

EMC AutoStart with Microsoft Exchange Server 2010

Applied Technology

Abstract

This white paper outlines EMC[®] AutoStart[™] integration with Microsoft Exchange Server 2010 database availability groups and synchronous storage array-based replication, including SRDF[®]/Synchronous and MirrorView[™]/Synchronous.

April 2010

Copyright © 2010 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h7097

Table of Contents

Executive summary	4
Introduction	4
Audience	4
EMC AutoStart overview	5
Exchange Server 2010 database availability group overview.....	6
AutoStart integration with Exchange Server 2010	7
Shared disk device for Windows data source	7
Mount point data source	7
EMC SRDF/S mirroring data source	7
EMC MirrorView/S mirroring data source.....	7
Configuration requirements.....	9
Log and database file placement	9
Storage presentation	9
Replication.....	9
Virtualization support.....	9
AutoStart installation.....	10
DAG and database copy creation.....	10
Required licenses for the Exchange Server 2010 module	12
Creating AutoStart Exchange Server 2010 instances.....	12
Node resource group.....	13
Database resource group	13
AutoStart Exchange Server 2010 Instance Creation Wizard.....	13
DAG and database copy management considerations	20
AutoStart Exchange Server 2010 shell commands	20
Manually moving active database copies between nodes	21
Controlling automatic database copy activation.....	22
Adding or removing nodes for an existing AutoStart Exchange instance	23
Adding or removing databases for an existing AutoStart Exchange instance	24
Querying or controlling replication.....	25
Unplanned failover when node majority is lost	26
Conclusion	27
References	27

Executive summary

Microsoft Exchange Server 2010 introduces a new feature called database availability groups (DAG), which provide high availability (HA) and site resiliency for Exchange-specific server and storage components. DAGs replace previously existing Exchange and Windows failover clustering-based technologies used for HA and site resiliency, such as single copy clusters (SCC), continuous cluster replication (CCR), and standby continuous replication (SCR). The removal of single copy clusters in particular impacts native HA and disaster-recovery (DR) solutions previously available with Microsoft Exchange Server via storage-based technologies, including geographically dispersed clusters that integrate with Windows failover clustering.

In order to help address storage-based HA and DR technologies, Microsoft Exchange Server 2010 includes an application programming interface (API) to integrate third-party replication solutions into the DAG framework. When enabled, third-party replication support will disable the native network-based log shipping mechanism used by DAGs. Synchronous replication technologies can then be used to protect the Exchange database copies specified within the Exchange Server 2010 environment.

EMC[®] AutoStart[™], an HA cluster product that provides heterogeneous application infrastructure monitoring and automated restart, has been enhanced with version 5.3 SP5, to include an Exchange Server 2010 module that integrates with the DAG third-party replication API. The AutoStart Exchange Server 2010 module enables the ability to use both shared local storage as well as synchronously replicated storage with SRDF[®]/S and MirrorView[™]/S as database “copies” within a DAG. The ability to utilize shared local storage as well as synchronously replicated remote copies helps to enable HA and site-resiliency functionality similar to SCC and geographically dispersed cluster capabilities available with previous versions of Exchange.

AutoStart integration with Exchange Server 2010 offers several key features for organizations that include:

- Synchronous disaster recovery capabilities that enable zero data loss recovery point objectives (RPO)
- Shared storage capabilities that decrease the number of database copies required to achieve server-based HA
- The ability to leverage storage replication with Exchange Server 2010 and maintain a common DR/business continuity (BC) strategy across all applications and platforms in an environment.
- Maintaining existing investments in storage-based replication technologies, processes, and procedures for managing HA and DR/BC for servers and storage in a data center.
- Resource preservation provided by AutoStart’s implementation returns host cycles and network bandwidth that would normally be used by native DAG replication.

Introduction

This white paper will provide an overview of AutoStart 5.3 SP5 and its integration with Exchange Server 2010 database availability groups. Sections will include a high-level overview of AutoStart, an overview of Exchange Server 2010 DAGs, and specific integration capabilities when combining AutoStart and storage array-based replication using SRDF/S and MirrorView/S into the DAG framework. Additionally, supported configurations and installation requirements will be outlined, and some general processes and procedures for managing an environment configured with AutoStart and Exchange Server 2010 will be discussed.

Audience

This white paper is intended for Microsoft Exchange database administrators and storage architects who want to understand how storage-based replication can be used within the Exchange Server 2010 database availability group framework.

EMC AutoStart overview

AutoStart is an HA cluster product that provides heterogeneous application infrastructure monitoring and automated restart for UNIX, Linux, and Windows environments and geographically dispersed data centers. AutoStart continuously gauges replication status (via SRDF/S or MirrorView/S, for example) and maintains availability in the event of network, server, or application failures with both local and remote failover across physical and virtual or mixed environments. AutoStart provides four primary capabilities:

- **Business continuity**—AutoStart provides HA and DR through clustering for local and long-distance environments.
- **Automated restart**—AutoStart automates application restart processes on a local or remote server in the event of a planned or unplanned outage. AutoStart provides rapid and efficient failback of services, applications, and data to maintain business continuity.
- **Integration with replication solutions**—AutoStart works jointly with SRDF/Synchronous, and MirrorView/Synchronous along with SRDF/Asynchronous and and MirrorView/Asynchronous. However, since Microsoft only supports synchronous remote replication, this paper will focus on SRDF/S and MirrorView/S. AutoStart seamlessly transfers control of storage resources to a remote facility in the event of an application restart and provides automatic or operator-initiated failover.
- **Monitoring**—By monitoring servers, networks, and data parameters, AutoStart can notify administrators to remedy issues and maintain uptime.

In addition to these features, AutoStart directly integrates with key applications such as Oracle, SQL Server, and various Exchange Server versions. For the purposes of this paper, we will focus on the Exchange Server 2010 module, which integrates with the DAG third-party replication API.




Figure 1. High-level depiction of AutoStart providing server- and storage-based HA and DR/BC

Exchange Server 2010 database availability group overview

Database availability groups (DAGs) are a new feature available with Exchange Server 2010. DAGs replace all previous HA and site-resiliency solutions that existed for Exchange 2007 including SCC, CCR, and SCR. A DAG is a group of up to 16 Exchange Server 2010 mailbox servers that can host independent copies of mailbox databases. Mailbox databases are replicated continuously from an active mailbox copy to one or more passive mailbox copies. Passive mailbox copies are hosted by separate servers within the DAG, and are updated by Exchange Server-based log shipping and log replay. Figure 2 provides a logical view of a DAG with multiple active and passive database copies.

