

# EMC<sup>®</sup> Storage Integrator for Oracle VM Version 3.3.1

Release Number 3.0

## Product Guide

P/N 300-015-428

REV 03

APRIL 2015

**EMC<sup>2</sup>**

Copyright © 2015 EMC Corporation. All rights reserved. Published in the USA.

Published April 2015

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

This guide presents the following topics:

◆ Revision history .....	4
◆ Product description.....	4
◆ System prerequisites .....	6
◆ Installation .....	7
◆ Upgrade ESI for Oracle VM.....	10
◆ Use ESI for Oracle VM.....	12
◆ Getting help.....	25

# Revision history

Table 1 presents the revision history of this document.

**Table 1 Revision history**

Revision	Date	Description
03	April 2015	EMC Storage Integrator for Oracle VM version 3.3.1 Product Guide (release 3.0)
02	October 2014	EMC Storage Integrator for Oracle VM version 3.3 Product Guide (release 2.0) <b>Note:</b> Release 1.0 included release notes but not a product guide.

## Product description

EMC® Storage Integrator (ESI) for Oracle VM version 3.3.1 is a plug-in that enables Oracle VM (OVM) to discover and provision EMC storage arrays. The ESI plug-in is built upon the Oracle Storage Connect (OSC) framework. The framework provides a set of storage discovery and provisioning application programming interfaces (APIs) that enhance the ability to manage and provision storage devices in an OVM environment.

**Note:** For more information about the OSC framework and APIs, visit the Oracle website.

ESI for Oracle VM supports the following storage management and provisioning operations for EMC storage arrays:

- ◆ LUN discovery and inquiry
- ◆ Thin LUN/MetaLUN creation and removal
- ◆ Auto-provisioning (LUN masking)
- ◆ Register and refresh arrays
- ◆ Clone device creation and removal
- ◆ Snap device creation and removal

ESI for Oracle VM uses EMC SMI-S Provider or the Representational State Transfer (REST) interface to do the actual storage management work. SMI-S Provider supports the Storage Networking Industry Association (SNIA) Storage Management Initiative (SMI), an American National Standards Institute (ANSI) standard for storage management.

For storage provisioning and management operations, the ESI plug-in uses the SMI-S Provider interface for EMC VNX®, EMC VMAX®, and EMC VMAX3™ storage arrays. The plug-in uses the REST interface for EMC XtremIO™ storage arrays.

**Note:** The MetaLUN feature is not supported for VNX and XtremIO arrays.

## Interoperability differences between OSC and EMC Unisphere

EMC Unisphere® and EMC Solutions Enabler (SE) are the traditional storage provisioning and management interfaces for EMC storage arrays. Unisphere is a fully-functional storage provisioning tool for VMAX, VMAX3, and VNX arrays, and OSC simplifies storage array management for other array vendors. OSC uses a simplified management object model, but it does not exactly fit into a VMAX, VMAX3, or VNX object model. Some interoperability issues occur for mixed use of the EMC plug-in and Unisphere. In general, storage objects created by OSC should be consistently managed only by OSC. Similarly, storage objects (such as masking view, storage group, and so on) that are created by Unisphere should be managed only by Unisphere. Some interoperability issues occur when users attempt to manage the same storage objects with both the EMC OSC plug-in and Unisphere.

Unisphere provides all aspects of VMAX, VMAX3, and VNX management, including directors, disks, devices, pools, auto-provisioning, and so on. OSC provides a subset of functionality. The EMC XtremIO storage management application also provides easy-to-use tools for performing most of the cluster operations. OSC uses the XtremIO RESTful API, which has an HTTP-based interface for storage provisioning operations. [Table 2](#) shows the differences between OSC, Unisphere, and XtremIO terminology.

**Table 2 OSC Unisphere, and XtremIO terminology matrix**

OSC terms	Unisphere terms	XtremIO terms
Storage server or storage array	VMAX, VMAX3, or VNX array	XtremIO array
Storage element	VMAX, VMAX3, or VNX device	XtremIO volume
Volume group	<ul style="list-style-type: none"> <li>Thin pool (for thin device creation only)</li> <li>Storage resource pool (for VMAX3 arrays)</li> </ul>	XtremIO cluster shown as a volume group with the volume group named with the cluster name
Access group	<ul style="list-style-type: none"> <li>Initiator group (with default prefix <b>osc_</b> or a user-defined prefix) for VMAX and VMAX3 arrays</li> <li>Storage group (with default prefix <b>osc_</b> or a user-defined prefix) for VNX arrays</li> </ul>	Initiator group (with default prefix <b>osc</b> or a user-defined prefix)
Present or unpresent	LUN masking, to make a LUN visible or invisible to a host	LUN mapping allows the ports of the host to access the volume

### LUN masking

LUN masking in OSC differs significantly from Unisphere, and interoperability cannot be implemented between them.

The basic guideline for LUN masking is that OSC and Unisphere should never interfere with each other. Do not use Unisphere to modify OSC-managed masking groups and views, and do not use OSC to modify masking groups and views managed by Unisphere. Violating these guidelines may cause unexpected errors.

To distinguish OSC-managed masking groups and views from masking groups and views managed by Unisphere, a default **osc\_** prefix (or any user-defined prefix) is added to the name of OSC-managed objects. Never modify OSC-managed objects with Unisphere. OSC blocks the association of any OSC-managed initiator groups, storage groups, and port groups to masking views managed by Unisphere. Any modifications performed outside of OSC may cause unexpected failures.

To specify an access group prefix other than the default **osc** prefix, type the following command in the configuration file:

```
AccessGroupPrefix = User_Defined_Prefix  
/opt/storage-connect/plugins/emc/isa/isa.conf
```

For example, to define the **mycomp** prefix, type the following line in the configuration file:

```
AccessGroupPrefix = mycomp
```

---

#### Notes:

- ◆ To avoid operation failures, configure the prefix before you configure the OSC usage, and do not change the configuration after you create the access group.
- ◆ If the OVM manager manages multiple servers, the **isa.conf** file must be consistent across all the servers. Configure the same value for **AccessGroupPrefix** in all of the managed servers.
- ◆ [“Edit the ESI for Oracle VM configuration file”](#) provides more details about how to edit the configuration file.
- ◆ “Known problems and interoperability for OSC with Unisphere and Solutions Enabler” in the *EMC Storage Integrator for Oracle VM Version 3.3.1 Release Notes* provides more details about OSC and Unisphere interoperability.
- ◆ The EMC XtremIO Management Application is used to perform storage management and provisioning operations on XtremIO arrays. The other restrictions in these notes also apply to XtremIO arrays.

## System prerequisites

[Table 3](#) lists all the configurations supported for version 3.0 of ESI for Oracle VM version 3.3.1.

**Table 3 Environment and system requirements**

Type	Name	Supported version
Oracle environment	Oracle VM Manager and Oracle VM Server	3.3.1
Array management server	EMC SMI-S Provider	<ul style="list-style-type: none"> <li>• 4.6.2.9 (VMAX and VNX)</li> <li>• 8.0.1, 8.0.2 (VMAX3)</li> </ul>

**Table 3 Environment and system requirements**

Type	Name	Supported version
Disk array (VMAX, VMAX3, VNX, or XtremIO)	VMAX arrays	EMC Enginuity™ 5874
	VMAX3 arrays	EMC Enginuity 5977
	VNX arrays	Fibre Logic Array Runtime Environment (FLARE)/OE 5.32 and 5.33
	XtremIO arrays	Version 3.0.2 build 14 or later

# Installation

This section describes how to install, uninstall, and deploy ESI for Oracle VM, and also includes details on how to edit the configuration file to define runtime behavior for the plug-in.

