

Cluster configuration on Database Servers

Table of Contents

1. PREREQUISITES BEFORE SETTING UP CLUSTER.....	3
2. INSTALLING CLUSTER PACKAGES	3
3. CLUSTER CONFIGURATION	4
3.1 CREATE NEW CONFIGURATION	5
3.2 CUSTOM CONFIGURE MULTICAST	6
3.3 USE A QUORUM DISK	7
3.4 ADDING A MEMBER TO A CLUSTER.....	7
3.5 CONFIGURING FENCE DEVICES	9
3.6 CONFIGURE FENCE FOR THE NODE	10
3.7 ADDING CLUSTER RESOURCES.....	11
3.7.1 Adding IP Address to the Resources	11
3.7.2 Adding GFS to Resources	12
3.7.3 Adding PostgreSQL Script to Resources.....	13
3.7.4 Adding Service to Cluster	14
3.7.5 Adding Shared Resources	16
3.7.6 Propagating the updated configuration to other nodes	17
3.7.7 Starting of services.....	17
3.7.8 Stop Cluster Services	18

Table of Figures

FIGURE 1: CLUSTER CONFIGURATION	5
FIGURE 2: NAMING THE NEW CLUSTER CONFIGURATION	6
FIGURE 3: ADDING A MEMBER TO A NEW CLUSTER.....	8
FIGURE 4: FENCE DEVICE CONFIGURATION.....	9
FIGURE 5: CONFIGURING FENCE DEVICE TO NODE	10
FIGURE 6: RESOURCE CONFIGURATION	12
FIGURE 7: RESOURCE CONFIGURATION (CONTD...)	13
FIGURE 8: RESOURCE CONFIGURATION (CONTD...)	14
FIGURE 9: ADDING A NEW SERVICE TO CLUSTER	15
FIGURE 10: SERVICE MANAGEMENT	16
FIGURE 11: CLUSTER MANAGEMENT.....	18

1. Prerequisites before setting up cluster

- Cluster should have an IP address and it should be in the same network as that of the Database servers.
- There should be SAN with 5 TB capacity which should be mapped on database servers, It can be checked through **fdisk -l** command from the terminal. e.g. /dev/sdb.

Convert the provided SAN space in to LVM partition by using below commands

```
pvccreate <Enter the visible LUN space>  
vgcreate -s 5T <Volume Group Name> <Physical Volume Name>  
lvcreate -L 5T -n <Logical Volume Name> <Volume Group Name>
```

Then convert this Logical Volume partition into gfs file system by using the below command

```
gfs_mkfs -p lock_dlm -t <ClusterName>:<FileSystemName> -j 2  
<BlockDevice Path>
```

2. Installing cluster packages

Use SSDG Solution DVD1 and switch to CentOS directory, all rpm's are stored in this directory. The rpm packages required for cluster installation are listed below.

Use **rpm -ivh <package1, package2, ... ,packageN>** command.

rgmanager-2.0.52-6.el5.centos.x86_64.rpm
openais-0.80.6-16.el5.x86_64.rpm
cman-2.0.115-34.el5.x86_64.rpm
perl-Net-Telnet-3.03-5.noarch.rpm
perl-XML-LibXML-1.58-6.x86_64.rpm
perl-XML-LibXML-Common-0.13-8.2.2.x86_64.rpm
perl-XML-Namespacesupport-1.09-1.2.1.noarch.rpm
perl-XML-SAX-0.14-8.noarch.rpm
pexpect-2.3-3.el5.noarch.rpm
lvm2-cluster-2.02.56-7.el5.x86_64.rpm
gfs-utils-0.1.20-7.el5.x86_64.rpm
gfs2-utils-0.1.62-20.el5.x86_64.rpm

Cluster Configuration Manual

net-snmp-5.3.2.2-9.el5.x86_64.rpm
net-snmp-libs-5.3.2.2-9.el5.x86_64.rpm
Cluster_Administration-en-US-5.2-1.el5.centos.noarch.rpm
cluster-snmp-0.12.1-2.el5.centos.x86_64.rpm
modcluster-0.12.1-2.el5.centos.x86_64.rpm
lm_sensors-2.10.7-9.el5.x86_64.rpm
cluster-cim-0.12.1-2.el5.centos.x86_64.rpm
tog-pegasus-2.9.1-2.el5.x86_64.rpm
system-config-cluster-1.0.57-3.noarch.rpm
ipvsadm-1.24-10.x86_64.rpm

3. Cluster Configuration

To configure the **installed** cluster follow the steps given below

Step 1: Go to Top Left Menu on the Desktop of the server

Step 2: Then **System** → **Administration** → **Server Settings** → **Cluster Management**.

You will get the below screen:

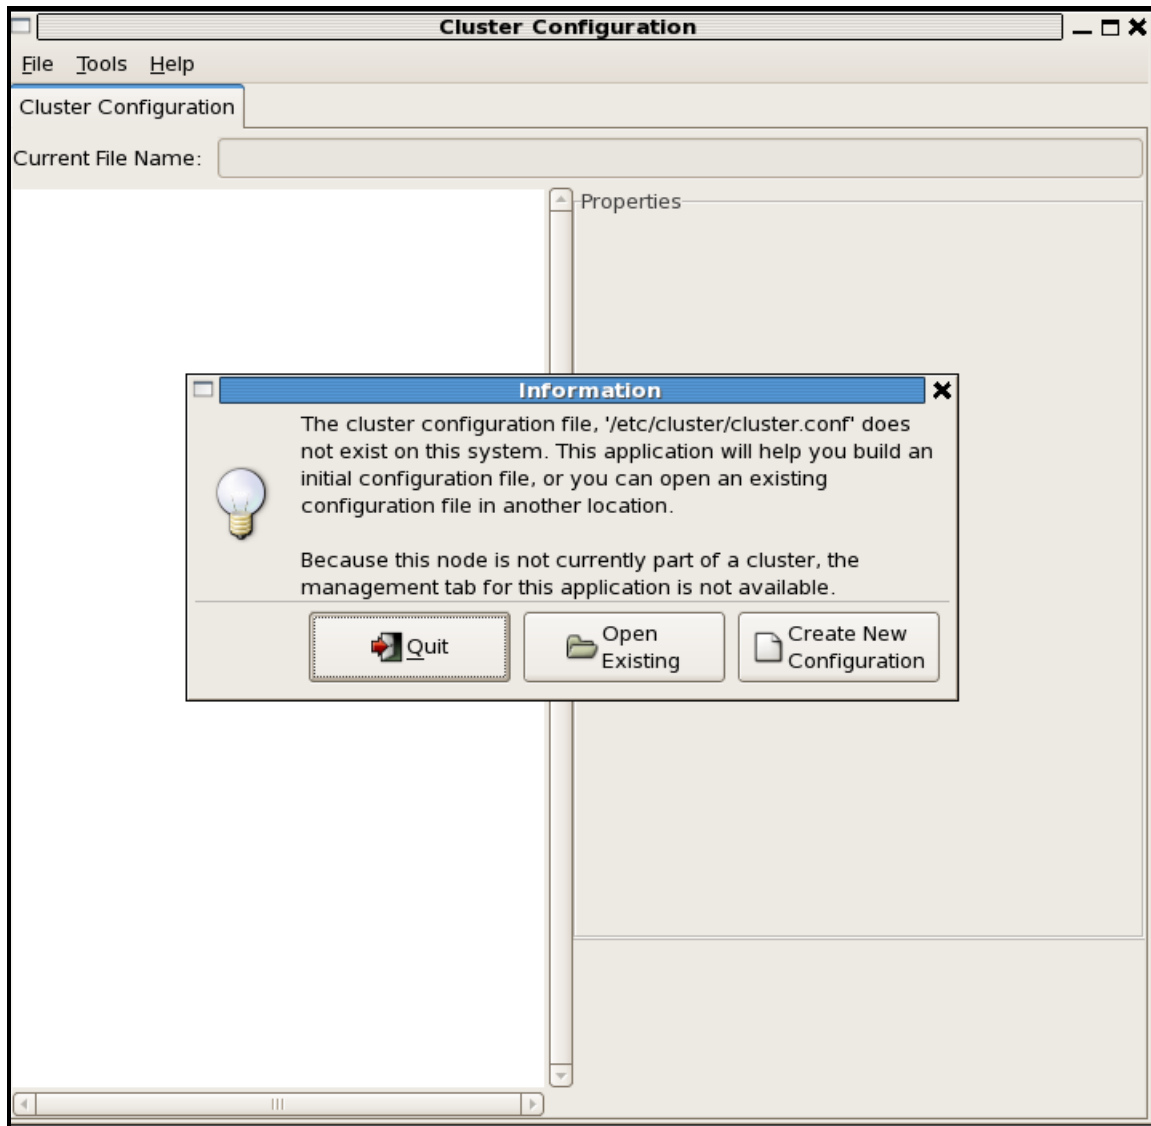


Figure 1: Cluster Configuration

Step 3: Click on the **Create New Configuration** button.

3.1 Create New Configuration

Now one needs to create a new cluster. When we start creating a new cluster, we get the screen shown in Figure below asking for cluster name. The **New Configuration** dialog box provides a text box for cluster name and the following checkboxes: **Custom Configure Multicast** and **Use a Quorum Disk**. In most circumstances you only need to configure the cluster name.

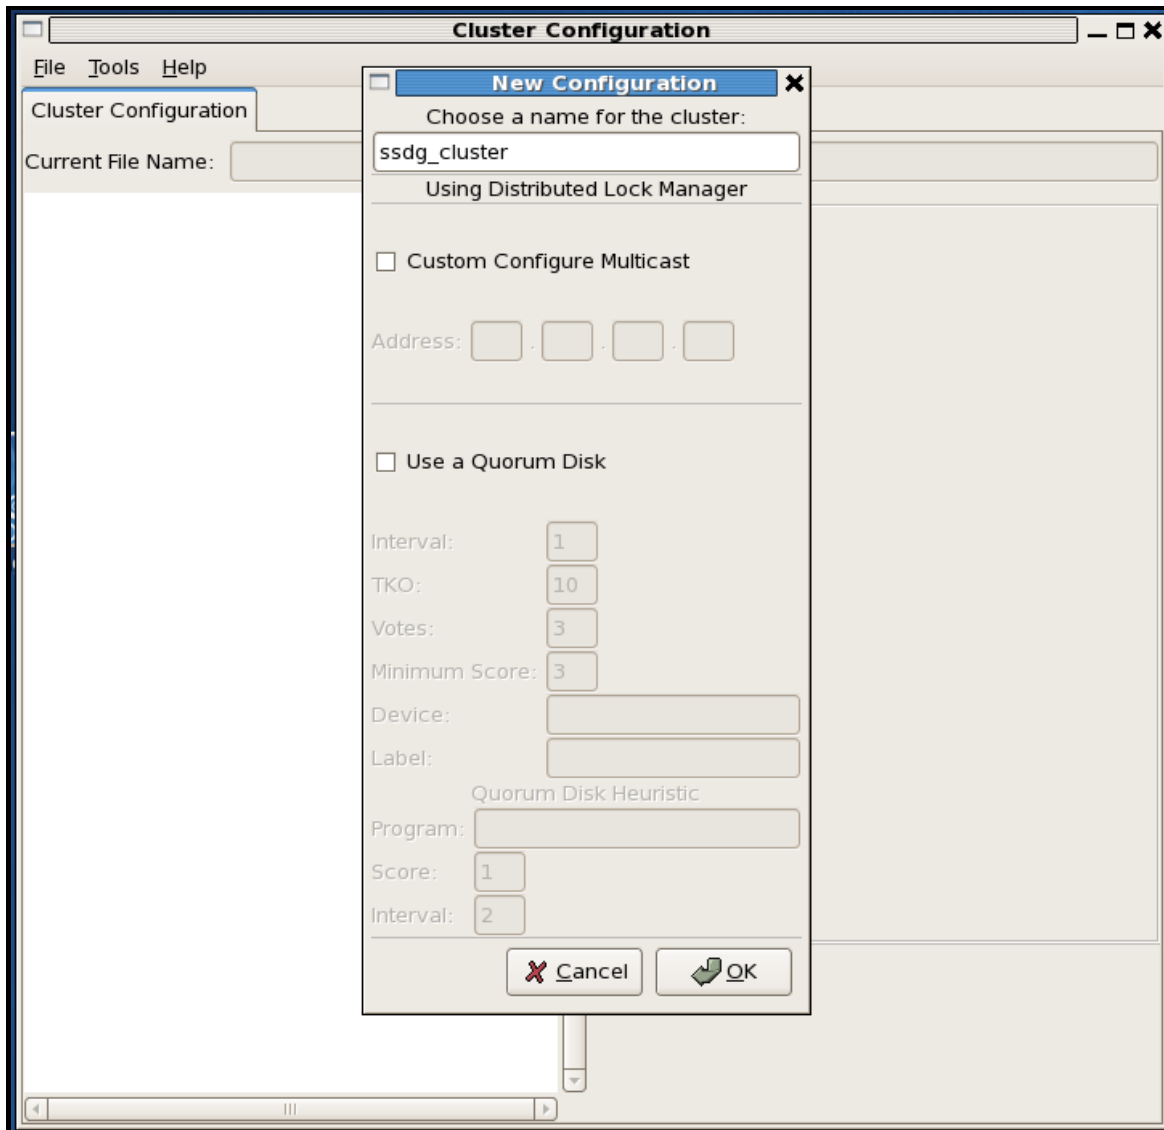


Figure 2: Naming the new cluster configuration

- Specify a cluster name such as `ssdg_cluster` for the cluster.

Note: Cluster Name should be the same as given at the time of GFS file system creation. **Also choose the cluster name carefully.** The only way to change the name of a Red Hat cluster is to create a new cluster configuration with the new name.

3.2 Custom Configure Multicast

Cluster software chooses a multicast address for cluster management communication among cluster nodes. If you need to use a specific multicast address, click the **Custom Configure Multicast** checkbox and enter a multicast address in the **Address** text boxes. If you do not

Cluster Configuration Manual

specify a multicast address, the Red Hat Cluster software (specifically, **cman**, the Cluster Manager) creates one. It forms the upper 16 bits of the multicast address with 239.192 and forms the lower 16 bits based on the cluster ID.

Note: If you do specify a multicast address, you should use the 239.192.x.x series that **cman** uses. Otherwise, using a multicast address outside that range may cause unpredictable results.

3.3 Use a Quorum Disk

We will leave the **Use a Quorum Disk** checkbox unchecked, as not required in our scenario.