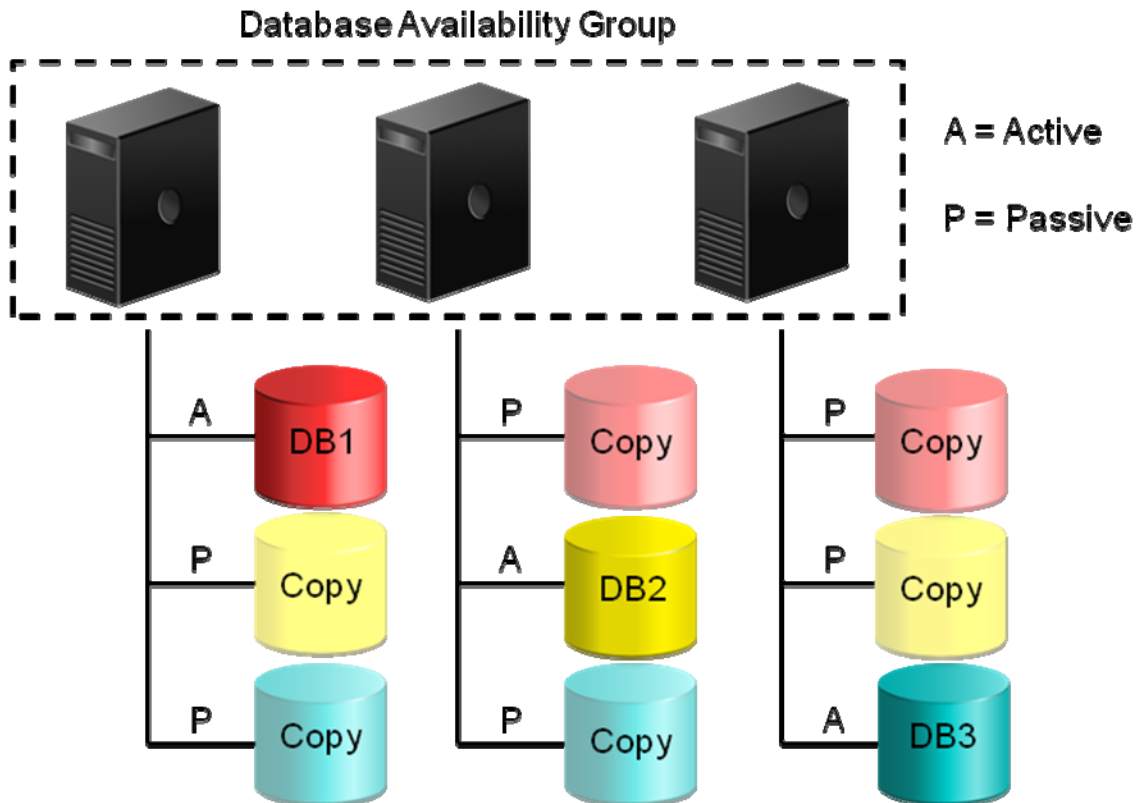


Figure 2. Logical view of a DAG and multiple database copies

DAGs rely on a subset of functionality provided by Windows failover clustering in order to help provide registration and coordination between its members. At its core, a DAG is a node majority based cluster that can include an optional file share witness. The actual quorum model used by a DAG will change automatically depending on the number of server members. If there are an odd number of servers within a DAG, a node majority quorum model will be used. If there is an even number of servers within a DAG a node and file share majority quorum model will be used.

In addition to failover clustering, Exchange Server 2010 introduces a new component called “active manager” that is used to provide Exchange application-awareness into the DAG environment. Active manager replaces what was previously available within Windows failover clustering where Exchange-specific cluster resources would monitor Exchange service availability.

The active manager runs on all mailbox servers configured as members within a DAG. There are two active manager roles, the primary active manager (PAM) and the secondary active manager (SAM). One server operates as the PAM role, specifically the node that owns the default cluster group (which maintains

the DAG node name and IP address(es)). The PAM is the decision maker for reacting to failures and determining which DAG member will operate with the active database copy. The PAM also acts as the SAM for the server on which it runs. The SAM provides detection of local database and information store failures and runs on the rest of the DAG nodes. The SAM reacts to failures by requesting the PAM to initiate a failure for a given database copy or store process.

By combining multiple mailbox database copies, node majority clustering and active manager monitoring for database and process health and availability, Exchange Server 2010 DAGs allow for database, server and storage HA and site resiliency. Additional information regarding native Exchange server 2010 DAG functionality can be found on TechNet under the Exchange Server 2010 heading “High Availability and Site Resilience (<http://technet.microsoft.com/en-us/library/dd638137.aspx>).

AutoStart integration with Exchange Server 2010

Exchange Server 2010 offers an API to allow synchronous replication vendors to integrate with the DAG framework. AutoStart 5.3 SP5 has been updated to include an Exchange Server 2010-specific module to integrate with this API. AutoStart provides support for specific storage-based technologies based on components called data sources. In combination with the AutoStart Exchange Server 2010 module, the following AutoStart data sources are supported:

Shared disk device for Windows data source

A shared disk device for Windows data source is a common LUN that can be shared by multiple Windows hosts. Only one Windows host will have read/write access (defined as “attached” within AutoStart) to the disk at a given time. When a node does not have access to a disk, the device is marked as not ready by the AutoStart filter driver. The not ready state will prevent unwanted access to the LUN for the nodes where the resource is detached. To use the shared disk device for the Windows data source, the disk must be created as a Basic disk. Either MBR or GPT partitioning can be used.

Mount point data source

A mount point data source is simply a LUN that is mounted as a folder within a directory structure. AutoStart supports the use of mount points in combination with the shared disk device data source. The shared disk data source will represent the root of the mount point, while the mount point data source represents the LUN and folder mounted within the root. The startup sequence will be such that the shared disk data source comes online first, followed by the mount point data source.

EMC SRDF/S mirroring data source

SRDF/S provides for synchronous replication between LUNs between multiple Symmetrix® storage arrays. Microsoft support for third-party replication requires that the replication mechanism be synchronous. This limits Microsoft support to SRDF/S.¹ A mount point or shared disk data source should be configured above the SRDF/S data source to provide access control.

EMC MirrorView/S mirroring data source

CLARiiON® MirrorView/S mirroring provides for synchronous or asynchronous replication between LUNs within multiple CLARiiON storage arrays. Microsoft support for third-party replication requires that the replication mechanism be set to synchronous mode. This limits Microsoft support for MirrorView/S to MirrorView/S.¹ A mount point or shared disk data source should be configured above the MirrorView/S data source to provide access control.

¹ Asynchronous replication is not blocked from being configured. A warning will be displayed should SRDF/A or MirrorView/A replication be configured for a data source. An administrator must select and confirm that they will be running in a configuration that Microsoft does not support.

Much of the DAG framework remains the same for when third-party replication is enabled. The DAG will continue to operate as a node majority cluster. AutoStart will support the maximum number of nodes within a DAG of 16 and the maximum number of 100 databases per mailbox server.

