Help Documentation

# Configuring SmarterMail for Failover

## Who Should Use This

This document is intended for use by administrators deploying SmarterMail in high-volume environments and/or for organizations that want to ensure maximum uptime. It provides minimal system requirements and considerations for deploying SmarterMail in a failover environment. Note: Failover requires activation of SmarterMail Enterprise. For licensing information for this product, contact the SmarterTools Sales Department .

## Failover Overview

SmarterMail Enterprise allows organizations to decrease the likelihood of service interruptions and virtually eliminate downtime by installing SmarterMail on a hot standby that is available should the primary mail server suffer a service interruption. For businesses that use their mail server as a mission-critical part of their operations, failover functionality ensures that the business continues to communicate and that productivity remains at the highest levels possible, even if there is a primary server failure.

## Understanding How Failover Works

The main components of failover functionality are a primary server that acts as the default SmarterMail server and manages the licensing of the server cluster and a secondary server that remains connected and available in a "hot standby" mode until the primary server experiences problems with network access or system hardware.

If the primary server fails, SmarterMail can be configured to automatically enable the secondary server. When this occurs, the secondary server takes over responsibility for processing background threads and supporting all email functionality. This server will remain in active status until another failure occurs or the primary mail server comes back online.

The initial set up of SmarterMail's failover functionality entails system administrators manually disabling both the node and SmarterMail service on the primary server and then starting the node and SmarterMail service on the hot standby. However, system administrators can easily use third-party monitoring systems and script an automated failover and recovery strategy as needed. An example of this is provided at the end of this document.

## Minimal System Requirements

- A minimum of two servers running Microsoft Windows Server 2008 R2 or higher. (Windows Server Core is not currently supported).
- Three IP addresses

• Both servers must have their server times synchronized

• NFS/SMB share for mail and system files. We recommend that the share is running on a NAS/SAN that is configured as RAID 10

## Adding Network Load Balancing to Your Servers

Note: This needs to be performed on each server that will be used in the failover environment.

• Open the server manager console

• Right click on Features in the tree view and select Add Features

• Check the box next to Network Load Balancing and select Next

• Click Install

• Once the installation finishes, click Close

## Configuring the Load Balanced Cluster for Use with Failover

• Navigate to Start -> Administrative Tools -> Network Load Balancing Manager

• Click the Cluster menu item and select New

• In the New Cluster: Connect window, type the IP of your primary server in the Host: text box and select New

• When the Interface Name and Interface IP appear, select the Interface Name and click Next

• Since this is the primary node, ensure the host Priority is set to 1

• In the New Cluster: Host Parameters window, confirm the IP address and Subnet mask are correct and change the initial host state to Stopped . This is to prevent any issues with connectivity if a machine randomly reboots or suffers from a hardware failure. If all nodes are set to Started for their initial host state, traffic will be split between the two (or more) machines. Note: Monitoring software can be used to execute scripts that will start and stop hot standbys in the event of a failure and recovery. If you are not executing scripts via monitoring software then all failover will need to be handled manually.

• Click Next

• In the New Cluster: Cluster IP Addresses window, click Add and enter in your cluster IP address and the same subnet mask as in Step 6

• Select Next

• In the New Cluster: Cluster Parameters window, confirm the IP address and subnet mask, then enter a Full Internet Name , though this is optional

• Ensure the cluster operation mode is set to Multicast

• Click Next

• In the New Cluster: Port Rules window, click Edit

• If you want you can restrict the cluster IP to work on an individual port or across a port range. You can also simply allow the cluster IP to work across all ports on the server

• Ensure your port rules are set to Single Host in the Filtering Mode section

• Click OK

• Verify your settings and click Finish to complete the setup

## Joining Additional Nodes to the Cluster

• From the secondary server navigate to Start -> Administrative Tools -> Network Load Balancing Manager

• Click the Cluster menu item and select Connect to Existing . Note: the existing cluster will need to be running before a secondary node can be added

• In the Connect to Existing: Connect window, enter the IP address of your existing cluster as the Host and click Connect

• Select the existing cluster that appears in the Clusters section and click Finish

