Achieving Higher Availability with Multiple Schedulers & Clustering

©2024 Software Pursuits, Inc.

Table of Contents

Introduction

The demands for data availability continue to increase. A business experiences increased cost when employees are unable to access critical files or when systems are not accessible. SureSync provides a powerful set of file replication and synchronization features to make data more accessible.

This document discusses Windows File Server Failover Clusters and the method for synchronizing data stored in these highly available file servers. SureSync database availability and Scheduler redundancy are also discussed.

Contact Information

If you need further information about SureSync or need clarification on anything within this document, please contact our support group and they will be happy to assist you.

Software Pursuits, Inc.
140 Chestnut Ln
San Mateo, CA 94403

Phone: +1-650-372-0900
Fax: +1-650-372-2912

Sales e-mail: sales@softwarepursuits.com
Support e-mail: support@softwarepursuits.com

Technical support is available between 7:00AM and 4:00PM PST Monday through Friday.

Tested Architecture

SureSync has been tested interacting with a Windows 2019 File Server Failover Cluster. This testing includes real-time monitors and file locking.

The full SureSync application was installed on a machine outside of the cluster. The Communications Agent was installed in the cluster as a Generic Service.

UNC Path Synchronization on a Schedule

Performing a synchronization to a cluster via UNC path on a Scheduled basis requires no additional SureSync configuration. Simply define the cluster UNC path as you would any other UNC path in a Job. The cluster handles virtualizing the server name and all works as expected.

Using the Communications Agent

The Communications Agent provides a powerful set of performance enhancing functionality. This functionality includes:

- **Real-Time Monitors**
  Real-Time processing increases data availability by continuously monitoring data for

updates. By processing files immediately after they change, SureSync can avoid folder scans that are necessary for scheduled jobs.

- **Remote Differential Compression (RDC)**
  RDC copies only the changes made to a file, reducing synchronization time and decreasing the consumption of network resources.
    - **RDC to Non-Windows Machines**
      If you map a non-Windows drive to a drive letter on a Windows machine you can process delta copies to that non-Windows machine. This in turn provides bandwidth savings between any machines, regardless of the OS.
- **File Compression**
  Transmit files in a smaller package across your network by compressing them. File compression is particularly useful on slow network connections.
- **TCP/IP Transmissions**
  Transmissions across the Internet give you more options on how you connect to remote machines. Now it is even easier to keep office branches connected to the information they need.
- **Encryption**
  Ensure your files are being transmitted securely by encrypting each file transmission. Multiple FIPS certified encryption algorithms are available. This is an especially useful and necessary option when transmitting sensitive data across the Internet.
- **Change Journal Support for Schedules**
  When a Schedule is configured to run a Job using the Communications Agent Add-On, SureSync uses the NTFS Change Journal to minimize the amount of scanning performed. Eliminating much of the scanning creates a significant performance improvement for Schedules.

When using the Communications Agent with a file server cluster, you must configure the Communications Agent service to operate as a Generic Service within the cluster.

High Availability for SureSync & the Database

If you are attempting to create a highly available SureSync environment, the SureSync machine and the machine hosting the SureSync database can be single points of failure. For example, a reboot of the SureSync machine to apply Windows updates will result in downtime of the synchronization environment for the length of the reboot. In many environments, this is an acceptable amount of downtime.

In environments where higher levels of availability are desired, it is recommended that the SureSync database be a SQL database stored on a clustered Microsoft SQL server. This makes the database highly available.

You can install SureSync on multiple servers and run multiple Schedulers to provide an additional level of redundancy / availability.

Example Architecture

The full SureSync application (where the SureSync Desktop is installed) should not be run inside a cluster. Instead, SureSync should be installed on a machine outside of the cluster and the Communications Agent ran within the cluster. For the Communications Agent to function properly, it must be configured as a Generic Service within the cluster.

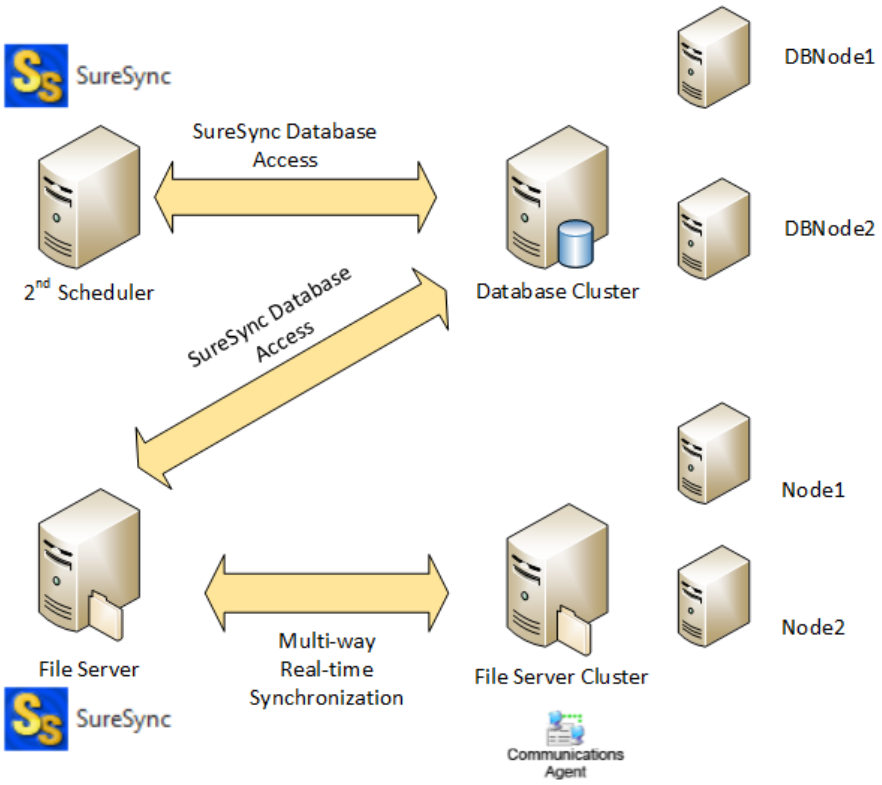

Expected Behavior on File Server Failover

If you are running a Real-time Monitor, when the primary node in the cluster goes offline, failover will occur. When failover occurs, the path that resides on the cluster will temporarily go offline and then recover. The cluster takes a few seconds to complete the failover process and during that time the path will be offline.

If you are running a Schedule to a clustered file server during a failover the path will be dropped, and the Schedule will be rescheduled to its next execution time. Retries can be configured on the Scheduling tab if you would like the Schedule to run again right away when this type of error is encountered.

The SureSync Database

The SureSync database is required to be available for SureSync to operate. If the database machine goes offline, the synchronization process will stop until the database is available again. The database is the first item to consider when attempting to increase SureSync availability.

The recommendation for environments where high availability is a requirement is to use a SQL database hosted within a SQL cluster. The SQL cluster makes the database highly available.

Using Multiple Schedulers

The Scheduler is the Windows service within SureSync responsible for launching Schedules and Real-Time Monitors at their appropriate times. In most environments, a single Scheduler is used. The Scheduler service is generally installed on the same machine as the full SureSync installation.

In highly available environments, Windows maintenance and other hardware issues can result in synchronization downtime. This downtime can be minimized with the use of multiple Schedulers. This requires at least two machines running the full SureSync application. The SureSync installation on each of the machines will open the same database and a Scheduler will be installed on each machine.

Step 1 – Select Machines and Install SureSync

Select two machines that will run SureSync Scheduler services. Launch SureSync8Setup.exe and install the full SureSync application using the “SureSync” option on the “Which features of SureSync 8 do you want to install” panel.