A DAG is created in much the same way, with the exception of marking the DAG as having “thirdpartyreplication” enabled. Database copies are created within the DAG like with native Exchange replication; however, the term “copy” takes on a new meaning when AutoStart is used.

When a database copy is added to a DAG configured for third-party replication, no database seeding process will occur and the native log shipping mechanism otherwise used by Exchange will be disabled. A database copy under AutoStart will represent either a shared disk or a copy as replicated with SRDF/S or MirrorView/S. A copy under shared disk functionality will represent the same physical copy (as protected by RAID 1, RAID 5, or RAID 6 within the storage array) of the database across multiple nodes. The only true “copy” in a physical sense will be the copy synchronously replicated by SRDF/S or MirrorView/S to an alternate Symmetrix or CLARiiON storage array, respectively. Figure 3 provides a logical overview for how AutoStart can allow for shared storage as well as remote synchronous storage management within a DAG.

The active managers will also continue to provide failure detection for databases and Exchange-specific processes within a DAG where third-party replication is enabled. While the active managers continue to monitor the environment, they no longer control which database copy should be active. When a failure is detected, the active manager will alert AutoStart, which essentially becomes an extension to the PAM, and will then control where databases will become active. By default AutoStart will first move active database copies to nodes with shared disk device access before moving resources to remote nodes that maintain copies as replicated by SRDF/S or MirrorView/S. AutoStart will also continue to monitor the environment and can initiate movement of active databases between nodes independent of the active managers.

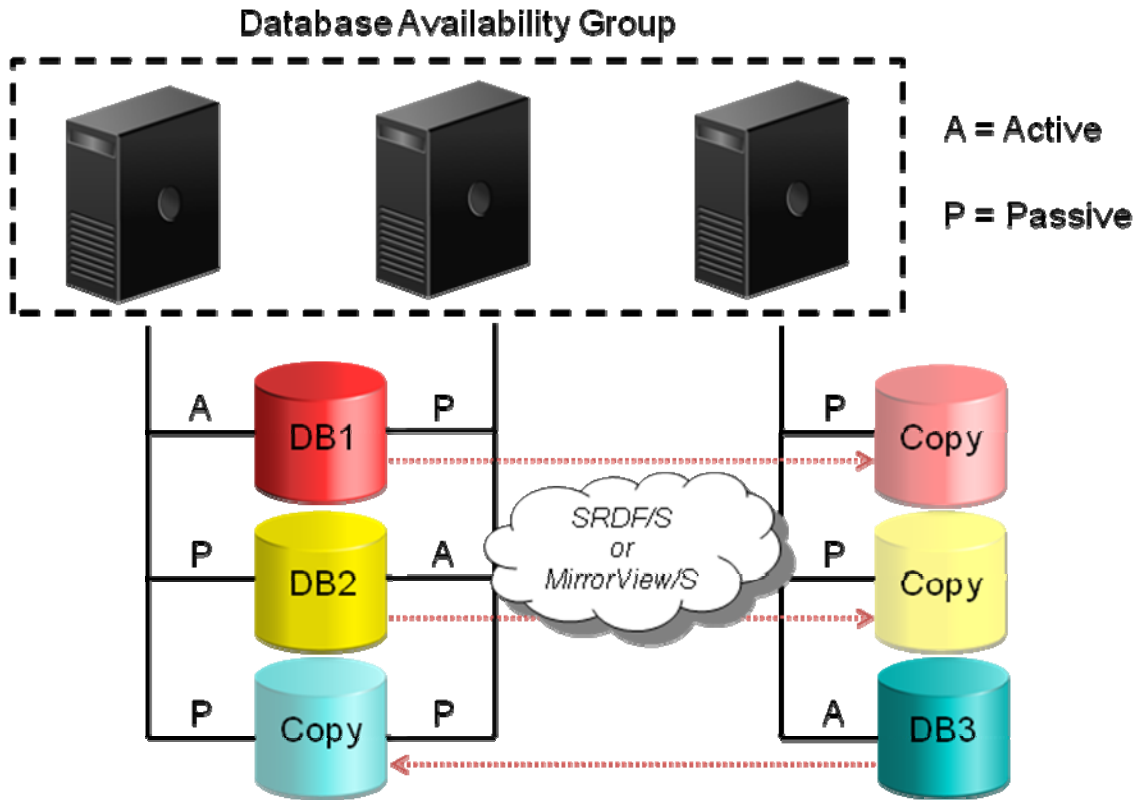


Figure 3. High-level view of AutoStart integration into the DAG framework

Configuration requirements

Prior to configuring the Exchange Server 2010 DAG and AutoStart environment, several requirements should be met. The following headings outline these requirements:

Log and database file placement

Exchange Server 2010 has removed the concept of a storage group as existed for previous Exchange versions and now supports up to 100 discrete database instances (database and log pairs) per server. For support with AutoStart, each database instance must reside on dedicated LUNs as presented to the Windows operating system. AutoStart performs the movement of active databases at the LUN level, therefore to maintain database level granularity for moves between servers, a database must reside on its own LUNs. A database file and log directory combination for a given database can share a LUN or be on separate LUNs.

Storage presentation

AutoStart requires that all LUNs to be used for either shared disk access or replication be zoned, mapped, and masked to the nodes participating in the DAG. As stated previously, shared disk access requires that the LUN be configured as a basic disk with either MBR or GPT partitioning. For the initial configuration of the shared disk resources, each node must mount its respective LUNs to the same mount points on all servers in the DAG.

Replication

For SRDF/S or MirrorView/S data sources to be configured, the replication pairs must first be defined prior to adding the data source within AutoStart. With SRDF/S this entails creating the SRDF/S R1 and R2 relationship. Additionally for SRDF/S, a device group (DG) or composite group (CG) must be created either locally on each appropriate node within the DAG, or configured within Group Name Services (GNS) as maintained globally in the Symmetrix. The requirement to manually create a DG or CG may be changed in a future release of AutoStart; please check release notes for details. Also for SRDF/S, it is recommended to present two unique gatekeepers on all nodes for each configured data source.

As previously noted, for Microsoft support of the third-party replication environment, it must be synchronous replication, thus either SRDF/S or MirrorView/S

Virtualization support

The AutoStart Exchange Server 2010 module and its supported data sources can be installed and run in both VMware ESX and Windows 2008 R2 Hyper-V based virtual machines. AutoStart Exchange Server 2010 instances can contain all virtual machines or contain a mix of virtual and physical machines running within the Exchange Server 2010 DAGs.

For AutoStart support in an ESX server environment, virtual machines must have their database and gatekeeper LUNs configured as raw device mappings (RDM).

For AutoStart support in Windows 2008 R2 Hyper-V environments, virtual machines must have their database and gatekeeper LUNs configured as pass-through disks. Additionally, a special setting is required to be set per virtual machine from the Hyper-V parent partition in order to allow SRDF/S or MirrorView/S data sources to work within the virtual machines. By default, SCSI commands are filtered with Hyper-V, which may prevent important commands used for communication with Symmetrix and CLARiiON arrays from being processed.