## Install ESI for Oracle VM

Before you run the installation program, confirm that the system configuration is listed in the “[System prerequisites](#)” and verify that you have root access to the Oracle VM server.

To install ESI for Oracle VM:

1. Download the installation package file from EMC Online Support (<https://support.emc.com>).
2. Log on as the root user to the Oracle VM server.
3. Type the following command to install the package, as shown in [Figure 1](#):

```
# rpm -ivh ./emc-osc-isa- 3.0.0.7-1.e16.x86_64.rpm
```

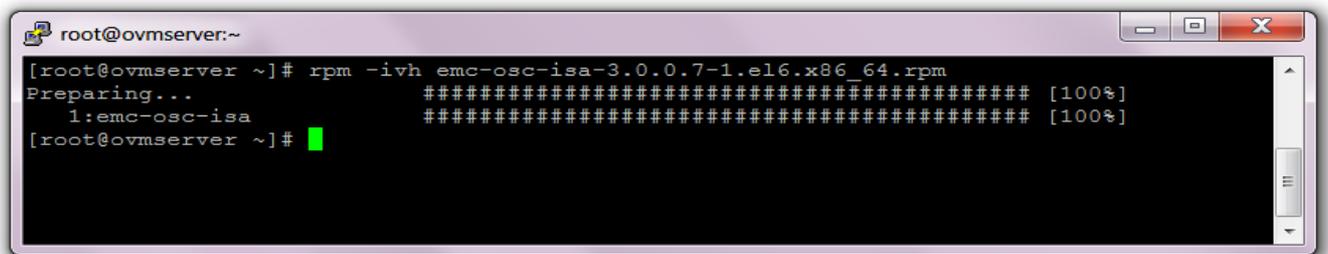
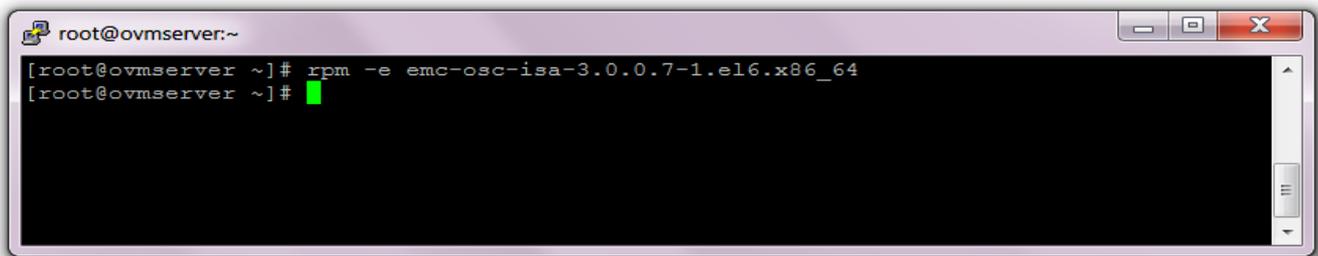


Figure 1 Install ESI

## Uninstall ESI for Oracle VM

To uninstall ESI for Oracle VM, type the following command to remove the package, as shown in [Figure 2](#):

```
# rpm -e ./emc-osc-isa-3.0.0.7-1.e16.x86_64
```



```

root@ovmsserver:~
[root@ovmsserver ~]# rpm -e emc-osc-isa-3.0.0.7-1.el6.x86_64
[root@ovmsserver ~]#

```

Figure 2 Uninstall ESI

## Deploy ESI for Oracle VM

ESI for Oracle VM is deployed to an Oracle VM environment and is installed on an Oracle VM Server. It requires the EMC SMI-S Provider server to provision and manage VMAX arrays. The provider server runs on a standalone Windows or Linux management host. *EMC SMI-S Provider Release Notes* on EMC Online Support provide installation and configuration instructions.

The SMI-S Provider server is based on Solutions Enabler, so it requires the same setup and connectivity to the VMAX array, such as zoning and configuring gatekeeper devices. The OVM server (where the ESI plug-in is installed and runs) requires TCP/IP network connections to the SMI-S Provider server and connectivity to the storage array. The SMI-S Provider server cannot be installed on OVM servers because OVM servers have a restricted environment.

For the XtremIO array, the REST server that runs on the XMS array enables you to control and manage the array by providing GUI and RESTful API services to clients with all communications, by using standard TCP/IP connections.

## Edit the ESI for Oracle VM configuration file

ESI for Oracle VM uses **isa.conf** to define its runtime behavior. To change the default behaviors, manually create a plain text file named **isa.conf** under **/opt/storage-connect/plugins/emc/isa/** on the Oracle VM server after installation.

Permissions of the **isa.conf** file must be readable by the root user. Any line starting with **#** will be ignored by ESI for Oracle VM. [Table 4](#) describes the ESI plug-in runtime settings.

**Table 4 isa.conf properties**

Properties	Values	Descriptions
AccessGroupPrefix	A string composed of alphanumerical characters, “_” or “.”  <b>Note:</b> For example: <b>osc1_</b>	User defined prefix for OSC-managed initiator group, port group, device group, and masking view on VMAX, VMAX3, VNX, and XtremIO arrays. This prefix makes it easier to differentiate between OSC-managed groups and other managed groups.
AutoMetaEnabled	<b>True</b> or <b>False</b>	A flag to indicate whether a VMAX auto meta is enabled or disabled.  <b>Note:</b> VNX, VMAX3, and XtremIO arrays do not support meta devices. Devices up to 16 TB in size can be created on these arrays. These arrays handle the creation of maximum size LUNs.
LogLevel	debug	Set the verbose logging flag for troubleshooting.
VolumeGroupPrefix	A prefix or the exact pool name. Multiple comma-separated values can be specified.  <b>Note:</b> For example: <b>esi, osc_pool1</b>  <b>Note:</b> To apply these configuration changes on the configured arrays, remove and re-register the arrays.	Filters and lists only the desired volume groups in the storage array.  <b>Note:</b> This configuration is only applicable for VMAX and VNX arrays and will not affect VMAX3 and XtremIO arrays.

# Upgrade ESI for Oracle VM

To upgrade the ESI plug-in from version 2.0 to 3.0, first uninstall the version 2.0 plug-in and then install the version 3.0 plug-in.

**Note:** To upgrade ESI for Oracle VM from version 3.2.8 to 3.3.1 of the Oracle VM server, refer to the Oracle migration guide. Contact Oracle support if you encounter any issues during the upgrade.

To upgrade the storage servers from version 2.0 to version 3.0:

**Note:** When you perform the upgrade, also back up the configuration file so that it can be restored after you complete the upgrade.

1. Navigate to **Storage > SAN servers > Edit**.
2. Select the EMC storage plug-in from the drop-down list.
3. Click **OK**.

After you successfully upgrade ESI for Oracle VM, you must rediscover the OVM server from the **Servers and VMs** tab for all the OVM servers where ESI for Oracle VM is installed.