If you still need to use a quorum disk, click the **Use a Quorum disk** checkbox and enter quorum disk parameters. The following quorum-disk parameters are available in the dialog box if you enable **Use a Quorum disk**: **Interval**, **TKO**, **Votes**, **Minimum Score**, **Device**, **Label**, and **Quorum Disk Heuristic**.

When you have completed entering the cluster name and other parameters in the **New Configuration** dialog box, click **OK**. Clicking **OK** starts the **Cluster Configuration Tool**, displaying a graphical representation of the configuration.

3.4 Adding a Member to a Cluster

To add a member to a new cluster, follow these steps:

1. Click **Cluster Node**.
2. At the bottom of the right frame (labeled **Properties**), click the **Add a Cluster Node** button. Clicking that button causes a **Node Properties** dialog box to be displayed. The **Node Properties** dialog box presents text boxes for **Cluster Node Name** and **Quorum Votes** (refer to *Figure 3, "Adding a Member to a New Cluster"*).
3. At the **Cluster Node Name** text box, specify a node name. The entry can be a name or an IP address of the node on the cluster subnet.
4. Specify a node name as **SSDGDB1 in Cluster Node Name** Text Box.
5. Optionally, at the **Quorum Votes** text box, you can specify a value; however in most configurations you can leave it blank. Leaving the **Quorum Votes** text box blank causes the quorum votes value for that node to be set to the default value of **1**.
6. Click **OK**.

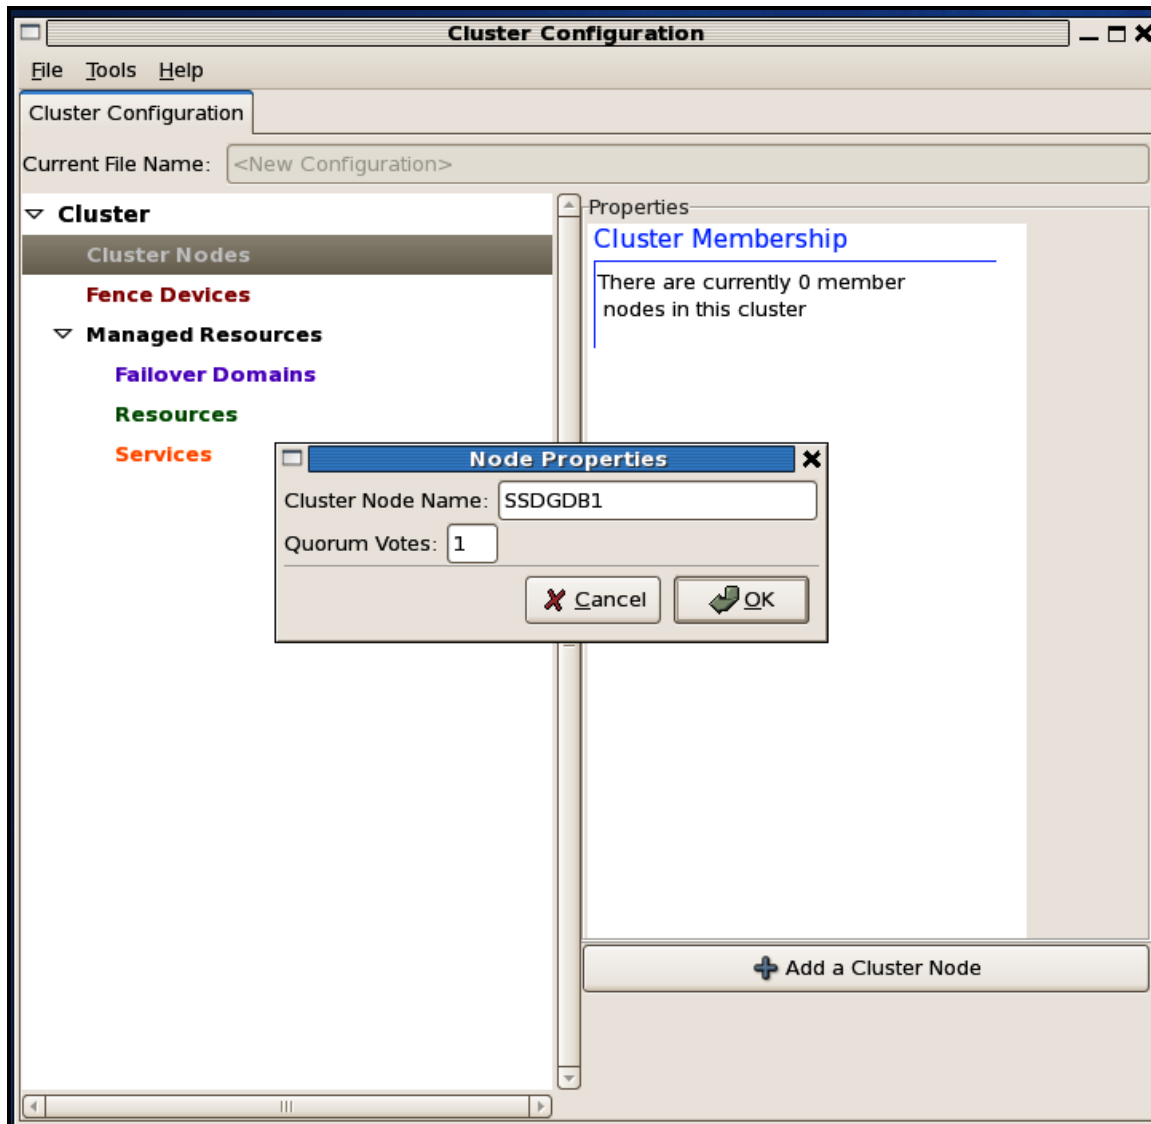


Figure 3: Adding a Member to a New Cluster

Note: Repeat the above steps to add another node with node name SSDGDB2 .

- Save cluster configuration changes by selecting File => Save.

Note: Other node must be on the same subnet as the node from which you are running the Cluster Configuration Tool and entry of the same should be in the /etc/hosts file of each cluster node.

3.5 Configuring Fence devices

Configuring fence devices for the cluster consists of selecting one or more fence devices and specifying fence-device-dependent parameters (for example, name, IP address, login, and password).

To configure fence devices, follow these steps:

1. Click **Fence Devices**. At the bottom of the right frame (labeled **Properties**), click the **Add a Fence Device** button. Clicking **Add a Fence Device** causes the **Fence Device Configuration** dialog box to be displayed (refer to Figure 4, “*Fence Device Configuration*”).