A bypass of this filtering is provided with Windows Server 2008 R2 Hyper-V. Full SCSI pass-through must be enabled to allow for appropriate EMC Solutions Enabler discovery options to occur from the virtual machine. Microsoft supports allowing full pass-through of SCSI commands as referenced at [http://technet.microsoft.com/en-us/library/dd183729\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd183729(WS.10).aspx). EMC recommends allowing full SCSI command pass-through only for those virtual machines where it is necessary. For more information, including steps for enabling this setting, please refer to the white paper *EMC Symmetrix with Microsoft Hyper-V Virtualization* available on <http://Powerlink.EMC.com>.

AutoStart installation

Prior to configuring the DAG environment, it is recommended to first install AutoStart. The AutoStart install will require a reboot that may otherwise impact the DAG configuration including the availability of active database copies. When installing AutoStart, ensure to include the Exchange Server 2010 module, the AutoStart console, and the appropriate data sources on each mailbox database server participating in the DAG. The AutoStart filter driver, which includes the shared disk and mount point data sources, is required to be installed. During the installation, ensure all nodes to be running within a DAG are within the same AutoStart domain.

AutoStart also requires an appropriate version of Solutions Enabler to be installed on each node to operate within the DAG. If MirrorView/S is to be used, an appropriate version of Navisphere[®] must be available in the environment. Please see the *EMC AutoStart Compatibility Guide* for the appropriate versions of these required applications on EMC Powerlink[®].

DAG and database copy creation

Prior to configuring AutoStart, the Exchange Server 2010 DAG environment needs to be configured. The first configuration step is to define the DAG itself. For use with AutoStart, the DAG must be created with an attribute named “thirdpartyreplication.” This attribute cannot be set or enabled from the Exchange Management Console (Exchange console); the Exchange shell command “new-databaseavailabilitygroup” must be used. The following is an example of how to use the new-databaseavailabilitygroup command to create a DAG with thirdpartyreplication enabled.

```
new-databaseavailabilitygroup -name DAG1 -databaseavailabilitygroupipaddresses 10.0.0.185 -thirdpartyreplication Enabled -witnessdirectory c:\DAGFSW -witnessserver server185
```

A DAG can be queried with the “get-databaseavailabilitygroup” command, as shown in Figure 4, to ensure that the thirdpartyreplication attribute has been enabled.

```

Machine: LICOC212.MMEX2010.local
[PS] C:\Windows\system32>get-databaseavailabilitygroup | fl

RunspaceId           : f6a619c5-fe5f-49c8-ae7
Name                  : DAG1
Servers               : <LICOC113, LICOC213, E
WitnessServer        : LICOC219.MMEX2010.LOC
WitnessDirectory     : c:\DAGFSW
AlternateWitnessServer :
AlternateWitnessDirectory :
NetworkCompression   : InterSubnetOnly
NetworkEncryption    : InterSubnetOnly
DatacenterActivationMode : Off
StoppedMailboxServers : <>
StartedMailboxServers : <>
DatabaseAvailabilityGroupIpv4Addresses : <10.243.157.185>
OperationalServers   :
PrimaryActiveManager :
ThirdPartyReplication : Enabled
ReplicationPort       : 0
NetworkNames          : <>
AdminDisplayName     :
ExchangeVersion       : 0.10 (14.0.100.0)
DistinguishedName     : CN=DAG1,CN=Database Av
Identity              : DAG1
Guid                  : 42448d1f-281a-49c0-a31
ObjectCategory        : MMEX2010.local/Configu
ObjectClass            : <top, msExchMDBAvailab
WhenChanged           : 3/9/2010 2:10:25 PM
WhenCreated           : 3/9/2010 2:10:25 PM
WhenChangedUTC        : 3/9/2010 10:10:25 PM
WhenCreatedUTC        : 3/9/2010 10:10:25 PM
OrganizationId        :
OriginatingServer     : LICOC180.MMEX2010.local

```

Figure 4. Query of a DAG to see if third-party replication is enabled

It should be noted that once a DAG is created with `thirdpartyreplication` enabled, it cannot be disabled without removing the DAG from the configuration and re-creating the DAG object.

Once the DAG is created, servers running the mailbox database role must be added. This can be accomplished from the Exchange console, or with the “Add-DatabaseAvailabilityGroupServer” shell command, for example:

```
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer LICOC212
```

After the appropriate servers are added, database copies, for active databases owned by servers within the DAG, need to be specified. In the context of AutoStart, database copies should be specified for servers that have access to either a shared LUN, as managed by the shared disk device data source, or access to a LUN as replicated from the active database copy by either SRDF/S or MirrorView/S. Database copies can be added with either the Exchange console or with the Exchange shell. When adding a database copy using either method, the node adding the database copy is expected to have access to mount points that match the same mount points available for the active database.

To add a database copy to a node, the following Exchange shell command can be run:

```
Add-MailboxDatabaseCopy -Identity DB2 -MailboxServer licoc113
```

As previously noted, when a database copy is added to a DAG configured for third-party replication, no database seeding process will occur and the native log shipping mechanism otherwise used by Exchange will be disabled. The database copy status will also be marked as “NonExchangeReplication” when

viewed from the Exchange console, as depicted in Figure 5, or when viewed from the “get-mailboxdatabasecopystatus” shell command