To rediscover the OVM server:

**Note:** Stop the **ovs-agent** service on the OVM server while uninstalling and installing the OSC plug-in RPMs.

1. In **Servers and VMs**, click **Rediscover Server**, as shown in [Figure 3](#).

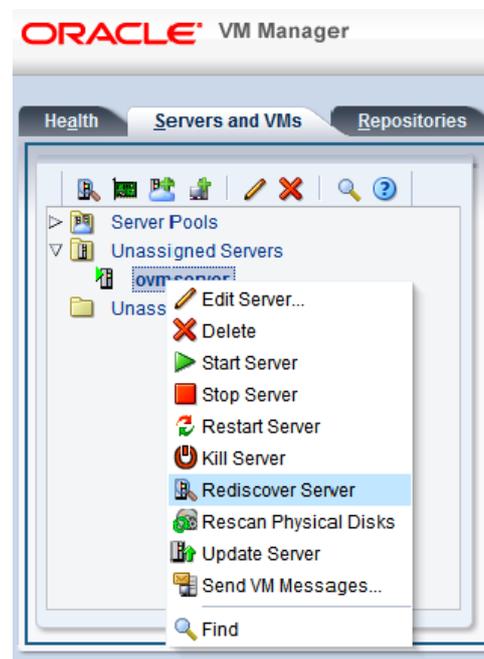
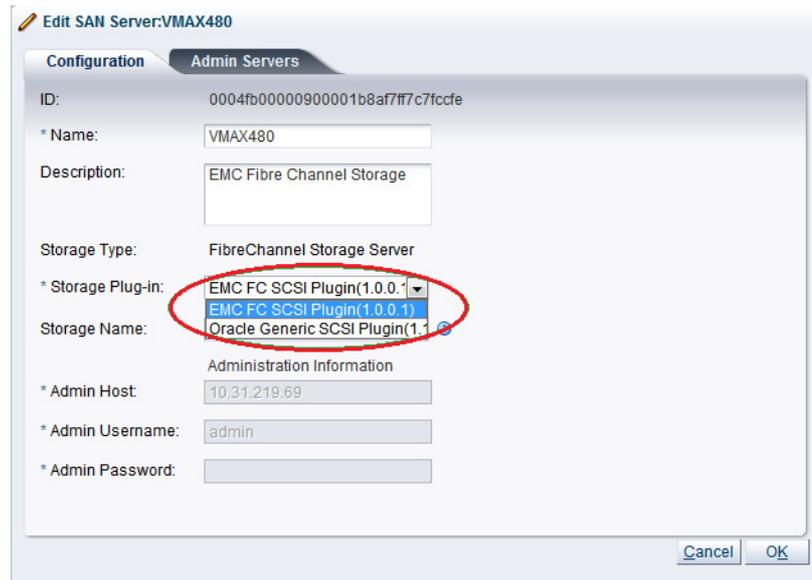


Figure 3 Rediscover the OVM server

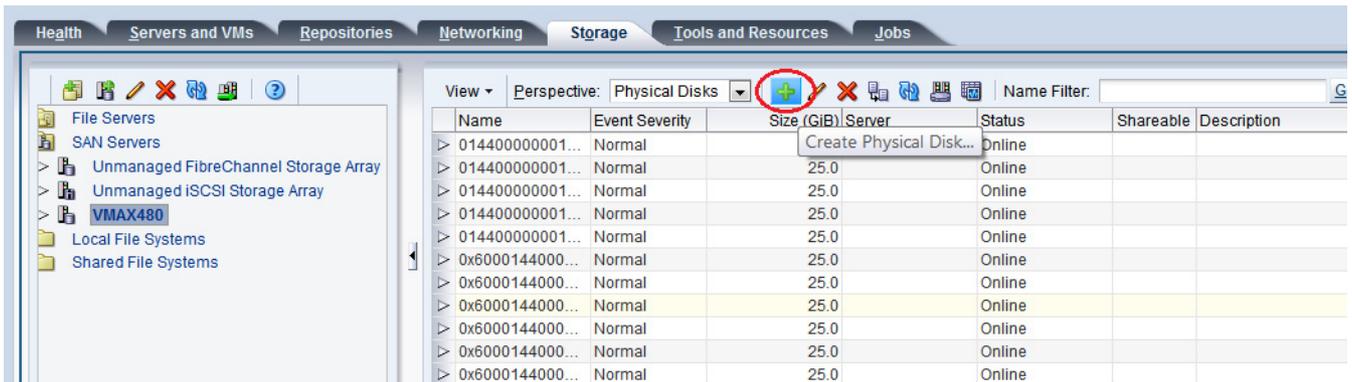
2. In **Configuration**, select the plug-in that you are upgrading, and then click **OK**, as shown in [Figure 4](#).



**Figure 4** Select the plug-in to be upgraded

In **Storage**, the green + confirms that the plug-in was successfully upgraded, as shown in [Figure 5](#).

**Note:** You will see a red + if the upgrade was unsuccessful. Repeat the steps until you are successful.



**Figure 5** Upgrade confirmation

# Use ESI for Oracle VM

This section describes how to use ESI for Oracle VM, including the following topics:

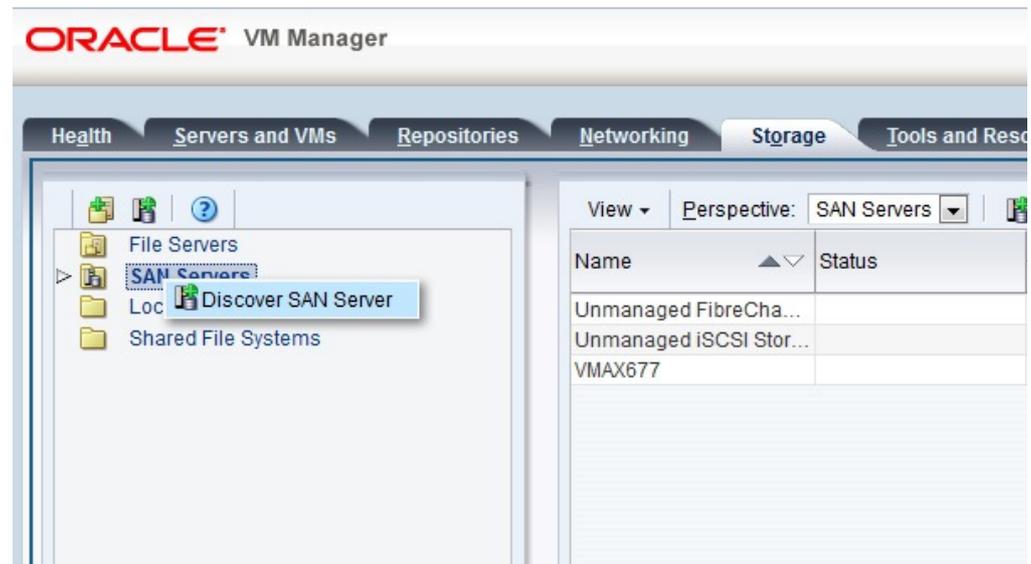
- ◆ Register a storage array
- ◆ LUN discovery and inquiry
- ◆ Create and remove thin LUNs
- ◆ Auto-provisioning (LUN masking)
- ◆ Create access groups
  - Create access groups
  - Edit access groups
    - Add or remove initiators from an access group
    - Start or stop presenting physical disks
- ◆ Create snapshots
- ◆ Create clones

## Register a storage array

You must install ESI for Oracle VM on an OVM server and then discover the plug-in through OVM Manager before a storage array can be registered.

To register and refresh a SAN storage array:

1. Click **Storage > SAN Servers > Discover SAN Server**, as shown in [Figure 6](#).



**Figure 6** Rediscover the server

2. In the **Access Information (if required)** dialog box, enter the following details, as shown in [Figure 7](#), and then click **Next**:
  - **Name**—User defined name for an EMC storage array.