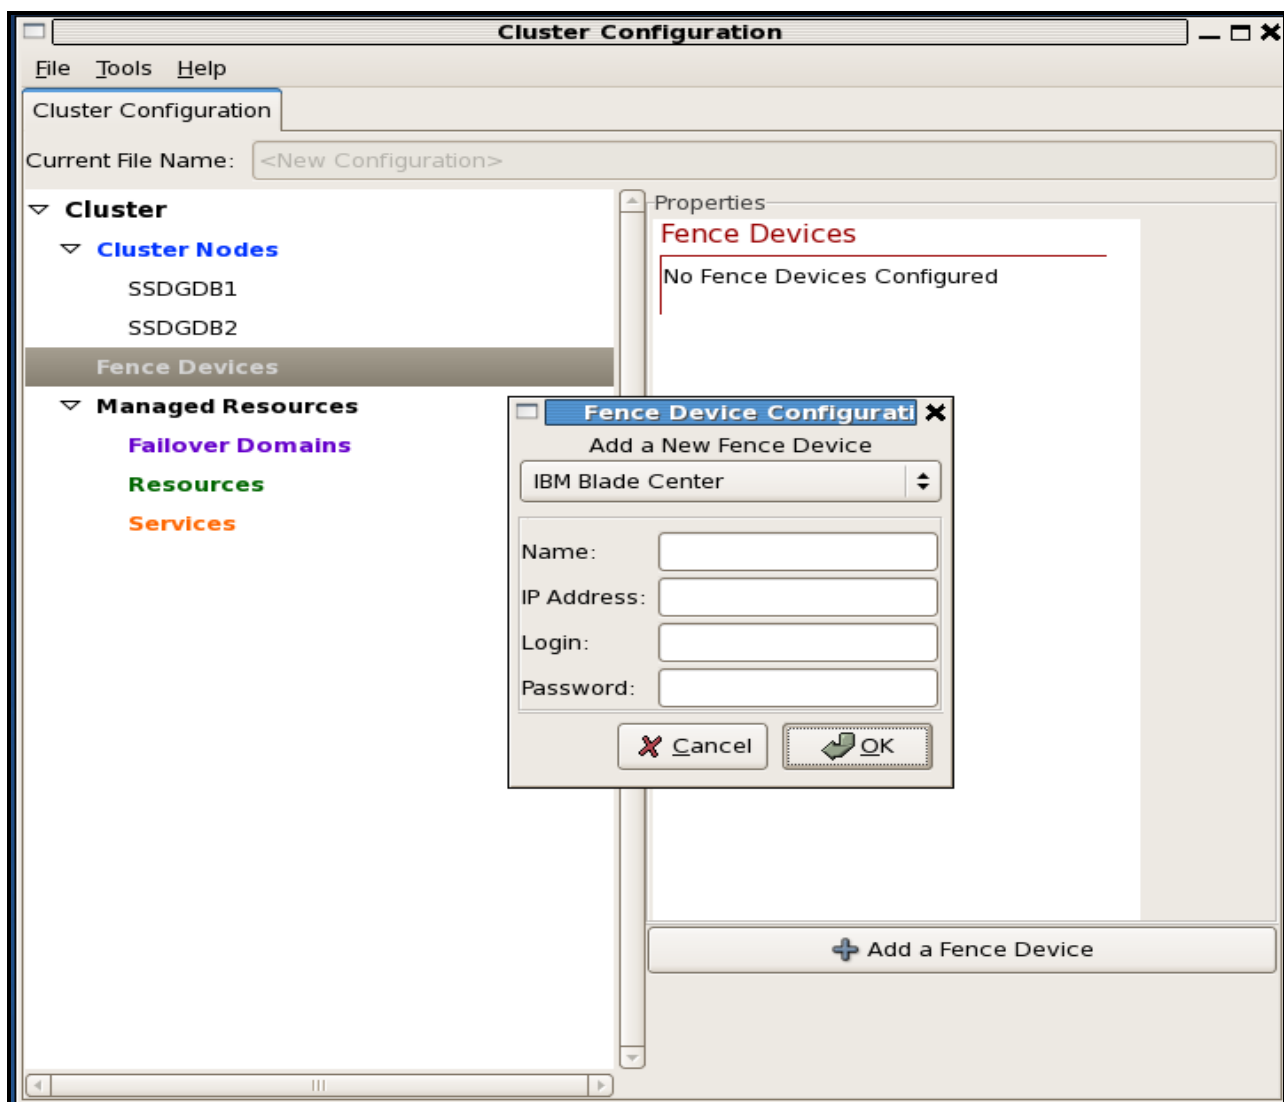


Figure 4: Fence Device Configuration

Cluster Configuration Manual

2. At the **Fence Device Configuration** dialog box, click the drop-down box under **Add a New Fence Device** and select the type of fence device to configure.
3. Specify the information in the **Fence Device Configuration** dialog box according to the type of fence device.
4. Click **OK**.
5. Choose **File => Save** to save the changes to the cluster configuration.

3.6 Configure Fence for the node

We need to add the fence device to the both nodes that we have created earlier (SSDGDB1 and SSDGDB2). Adding Manual Fence to the Node1 (SSDGDB1) is illustrated below in Figure 5, You may select your respective fence device.