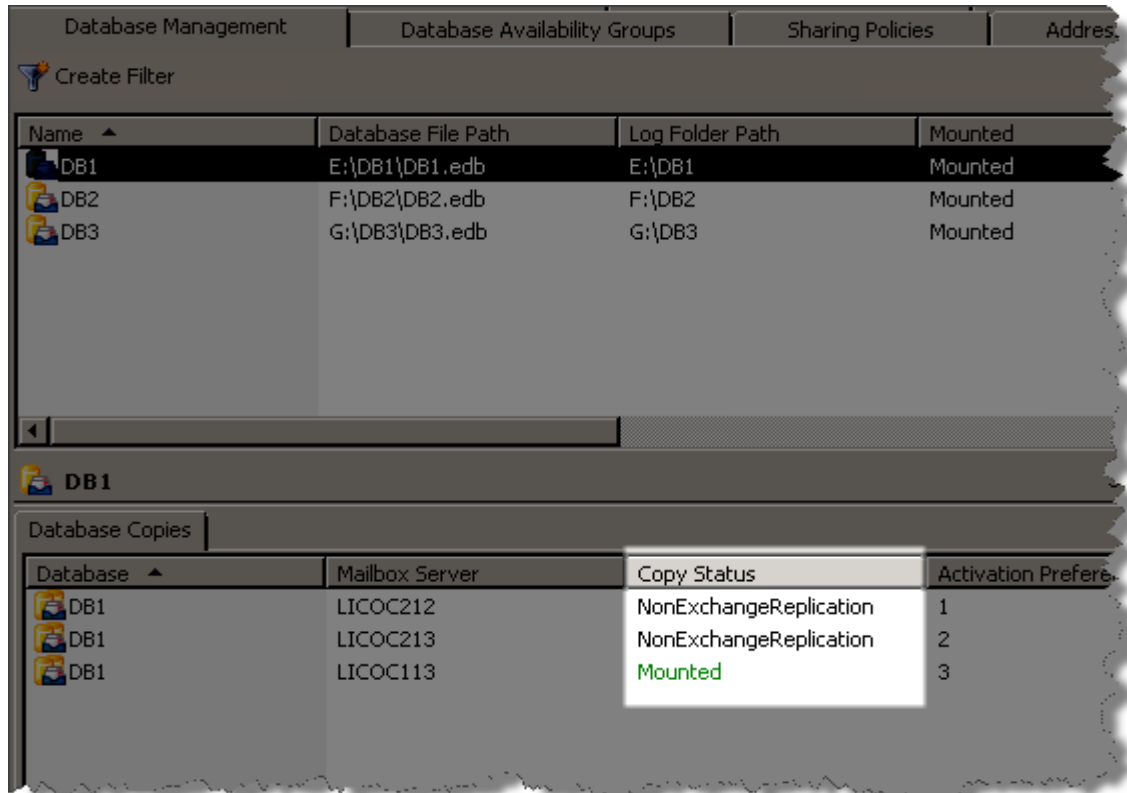


Figure 5. Database copy status when viewed from a DAG with third-party replication enabled

Once these steps have been completed, it is now possible to configure AutoStart to manage the environment.

Required licenses for the Exchange Server 2010 module

Before configuration of the AutoStart Exchange Server 2010 module, make sure that the necessary AutoStart licenses are installed for the configuration. The AutoStart Exchange Server 2010 module will require a new license key. There is no AutoStart-provided upgrade path from Exchange 2007 to Exchange Server 2010. There must be one license for each of the following products, for each node in the configuration. The licenses you must have are:

- One EMC AutoStart 5.3 SP5 license per node
- One EMC AutoStart Exchange Server 2010 module license per node
- If you are using SRDF/S mirroring, then one EMC AutoStart SRDF/S data source license per node
- If you are using MirrorView/S mirroring, then one EMC AutoStart MirrorView/S data source license per node

Creating AutoStart Exchange Server 2010 instances

The AutoStart Exchange Server 2010 module instance ensures high availability for the Microsoft Exchange Server 2010 database copies belonging to a particular DAG. The instance will contain two forms of resource groups:

Node resource group

A node resource group is created for each DAG node and contains a listener process. This resource group ensures that the listener process always runs on the assigned DAG node. The listener process will receive failure notification from the Exchange Server 2010 active managers and will pass the notifications on to AutoStart for further action.

Database resource group

A database resource group is created for each Exchange Server 2010 database as selected by an administrator. This resource group will ensure that the appropriate objects associated with an Exchange database are in a healthy state. The resource group will also define the startup and shutdown order for the resources associated with a database and will maintain the preferred node list for a given database copy. Additionally the resource group is the object from which movement of a database between nodes is initiated. The elements of a database resource group will consist of:

- An InitDb utility process that prepares the Exchange Server 2010 database for mount operations.
- A list of data sources that host the Exchange Server 2010 database. This will include a share disk data source as well as the appropriate SRDF/S or MirrorView/S datasource supporting the physical database copy
- A MountDb utility process that controls the mounts for the Exchange Server 2010 database
- A DbProxy process that reflects the state (mounted/dismounted) of the Exchange Server 2010 database

Figure 12 shows a view of the node and resource groups, including the associated data sources and processes from within the AutoStart console.

AutoStart Exchange Server 2010 Instance Creation Wizard

Important: Prior to continuing, ensure the mailbox databases to be managed by AutoStart are in a state where they can be dismounted.

The creation of the AutoStart instance and the associated resource groups will dismount the active Exchange databases specified within the following process.

To create a new AutoStart Exchange Server 2010 module instance:

1. In the AutoStart console, expand the **Modules** tree.
2. In the **Modules** tree, right-click the AutoStart module for Exchange Server 2010 version 1.0, then select **Create Exchange Server 2010 Instance**, as shown in Figure 6.

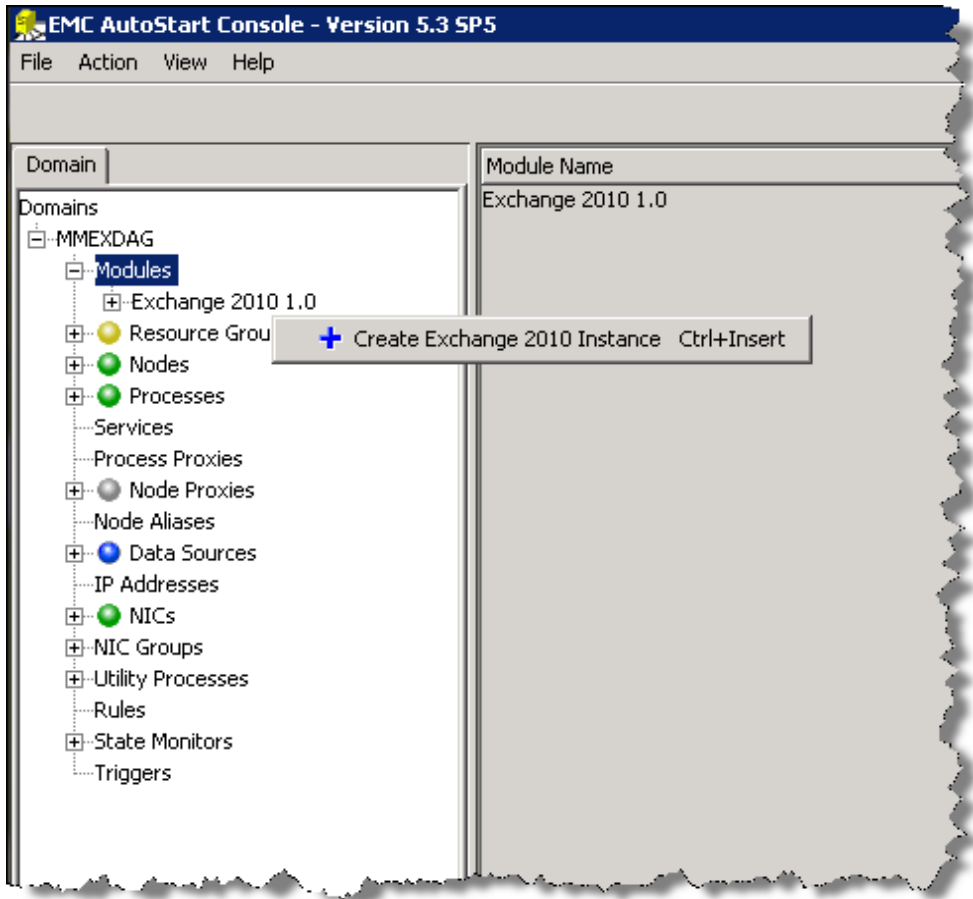


Figure 6. Create an Exchange Server 2010 instance

3. In the Instance Name field (Figure 7), enter a name and description for the new module instance. When considering the module instance represents a particular DAG, naming the instance with reference to the DAG name would be appropriate.

