Manual Configuration of InfoScale Disaster Recovery Setups for Applications

Arctera InfoScale for Windows 7.0, 7.0.1, 7.1, 7.2, 7.3, 7.3.1, 7.4, 7.4.1, 7.4.2, 8.0, 8.0.1, and 8.0.2

TechNote | 27 February 2025

Contents

About this document

This document describes the manual processes for configuring InfoScale Disaster Recovery (DR) setups for applications. It assumes that you have already configured the necessary Global Cluster Option (GCO) setup and the application service group at the primary site.

Configuring storage at the DR site

To configure storage, create dynamic disk groups and volumes using Veritas Enterprise Administrator (VEA) on any InfoScale cluster node at the DR site. The disk group and the volume name must be the same at both sites.

Creating disk groups

1. In the VEA tree view, right-click **Disk Groups** and click **New Dynamic Disk Group**.

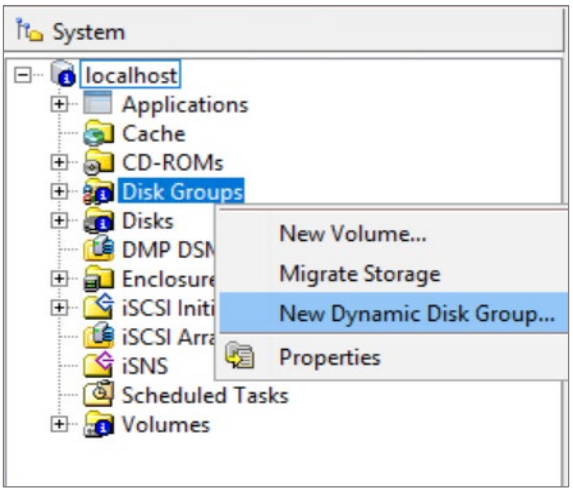

2. On the New Dynamic Disk Group panel, provide the following input:

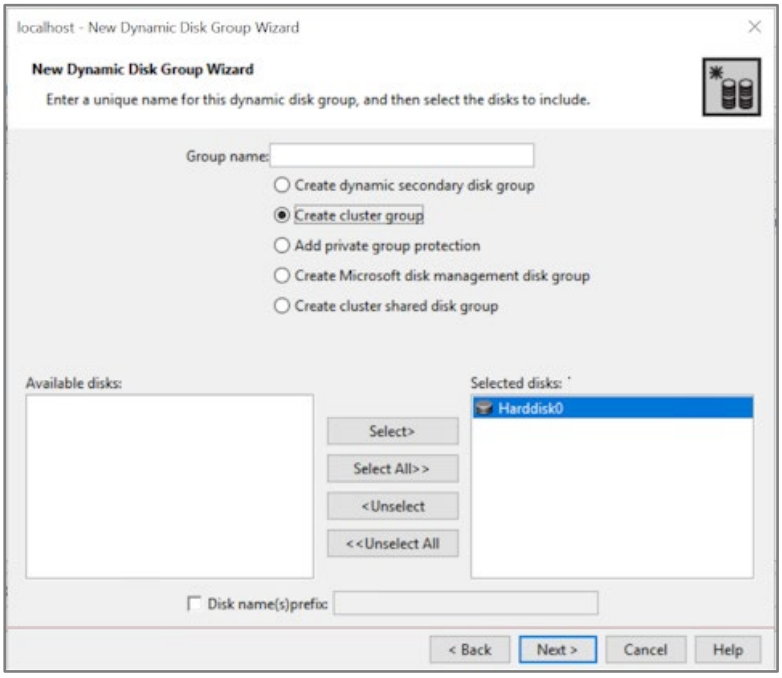

   a. In the **Group name** text box, specify a name for the disk group.
   b. Select **Create cluster group**.
   c. Select the required disks from the **Available disks** box and click **Select**.
   d. Click **Next**.