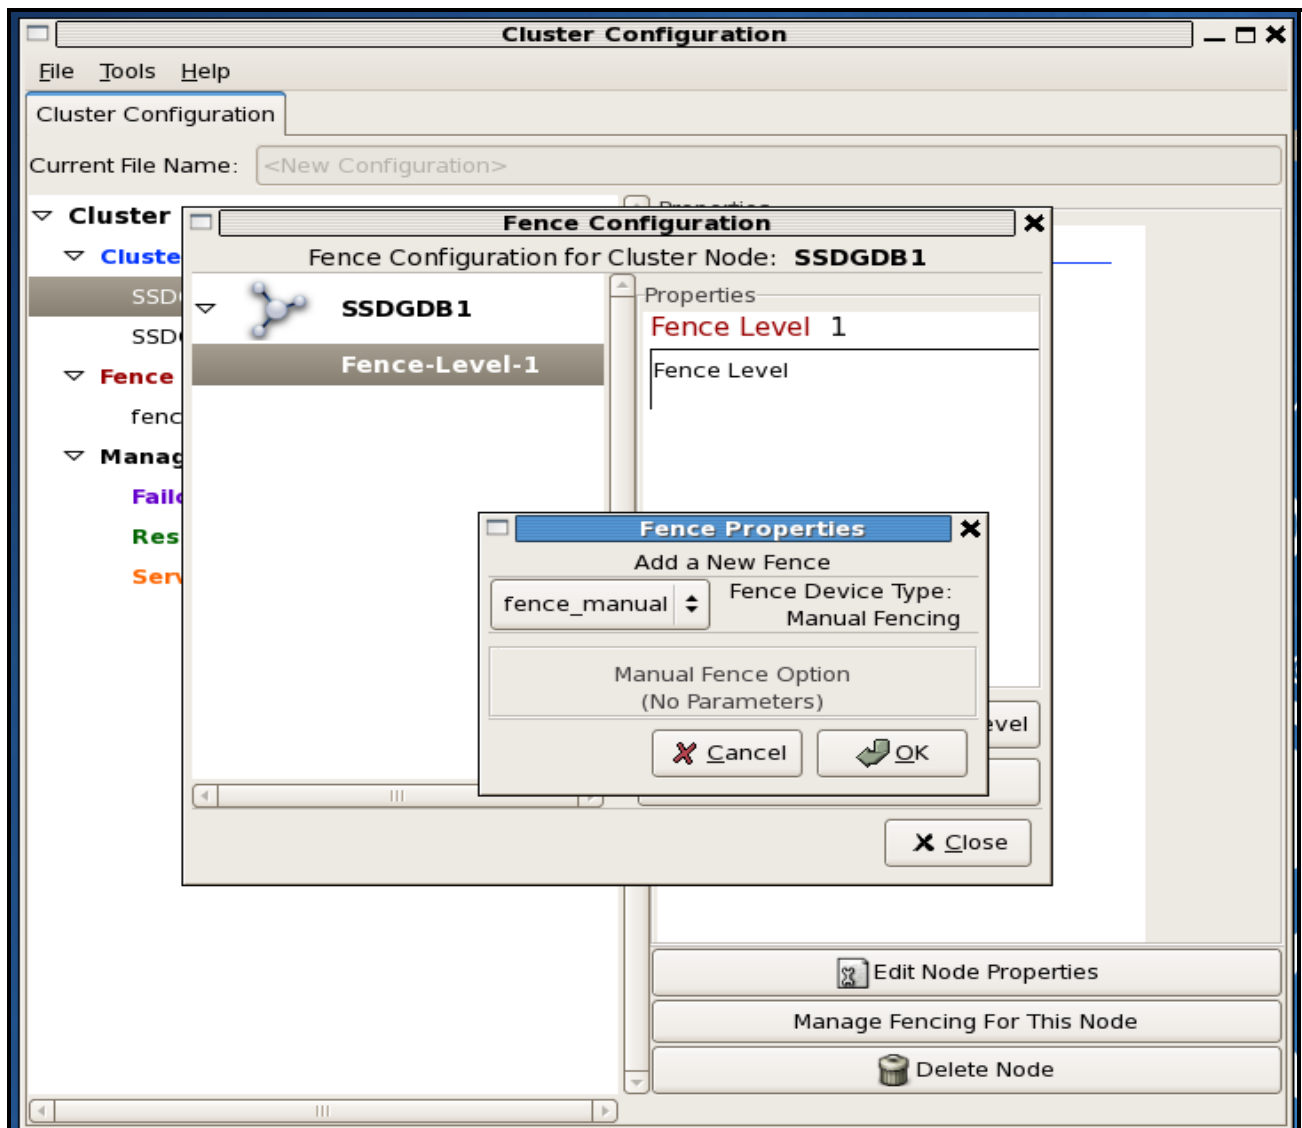


Figure 5: Configuring Fence Device to Node

Cluster Configuration Manual

Follow the below mentioned steps to configure Fence Device to the created Node.

1. Click the node that you added in the previous step i.e. SSDGDB1.
2. Now click on **Manage Fencing For This Node** button.
3. At the bottom of the right frame (below **Properties**), click **Manage Fencing For This Node**.
4. Clicking **Manage Fencing For This Node** causes the **Fence Configuration** dialog box to be displayed.
5. Click on **Add a New Fence Level** button. This causes a fence-level element to be displayed.
6. Click on **Fence-Level element**.
7. Click on **Add a New Fence** button to this Level.
8. Clicking **Add a New Fence** to this Level causes **Fence Properties** dialog box to be displayed.
9. At **Fence Properties** dialog box, click **Fence Device Type** drop-down box and select **fence device** for this node. Also, provide additional information required. An example for adding Manual Fence i.e. fence_manual to the Node SSDGDB1 is provided in Figure 5, You may select your respective fencing device.
10. At **Fence Properties** dialog box, click **OK**. Clicking **OK** causes a fence device element to be displayed below the fence-level element.
11. If you have configured all the fence levels and fence devices for this node, click **Close**.

*** Note: Repeat the above steps for node SSDGDB2.**

Choose File => Save to save the changes to the cluster configuration.

3.7 Adding Cluster Resources

Now we need to add Resources to our configured nodes SSDGDB1 and SSDGDB2.

On the **Resources** property of the **Cluster Configuration Tool**, click the **Create a Resource** button. Clicking the **Create a Resource** button causes the **Resource Configuration** dialog box to be displayed.

3.7.1 Adding IP Address to the Resources

- a) Click on the Resources. At the bottom of the right frame you will get **Create a Resource** button ((refer to Figure 6, “Resource Configuration”).
- b) Click Create a Resource.
- c) Choose IP Address from the drop-down menu.
- d) Enter the IP Address to be associated with the Postgresql Server service.
- e) Make sure that the **Monitor Link** checkbox is left checked.
- f) Click OK.