• In the main Network Load Balancing Manager , expand Network Load Balancing Clusters and right click on your Cluster (it may be the IP address of your cluster) and select Add Host to Cluster

• In the Add Host to Cluster: Connect window, enter the IP address of the secondary server in the Host: section and click Connect

• When the Interface Name and Interface IP appear, select the Interface Name and click Next

• In the Add Host to Cluster: Host Parameters window, confirm the IP address and subnet mask and ensure the Initial Host State is set to Stopped . As this is the second node you're adding to your cluster, the Priority should be set at 2

• Click Next

• Just as with the primary node, in the Add Host to Cluster: Port Rules window you have the ability to set this node to respond via specific ports or a port range. If you wish to set these rules, click Edit . Otherwise, click Finish to complete the setup

• Wait for the nodes to converge and, if necessary, stop the secondary sever by right clicking the second server's name, select Control Host -> Stop

## Configure a Shared Service Directory

• Using Network File Sharing (NFS) or Samba (SMB), create a shared directory named SmarterMail , preferably on a NAS or SAN. NOTE: We recommend that this shared directory be hosted on a server that utilizes a RAID 10 configuration for the data.

• Inside that new SmarterMail folder, create a Service folder

• Configure your permissions accordingly. If special permissions are required, configure the SmarterMail service to run with the proper credentials within the Windows Services console.

Note: When performing updates to the software, the credentials will need to be re-applied to the service

## Configuring a Fresh Installation of SmarterMail for Failover

• Install SmarterMail Enterprise on a server. This will be your hot standby. Leave all setup information as the default settings and after setup is complete, configure SmarterMail as an IIS site.

• Stop the SmarterMail service on the hot standby

• Edit the failoverConfig.xml file in the primary server's Service folder as follows:

• SharedSystemFilesPath - Set to the shared network shared system folder

• FailoverIPAddress - Set this to the IP address of the Network Load Balancer

• IsEnabled - Set this to True

• Save this file, then copy it to the hot standby's Service folder and replace the existing failoverConfig.xm l

• Copy over all folders, DAT and XML files from C:\Program Files (x86)\SmarterTools\SmarterMail\Service to the Service folder in the shared service directory you created

• Start the service on the hot standby server and verify that the paths are pointing to the network shared paths

• Activate your Enterprise key on the hot standby by logging into SmarterMail's management interface as the system admin and going to Settings -> Activation -> Licensing , then stop the SmarterMail service on the server

• Start the service on the primary server, then reactivate your Enterprise license key in the SmarterMail management interface

• After re-activating the license, go to Settings -> Bindings -> IP Address and bind all the ports to the load balancer's IP address and make sure no other IPs have any ports bound to them

• Both servers are now set up for failover. To verify this, when logged into the primary server as the system admin, go to Settings -> Failover Servers to view the servers that are part of the failover cluster

## Adding Failover to an Existing Installation of SmarterMail

Note: You will need to configure both servers for Network Load Balancing and set up a shared service directory. See the steps outlined in the Adding Network Load Balancing to Your Servers , Configuring the Load Balanced Cluster for Use with Failover , Joining Additional Nodes to the Cluster and Configure a Shared Service Directory sections earlier in this document for more information.

• Ensure the primary server is running the latest version of SmarterMail and that it is also configured as an IIS site. Ensure the IIS binding is pointing to your cluster IP address

• Install SmarterMail on a hot standby and configure it as an IIS site. Ensure the cluster node is stopped on the hot standby and ensure the IIS binding is also pointing to the cluster IP

• Stop the SmarterMail service on the hot standby

• Copy all of your mail data (located in C:\SmarterMail\ by default) to your shared service directory. If possible, use robocopy to do this because it will not result in any downtime for the mail service

• Once robocopy finishes, run it one more time. This second pass will only copy any new data

• Stop the SmarterMail service on the primary server

• Edit the failoverConfig.xml file in the primary server's Service folder as follows:

• SharedSystemFilesPath - Set to the shared network shared system folder

• FailoverIPAddress - Set this to the IP address of the Network Load Balancer

• IsEnabled - Set this to True

• Run the robocopy one more time to copy over any modified files and remaining spool e-mails