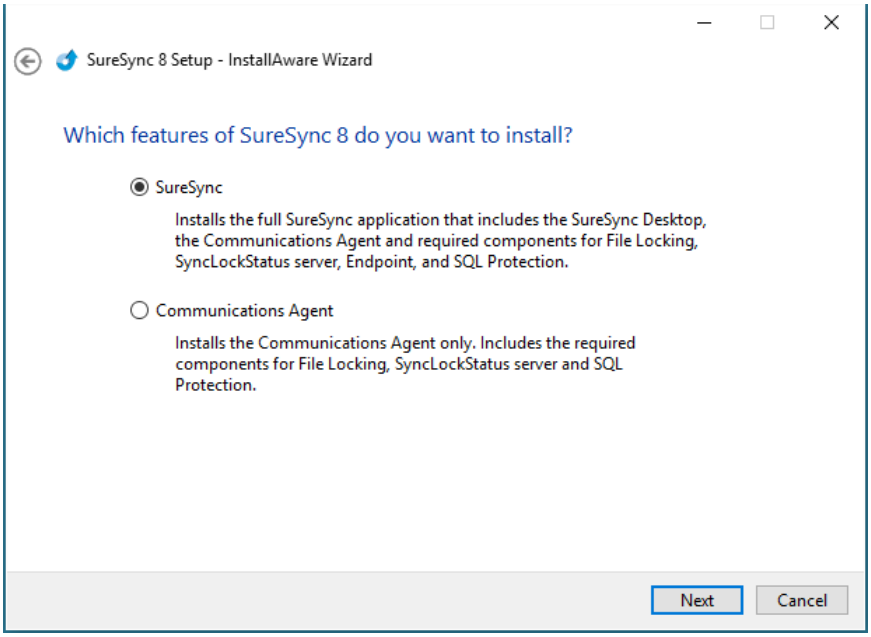

Step 2 – Install the Scheduler on the First Machine

> If SureSync is already deployed in your environment, you likely already have the product installed on a machine with a Scheduler running. In that case, you can move to installing the Scheduler on the second machine.

On the first machine, launch SureSync and open the SureSync database. To perform the installation, click on the “Home” button in the upper left-hand corner of the SureSync Desktop. Click on “Scheduler Services” to launch the Scheduler Services panel.

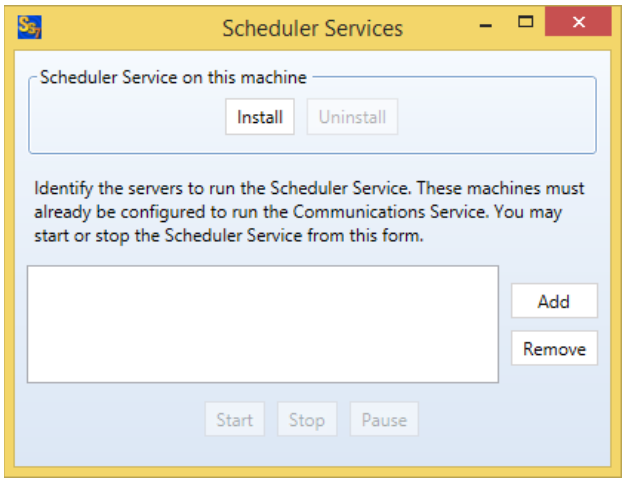

Click the “Install” button.

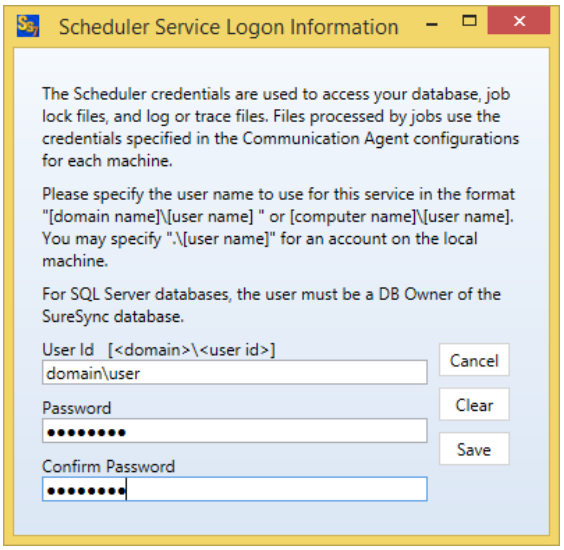

The user account the Scheduler service runs under must be a member of the local administrators group on the SureSync machine. If using SQL, the account must also be a DBOwner on the SureSync database.

On the installation panel, you will provide a User Id and password. The User Id must be provided in domain\user or machinename\user format. Enter the password for the account twice and click the “Save” button.

You will receive confirmation that the service has been successfully installed.

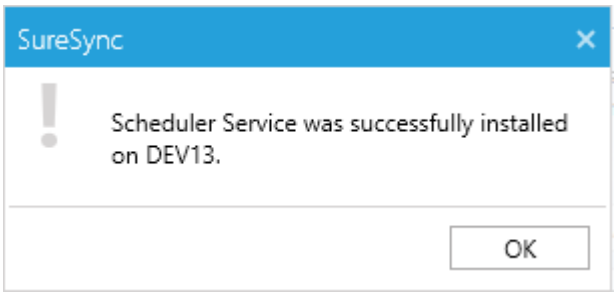

The Scheduler will now be listed on the Scheduler Services panel in a “(Stopped)” state.

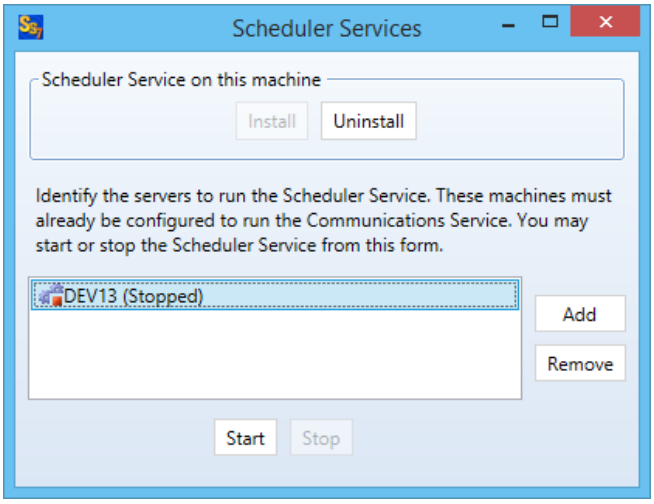

Click on the Scheduler and then the “Start” button to start the Scheduler.

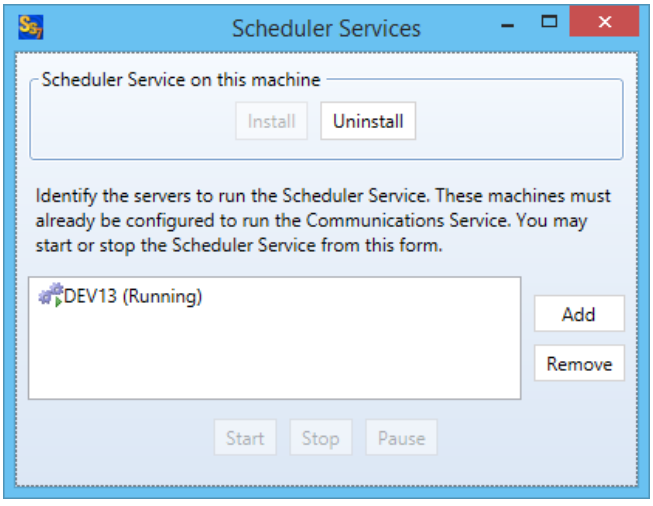

Step 3 – Install the Scheduler on the Second Machine

Launch SureSync on the second machine and open the shared database. This is accomplished by launching the SureSync Desktop, clicking on Home, Database, and Open Existing Database. Provide the same database details used on the first machine.