Figure 7. Instance Name field, representative of a DAG

4. Click **Next** to continue to the **Enter domain administrator credentials** dialog box (Figure 8).
5. Enter a Windows domain user account with sufficient privileges to perform Microsoft Exchange administrative operations against the DAG and associated databases:
 - a. In the **Domain** text box, type the name of the Windows domain.
 - b. In the **Userid** text box, type the Windows username.
 - c. In the **Password** text box, enter the password.
 - d. In the **Confirm Password** text box, re-enter the password.

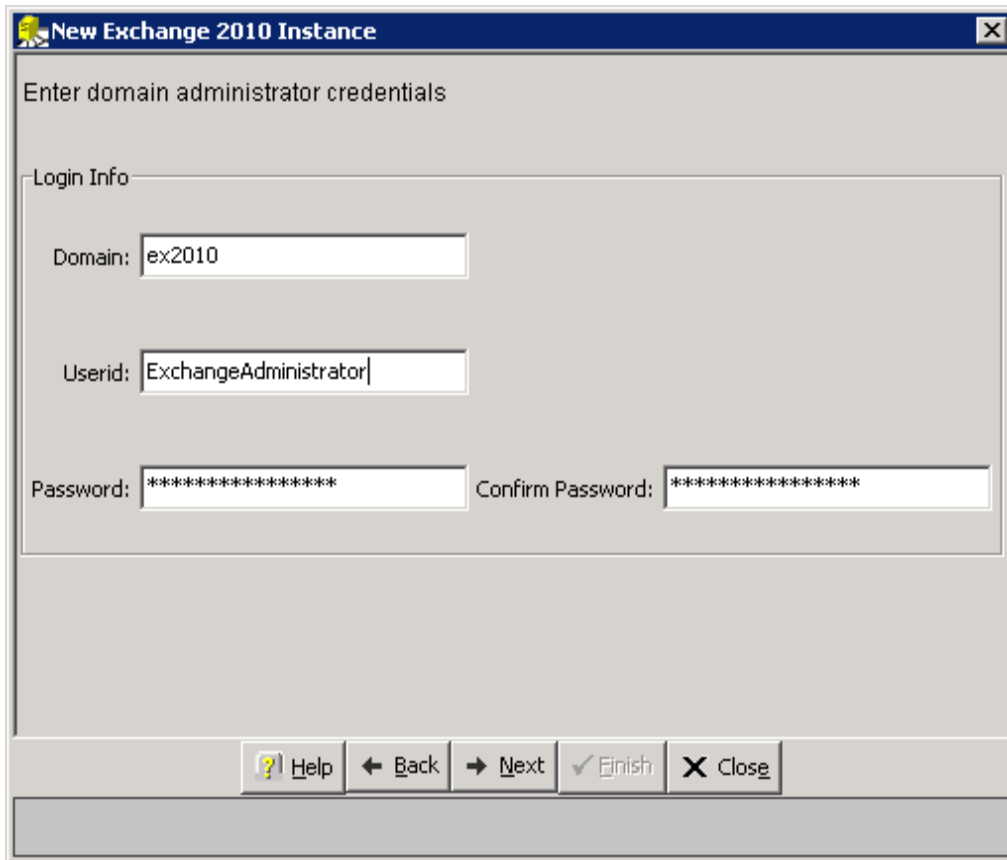


Figure 8. Credentials dialog box

6. Click **Next** to continue. The list of DAGs and their corresponding Exchange Server 2010 databases will be discovered and populated in the **DAG and Database list Configuration** box (Figure 9).
7. Choose the DAG to be associated with the AutoStart instance from the dropdown list. The list of databases associated with the DAG will be populated in the **Available Databases** box. Move the databases as required to be managed by AutoStart into the **Selected Databases** box using the left arrow button.
8. If the discovered storage is EMC CLARiiON , then additional array credentials for MirrorView/S data sources will be required. In this case a dialog box **Enter CLARiiON Login credentials** will appear. Enter the username and password for the CLARiiON array and click **OK**.