• Copy over all folders, DAT and XML files from C:\Program Files (x86)\SmarterTools\SmarterMail\Service to the Service folder in the shared service directory you created

• Edit the domainlist.xml file in the shared Service folder and change the path of your domains to match the new NFS\SMB path. (For example, \\NAS01\SmarterMail\Domains\mydomain.com)

• Edit the mailconfig.xml file and replace any instances of the old physical path's with your new network location for SmarterMail. (For example, if all of your data was hosted on E:\Smartermail, you would then perform a find and replace for all instances of E:\Smartermail to \\NAS01\Smartermail).

• On the primary server, go to Start -> Administrative Tools -> Network Load Balancing Manager and stop the cluster node, then start the NLB on the secondary node

• Start the SmarterMail service on the hot standby

• Access SmarterMail's web interface at the cluster IP and sign in as the system admin

• Activate your Enterprise key on the hot standby by going to Settings -> Activation -> Licensing

• Verify that the data and settings are being picked up from the shared Service directory

• Stop the SmarterMail service on the hot standby and stop the secondary cluster node

• Start the cluster node and the SmarterMail service on the primary server

• Sign into the web interface on the primary server and re-activate the Enterprise license key by

going to Settings -> Activation -> Licensing

• Verify mail data and settings are being accessed from the shared service directory

## Scripting Failover

Below is an example of a PowerShell script that can be created to automate the SmarterMail failover process. You can utilize a third party monitoring product such as PRTG or SolarWinds (though there are many others) to execute this script when a failure is detected.

## Prepping PowerShell on the Servers

The servers will need to be configured to run remote scripts and accept remote PowerShell sessions. Therefore, on each server, run the following commands within an elevated PowerShell console:

• Set-ExecutionPolicy RemoteSigned - Press Y to accept
• Enable-PSRemoting -force

## Sample Script - Stop a Primary Server and Start the Hot Standby

In the scripts below, replace the "WAN" variable called in the –hostname parameter with the name of your interface. This can be obtained by opening a PowerShell console on the server and typing Get-NlbClusterNodeNetworkInterface . Also replace Server01 and Server02 with the NetBIOS names of your servers.

```
$StopPrimary = New-PSSession -ComputerName Server01 Invoke-Command -Session
$StopPrimary -ScriptBlock { Import-Module NetworkLoadBalancingClusters ;
Stop-nlbclusternode -HostName Server01 -InterfaceName "WAN" ; import-module
WebAdministration ; stop-webapppool SmarterMail; set-service -computerName
Server01 -name mailservice -status stopped ; remove-pssession Server01}

$StartSecondary = New-PSSession -ComputerName Server02 Invoke-Command -
Session $StartSecondary -ScriptBlock { Import-Module
NetworkLoadBalancingClusters ; Start-nlbclusternode -HostName Server02 -
InterfaceName "WAN" ; set-service -computerName Server02 -name mailservice
-status running ; import-module WebAdministration ; start-webapppool
SmarterMail ; remove-pssession Server02 }
```

## Sample Script - Stop the Hot Standby and Re-start the Primary Server

These scripts can be used to bring the primary server back online and stop the hot standby after your monitoring software issues an all-clear.

```
$StopSecondary = New-PSSession -ComputerName Server02 Invoke-Command -
Session $StopSecondary -ScriptBlock { Import-Module
NetworkLoadBalancingClusters ; Stop-nlbclusternode -HostName Server02 -
```

```
InterfaceName "WAN" ; import-module WebAdministration ; stop-webapppool
SmarterMail; set-service -computerName Server02 -name mailservice -status
stopped ; remove-pssession Server02}

$StartPrimary = New-PSSession -ComputerName Server01 Invoke-Command -
Session $StartPrimary -ScriptBlock { Import-Module
NetworkLoadBalancingClusters ; Start-nlbclusternode -HostName Server01 -
InterfaceName "WAN" ; set-service -computerName Server01 -name mailservice
-status running ; import-module WebAdministration ; start-webapppool
SmarterMail ; remove-pssession Server01 }
```