection group. The placement of VMs onto a Tintri VMstore array enables tangible benefits that may not be available with local disk or other storage media. A subset of benefits includes:

- Tintri snapshots
- Tintri ReplicateVM
- Tintri cloning
- Tintri SyncVM which enables full VM, individual virtual disk, as well as granular folder and file recovery
- Tintri Quality of Service (QoS)
- Tintri individual VM analytics

### **Stale Folders**

By design, the Tintri SRA does not delete or overwrite the contents of a Tintri service group destination folder. Recovery plan execution of failover and failback operations will result in the creation of a new destination folders.

**Note:** The recovery plan test operation does not create stale folders. Folders used for recovery plan test operations are deleted at the point in time when the cleanup task is executed.

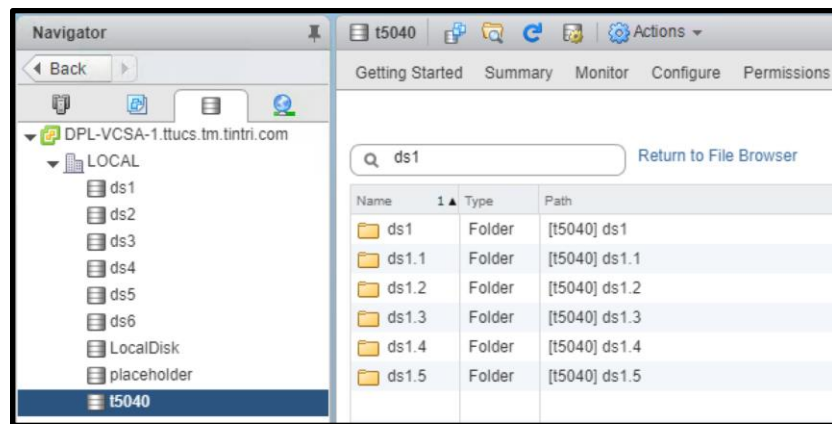


Figure 55 - Stale folders

In the example provided, the datastore named “ds1” was originally backed by the folder path “/tintri/ds1”. At the point in time when recovery plan failback and reprotect operations were performed, a folder named “ds1.1” was automatically created and utilized by SRM. At that point in time the “ds1” folder became stale. Subsequently, repeated recovery plan failover and failback operations were performed resulting in a number of stale folders.

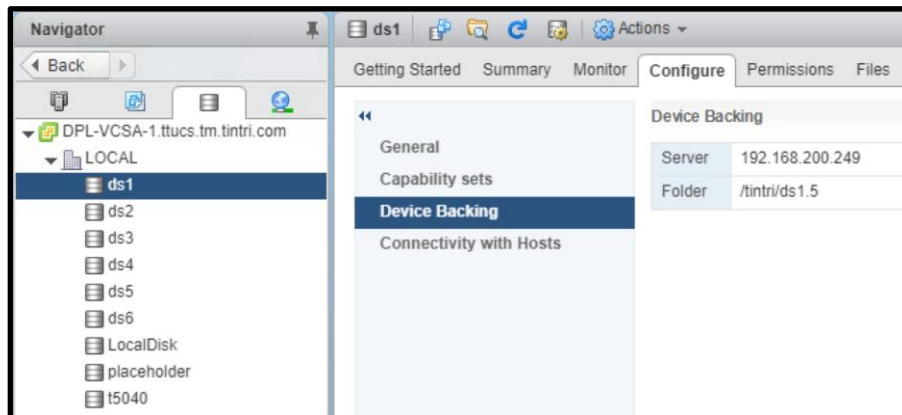


Figure 56 - Datastore "ds1" device backing

In the example provided, the datastore “ds1” is backed by the “/tintri/ds1.5” folder. In effect, all prior versions of the “ds1” folder have become stale and are effectively orphaned.

Within reason, the VM data in stale destination folders is considered valuable. In the event of an SRM infrastructure failure, unregistered VMs residing in a stale folder can potentially be added to inventory and powered on. VMs in stale folders are representative of the time at which they became stale.

When repeated recovery plan failover and failback operations occur, Tintri VMstore arrays used in conjunction with SRM may accumulate a significant number of stale folders. Stale folders consume array capacity, and a large number of stale folders may not be desired.

## Conclusion

The Tintri SRA enables tight, full featured integration between Tintri VMstore arrays and VMware SRM. Initial configuration and deployment is simple and straightforward. Array based replication leverages Tintri ReplicateVM and offloads the replication workload from the hypervisor. Recovery plan test, failover, reprotect, and failback operations execute flawlessly.

## References

## Tintri

Tintri reference documentation is available on the support portal at [support.tintri.com](http://support.tintri.com). Accessing the Tintri support portal requires login credentials.

- Tintri Storage Replication Adapter
  - SRA download
  - *Tintri SRA Plugin Release Notes*
  - *Tintri OS and External Compatibility Guide*
- *Tintri VMstore All Flash/Hybrid System Administration Manual*

## VMware

VMware Site Recovery Manager documentation is available at [docs.vmware.com/en/Site-Recovery-Manager](https://docs.vmware.com/en/Site-Recovery-Manager).

- *Site Recovery Manager Compatibility Matrices*
- *Site Recovery Manager Installation and Configuration*
- *Site Recovery Manager Administration*

The VMware SRM download portal is at <https://my.vmware.com>.

© 2017 Tintri, Inc. All rights reserved. Tintri, Tintri VMstore, Tintri Global Center, ReplicateVM, SecureVM, and SyncVM are trademarks of Tintri, Inc., and may be registered in the U.S. Patent and Trademark Office and in other jurisdictions. All other marks appearing in this publication are the property of their respective owners.

Tintri believes the information in this document is accurate as of its publication date. The information in this publication is provided as is and is subject to change without notice. Tintri makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose.

