



Configuring Pervasive Postgres™ with Microsoft Cluster Service

Revision 1.0

August 2005

Pervasive Software Inc.

Pervasive Software Inc.
12365 Riata Trace Parkway B
Austin, Texas 78727
Telephone: 800-287-4383/ 512-231-6000
Fax: 512-231-6010
Internet: <http://www.pervasive-postgres.com>

© 2005 Pervasive Software Inc., Pervasive, Pervasive.SQL, Btrieve, and the Pervasive logo are registered trademarks of Pervasive Software Inc. All other product names are trademarks of their respective companies. All rights reserved worldwide.

Microsoft Cluster Service is supported by the Postgres on the following configurations:

- Windows 2003 Advanced Server
- Windows 2000 Advanced Server with SP2 or higher
- Windows NT 4 Server with SP6a or higher

This section assumes that you are familiar with the installation and configuration of Cluster Service and need only the information required to add the database services to a Cluster Service group. These instructions work equally well for Pervasive Postgres and the standard community PostgreSQL.

If you are not familiar with Cluster Service technology, please refer to the Microsoft documentation for how to install Cluster Service, verify it is working correctly, and perform tasks with it. You may also find the following links beneficial:

- <http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspx>
- <http://www.microsoft.com/windows2000/technologies/clustering/default.asp>
- http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/WindowsNetServer/proddocs/SCCon_BP.asp
- <http://www.microsoft.com/ntserver/ProductInfo/Enterprise/clustering/ClustArchit.asp>.

How to Proceed

This section is organized in the manner in which you should proceed to add the Pervasive Postgres services to a cluster group. That is, this section gives you the following suggested approach:

- Verify that Cluster Service is functioning correctly
- Add a cluster group for Pervasive Postgres
- Install Pervasive Postgres on the cluster nodes
- Add the Pervasive Postgres resources to the cluster group
- Establish the Pervasive Postgres database(s) on the cluster shared disk and verify connection from a Pervasive Postgres client.

Verify Cluster Service Functioning Correctly

It is essential that Cluster Service be functioning correctly before you add the Pervasive Postgres services to a cluster group. Refer to the Microsoft documentation for how to install Cluster Service, verify it is working correctly, and perform tasks with it.

We recommend that you select servers and disk subsystems from the Microsoft hardware compatibility list (HCL) for clusters.



Note

If Cluster Service is not set up properly, you will be unable to get your Pervasive Postgres resources to work properly.

Verify Network Client Communication

Before you add Pervasive Postgres to a cluster environment, we recommend that you verify network client communication with the cluster. A network client must be able to communicate with a cluster before and after a failure of a cluster node (failover). The following task lets you verify this.



To verify network client communication and cluster failover

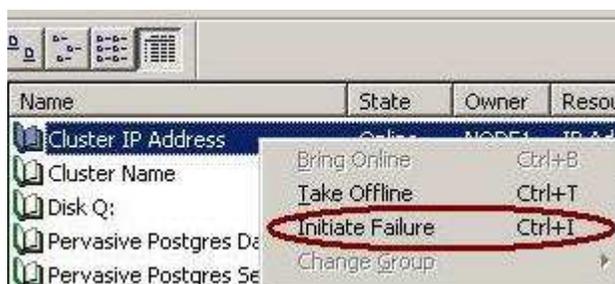
1. Verify that you can connect to the IP address of the cluster. From a command prompt, execute the following command: ping *Network Name*. For example, ping MyCluster.

If the command returns something like Reply from *IP address*, then you can connect to the IP address. For example, the following reply shows a successful connection:

```
Reply from 10.10.10.1: bytes=32 time<10ms TTL=128
```

Check that the IP address in the reply is the correct address of the cluster.

2. In the Cluster Administrator, right-click on the Network Name, then click **Initiate Failure**. Repeat this action until a failure occurs. The number of times you must repeat the action depends on the threshold setting.



This fails the node currently controlling the shared disk of the cluster. A surviving node in the cluster becomes active and assumes control of the shared disk.



Note

Ensure that the failover occurred and that a surviving node is now controlling the shared disk. If you cannot initiate a failure, then your clustering environment may not be set up

correctly and you will be unable to get your Pervasive Postgres resources to work properly.

3. On the client computer used for step 1, verify that you can still connect to the IP address of the cluster by repeating the ping *Network Name* command.

The cluster is functioning correctly if you can execute a command on the cluster from a client computer before and after a failure. If you cannot, check the setup of the Cluster Service and the resources.

Add a Cluster Group for Pervasive Postgres

Your next task is to create a cluster group to which you will later add the Pervasive Postgres resources. You create a cluster group with the Microsoft Cluster Administrator.

