

Deploying SAS[®] Metadata Server to a Windows Server 2003 Failover Cluster

[Overview](#)

[Assumptions and Limitations](#)

[Step 1: Pre-Installation](#)

[Step 2: Installation](#)

[Step 3: Configure SAS Group to Microsoft Cluster Service](#)

[Step 4: Point to the "Clustered" Metadata Server](#)

[Repository Back-up Considerations](#)

Overview

The following information describes the environment, installation and configuration requirements needed to enable the SAS Metadata Server to execute within a Microsoft Cluster Service failover scenario. Within the Microsoft Cluster Service environment the SAS Metadata server is defined as a "Generic Service" resource type. Defining the SAS Metadata Server as a "Generic Service" resource type allows the server to be installed and configured using the standard SAS installation and configuration utilities provided with the product. Additionally, to mitigate changes to the infrastructure, the SAS Metadata server components execute as "cluster-unaware", which allows the server to utilize the existing installation and configuration framework.

Assumptions and Limitations

- SAS Metadata Server deployed to Microsoft failover clusters is supported with SAS 9.1.3 SP4 (or later) on Microsoft Windows Server 2003.
- The SAS 9.1.3 Metadata Server has internal provisions for detecting, handling and recovering from fault conditions. In these conditions the Microsoft Cluster Service failover may not be aware of internal recovery performed by the SAS Metadata Server, and therefore the full capabilities of MSCS may not be exploited.
- Clients actively connected to the SAS Metadata Server during cluster failover processing will have repository information preserved at the last "update transaction" boundary. Client connection state information may not be preserved across a cluster failover boundary.
- A clustered SAS 9.1.3 Metadata Server appears as a single-system image to all external clients.
- You are installing a new instance of the SAS Metadata Server and will work with the standard plan file *Metadata Server* included in the installation media.
- If you are installing other products and components, it is recommended that you install and configure the SAS Metadata Server first following these instructions. Once you have confirmed

that you have a fully-functioning SAS Metadata Server running in your cluster environment, you can install the other products.

- While these steps will allow you to provide failover, the use of Microsoft Cluster Services to provide load-balancing of the SAS Metadata Server is not possible.

Step 1: Pre-installation

- Before you install Metadata Server, be sure that you have completed the [Pre-Installation Checklist](#) for Windows so the proper domain IDs have been established for administering and running Metadata Server.
- Also, it is important to make sure SAS General Server user account (generally, SASSRV) has been granted the following rights, either directly or through its group memberships:
 - “Act as part of the operating system.”
 - “Logon as batch job.”

The SAS Metadata Repository and configuration information must reside on a disk that is shared between the cluster nodes. A good choice for the repository disk drive letter might be **R:**. For more suggestions on a drive letter assignment for this disk, see Microsoft Knowledge Base article [318534](#), “*Best Practices for Drive-Letter Assignments on a Server Cluster*.”

Step 2: Install and Configure SAS Metadata Server on Cluster Nodes

Since there is not a specific, cluster-aware install program, the next step is to install the SAS Metadata Server on both nodes using the standard installation mechanism.

- Start SAS Software Navigator and select **Advanced** for the deployment.
- Click **Select a standard plan** from the Select a deployment plan page and use the drop-down box to select **SAS Metadata Server**.
- At the end of the set-up operation, the SAS Configuration Wizard is launched. The key to directing both nodes to the shared repository is to change the default drive letter (**C:**) to the drive letter that you selected for the shared repository disk (say **R:**) as shown in the following display.



- Specify **Run as services (recommended)** on the *SAS Server Configuration Options* panel.
- On the *Enter SAS Metadata Server Information* page, enter the physical node name, in the field *SAS Metadata Server Host Name*. This should be the fully qualified physical node name. For example, in the diagram on page 6, Step 4, the physical node name would be **node1.OrionGold.com**.
- Complete the *Configuration Wizard* and proceed with the **instructions.html** document produced by the *SAS Software Navigator*.
- Complete the installation of the SAS Metadata Server on the second physical node, using the standard installation methodology. During installation, when prompted for the *SAS Metadata Server Host Name*, use the fully qualified physical node name of the second node. For example, in the diagram on page 6, Step 4, the physical node name would be **node2.OrionGold.com**.