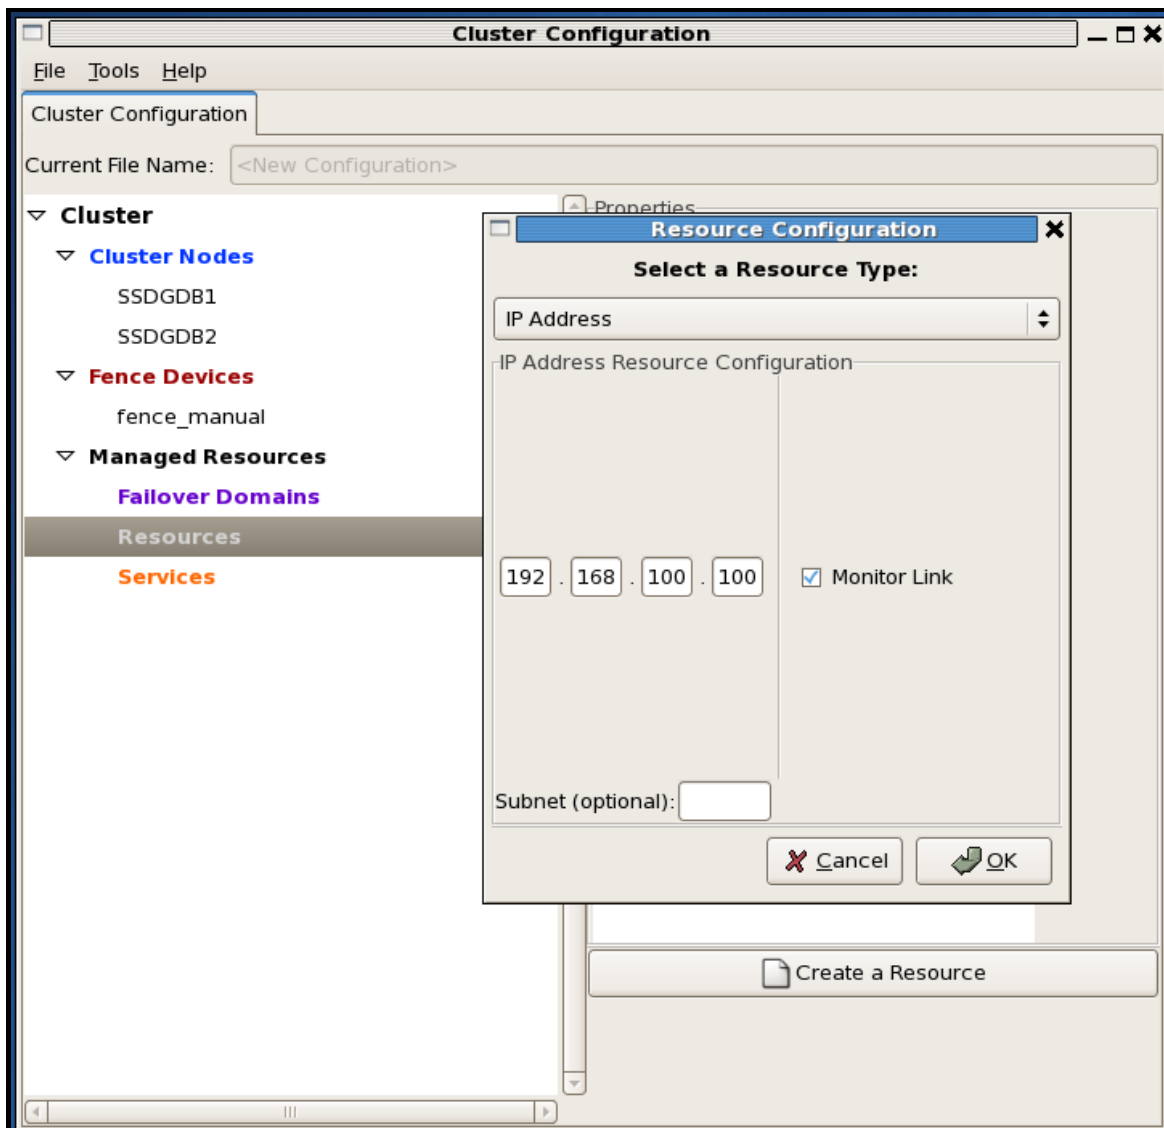


Figure 6: Resource Configuration

3.7.2 Adding GFS to Resources

- a) Click on **Resources**.
- b) Click **Create a Resource** button.
- c) Choose GFS from the drop-down menu.
- d) Enter Name of GFS File System provided at the time of creation.
- e) Enter Mount Point to be associated with the PostgreSQL Server service.
- f) Enter Device Name where SAN storage has been mounted (should be the same as given at the time of GFS file system creation).
- g) Choose file system type as **GFS** from drop down menu.
- h) **Click OK.**

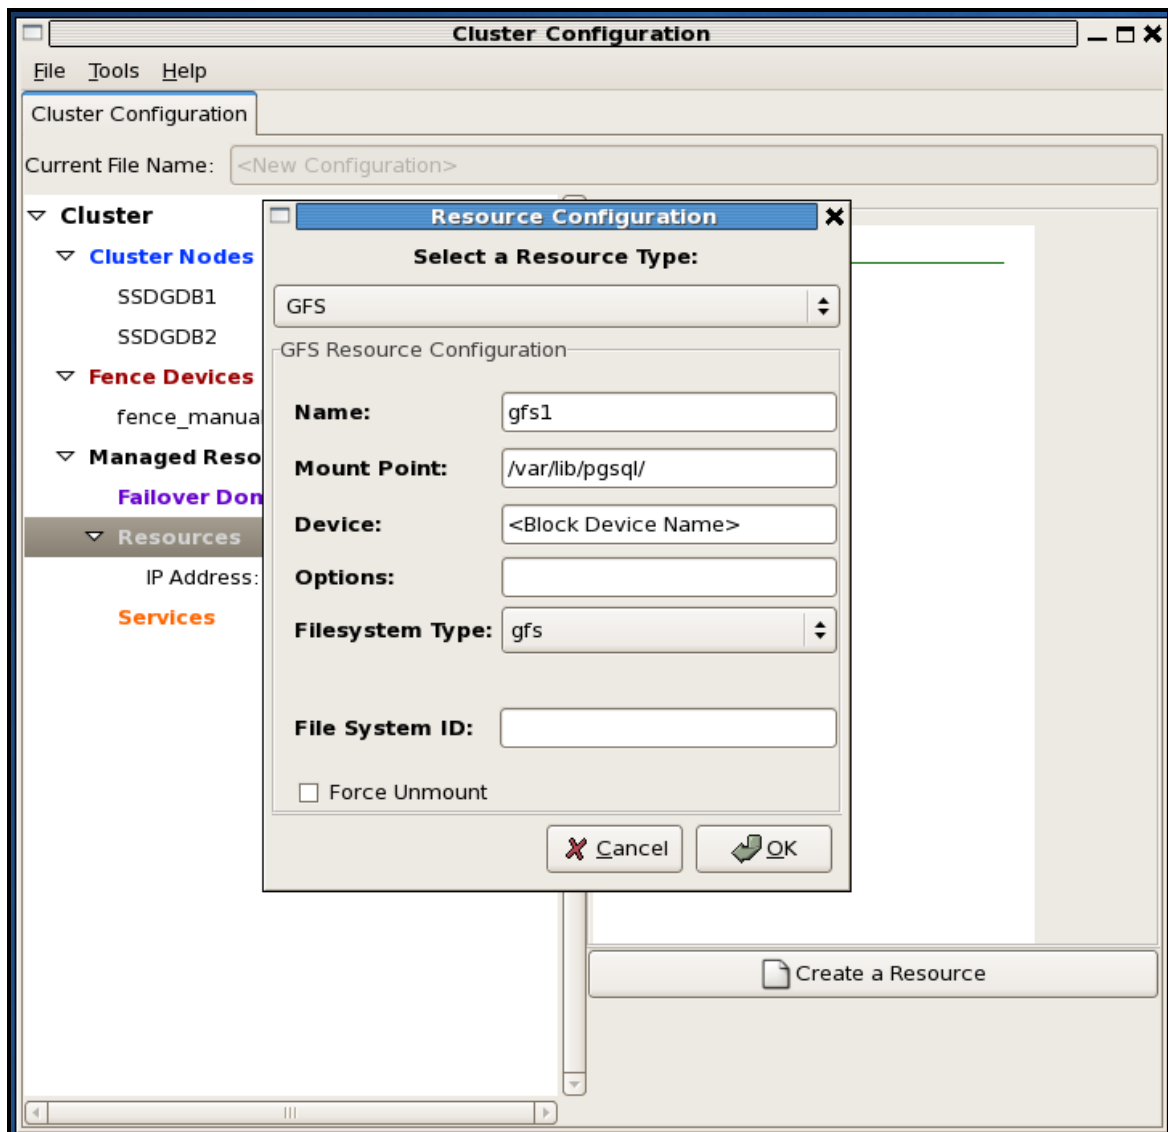


Figure 7: Resource Configuration (Contd...)

3.7.3 Adding PostgreSQL Script to Resources

- Select **Resources** tab
- Click **Create a Resource**. The Resources Configuration properties dialog box is displayed.
- Select **Script** form the drop down menu.
- Enter a Name to be associated with the PostgreSQL Server service.
- Specify the path to the PostgreSQL Server init script (i.e. **/etc/init.d/postgresql**).
- Click OK.