Please refer to the Microsoft Cluster Service documentation for how to add a cluster group. The steps to add the Pervasive Postgres resources to a cluster group are explained in *Add Pervasive Postgres Services to Cluster Group*.

All resources within a cluster group fail together. That is, if one resource within a group fails on a node, the remaining resources are taken offline. All resources within the group are then transferred to a surviving node and brought back online. For this reason, we suggest that you create a separate cluster group for Pervasive Postgres. This is not mandatory, but is a good administrative method.

At a minimum, ensure that a cluster group for Pervasive Postgres contains the following resources.

- **IP Address.** This is the IP address of the cluster. The Pervasive Postgres clients (libpq) connect to this IP address.
- **Network Name.** This is the name of the cluster that users see on a network. The name is of your choosing. For example, you may want to name the cluster "Pervasive Postgres Cluster" or something to that effect.
- **Physical Disk.** This is the disk subsystem that contains the Pervasive Postgres database. The cluster nodes share this disk, but only one node at a time controls the disk.

Failback

In addition to the required minimum resources (IP Address, Network Name, and Physical Disk), you may want to enable failback for the group. Failback is a feature of clustering that specifies a preferred controlling node within the cluster.

For example, suppose your cluster contains two nodes, Server A, the preferred controlling node, and Server B. If Server A fails, Server B assumes control. When Server

A comes back online, control passes back to it automatically. This automatic transfer of control is referred to as failback (clustering automatically "fails back" to the preferred node).

You specify failback on the properties dialog of the group name.



Note

When the cluster node fails back to the preferred server, a Pervasive Postgres client loses connection with the database engine. No automatic reconnection occurs. Your application must reconnect the client to the Pervasive Postgres database or you must restart the application.

Pervasive Postgres does not maintain the transaction state when the failback occurs. The transaction state does not transfer to the preferred server. Transactions are automatically rolled back to the state before the transaction began.

Install Pervasive Postgres on the Cluster Nodes

You install Pervasive Postgres executables, libraries, and similar on each cluster *node*. Do **not** install Pervasive Postgres on the cluster shared disk. Choose identical options for each installation.

The shared disk is where the Pervasive Postgres database(s) reside. That is, PGDATA should be on a directory from the shared disk. Do **not** use the individual cluster nodes' local disks for the data directory.

Add Pervasive Postgres Services to Cluster Group

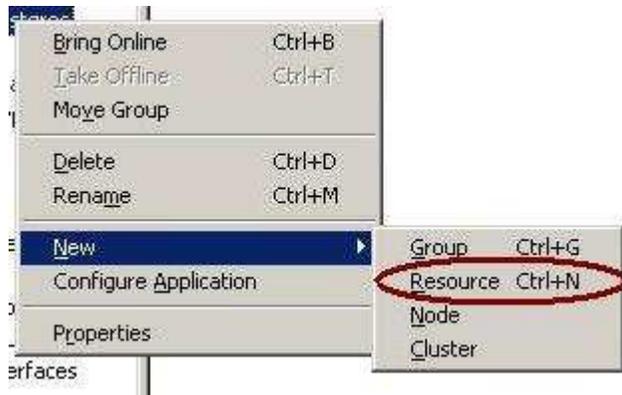
Your next task is to add the Pervasive Postgres services to the cluster group.