- **Storage Plug-in**—Select **EMC FC SCSI Plugin**.
- **Plug-in Private Data**—Storage array ID, or use the last several digitals of the storage array ID, if they uniquely identify the storage array. For example, both **SYMMETRIX+000195900311** or **311** are valid, but **SYMMETRIX+00019** is invalid. **XtremIO** is valid for an XtremIO array. This field is required in case there are multiple storage arrays managed by EMC SMI-S Provider.
- **Admin Host**—Host name or IP address of the host where SMI-S Provider is installed. For the XtremIO array, specify the IP address used to connect to the XMS server.
- **Admin Username**—Username to connect to the SMI-S Provider. The default value is **admin**.
- **Admin Password**—Password to connect to the SMI-S Provider. The default value is **#1Password**. For the XtremIO array, specify the password used for the admin user.

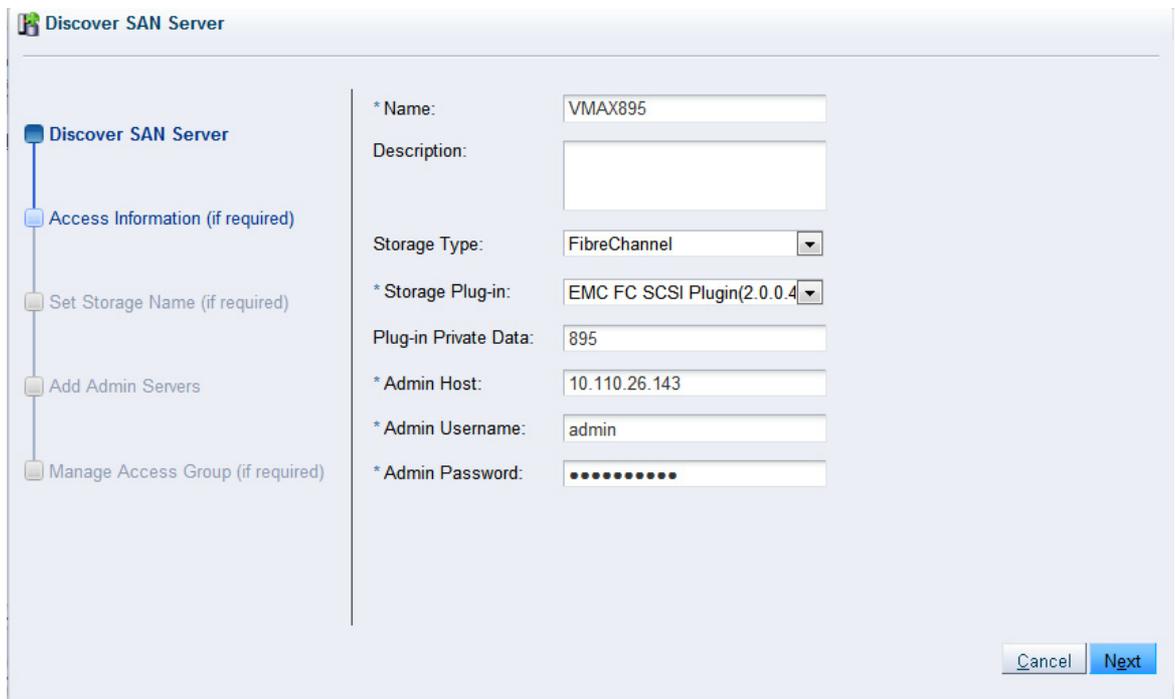
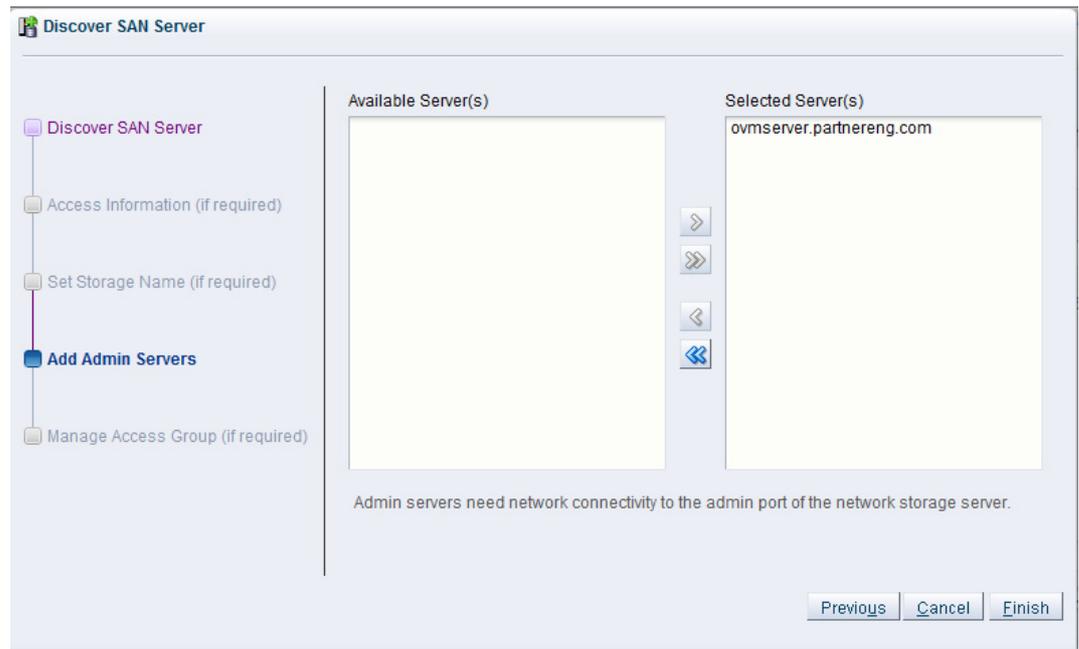


Figure 7 SAN Server registration for VMAX, VMAX3, and VNX arrays

3. In the **Add Admin Servers** dialog box, select the OVM server to be used for the storage operations from the **Available Servers(s)** list.

- Click the arrow to move the OVM server to the **Selected Servers(s)** list, and then click **Finish**, as shown in [Figure 8](#).



**Figure 8** Select the OVM server

- Click **Storage > SAN Servers >** the selected server > **Refresh** to refresh the storage array that you added, as shown in [Figure 9](#).



**Figure 9** Refresh the server

## LUN discovery and inquiry

After registering a storage array, refresh the storage array to list available storage access groups, volume groups, and devices.

Inquiry capabilities of a storage device include size, device World Wide Name (WWN), status, LUN ID, associated volume group, and access group, if any.

Only thin Fixed Block Architecture (FBA) devices are discoverable. Business Continuance Volumes (BCVs), clones, and R1 local devices are discoverable if they are also thin FBA devices. Count key data devices (CKDs), data devices (DATADEVs), save devices (SAVEDEVs), R2 target devices, and gate keeper (GK) devices are not discoverable.

The format of extra information associated with a storage device is as follows:

```
[UNMANAGED BCV/R1] LUN_ID [(A MIRROR/CLONE OF LUN_ID: COPY_STATE)]
```

**Note:** Items in square brackets ([ ]) indicate terms that are specific to special configurations and are not applied to every LUN.

where the following values are used:

- ◆ **Unmanaged**—LUN is not managed by OSC access group
- ◆ **BCV**—LUN is a BCV device
- ◆ **R1**—LUN is an R1 device
- ◆ **LUN\_ID**—VMAX, VMAX3, VNX, or XtremIO device number
- ◆ **A mirror/clone (or snapshot) of LUN\_ID: COPY\_STATE**—LUN is a replica target device (mirror or clone) with the copy state information

**Note:** XtremIO volumes with a Network Address Authority (NAA) name and page83 ID are listed by Oracle VM. The ESI plug-in ensures that the page83 ID is set for all volumes, clones, and snapshots created by Oracle VM.