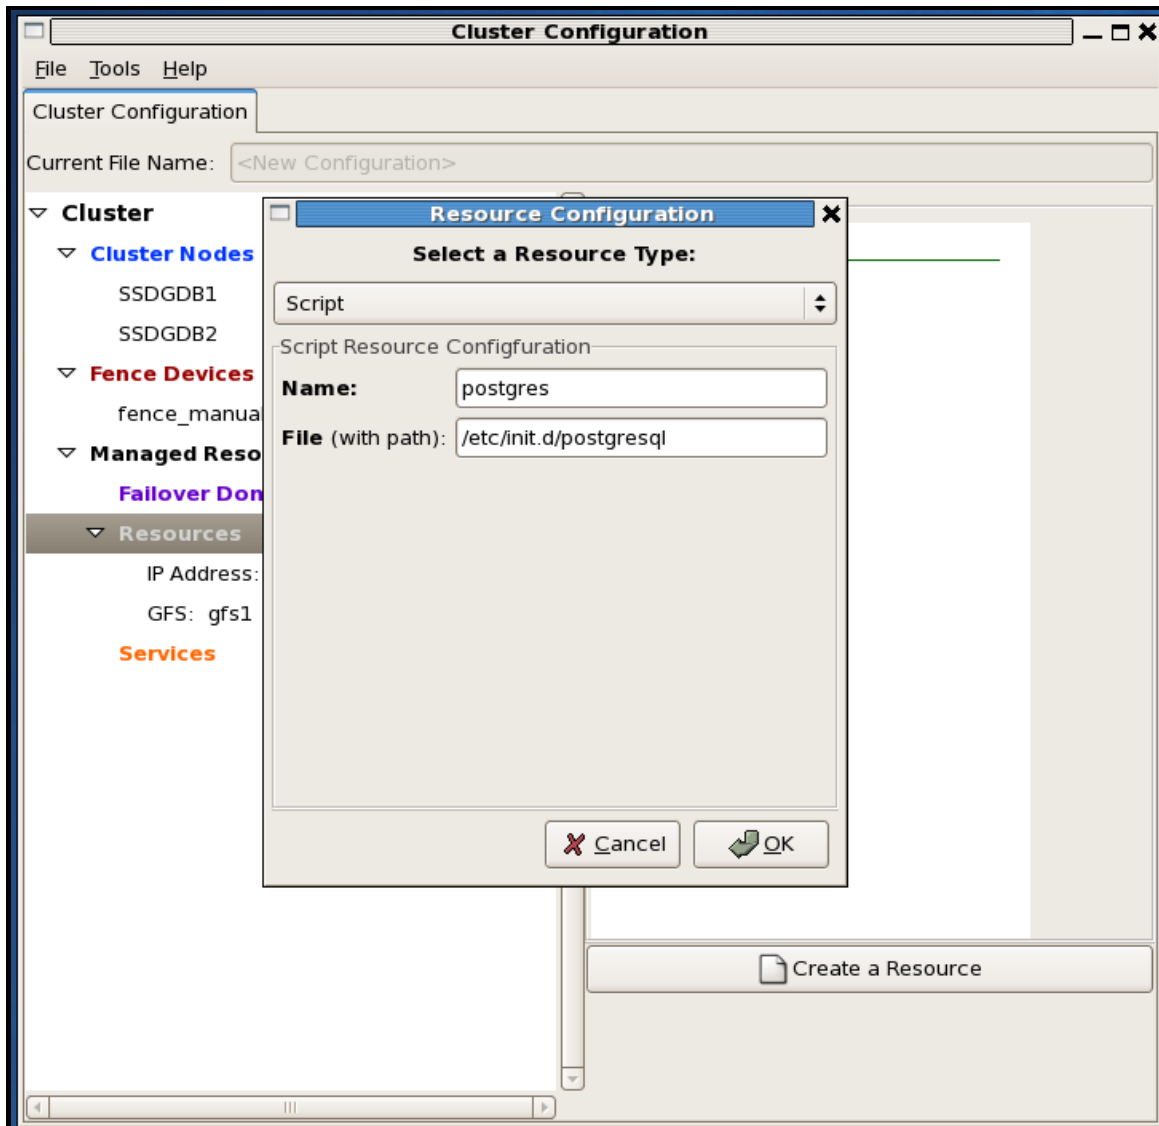


Figure 8: Resource Configuration (Contd...)

- g) Choose **File** -> **Save** to save your changes

3.7.4 Adding Service to Cluster

To add a cluster service to the cluster, follow the below mentioned steps:

- Click **Services**.
- Click **Create a Service** button. This causes the **Add a Service** dialog box to be displayed.
- At the Add a Service dialog box, type the name of the service in the Name text box as *ssdg_cluster_service*.

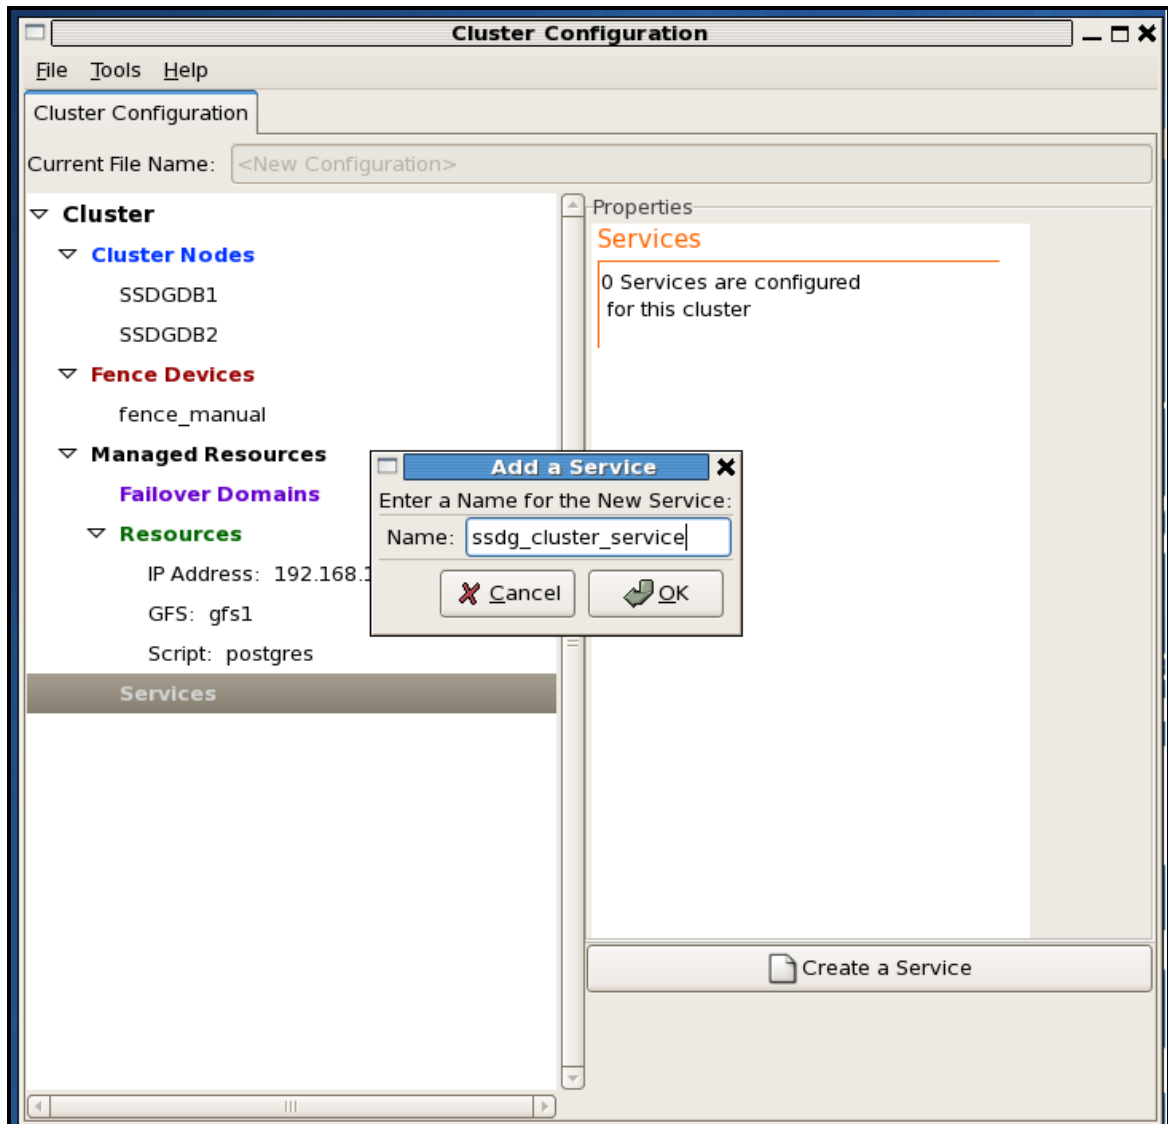


Figure 9: Adding a New Service to Cluster

d) Click **OK**.