Repeat the steps to install the second Scheduler service from the second machine. The result will be a Scheduler Services panel that shows two Schedulers that are both running.

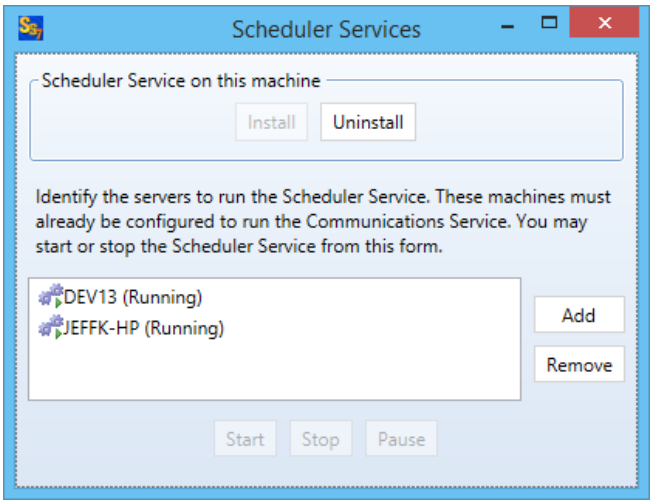

Understanding How the Schedulers are Used

Multiple Schedulers allow you to:

- Distribute load across multiple SureSync machines from a shared database
- Achieve a higher level of availability for the synchroniation / replication process

Two options on the Options tab of a Schedule or Real-Time Monitor determine how the Scheduler will act upon the item.

These options are:

- Computer that should excute this Schedule / Real-Time Monitor
- If the specified server is unavailable, use any server running the Scheduler Service

These options are highlighted in the screenshot below.

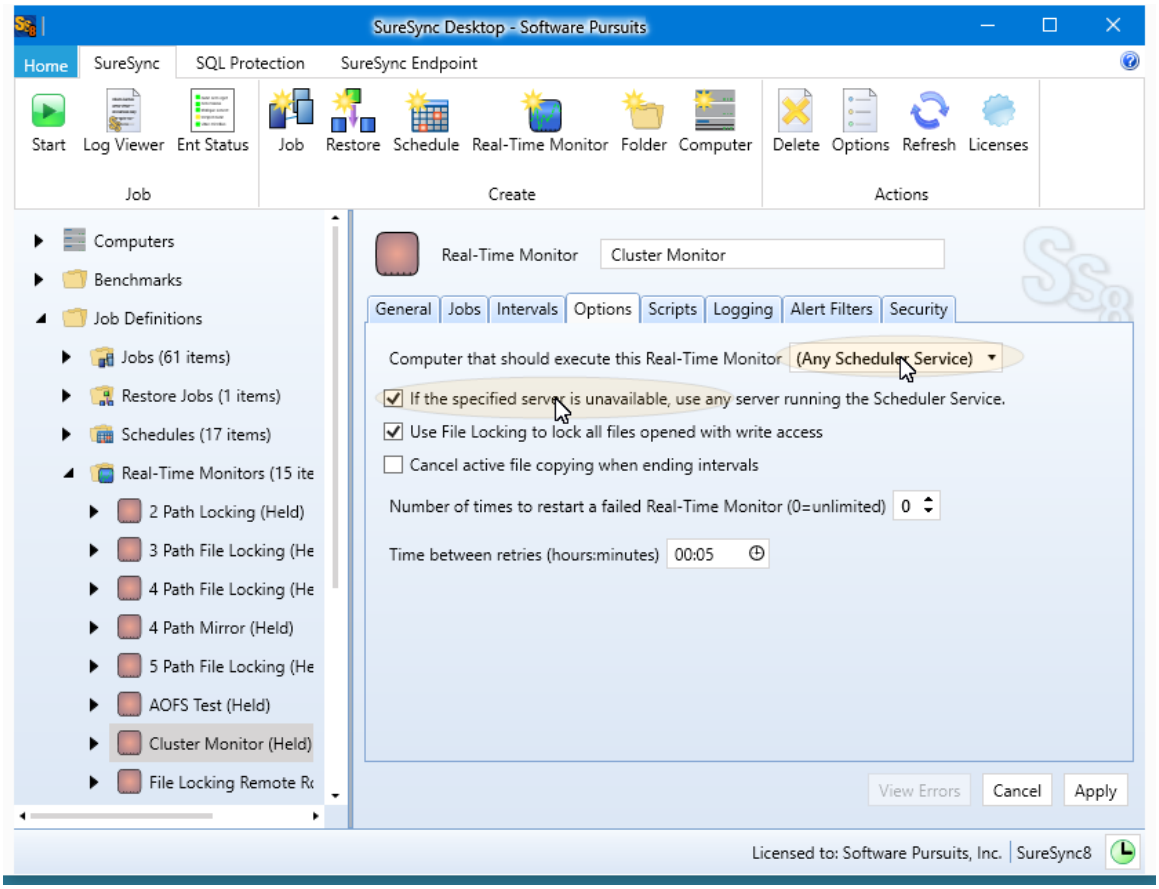

The “Computer that should execute this Schedule / Real-Time Monitor” option defaults to “(Any Scheduler Service)”. Configured this way, the Scheduler on either machine can execute the Schedule or Real-Time Monitor.

The SyncFiles.exe process responsible for performing the synchronization actions for the Schedule or Real-Time Monitor in question is executed on the machine where the Scheduler is running. With the “(Any Scheduler Service)” option selected, the process could end up running on either server.

Selecting a specific Scheduler server from the drop-down menu will control where SureSync first attempts to launch the Schedule or Real-Time Monitor.

The “If the specified server is unavailable, use any server running the Scheduler service” option is important for availability. With this option enabled, if the Scheduler defined is not available then another Scheduler will be used.

If the “If the specified server is unavailable, use any server running the Scheduler service” option is not enabled then the Scheduler will only execute on the machine defined in the “Computer that should execute this Schedule / Real-Time Monitor” option.

How Multiple Schedulers Increase Availability

If you have multiple Schedulers installed and the “If the specified server is unavailable, use any server running the Scheduler service” option is enabled, you get an availability benefit. Assume a Real-Time Monitor is running on Server1 and Server1 is rebooted. The Real-Time Monitor will go to the “Waiting to start” state and another Scheduler will start the Monitor on the other machine.

Limitations of Multiple Schedulers

The amount of time that it takes for the second Scheduler to detect the need to start the Schedule or Real-Time Monitor that was running on the first machine can take up to 5 minutes. The process is not immediate. However, multiple Schedulers do provide a higher level of availability.

With multiple Schedulers, there is no “fallback.” If Server1 running a Real-Time Monitor is rebooted, Server2’s Scheduler will detect the issue and start the Monitor. The Monitor will stay running on Server2. It will not go back to Server1 until the Monitor is restarted, assuming Server1 is defined as the “Computer that should execute this Real-Time Monitor” on the Options tab.

Cluster Basics for the Communications Agent

This section outlines some basic information about using the Communications Agent in a cluster.

Understanding the Hard Drives the Communications Agent can Access

Windows will allow a generic service to operate in a Role on one and only one Cluster Node at a time. This Cluster Node is referred to as the Owner Node of the Role.

The Communications Agent will see the system volume (C:) of that Node. It will also see drives running in any Roles currently running on the same Node. In the screenshot below, there are two File Server Roles running on the same Owner Node as the Generic Service Role running the Communications Agent. The Communications Agent would see the drives running on both.