3. Follow the wizard prompts through the Finish page to complete the workflow.

Creating volumes

1. In the VEA tree view, right-click the disk group you just created and click **New Volume**.

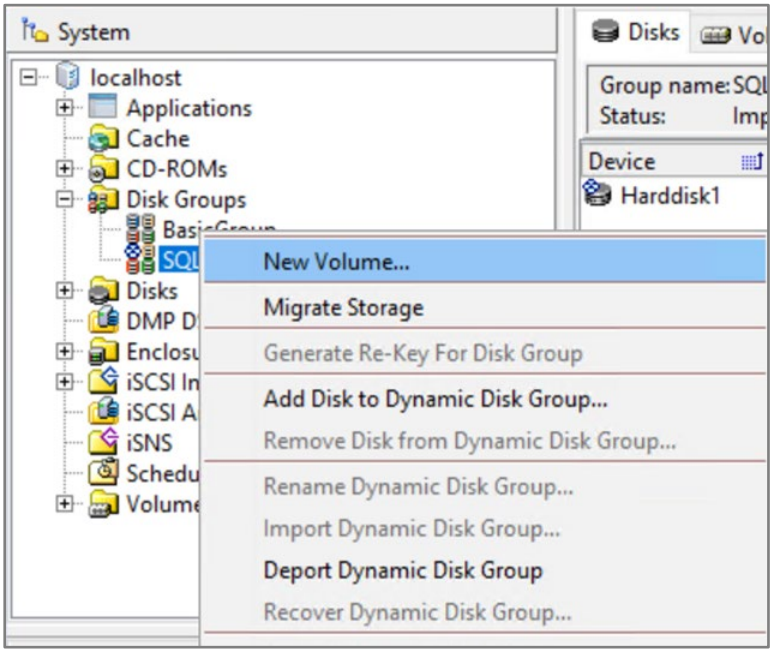

2. On the Assign Disks for Volume panel, provide the following input:

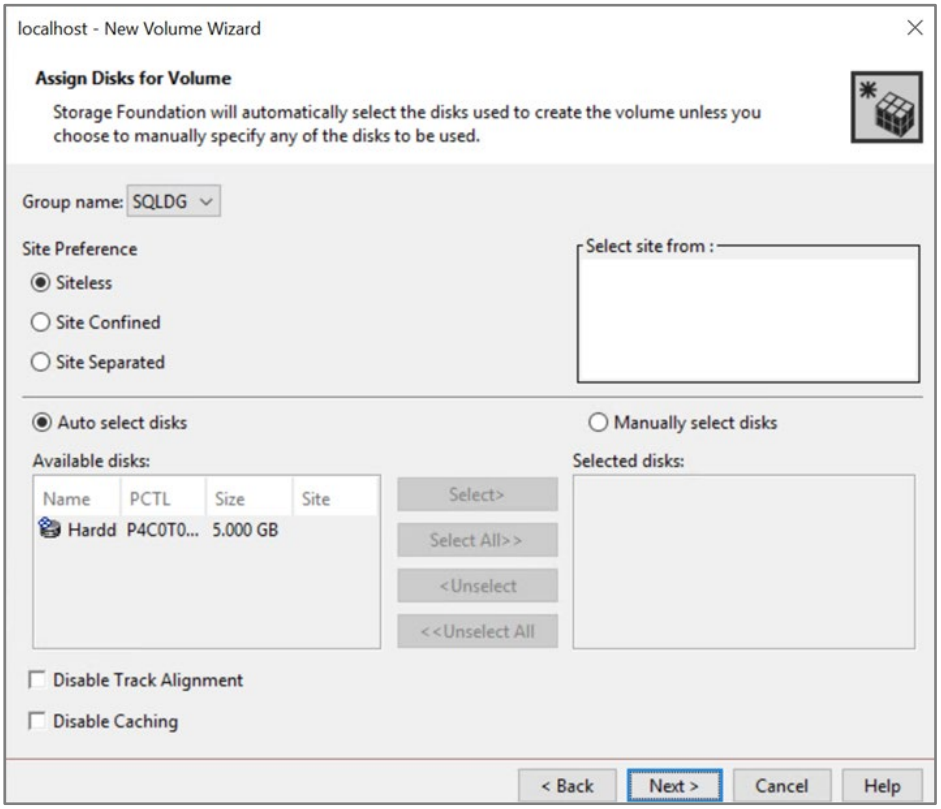

a. Select the disk group from the **Group name** drop-down list.
b. Select the appropriate **Site Preference** option.
c. Select the required disks from the **Available disks** box and click **Select**.
d. Click **Next**.