## Create and remove thin LUNs

Only thin devices can be created.

If the size specified is larger than the threshold of a standard device, a VMAX meta device is created as follows:

- ◆ If auto meta is disabled on the VMAX array, the ESI plug-in creates the required meta members first, then forms a concatenated meta device using the newly-created member devices. Striped meta devices cannot be created in this way.
- ◆ If auto meta is enabled on the VMAX array, add the following line in the configuration file `/OPT/STORAGE-CONNECT/PLUGINS/EMC/ISA/ISA.CONF` on the Oracle VM server:

```
AutoMetaEnabled = true
```

---

**Notes:**

- VNX, VMAX3, and XtremIO arrays do not support meta devices and the **AutoMetaEnabled** configuration parameter is not applicable for these arrays. You can create maximum-sized LUNs in these arrays.
  - [“Edit the ESI for Oracle VM configuration file”](#), provides more details about how to edit the configuration file.
  - The ESI plug-in depends on the VMAX auto meta feature to complete the device creation request. The physical disk size created by the auto meta feature may be greater than expected because the size must be a multiple of the **VMAX Auto Meta Member Size** attribute.
  - For VMAX arrays, you can also specify the meta member size using the **Extra Information** field of the **Create Physical Disk** dialog box. Use **meta\_member\_size=<Meta Member Size in GB>** to define the meta member size.
  - If the **AutoMetaEnabled** parameter is not set in the configuration file, the plug-in attempts to query whether the auto meta feature is enabled on the storage array when SMI-S Provider 4.6.2.9 or greater is used. If this query is unsuccessful, **AutoMetaEnabled** uses the default value of **FALSE**.
  - For VMAX arrays, if **meta\_member\_size** is specified while the VMAX auto meta is disabled, the ESI plug-in will always try to create a concatenated meta device according to the specified member size.
  - Volume group (storage pool) management occurs outside of OSC. Volume groups must be pre-configured before OSC usage.
  - Although BCVs and R1 devices are discoverable, they cannot be created through OVM. To remove these devices, follow the general Solutions Enabler guidelines for VMAX arrays. *EMC Solutions Enabler Symmetrix Array Control CLI Product Guide*, on EMC Online Support, provides more information.
- 

## Auto-provisioning (LUN masking)

To make a device accessible from Oracle VM server, you need to present it to an access group. The equivalent of an access group is an initiator group (in VMAX, VMAX3, and XtremIO) and a storage group (in VNX), which includes a list of host initiator WWNs.

After being presented, VMAX device information is updated in OVM. Access groups of a storage device are changed to the access group name. The device path becomes available on the Oracle VM server and the device is ready to use for OVM.

Ensure that you correctly configure zoning on the Fibre Channel switch before LUN masking. All host initiators must have Fibre Channel logins on the storage array to successfully add them to an access group. Adding an initiator to an access group in Oracle VM Manager requires a connection to the storage array with the Oracle VM server.

## Create access groups

To create and edit access groups:

1. Click **Storage > SAN Servers**, and then select the SAN storage array for the access group that you are creating.
2. Select a **Perspective of Access Groups**, and then click **+**, as shown in [Figure 10](#).

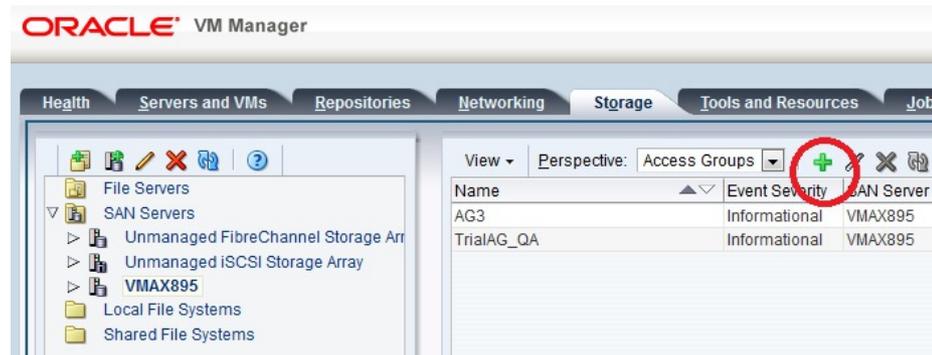


Figure 10 Storage tab

3. In the **Create an Access Group** dialog box, type an **Access Group Name** and a **Description**, and then click **Next**, as shown in [Figure 11](#).

**Note:** Do not include spaces in the **Access Group Name**. For example, use **Access\_Group\_1** or **AccessGroup1** instead of **Access Group 1**.

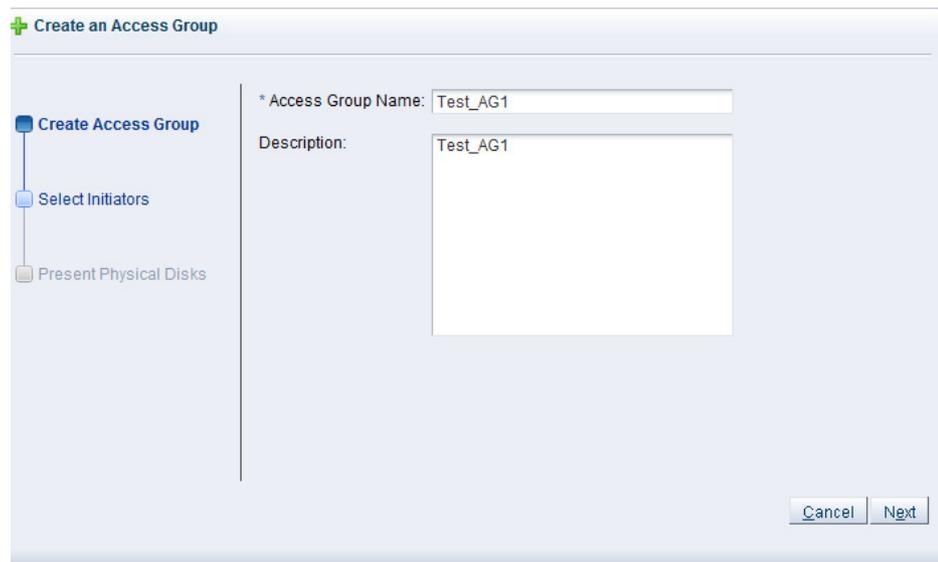


Figure 11 Create an access group

- In the **Select Initiators** dialog box, select an initiator from the **Available Storage Initiators** list.
- Click the arrow to move the initiator into the **Selected Storage Initiators** list, and then click **Next**, as shown in [Figure 12](#).