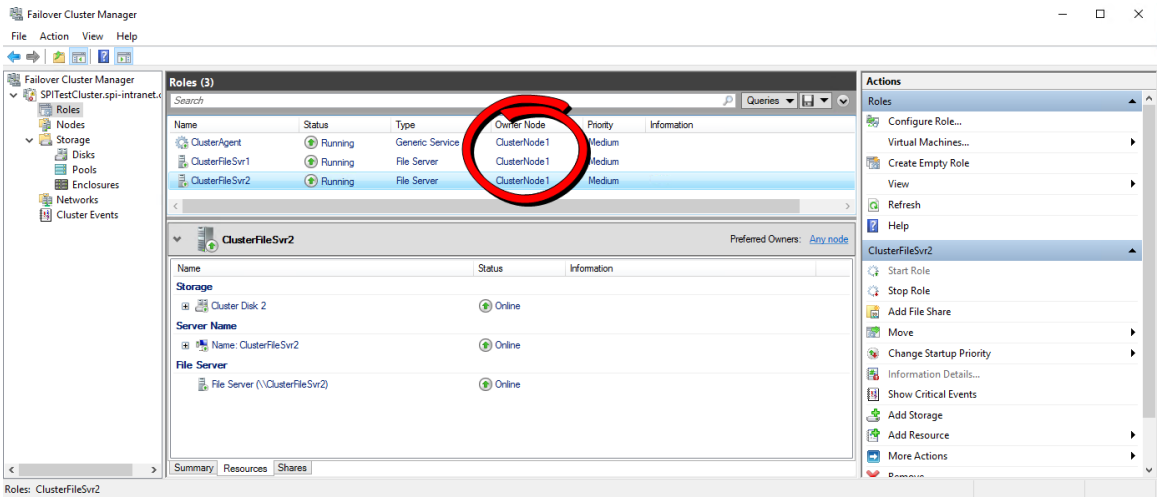

If a Role is moved to another Cluster Node, the Communications Agent will lose the ability to see those drives until the Role is back on the same node as the Communications Agent’s Generic Service.

Understanding Roles in a Cluster

In a cluster, you have the option of configuring a Generic Service as its own Role or as a Resource within another Role. Functionally, there is no difference. If the storage you’re wanting to access with the Communications Agent resides in a single Role, adding the Communication Agent’s Generic Service as a resource in the existing File Server Role helps ensure it’s always running on the same Node.

Often, there is a single File Server Role in a cluster. If there are multiple in your environment, it is important to keep in mind that the Communications Agent will not be able to see drives running in Roles on other Nodes. This is a Microsoft limitation with Generic Services.

Licensing in a Cluster

SureSync licensing requires a license for each server involved in the synchronization. This can cause some confusion with licensing a cluster. You might assume that each node in a cluster needs a license. This is incorrect.

SureSync will consume one server license and assign it to the Server Name\Client Access Point defined in the role running the Communications Agent as a Generic Service. The number of nodes in the cluster is transparent to SureSync.

Since a service can be run as a Generic Service in one and only one Role on the cluster, in effect only one license is required for the entire cluster.

Installing the Communications Agent in a New Role

When performing the installation, the Communications Agent must be installed on each node in the cluster. The service must already exist as an installed service on the nodes before you can configure the Communications Agent to run as a Generic Service.

> This guide is not intended to be a support resource for proper cluster configuration or maintenance. The guide's purpose is only to provide a conceptual overview of how the Communications Agent can be used in a cluster. Software Pursuits is unable to provide technical support related to the actual operation of a cluster. These questions should be directed to Microsoft.

To configure the Communications Agent to run as a Generic Service you should launch the Failover Cluster Manager. On the Actions menu located on the right side of the manager, click on the "Configure Role…" button.

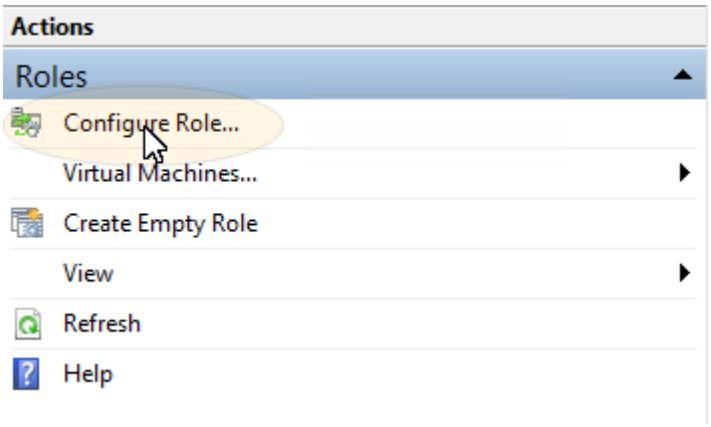

This will launch the High Availability Wizard.

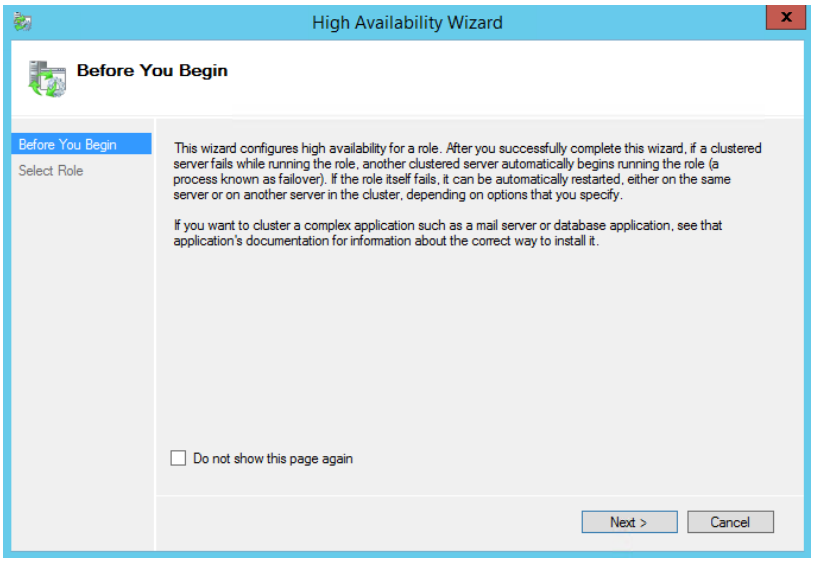

Click the “Next” button to continue.

Click on the “Generic Service” option.

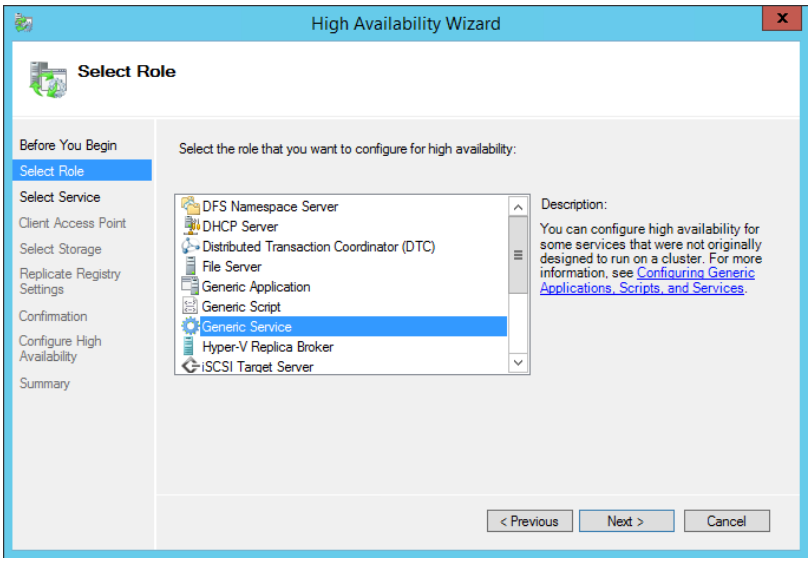

Click the “Next” button to continue.