Note: Use descriptive name that clearly distinguishes the service from other services in the cluster. Clicking **OK** causes **Service Management** dialog box to be displayed.

- e) Make sure that **Autostart This Service** checkbox is checked. This is checked by default. If **Autostart This Service** is checked, the service is started automatically when the clusters is started and running, (refer to Figure 10. Service Management).
- f) Make sure that **Run Exclusive** checkbox is checked. This sets a policy wherein the service only runs on nodes that have no other services running on them (refer to Figure 10. Service Management).

Cluster Configuration Manual

- g) At the upper right of the **Service Management** dialog box, there are three **Recovery Policy** options available:
- Check the **Relocate** option, as it makes sure that it does not restart the node where the service is currently located. It relocates the service before restarting (refer to Figure 10. Service Management).

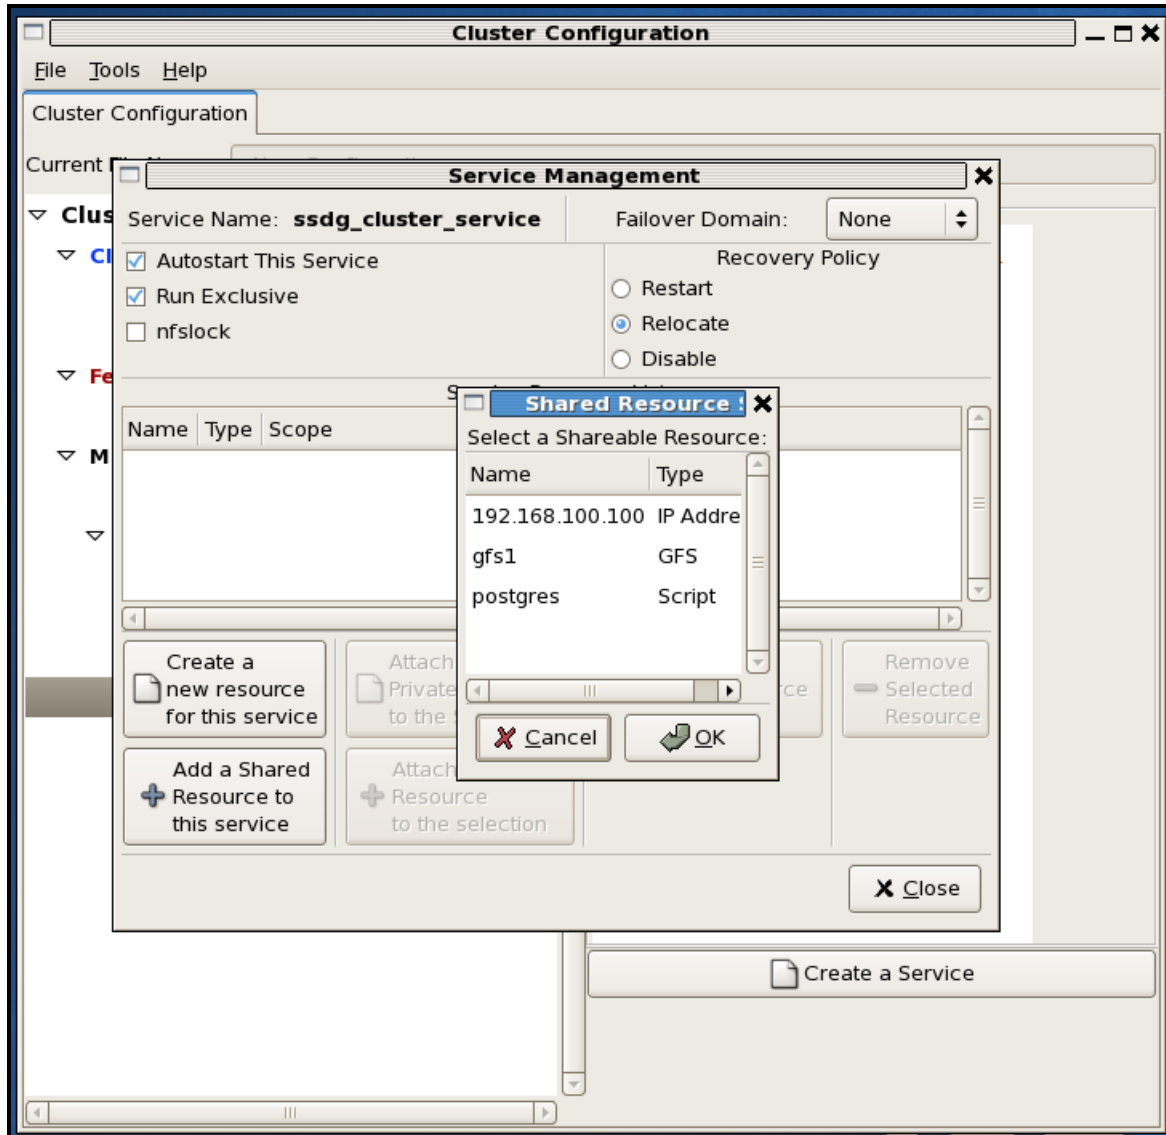


Figure 10: Service Management

3.7.5 Adding Shared Resources

To add Shared resources to the service `ssdg_cluster_service` that you have created, follow these steps:

Cluster Configuration Manual

- a) Click on the service `ssdg_cluster_service`. Clicking on the service displays the service management dialog (refer to Figure 10, Service Management).
- b) Click the **Add a Shared Resource to this service** button and choose all three shared resources created earlier i.e. IP Address, GFS file system and PostgreSQL script.
- c) When finished, click OK.
- d) Choose File => Save to save the changes to the cluster configuration.

3.7.6 Propagating the updated configuration to other nodes

- 1) Ping other nodes by hostname/node name given in the cluster configuration file.
- 2) Use the `scp` command to send the updated `/etc/cluster/cluster.conf` file from one of the existing cluster nodes to the new node.
e.g. `scp /etc/cluster/cluster.conf SSDGDB2:/etc/cluster/`

3.7.7 Starting of services

Start cluster software on all cluster nodes (including the added one) by running the following commands in this order:

Note: Start CMAN service on both nodes after interval of 20 seconds.

- a) `service cman start`
 - b) `service clvmd start`, if CLVM has been used to create clustered volumes
 - c) `service gfs start`
 - d) `service rgmanager start`
- 1) Click Cluster management GUI.
 - 2) At the Cluster Configuration Tool tab, verify that the configuration is correct. At the Cluster Status Tool tab verify that the nodes and services are running as expected. And Cluster configuration GUI must look like below screenshot (refer to Figure 11, Cluster Management).
 - 3) Make sure on which Node/Member the cluster is running
 - 4) Make sure the cluster service is present and is started.




Figure 11: Cluster Management

3.7.8 Stop Cluster Services

Stop the cluster software on the two running nodes by running the following commands at each node in this order:

- a) `service rgmanager stop`
- b) `service gfs stop`
- c) `service clvmd stop`, if CLVM has been used to create clustered volumes
- d) `service cman stop`

Note: Click **Send to Cluster** to propagate the updated configuration to other running nodes in the cluster, once nodes are part of the cluster.