Step 3: Configure SAS Group to Microsoft Cluster Service

Within the context of Microsoft Cluster Service, the term *group* refers to a combination of resources that are managed as unit for failover purposes (i.e. a set of things that need to be brought up together on the second node if a failure occurs on the first node). A *resource* is any individual hardware or software component within a group. To configure SAS Metadata Server as a group to Microsoft Cluster Service, you need to define four resources:

- A *Generic Service* resource that defines the running SAS Metadata Server. If the SAS Metadata Server was truly “cluster-aware”, the installation process would configure specific and custom resources when installing on a cluster.
- The physical disk where the repository is stored.
- A Network Name for the cluster.
- An IP Address for the cluster.

These resources should be defined within a new group that you create specifically for SAS Metadata Server. The following steps can be carried out either with a GUI or on the command line. The Cluster Administrator GUI, shown below, enables you to add a group and to add resources to the group by using GUI property panels.



Alternatively, you can invoke *cluster.exe* at a command line. The results of *cluster.exe* actions appear in the Cluster Administrator GUI immediately, just as though they had been completed by using the GUI itself.

The following series of commands provides an example showing how a group and its resources might be created. While the IP address, node names, and other names are shown for example only, the service and group names are recommended whether you use the GUI or commands.

Command	Purpose
<code>Cluster group "SAS" /create</code>	Create a group.
<code>Cluster group "SAS" /setowners:"NODE1", "NODE2"</code>	Set the possible owners of the group to be the two nodes of the cluster you have created.
<code>Cluster . res "Metadata" /create /group:"SAS"</code> <code>/type:"physical disk"</code> <code>Cluster . res "Metadata" /priv Drive="R:"</code> <code>Cluster . res "Metadata" /On</code>	Create a physical disk resource and bring it online. See Microsoft Knowledge Base Article 555312 " How to Create a Physical Disk Resource with Cluster.exe " for more information

Command	Purpose
<pre>Cluster . res "MS_IP" /create /group:"SAS" /type:"IP address" Cluster . res "MS_IP" /priv Network="Public" Address="10.10.10.118" SubnetMask="255.255.255.0" Cluster . res "meta" /create /group:"SAS" /type:"Network Name" Cluster . res "meta" /priv RequireDNS=1 Name="META" Cluster . res "meta" /adddep:"MS_IP"</pre>	Create IP address and network name resources.
<pre>Cluster . res "Metadata Server" /create /group:"SAS" /type:"Generic Service" Cluster . res "Metadata Server" /adddep:"Metadata" /adddep:"meta" Cluster . res "Metadata Server" /priv ServiceName="SAS Lev1 MS - MetadataSe" Cluster . res "Metadata Server" /On</pre>	Create the Metadata Server as a Generic Service resource and bring it online

The hierarchy of components created in the previous steps could be represented as:

```
SAS (Group)
  L->Metadata (Resource: Shared Disk)
  L->MS_IP (Resource: IP Address)
  L->Meta (Resource: Network Name)
  L->Metadata Server (Resource: Windows Service)
```

Another way of thinking about this group is to think of it as a *virtual server*. This virtual server has a specific network name and IP address and appears to the end user just like a physical server. The difference is that when the underlying physical server on which this virtual server is running encounters a problem, Microsoft Cluster Service will “move” the virtual server to the other physical server in the cluster. Since the clients communicate with the metadata server via the virtual server’s network name or IP address, the Metadata Server remains reachable in spite of the change in the underlying physical server.

Step 4: Establish the “Clustered” Metadata Server

During the Step 2 installation, each Metadata Server was installed on the physical nodes and there are two instances of the Metadata Server. In Step 3, the Cluster Resource was created using the **SAS Group** and the **Meta Resource: Network Name**. Before using the SAS Metadata Server as a clustered server, the SAS Metadata profiles need to be configured to use the *Cluster Resource Network Name*.