In the “Select Service” dialog, scroll down and select the “Software Pursuits SureSync 8 Communications Agent” service.

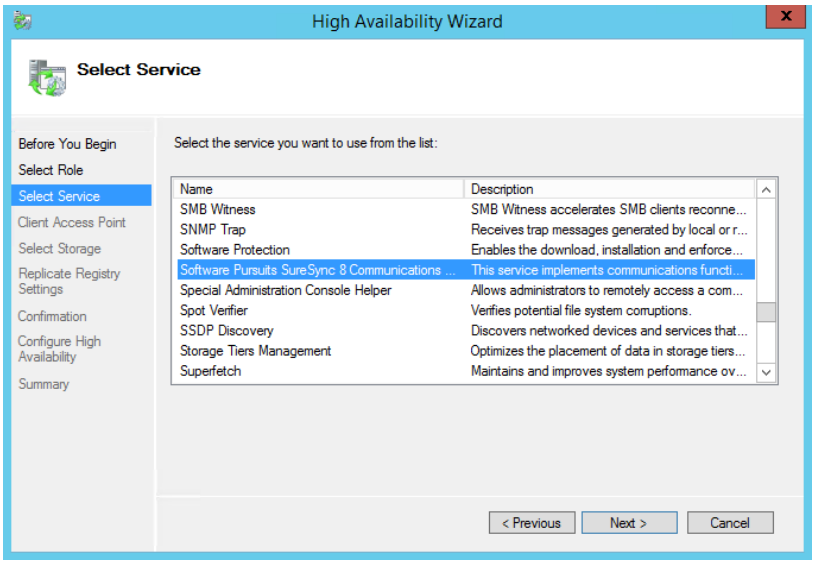

Click the “Next” button to continue.

The next panel is where you define the Client Access Point. This is the name that end users will use when accessing this clustered role.

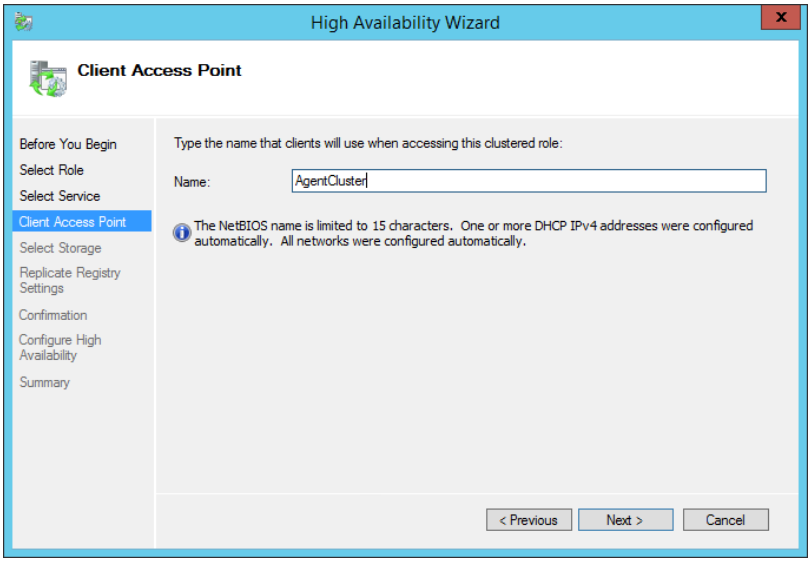

Click the “Next” button to continue.

The next wizard panel allows you to define storage that will be accessible to the Generic Service.

You do not need to define any storage specific to this Role. The Generic Service will see the storage assigned to all other Roles running on the same Owner Node.

If you want to add storage to tis Role, check the boxes for the storage you want to make available. Additional storage can be added later.

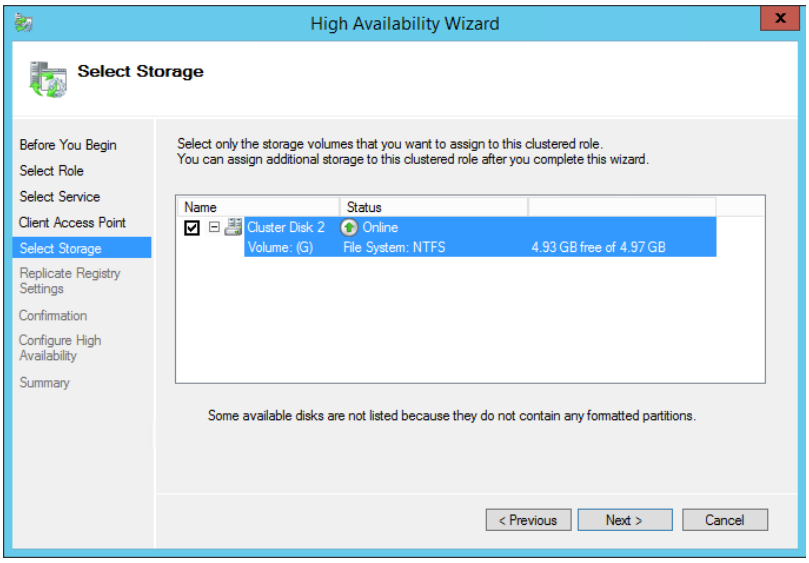

Click the “Next” button to continue.

The next panel allows you to define the replication of registry keys for the Generic Service. This does not need to be configured for the Communications Agent.

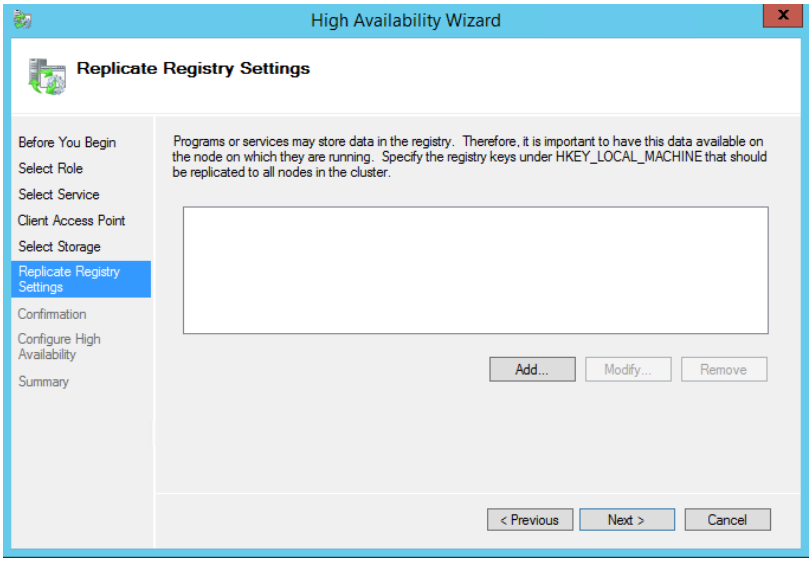

Click the “Next” button to continue.

The next panel provides summary information about the configuration that has been performed in the wizard.