3. On the New Volume panel, provide the following input:

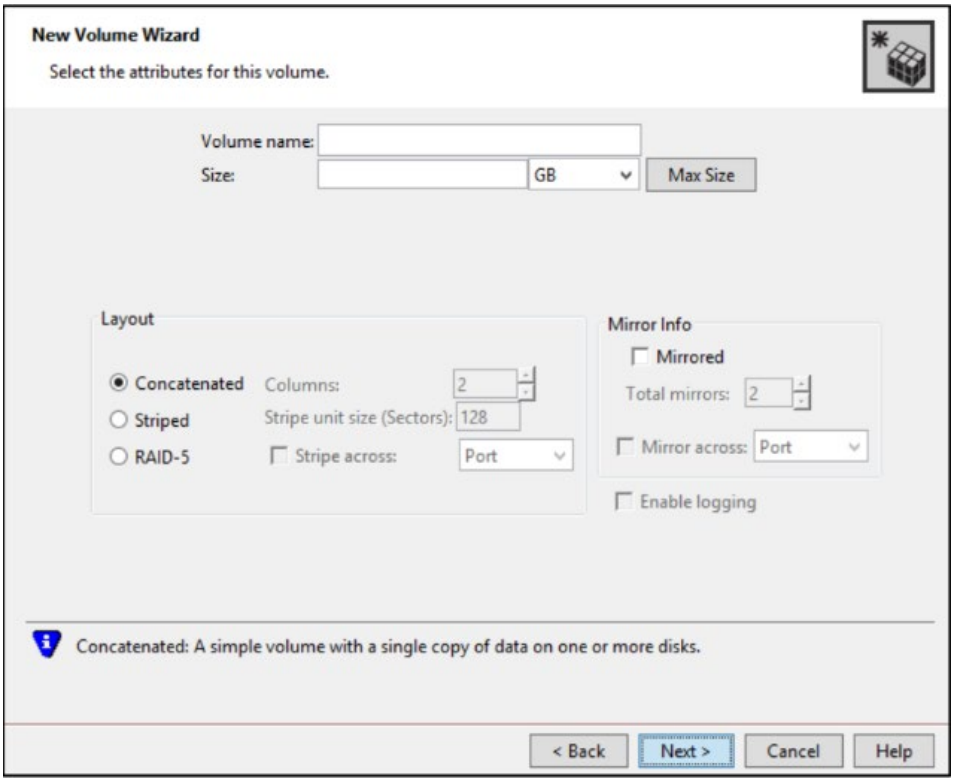

   a. In the **Volume name** text box, specify a name for the volume.
   b. Specify the size of the volume.
   c. Select the layout of the volume.
   d. In the **Mirror Info** section, select **Mirrored** and use the **Total Mirrors** field to specify the number of disks on which you want to mirror the volume data.
   e. Click **Next**.

4. On the Add Drive Letter and Path panel, select the drive letter and click **Next**.

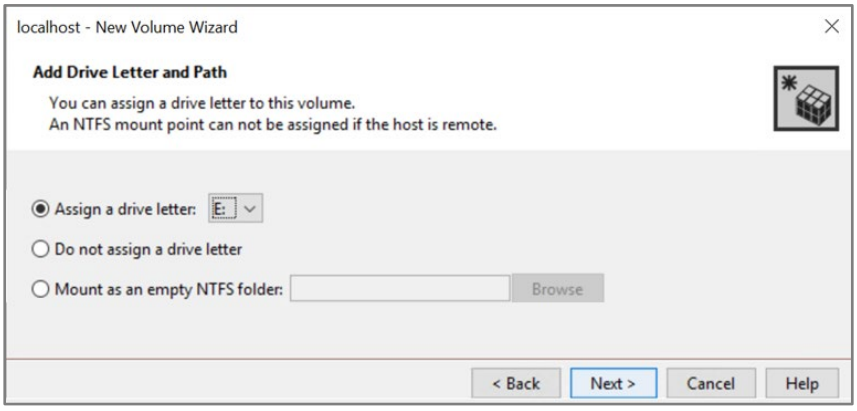

5. Follow the wizard prompts through the Finish page to complete the workflow.

Configuring replication

To configure data replication, you must first configure the VxSAS service and then configure the replicated data set (RDS) on all the cluster nodes at both sites.

For details about these configuration tasks, refer to the *Volume Replicator Administrator's Guide* applicable to your InfoScale version, for example: Volume Replicator 7.4 Administrator's Guide – Windows.

Changing the application service group and the replication service group configurations

Perform the following tasks at the primary site using the Java GUI:

1. For the Lanman resource in the application service group, set the **DNSOptions** attribute to **PurgeDuplicate**.

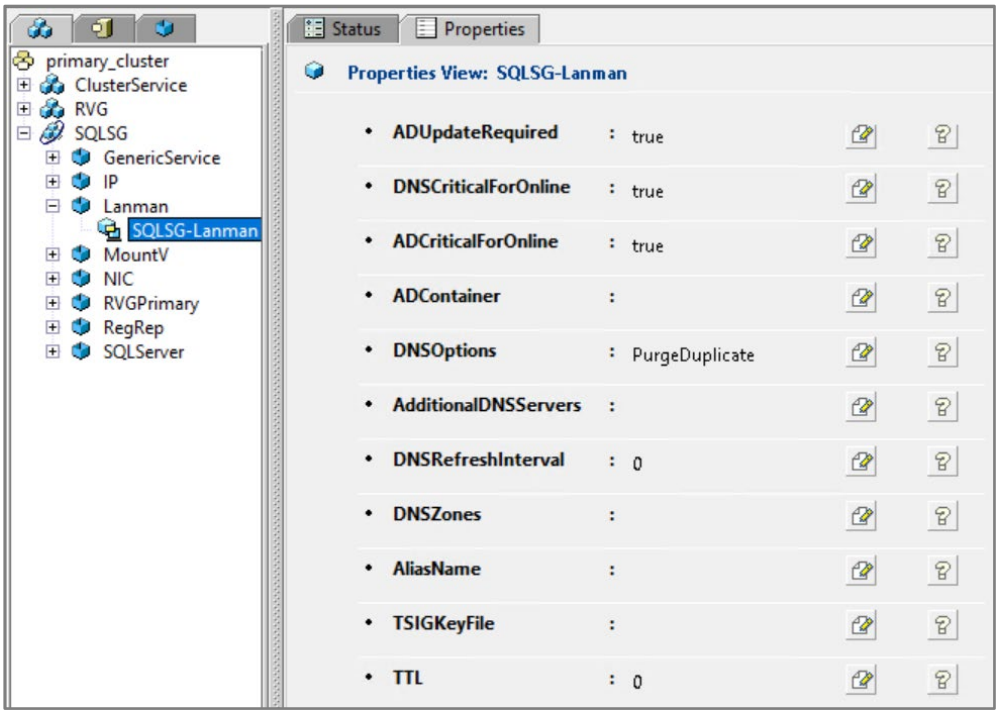

A sample Lanman resource in the `main.cf` (cluster configuration) file should look like:

```
Lanman SG-LanmanRes (
   VirtualName = LANMANPRI
   IPResName = SG-IP
   DNSUpdateRequired = 1
   ADUpdateRequired = 1
   DNSCriticalForOnline = 1
   ADCriticalForOnline = 1
   DNSOptions = { PurgeDuplicate }
   )
```

2. Create a replication service group with the same **SystemList** attribute value as that of the application service group.