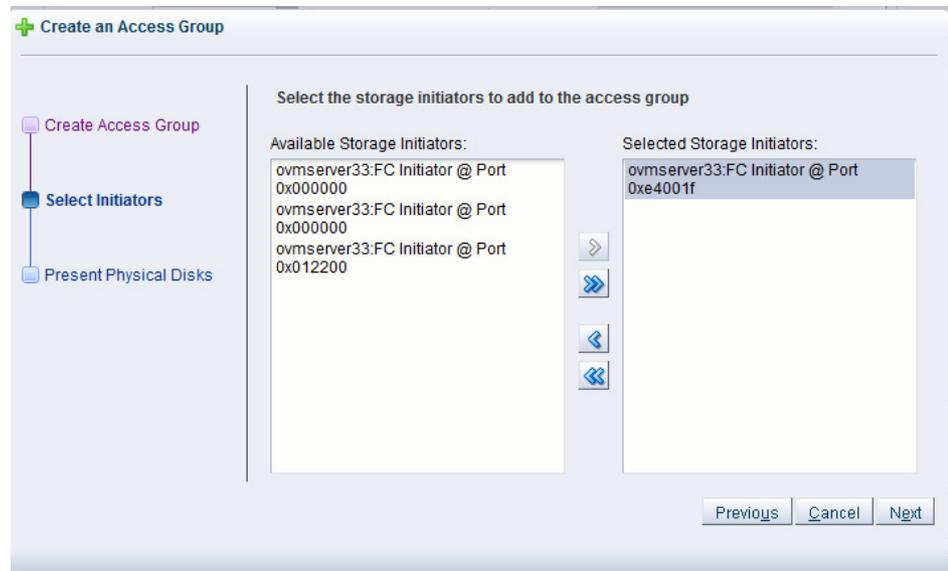


Figure 12 Select an initiator

- From the **Present Physical Disks** dialog box, select one or more physical disks from the **Available Physical Disks** list.
- Click the arrow to move the disk into the **Selected Physical Disks** list, and then click **Finish**, as shown in [Figure 13](#).

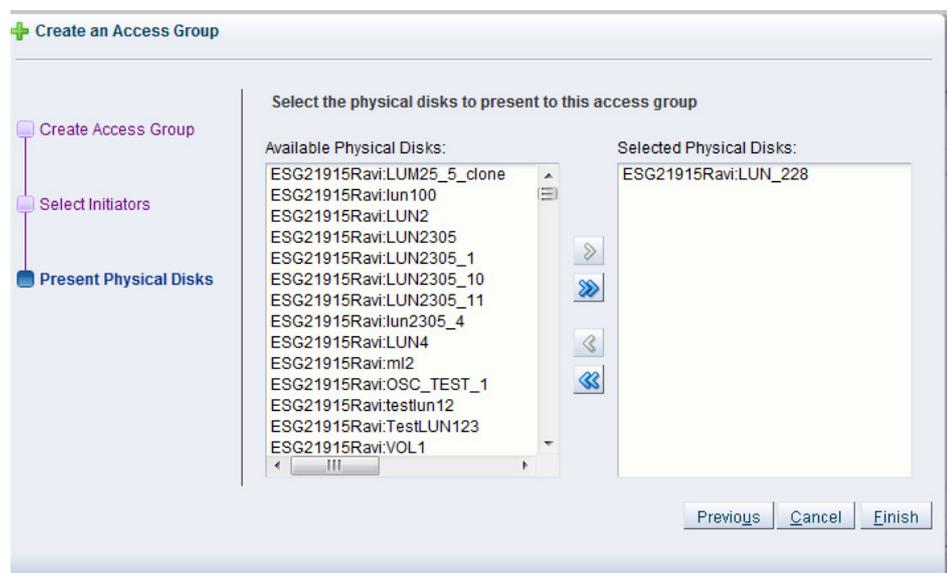


Figure 13 Present physical disks

## Add or remove initiators in an access group

To add an initiator to an access group:

1. Click **Storage Initiators** and select an initiator from the **Available Storage Initiators** list.
2. Click the arrow to move the initiator into the **Selected Storage Initiators** list, and then click **Next**, as shown in [Figure 14](#).

To remove an initiator from an access group:

1. Click **Storage Initiators** and select an initiator from the **Selected Storage Initiators** list.
2. Click the arrow to move the initiator into the **Available Storage Initiators** list, and then click **Next**, as shown in [Figure 14](#).

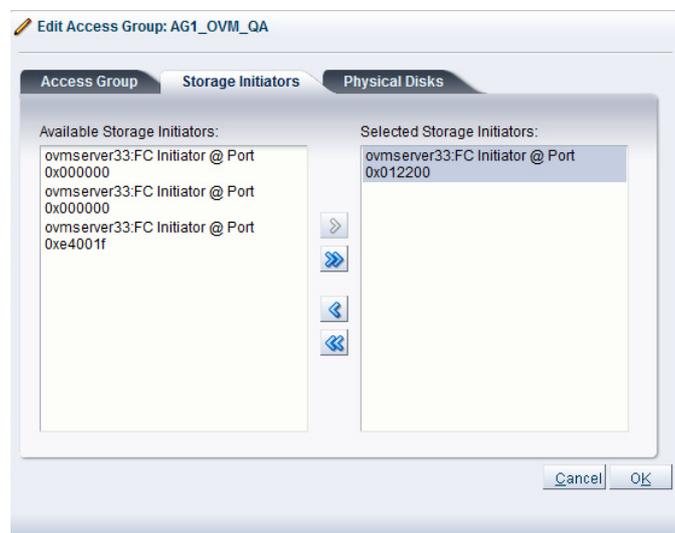


Figure 14 Add an initiator to an access group

## Start or stop presenting physical disks

To start a physical disk:

1. Select **Physical Disks** and select a physical disk from the **Available Physical Disks** list.
2. Click the arrow to move the initiator into the **Selected Physical Disks** list, and then click **OK**, as shown in [Figure 15](#).

To stop a physical disk:

1. Select **Physical Disks** and select a physical disk from the **Selected Physical Disks** list.

2. Click the arrow to move the physical disk into the **Available Physical Disks** list, and then click **OK**, as shown in [Figure 15](#).

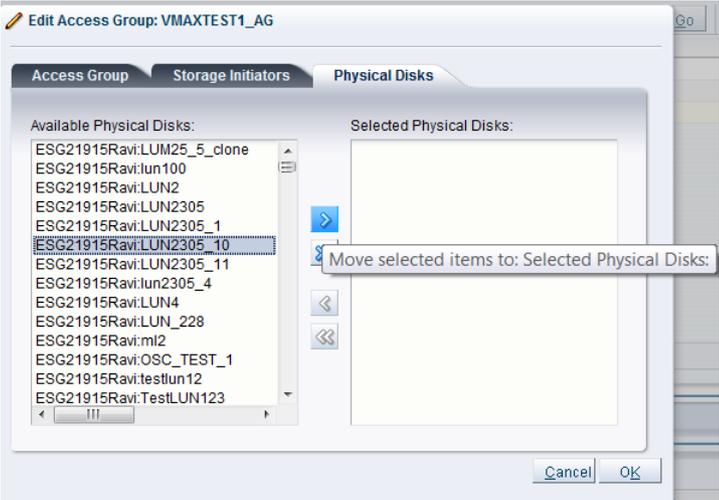


Figure 15 Start or stop presenting a physical disk

# Create clone devices

To create a clone device:

1. Click **Storage** and select **Perspective: Physical Disks**.
2. Select a **Volume Group**, right-click the storage element, and then select **Clone**, as shown in [Figure 16](#).