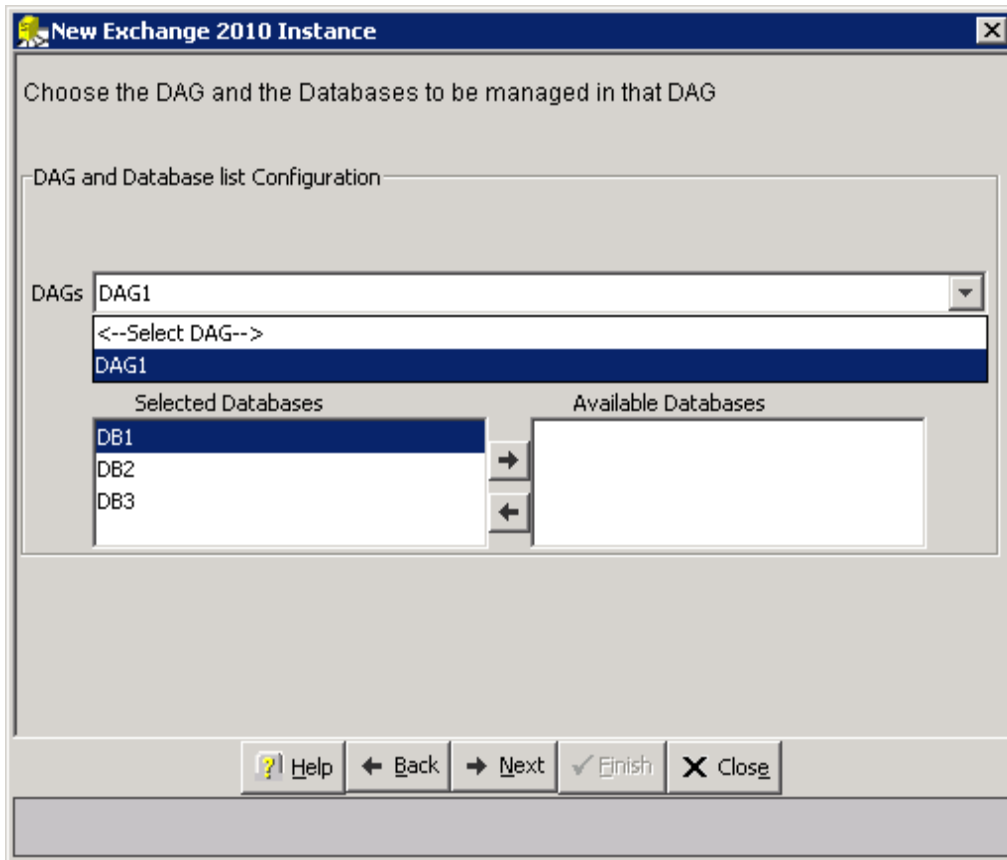


Figure 9. DAG and Database configuration list

9. Click **Next** to continue. The next dialog box will be used to configure data sources associated with the selected database from the previous screen. If AutoStart has been pre-configured with the necessary storage objects where the Exchange Server 2010 databases are residing, the wizard will perform an auto-mapping of the discovered storage to these pre-configured AutoStart data sources. This mapping will be displayed in the **Configure Data Sources** panel. If no pre-configured AutoStart data sources are found that map to the specified Exchange Server 2010 databases then the Auto Create option will be enabled. This feature will automatically create the required AutoStart data sources for supported storage arrays, as displayed in Figure 10.

If the AutoStart Exchange Server 2010 database is created on a SRDF/S LUN, the Auto Create will create a SRDF/S data source with the personality swap feature enabled.

Verify that the device mapping is correct. If any changes are required, move the required data sources from the **Available Data Sources** list to the **Selected Data Sources** list and click **Apply Selection**. If you manually add more than one data source to the Selected Data Sources list, you must sequence them so that they attach in the correct sequence when the Exchange Server 2010 instance starts up, and detach in the correct sequence during a shutdown. Click the up or down arrow to position a data source correctly in the Selected Data Sources list. For example, if you are using SRDF/S, the SRDF/S data source must be above the data source that holds the Microsoft Exchange Server 2010 databases to ensure that the SRDF/S data source attaches before the data source holding the Microsoft Exchange Server 2010 databases attaches.

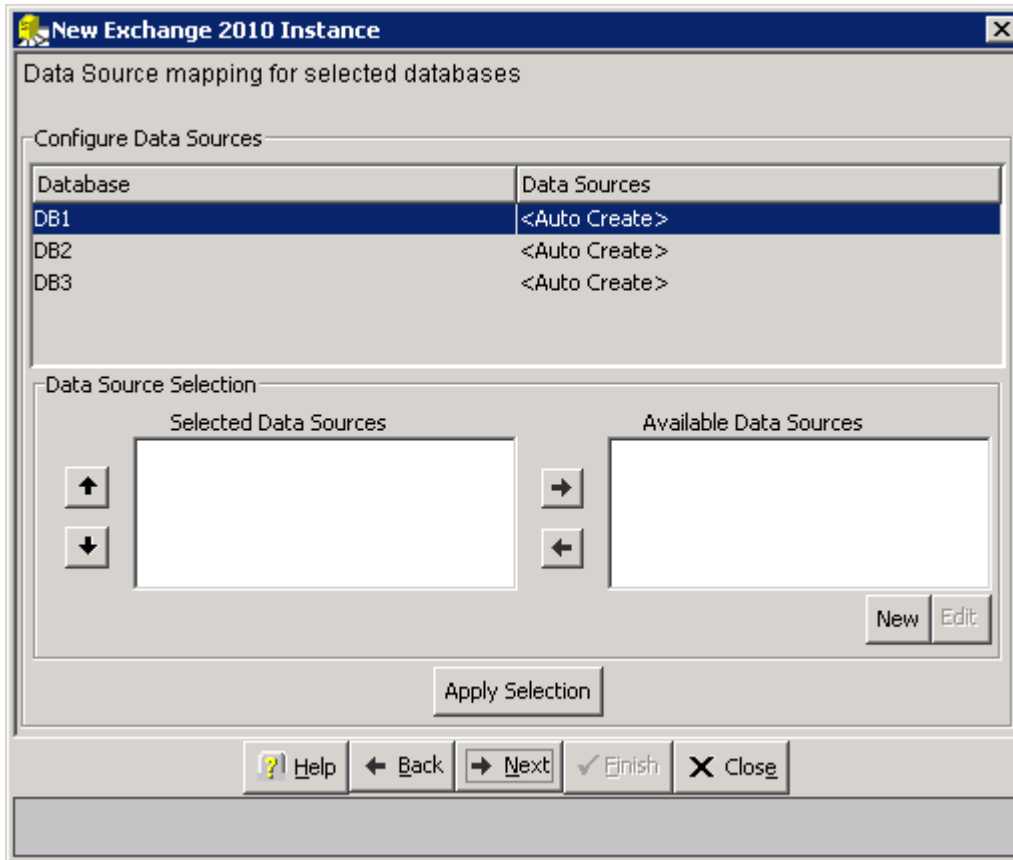


Figure 10. Data Source mapping and creation with Auto Create

10. Click **Next** to continue. The final screen (Figure 11) will allow for a notification e-mail address to be specified for specific events that occur against the Exchange Server 2010 instance. These settings are optional. Select the appropriate notification options and include an e-mail address. If notifications are enabled, ensure an SMTP server has been configured for the AutoStart domain from within the AutoStart console.