To add the Pervasive Postgres service to a cluster group

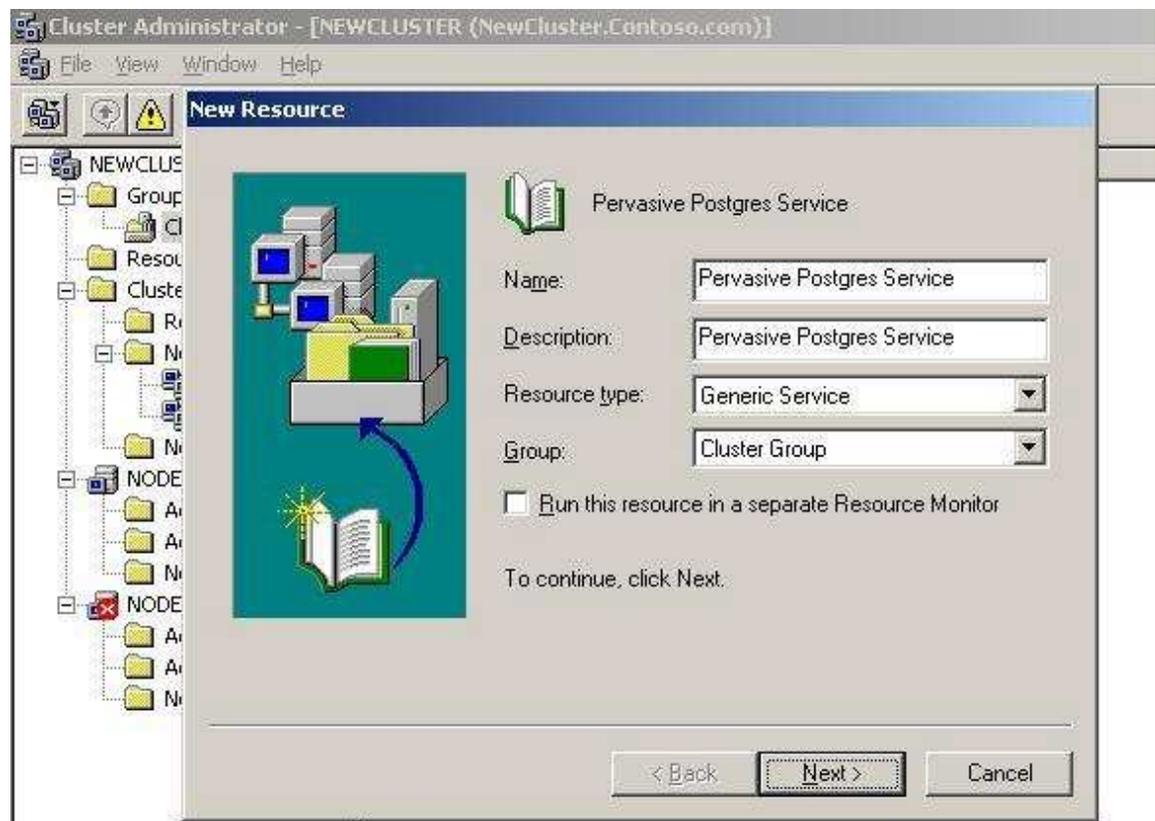
Perform this task on only one node in the cluster. All other nodes in your cluster share the information.

1. Ensure that Pervasive Postgres Server is installed on each cluster node and that the correct administrative permissions are set on each node.
2. In the Microsoft Cluster Administrator, right-click on the desired cluster group name, then click **New ▶ Resource**.