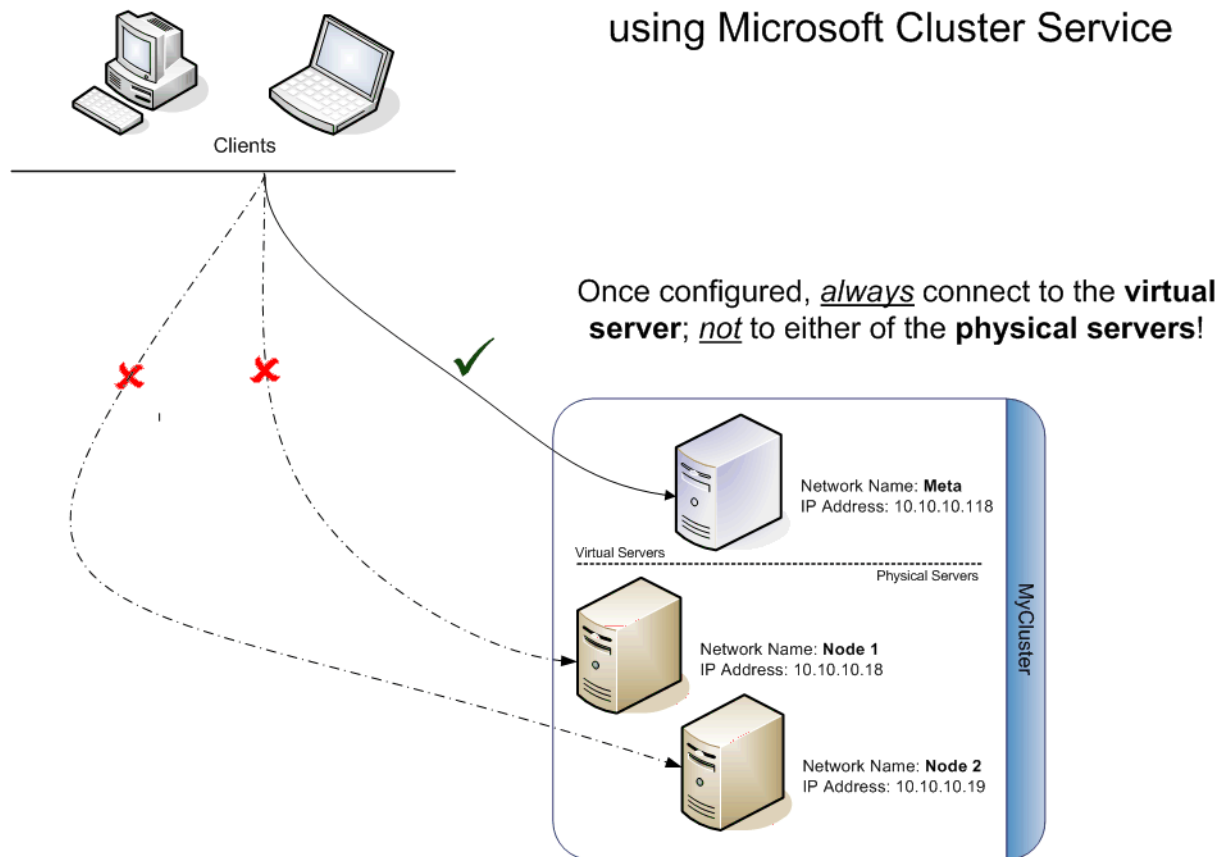
NOTE: The *Cluster Resource Network Name* or the *Cluster Resource IP Address* will be used for all SAS Metadata definitions and SAS client definitions to ensure that clients are using the cluster single instance image.

Update the Metadata Server definitions using the instances of SAS Management Console on each of the physical nodes. The metadata profiles currently point to the physical node machine name and/or IP address, for example **node1.OrionGold.com** and **node2.OrionGold.com**. Modify the definitions to use the *Cluster Resource Network Name*.

Perform the following steps on each node of the cluster:

- Launch SAS Management Console.
- The **Open a Metadata Profile** dialog box opens. Click **Edit**.
- On the **Metadata Profile** pane that asks you to specify a name for your metadata profile click **Next**.
- On the Connection Information panel change the Machine field so that it uses the resource network name that you specified when defining the group's resources or virtual server. For example, if the local name is **node1**, your domain name is **OrionGold.com** and the network name for the group is **meta**, you would change **node1.OrionGold.com** to **meta.OrionGold.com**.
- Click **Finish**.

Failover of the SAS ® Metadata Server using Microsoft Cluster Service



When installing any other SAS software, be sure to always use the virtual server's network name and IP address rather than that of any of the individual nodes when specifying the location of the

SAS Metadata Server. This is the only way to ensure you are taking full advantage of the fail-over features of the Microsoft Cluster Service.