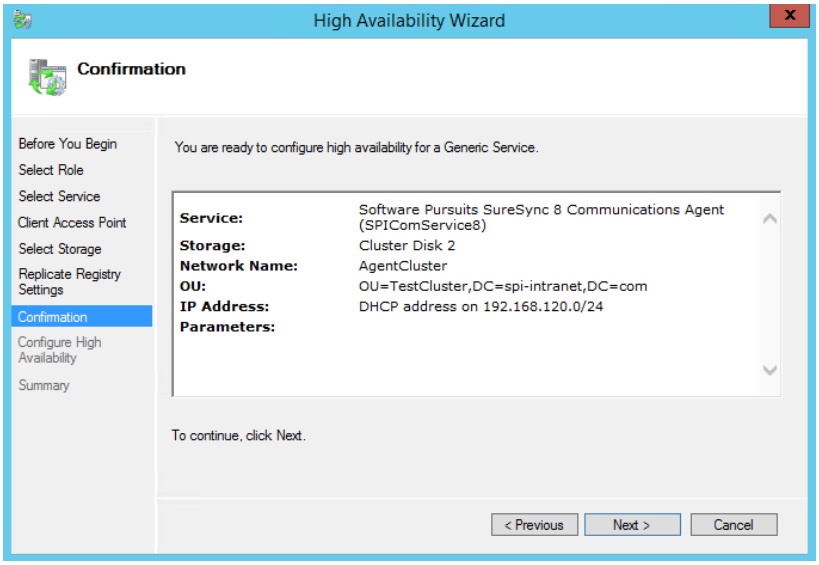

If all looks correct, click the “Next” button to continue.

The wizard will then move to the “Configure High Availabilty” phase and perform the configuration. When finished, a summary will be displayed.

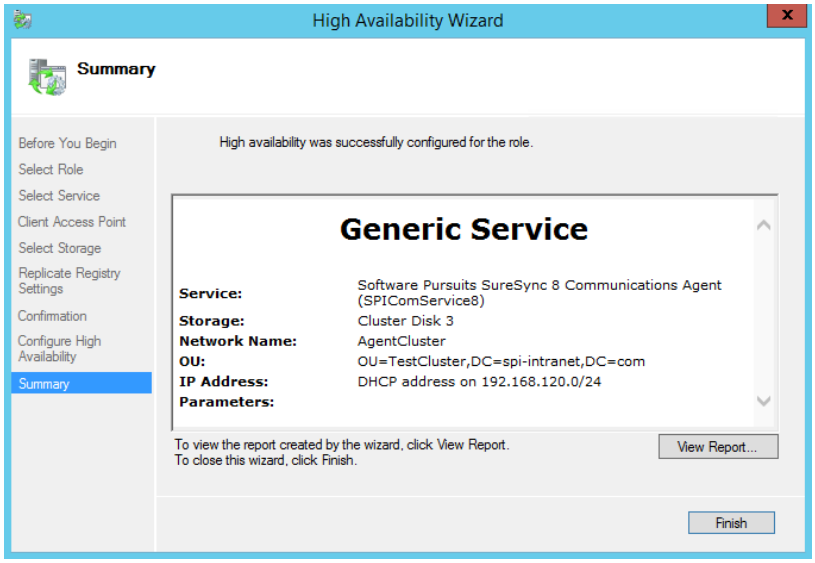

Click the “Finish” button to exit the wizard.

Additional configuration may be needed. If you’re going to allow end users to access files stored on storage within this Role, you would want to add a file share to the Role configuration.

Installing the Communications Agent in an Existing Role

When performing the installation, the Communications Agent must be installed on each node in the cluster. The service must already exist as an installed service on the nodes before you can configure the Communications Agent to run as a Generic Service.

This guide is not intended to be a support resource for proper cluster configuration or maintenance. The guide's purpose is only to provide a conceptual overview of how the Communications Agent can be used in a cluster. Software Pursuits is unable to provide technical support related to the actual operation of a cluster. These questions should be directed to Microsoft.

The Communications Agent can be added as a Generic Service Resource in an existing Role. For example, if you have an existing File Server Role.

To add the Communications Agent as a Generic Service, first click on the Role.

On the Actions panel, under the Role name is an option to "Add Resource." Click on "Add Resource" and a menu will show allowing you to select "Generic Service."

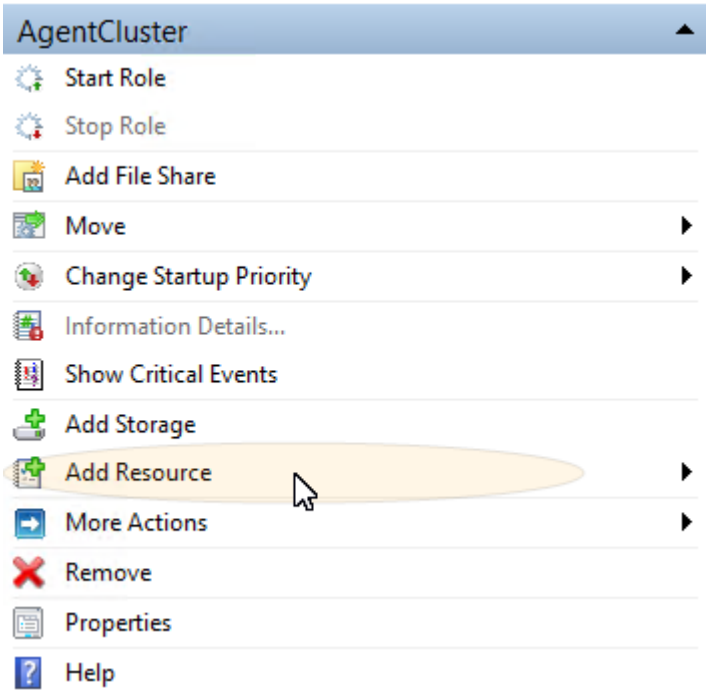

The New Resource Wizard will launch and display a list of available services on the machine. Scroll down and select the "Software Pursuits SureSync 8 Communications Agent" service.

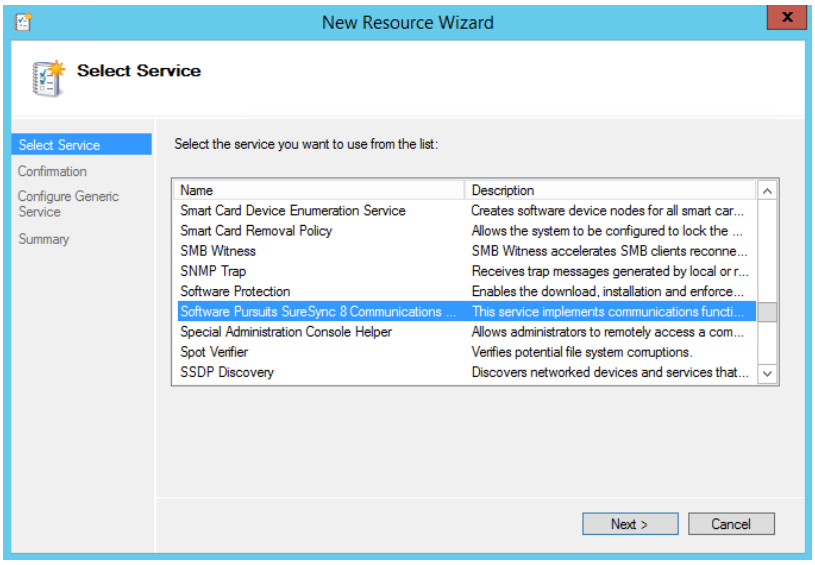

Click the “Next” button to continue.