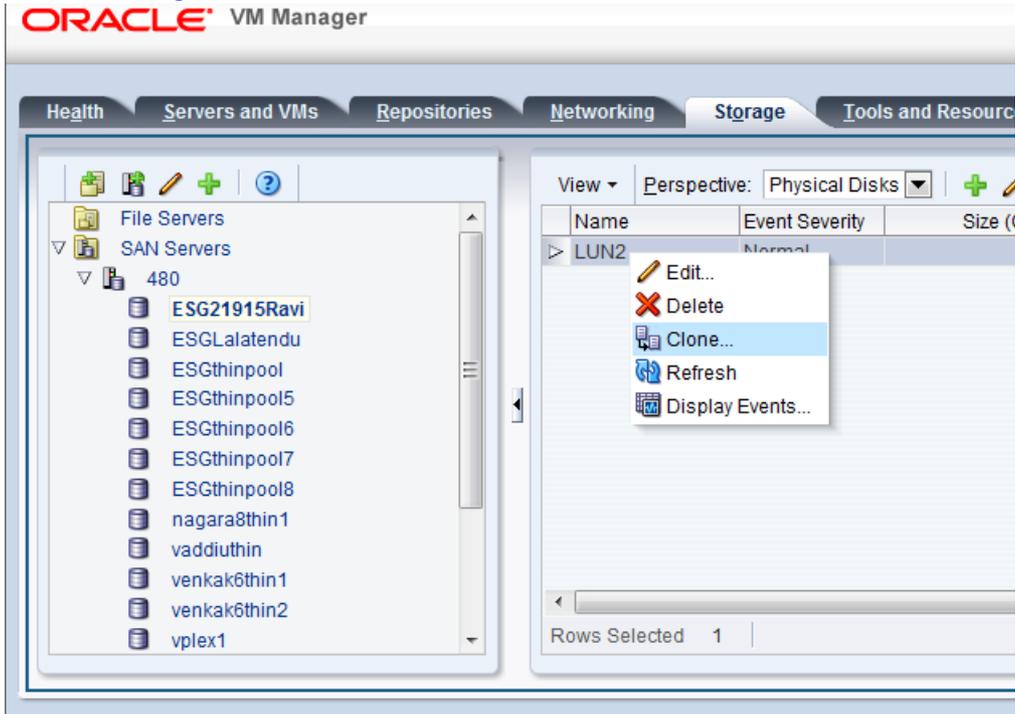
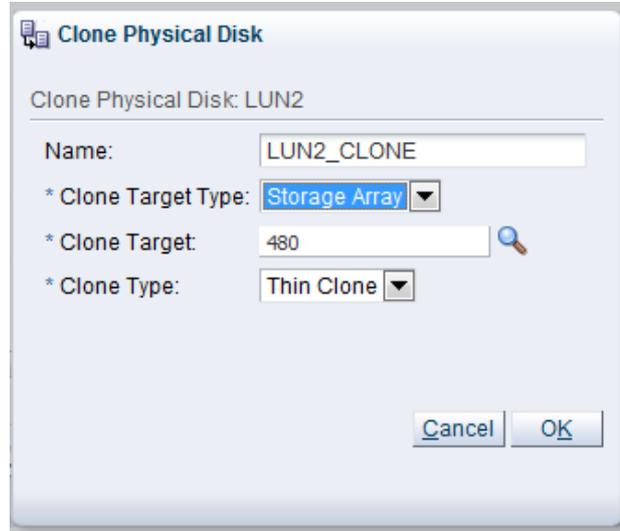


Figure 16 Selecting a clone

3. Type the clone name.
4. Select **Storage Array** as the clone target type.
5. Select the **EMC Storage Array** as the clone target.
6. Select **Thin Clone** as the clone type, and then click **OK**, as shown in [Figure 17](#).



**Figure 17** Specifying the clone details

After the clone is created, the clone becomes immediately accessible to its host, even while data copying is occurring in the background.

A source device can have multiple clones (different point-in-time copies) at the same time. However, the name for all clone devices is the same as the source device with a **\_clone** suffix. OVM does not provide a way to type a user-friendly name while creating a clone.

---

**Note:** Clones from snapshots are not supported, regardless of whether they are synchronous or asynchronous.

---

OSC clone functionality is limited to creation and removal. The advanced EMC TimeFinder® clone features such as Recreate and Restore are not available in OSC. Always create a clone for a new point-in-time copy.

---

**Note:** Because XtremIO supports only snapshots, the clones shown for XtremIO arrays are actually snapshots. OSC API **getCloneLimits()** and **getSnapLimits()** commands are not supported for XtremIO arrays.

---

## Create snapshots

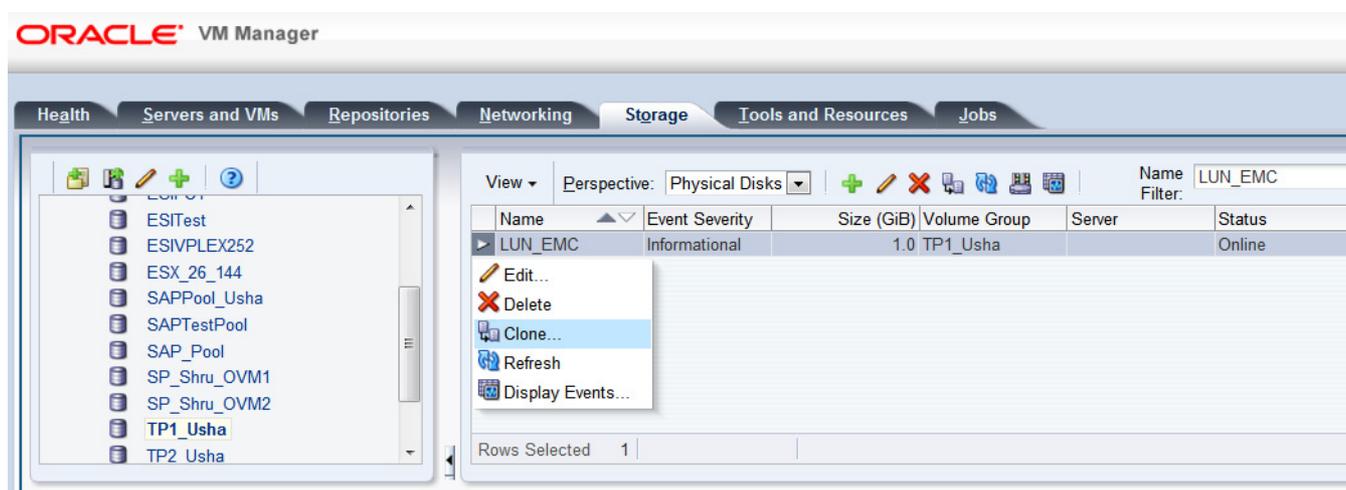
Snapshots that you create are listed under the volume group where the source storage element resides. Snapshots created with ESI can also be deleted, but cannot be restored. Snapshots cannot exist independently without a source. For that reason, source storage elements cannot be removed unless all associated snapshots are also removed.

For VMAX arrays with Enginuity version 5876 or above, the plug-in creates an EMC VP Snap. For earlier versions of Enginuity, the plug-in creates a TimeFinder Snap. Snapshot creation is extended from existing clone support by naming the replica with an **\_EMC\_SNAP** suffix. The ESI plug-in lists only OSC-managed snapshots and snapshots managed by Unisphere are filtered out.

VMAX3 arrays with Enginuity version 5977 use the new SnapVX snapshot technology, which is the only type of snapshot supported by the ESI plug-in.

To create a snapshot:

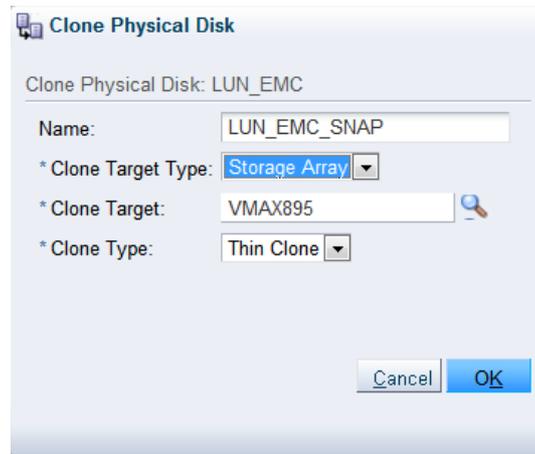
1. Click **Storage** and select a **Perspective of Physical Disks**.
2. Select a **Volume Group**, right-click the storage element, and then select **Clone**, as shown in [Figure 18](#).