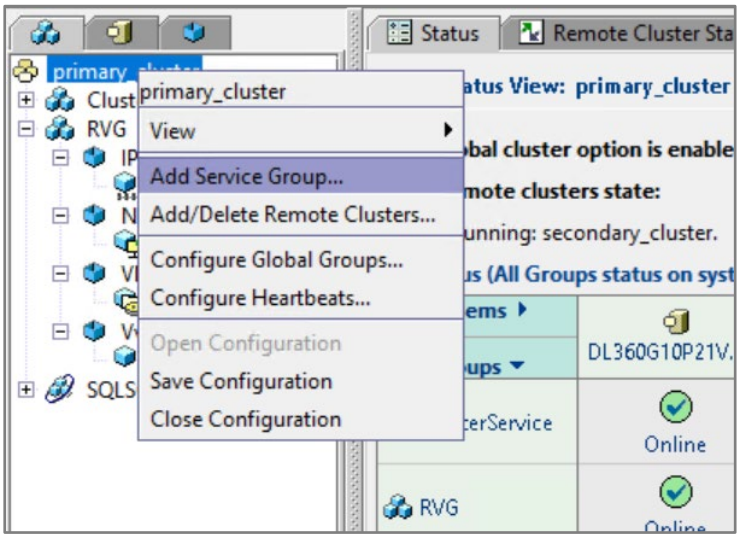

3. Add the required resources: **VMDg**, **IP**, **NIC**, and **VvrRvg**.

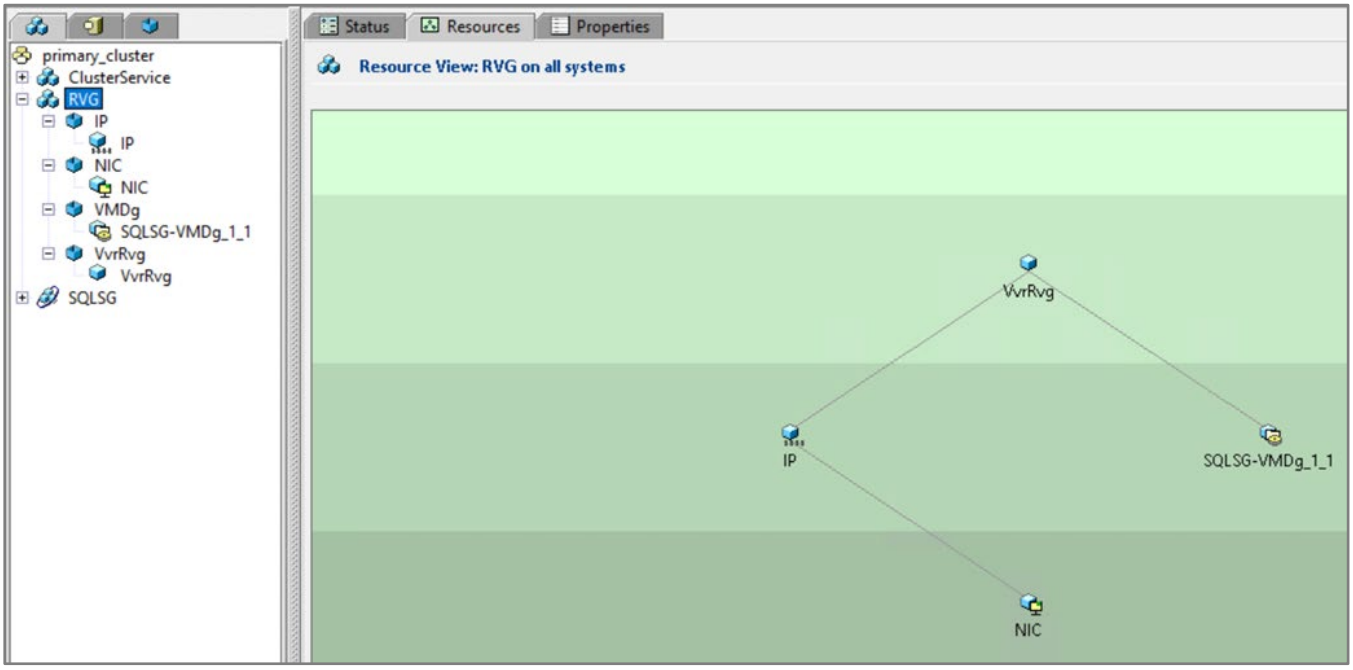

Adding resources

1. Copy the **VMDg** resource from the application service group to the replication service group and then delete it from the application service group.
2. On the replication service group, perform the following tasks:
   a. Add the **IP** resource and update the **Address**, **SubNetMask**, and **MACAddress** attributes. Add a **MACAddress** value for each system.
   b. Add the **NIC** resource and update the **MACAddress** attribute for each system.

c. Add the **VvrRvg** resource and update the following attributes:

i. **RVG**: Use the name given while configuring the RDS.
ii. **VMDgResName**: Use the resource name copied from the application service group to this service group.
iii. **IPResName**: Specify the IP resource name to be used for replication.

3. On the application service group, add the **RVGPrimary** resource and update the **RvgResourceName** attribute with the value of the **VvrRvg** resource name created in the replication service group.
4. Bring both the service groups online.

Setting up dependencies among service groups and resources

For the application to be taken offline and brought online gracefully during a failover, you must set the appropriate service group and resource dependencies.