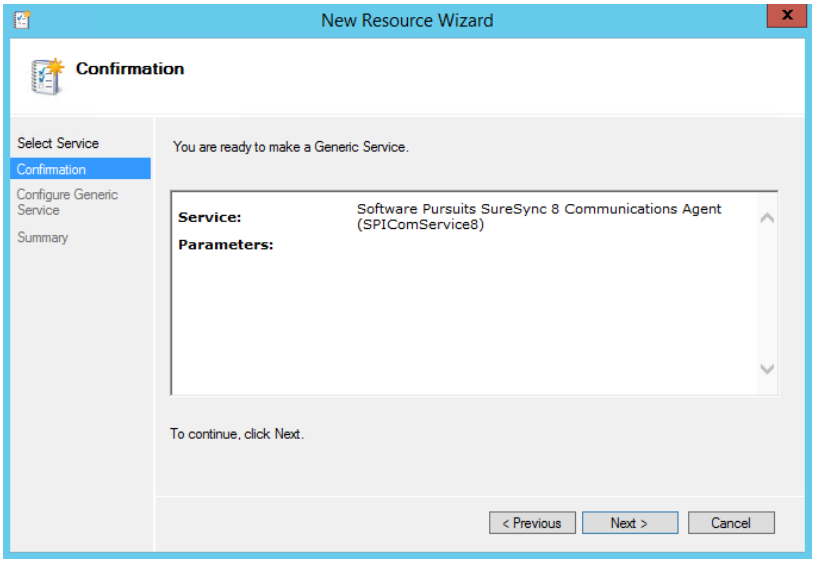

Click “Next” to continue.

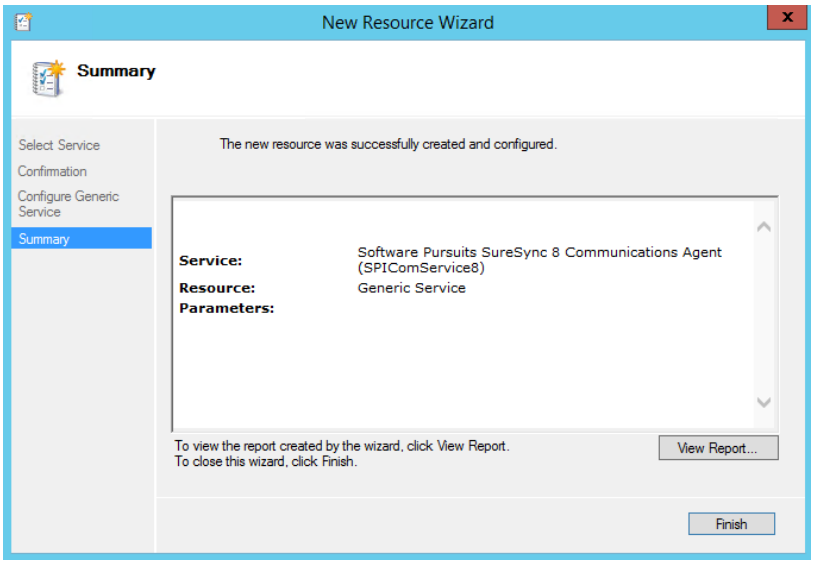

Click “Finish” to exit the wizard.

The Generic Service will be added to the cluster role but will be in an “offline” status. Before bringing the service online, some options should be configured.

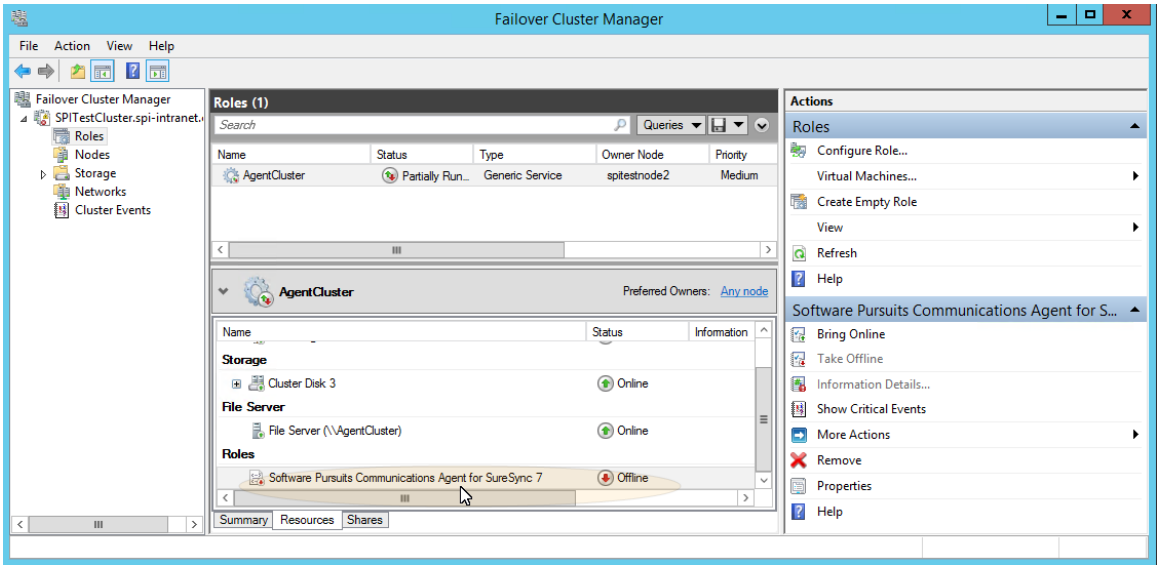

To perform the configuration, right click on the “Software Pursuits SureSync 8 Communications Agent” role and click “Properties.”

First click on the Dependencies tab. Click on the “Insert” button and under the Resource drop-down select “Name: [Computer Name]. Click “Apply.”

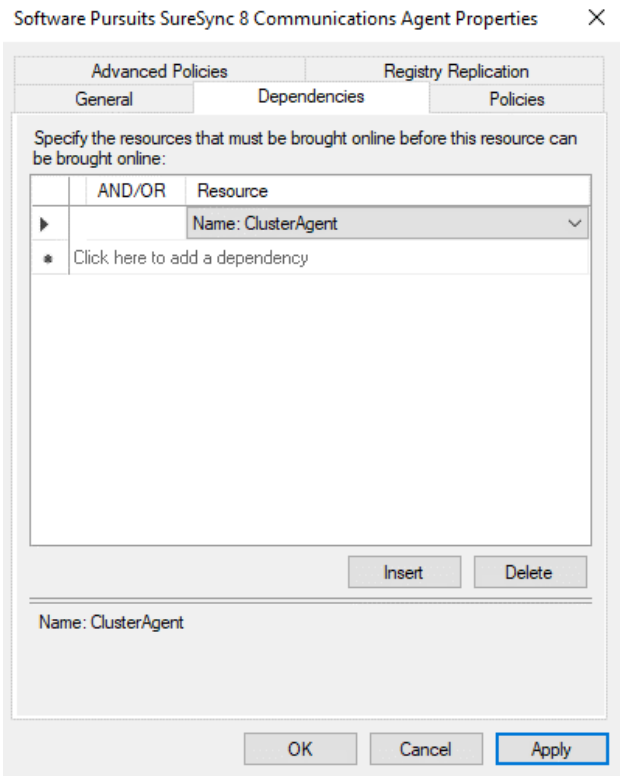

Finally, click on the “General” tab. Check the “Use Network Name for computer name” option. Click the “Apply” button.

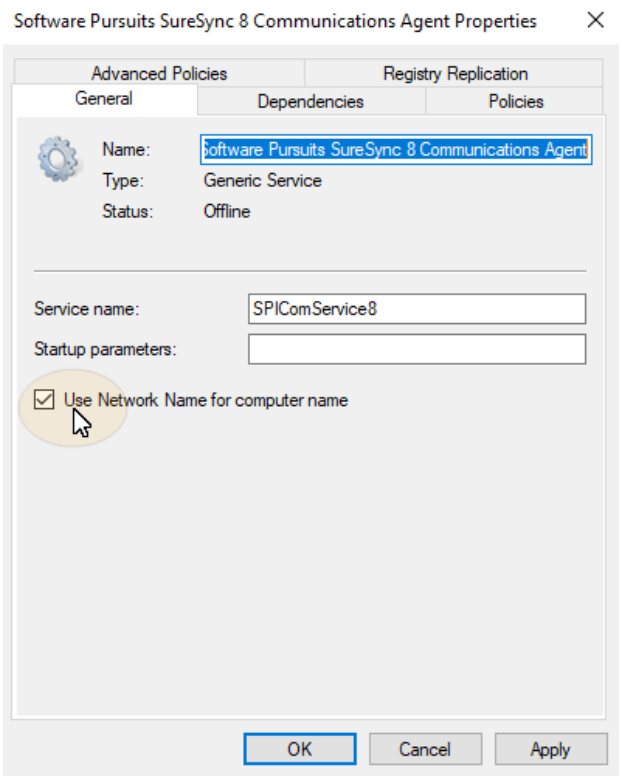

Click the “OK” button.