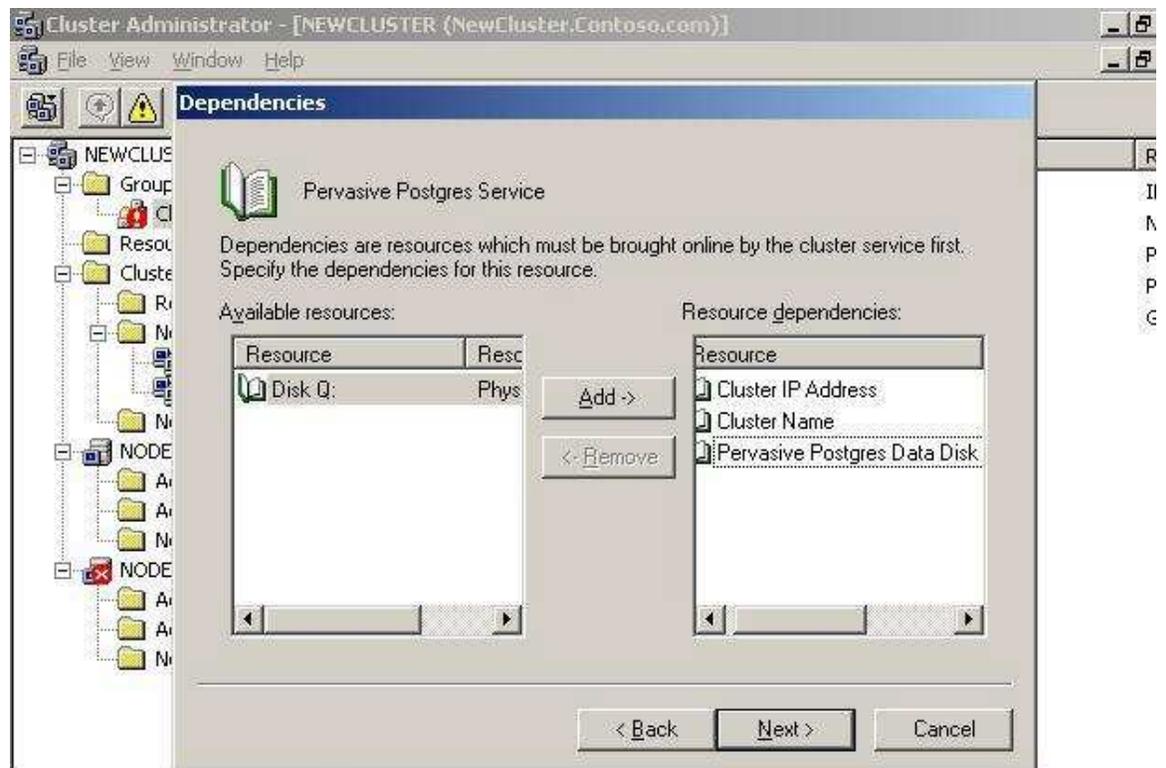


3. On the New Resources dialog, type in a name for the resource. The name is used only for display purposes within Cluster Administrator. Optionally, type in a description for the resource.
4. Click the list for Resource type then click **Generic Service**.
5. Click the list for Group then click the desired group.
6. Do *not* check the option **Run this resource in a separate Resource Monitor**.

The information should look something like the following example:

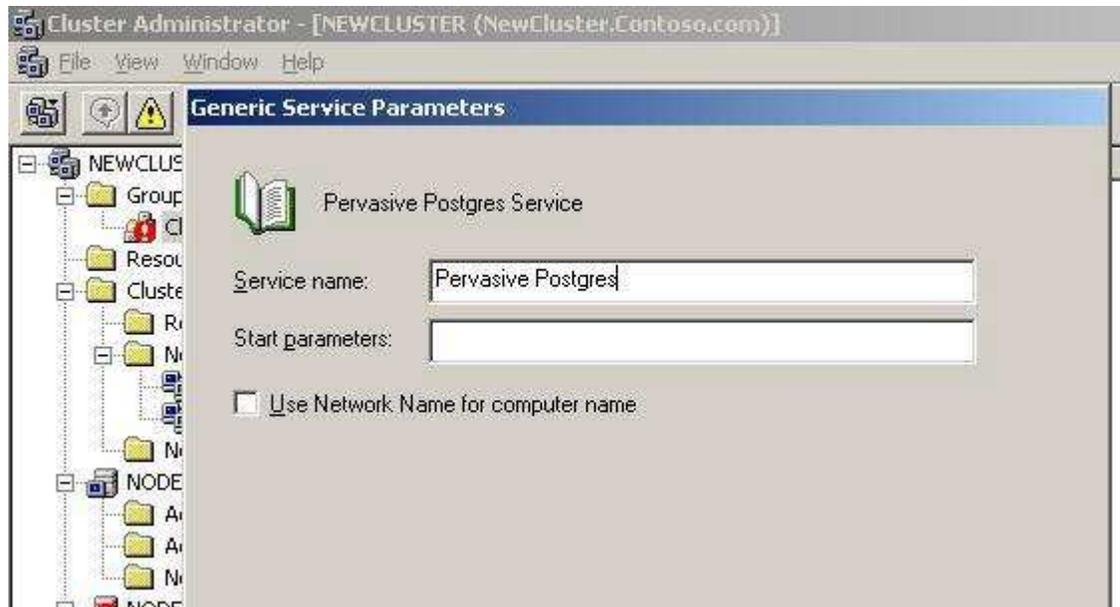


7. Click **Next**. The Possible Owners dialog appears.
8. Specify the nodes that can function as possible owners for the resource. (Add the nodes to the list of possible owners.)
9. Click **Next**. The Dependencies dialog appears.
10. Add the resources that must be brought online first. That is, the resources on which the new resource depends. Add the following dependencies to the list of resource dependencies:
 - o IP Address
 - o Network Name
 - o Physical Disk

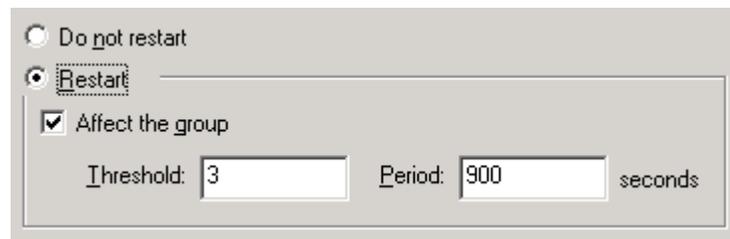


11. Click **Next**. The Generic Services Parameters dialog appears.
12. For **Service name**, type in **Pervasive Postgres**.
13. Leave the **Start parameters** blank.

The information must be the following:



19. Click **Finish** to add the new resource. A message appears stating that the resource was created successfully.
20. Click **OK**.
21. In the Cluster Administrator, right-click on the Pervasive Postgres resource you just added. Click **Properties**.
22. Click the **Advanced** tab.
23. Ensure that the **Restart** option and the **Affect the group** option are selected.



Optionally, you may set the Threshold and Period values to your choice.

24. Click **OK** (click Apply then OK is the Apply button is enabled).

You're Done!!