**Figure 18** Selecting a clone

3. Type the snapshot **Name**, which must end with **\_EMC\_SNAP**.
4. Select **Storage Array** as the **Clone Target Type**.
5. Select the EMC storage array as the **Clone Target**.
6. Select **Thin Clone** as the **Clone Type**.

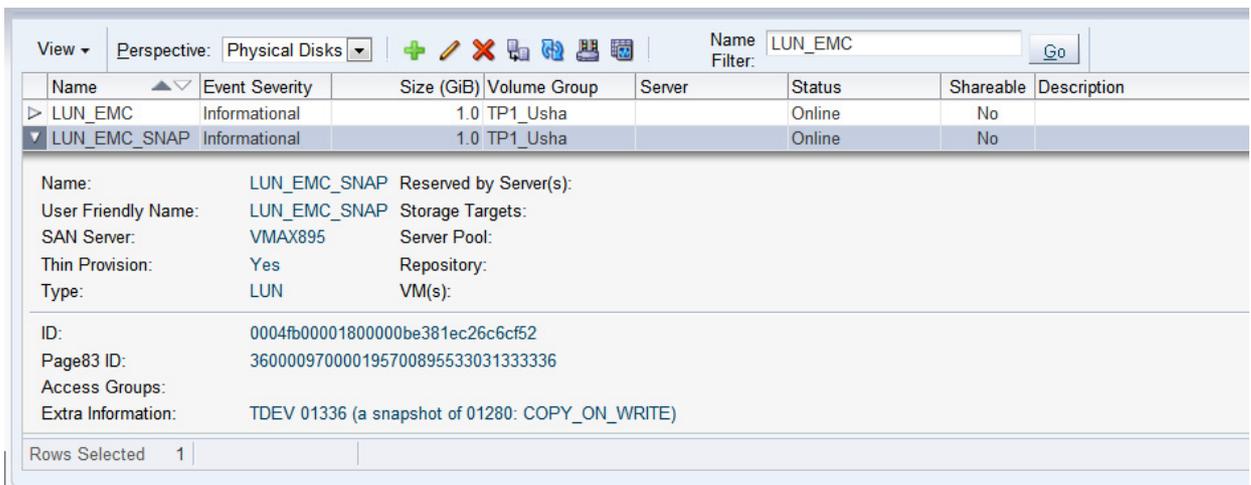
7. Click **OK**, as shown in [Figure 19](#).



**Figure 19** Specifying the clone physical disk details

8. Click the snapshot created under the selected volume group to view details for the following arrays:

a. VMAX and VNX arrays, as shown in [Figure 20](#).



**Figure 20** View snapshot details for VMAX and VNX arrays

b. VMAX3 arrays, as shown in [Figure 21](#).

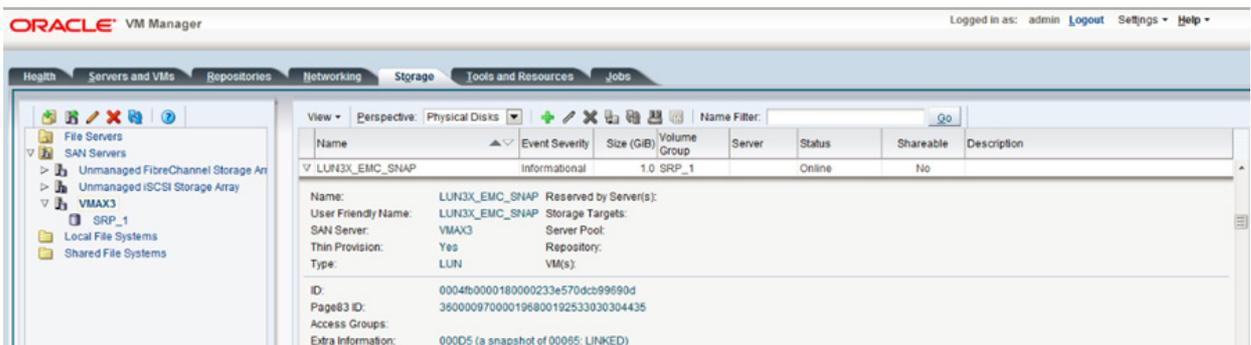


Figure 21 View snapshot details for VMAX3 arrays

c. XtremIO arrays, as shown in Figure 22.

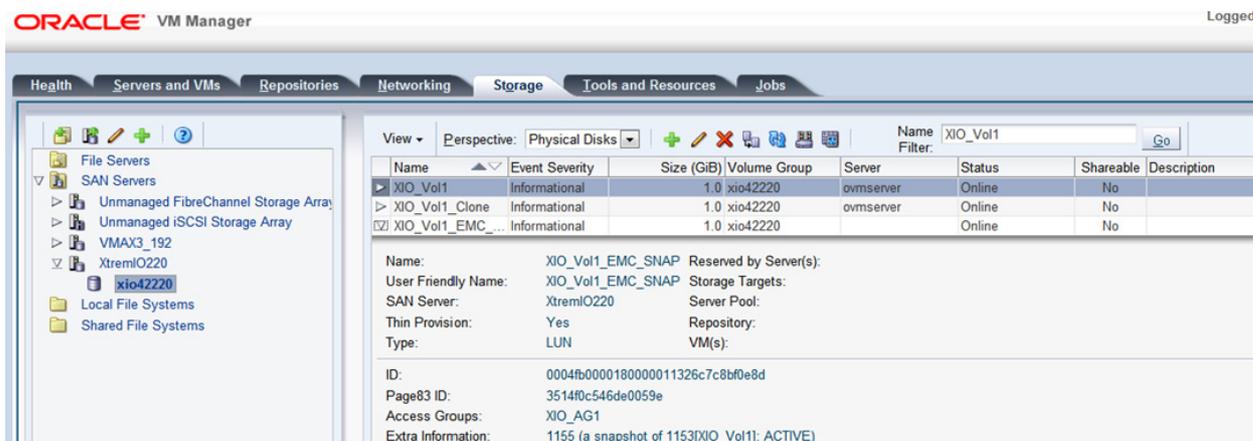


Figure 22 View snapshot details for XtremIO arrays

## Snapshot troubleshooting

The following problems can occur when you create snapshots:

- ◆ For VMAX arrays with Engenuity version 5876 or greater, snapshots will fail if the storage volume contains TimeFinder snap devices.
- ◆ For snapshot operations, the job description and job summary are displayed as **storage\_clone** instead of **storage\_snap**. This can be ignored.
- ◆ When auto meta is not enabled on the VMAX array, snapshots of a MetaLUN are not supported.
- ◆ Creating a snapshot of a MetaLUN fails for a VMAX array if **AutoMetaEnabled** is set to **FALSE** in the plug-in configuration file. If this occurs, an error message, such as the following example, is displayed:

```
Job Failed - Create Replica. error 1 : C:ERROR_CLASS_SOFTWARE
F:ERROR_FAMILY_NOT_SUPPORTED R:100002 Not Supported:
Creating composite volume targets with different size members.
```

- ◆ Cloning of a MetaLUN operation can fail if the **AutoMetaEnabled** parameter is changed in the plug-in configuration file or in the storage array after creation of the source MetaLUN. If this occurs, an error message, such as the following example, is displayed:

```
A configuration mismatch was detected among members of your
META device(s)
```

# Getting help

EMC support, product, and licensing information can be obtained as follows:

**Product information**—For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at:

<https://support.emc.com>

**Technical support**—Go to EMC Online Support and click **Service Center**. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.