To bring the Generic Service online, click on “Software Pursuits SureSync 8 Communications Agent” under the Role and on the Actions pane click the “Bring Online” button.

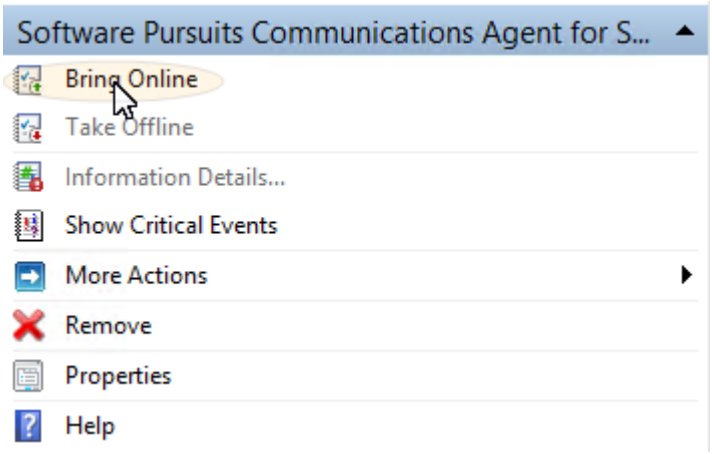

Configuring the Communications Agent Machine for the Cluster

Once the Communications Agent Generic Service has been defined in the cluster. The Communications Agent must be added into SureSync. This task is performed with the same procedure as adding any other Communications Agent.

To perform this configuration, launch the SureSync Desktop and click on the Computer button in the Ribbon bar.

The following wizard panel will be displayed:

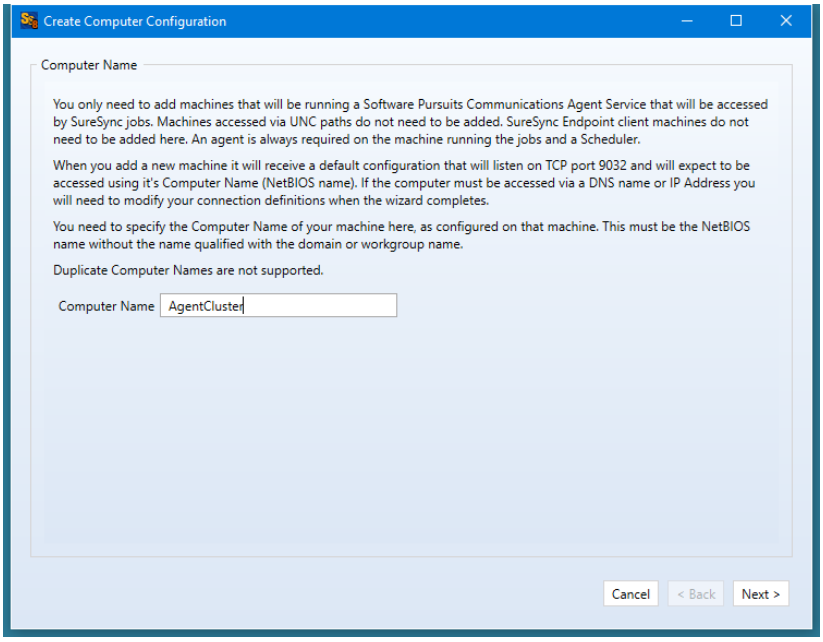

The name entered in the “Computer Name” field must be the name configured for the “Client Access Point” which was defined when configuring the cluster resource for the Communications Agent.

In this example, AgentCluster should be entered. Using an IP address or a random name for this field will result in the test failing.

Click “Next” to continue.

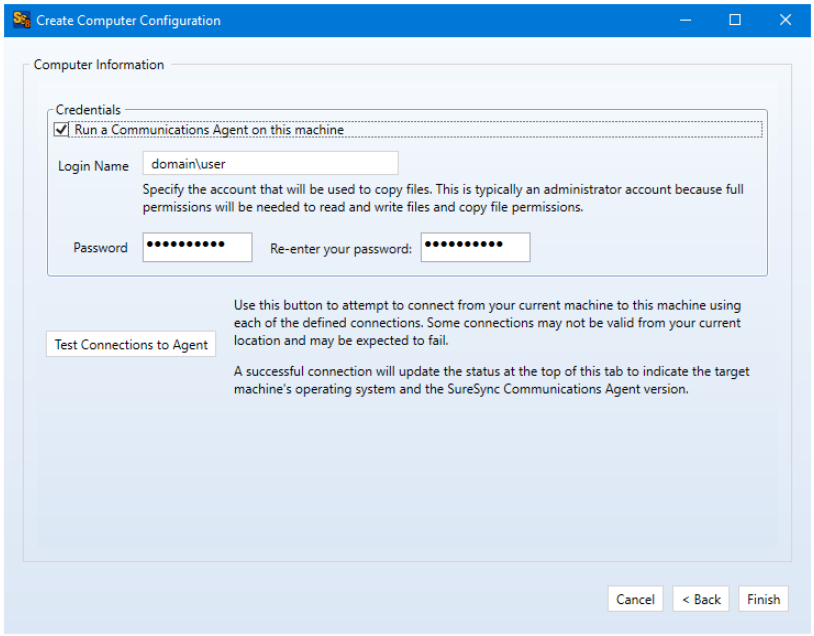

The “Computer Information” panel of the wizard is where you define the user account and password that the Communications Agent should use to access the files on the machine.

To ease initial configuration, the “Run a Communications Agent on this machine” option will be checked. Enter the account that should be used to access this Communications Agent in the “Login Name” field. The account must be in domain\user or machinename\user format.

Enter the password for the defined account twice in the password fields.

When a Communications Agent configuration is saved, a default connection for TCP port 9032 is created automatically. In most environments, only the default connection is used.

To test the connection, click the “Test Connection to Agent” button.

Click “Finish” to complete Communications Agent configuration.

Once the Communications Agent for the cluster has been defined, you can include paths on that Communications Agent as you would any other agent in your environment.

Updating the Communications Agent in a Cluster

When updating the Communications Agent in a cluster, it is recommended you run the setup on the passive node first. Once that node is updated then update the active node’s copy of the Communications Agent.

It is recommended that you run a test to the cluster agent after an upgrade. This is done by launching the SureSync Desktop. Click on the “Home” button, click on Communications Agents and finally click the “Configure Communications Agents” button. On the computers tab, select the cluster name from the list and click the “Test TCP Connection” button.

It is important that the cluster start the Communications Agent service. In a cluster, a Generic Service is set to the “Manual” startup type. The cluster manages starting the service on the appropriate node. This process also allows the cluster to substitute the cluster name for the NetBIOS computer name presented to the service. This cluster name substitution is critical to the Communications Agent functioning correctly in the cluster.

If the test of the connection returns a message about the request being received by the individual node’s computer name instead of the cluster name, bring up the Failover Cluster Manager. Click on Roles, select the Role running the Communications Agent as a generic service and click on Resources tab. Select the Communications Agent under Roles. Take the Communications Agent offline and bring it back online. This will trigger the cluster to stop and start the service which should resolve the name issue.