Validation

At this point, you should have a functioning SAS Metadata Server that leverages the Microsoft Cluster Service to provide failover. To validate this, install the SAS Management Console on a machine outside of the cluster. Configure your metadata profile to connect to the clustered Metadata Server using the virtual server's name and IP address. Confirm you can access the SAS Metadata Server. Now force failover to occur by terminating the SAS Metadata Server service running on the active node (or by powering down the node). Microsoft Cluster Service should detect that the node has failed and bring the SAS Metadata Server back up on the second node. Alternatively, you can use the Cluster Administrator UI to "move" the SAS group from one node to the other. In any case, if you can continue to bring up the SAS Management Console from this third machine *after* the switch, your configuration is correct.

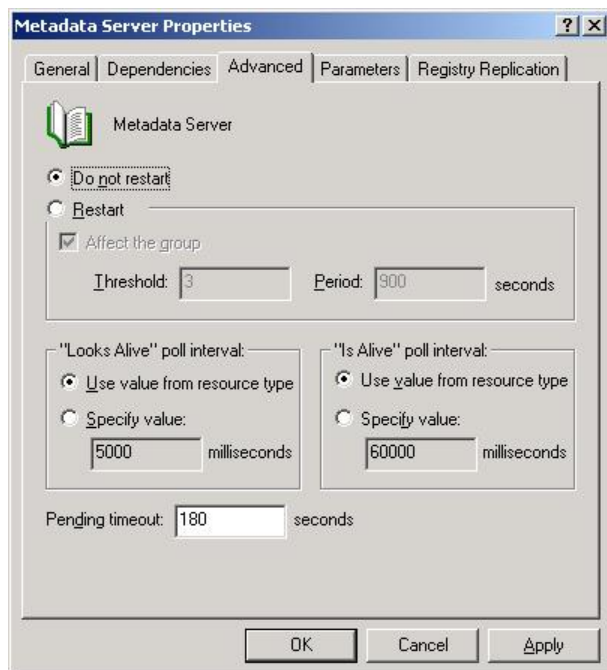
Repository Back-up Considerations

It is a recommended best practice to regularly back-up the metadata repository from the SAS Metadata Server. However, the SAS Metadata Server must be quiesced prior to performing a backup of the metadata repository. It is best to use Cluster Administrator from either node and take the group offline. Or, you can issue the equivalent command at a command prompt:

```
cluster . group "SAS" /offline
```

During installation, some *.bat* files are created on the shared disk for pausing, resuming, starting and stopping SAS Metadata server, but these will not work as expected under Microsoft Cluster Service. If a service pauses when it is defined as a resource under Microsoft Cluster Service, the Cluster Service will attempt to restart it. So, if you invoke the *pauseserver.bat* file it will appear to run normally (i.e. no error or warning messages would be generated), but the SAS Metadata Server will be immediately restarted on the second node and the command will not be effective.

If you must run the *.bat* files for some reason to control SAS Metadata Server you will need to re-configure Microsoft Cluster Service to avoid automatically restarting the service. The following image shows where to set the "**Do not restart**" property on the Advanced properties tab of the generic service resource by using Cluster Administrator.



The equivalent command is:

```
cluster . res "Metadata Server" /prop RestartAction=0
```

The command to restore MSCS so that it will restart the service is:

```
cluster . res "Metadata Server" /prop RestartAction=2
```

Similarly, stopping or pausing SAS Metadata Server via the Metadata Manager plug-in to the SAS Management Console will also result in an automatic restart unless "**Do not restart**" is selected in Cluster Administrator.

However, doing this will prevent you from taking advantage of the benefits of running the SAS Metadata Server with Microsoft Cluster Service. Therefore, it is strongly recommended that you preserve the default restart functionality and only use Cluster Administrator (or the equivalent commands) to stop the SAS Metadata Server by taking the SAS group Offline.

Copyright © 2007 SAS Institute Inc., Cary, NC, USA. All rights reserved. SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. Other brand and product names are registered trademarks or trademarks of their respective companies. This document is the confidential and proprietary property of SAS Institute Inc. It may contain approaches, techniques, and other information proprietary to SAS, and shall not be disclosed in whole or in part to third parties without the prior written consent of SAS. All Rights Reserved.