1. Set the application service group to have an online local hard dependency on the replication service group.

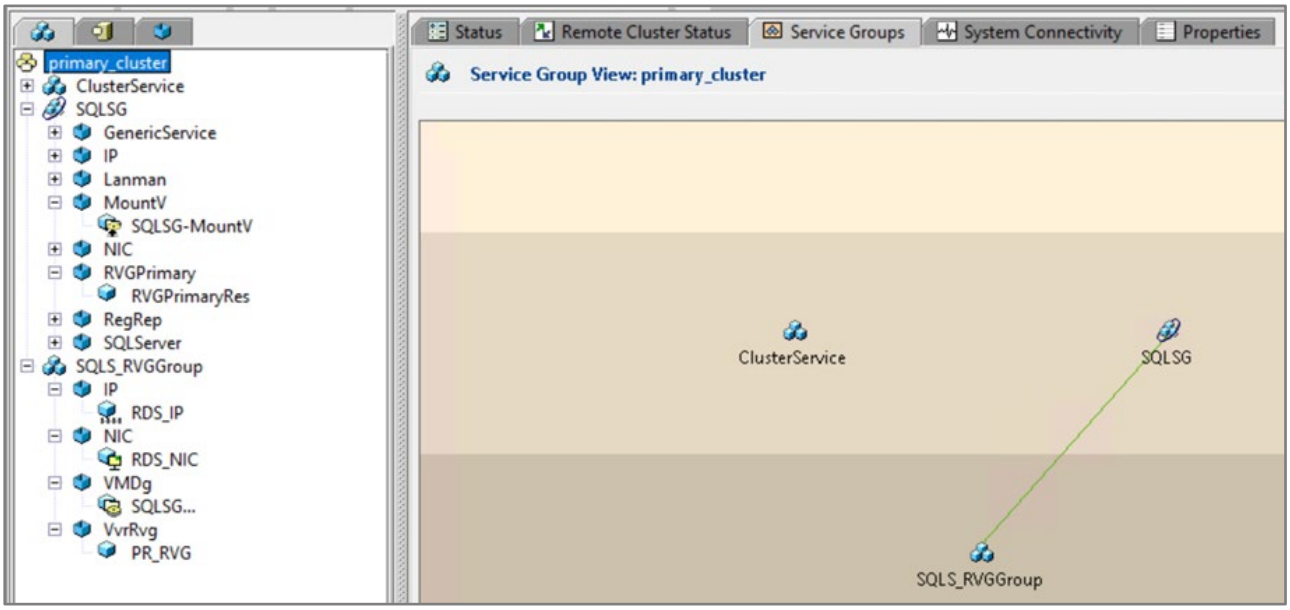

2. Set the MountV resource in the application service group to depend on the RVGPrimary resource.

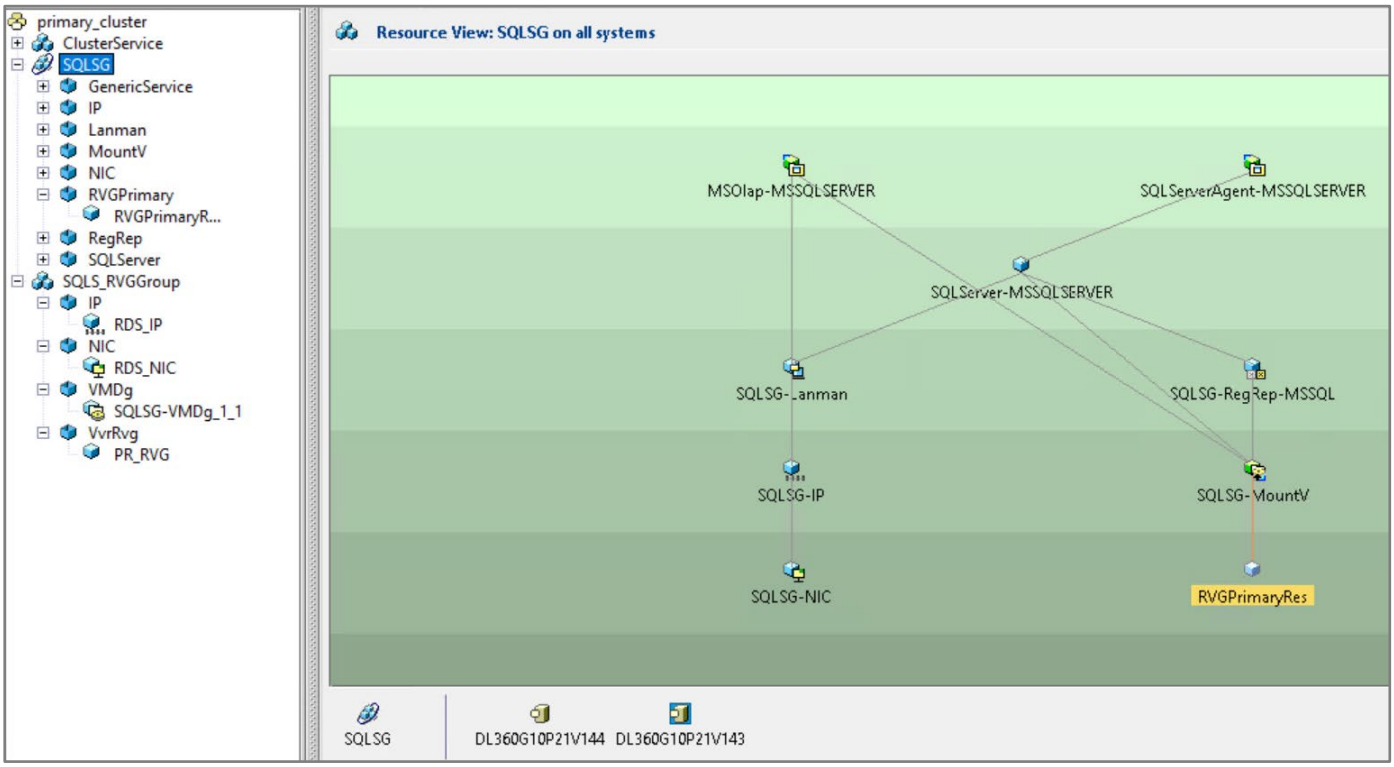

3. Set the resource dependencies in the replication service group as follows:

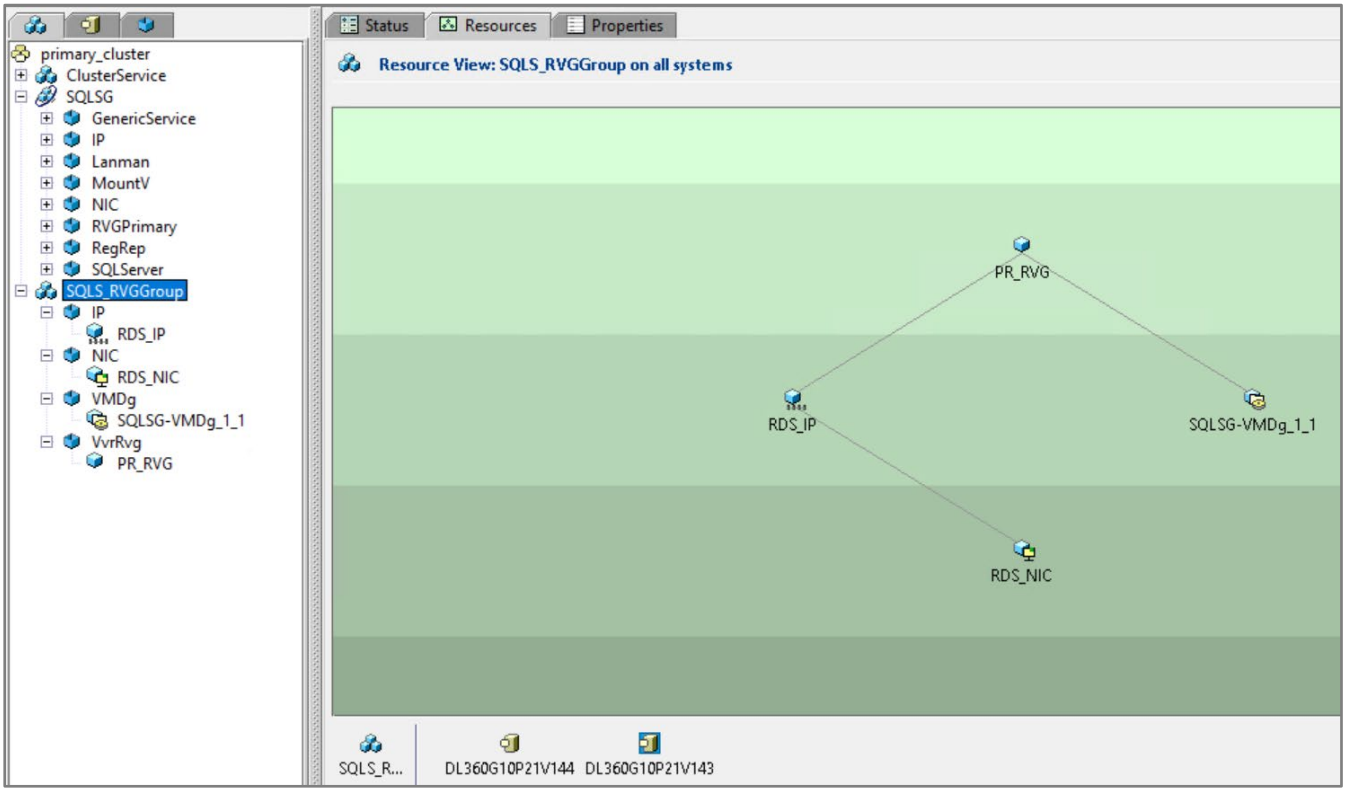

- The IP resource depends on the NIC resource.
- The VvrRvg resource depends on the IP resource.
- The VvrRvg resource depends on the VMDg resource.

4. Save the configuration for these changes to reflect in the `main.cf` file.

Configuring the corresponding service groups at the DR site

You can manually create the application service group and the replication service group at the DR site to correspond to those at the primary site. Alternatively, you can copy the service group details from the primary site cluster configuration to the DR site and then make only the necessary changes. The alternative method involves the following steps:

1. At the DR site, stop Cluster Server (VCS) by running the `hastop -all -force` command.
2. Copy the application service group details from `main.cf` at the primary site and add them to `main.cf` at the DR site.

   Make sure that you also copy over the lines related to the service group dependencies. Then, make the following changes:

   a. Update the SystemList attribute with the DR site cluster nodes.
   b. Update the IP resource attributes to include the MACAddress and the IP address details of each of the DR cluster nodes.
   c. Similarly, update the NIC resource attributes to include the MACAddress details of the each of the DR cluster nodes.

3. Copy the replication service group details from `main.cf` at the primary site and add them to `main.cf` at the DR site.

   Make sure that you also copy over the lines related to the service group dependencies. Then, make the following changes:

   a. Update the IP resource attributes to include the MACAddress and the IP address details of each of the DR cluster nodes.
   b. Similarly, update the NIC resource attributes to include the MACAddress details of the each of the DR cluster nodes.
   c. Update the DGGuid attribute of the VMDg resource with the DiskGroup ID, which you can identify by running the `vxdg list` command.

4. Run the `hacf -verify <location_of_main.cf>` command to ensure that the `main.cf` file does not have any syntax errors.
5. Start VCS by running the `hastart -all` command.
6. At both sites, configure the application service group as global.

In the Java GUI, right-click the application service group and click **Configure As Global**.

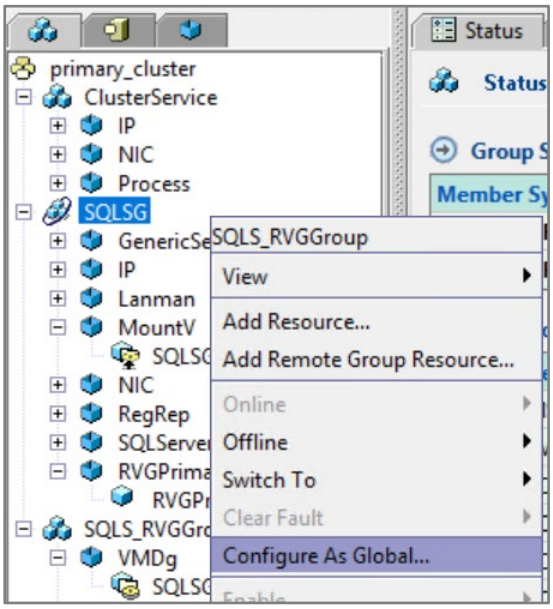

Alternatively, you can run the following commands sequentially to make the application service groups global:

a. On a primary cluster node, run:

```
haconf -makerw
hagrp -modify SQLSG ClusterList secondary_cluster 1 primary_cluster 0
hagrp -modify SQLSG ClusterFailOverPolicy Manual
haconf -dump
```

b. On a DR cluster node, run:

```
haconf -makerw
hagrp -modify SQLSG ClusterList secondary_cluster 0 primary_cluster 1
hagrp -modify SQLSG ClusterFailOverPolicy Manual
haconf -dump
```

7. From VEA, use the **File** > **User Manager** > **New User** option to add the Group Operator privilege on the service groups for the cluster nodes at both sites.

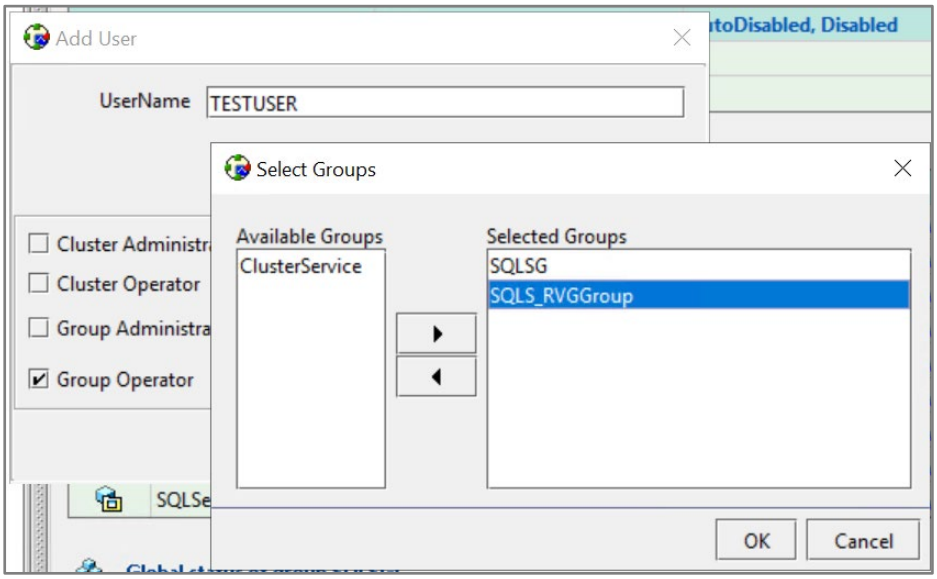

Note: The replication service group remains online at both sites and the application service group remains online only at the primary site.

Changing the replication settings

1. From VEA, launch **Replication Network**.
2. In the tree view, right-click each cluster node and click **Change Replication Settings** to update its IP address.

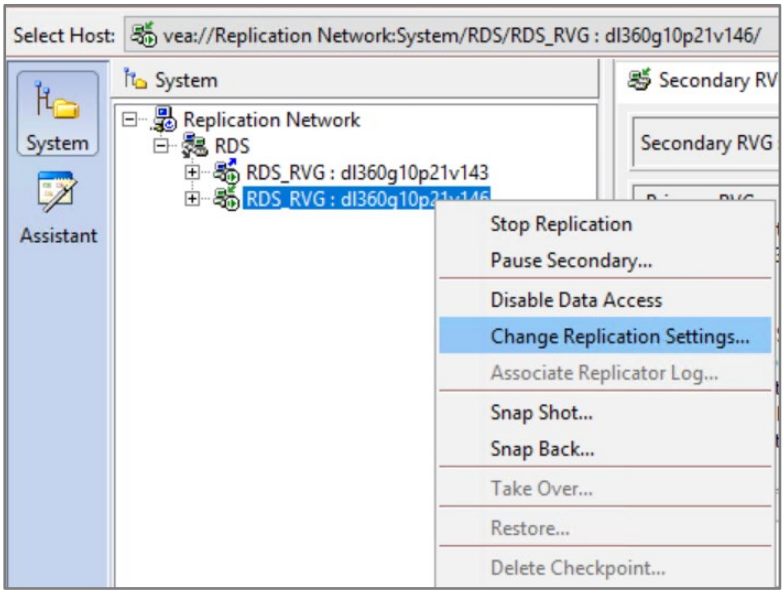

3. On the Change Replication Settings dialog box, use the **Primary side IP** and the **Secondary side IP** drop-down lists to select the appropriate IP addresses. These addresses are configured while adding the IP resource to the replication service group.

This page is intentionally left blank.