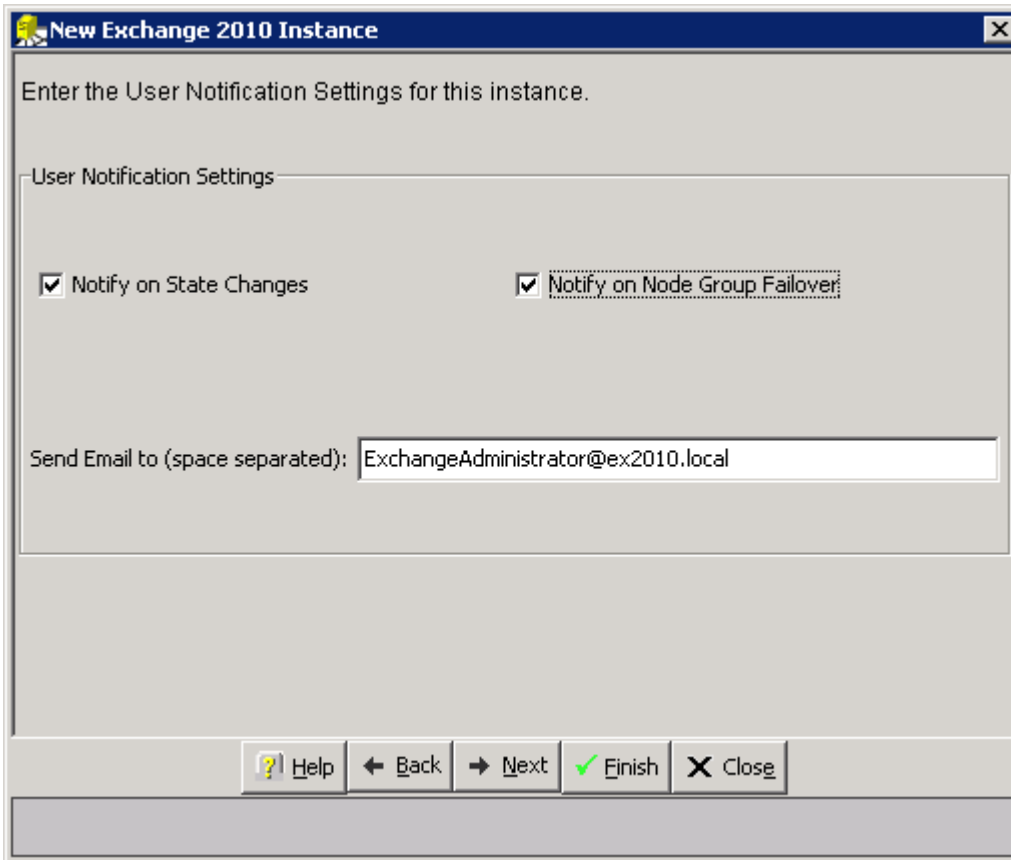


Figure 11. E-mail notification settings

11. Select **Finish** to complete the creation of the instance. As noted previously, the existing Exchange Server 2010 database copies will be dismounted as a part of the process to create the appropriate resource groups. Figure 12 gives an example of what an AutoStart Exchange Server 2010 instance looks like from the Resource View tab.

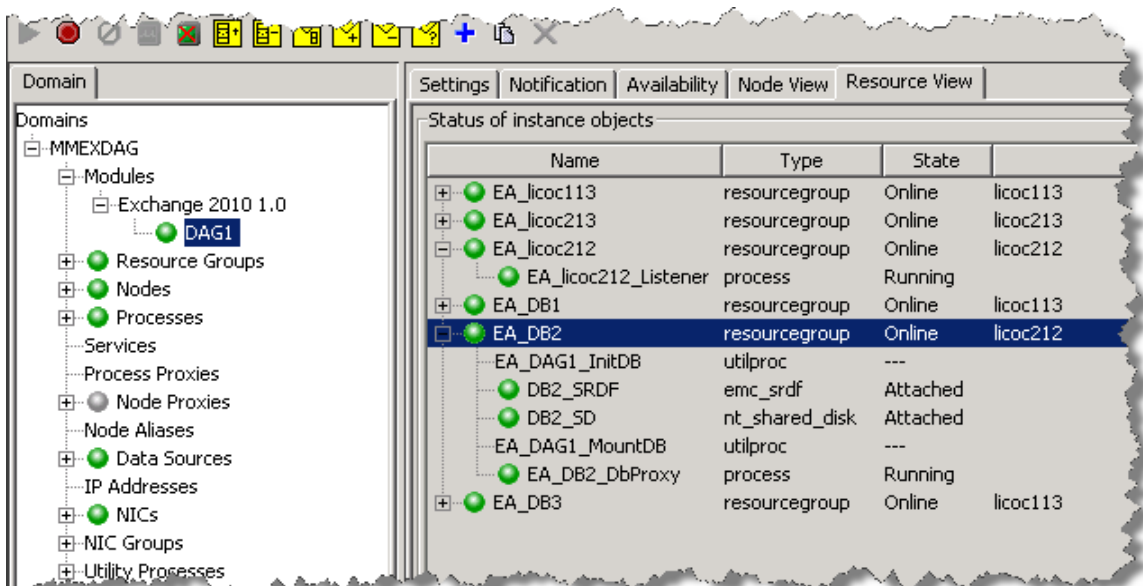


Figure 12. Exchange instance and associated resource groups

DAG and database copy management considerations

Because replication is managed by a third party, in this case SRDF/S or MirrorView/S by AutoStart, certain Exchange Management Console functions and Exchange shell commands will either not function as expected or provide limited information compared to when native replication is used. In other cases, a combination of Exchange commands and AutoStart commands need to be combined to perform specific maintenance tasks. The following sections outline the commands and scenarios for managing the AutoStart Exchange Server 2010 DAG environment.

AutoStart Exchange Server 2010 shell commands

Along with the AutoStart console, specific AutoStart Exchange Server 2010 shell commands are provided. To add the AutoStart Exchange commands for use within the Exchange shell, run *Add-PSSnapin AutoStart.Exchange*.

Note: When adding the AutoStart commands from a non-Exchange PowerShell window, first enter *Add-PSSnapin Microsoft.Exchange.Management**.

The AutoStart shell commands available are as follows:

Dismount-EMCActiveMailboxDatabase

Use the *Dismount-EMCActiveMailboxDatabase* cmdlet to bring an AutoStart resource group managing a Microsoft Exchange Server 2010 mailbox database offline. This operation will result in the Exchange database being dismounted and subsequently unavailable.

Syntax

Dismount-EMCActiveMailboxDatabase -Identity <Database ID Parameter>

Parameters

Parameter	Required	Description
<i>Identity</i>	Required	Use the <i>Identity</i> parameter to specify a name that represents the Exchange database that will be brought offline.

Example

Dismount-EMCActiveMailboxDataBase -Identity DB1

Mount-EMCMailboxDatabase

Use the *Mount-EMCMailboxDatabase* cmdlet to bring an AutoStart resource group managing a Microsoft Exchange Server 2010 mailbox database online on a particular host. This operation will result in the Exchange database being mounted to a specific host.

Syntax

Mount-EMCMailboxDatabase -Identity <Database ID Parameter> -Server <Target Host>

Parameters

Parameter	Required	Description
<i>Identity</i>	Required	Use the <i>Identity</i> parameter to specify a name that represents the Exchange database that will be brought online.
<i>Server</i>	Required	Use the <i>Server</i> parameter to specify where to bring the Exchange database online.

Example

Mount-EMCMailboxDatabase –Identity DB1 –Server LICOC213

Move-EMCActiveMailboxDatabase

Use the *Move-EMCActiveMailboxDatabase* cmdlet to relocate an online AutoStart resource group managing a Microsoft Exchange Server 2010 mailbox between nodes in the DAG. This operation will result in the Exchange database being accessible on the new host specified.

Syntax

Move-EMCActiveMailboxDatabase –Identity <Database ID Parameter> -Server <Target Host>

Parameters

Parameter	Required	Description
<i>Identity</i>	Required	Use the <i>Identity</i> parameter to specify a name that represents the Exchange database that will be relocated.
<i>Server</i>	Required	Use the <i>Server</i> parameter to specify where to bring the Exchange database online.

Example

Move-EMCActiveMailboxDataBase –Identity DB1 –Server LICOC113

Get-EMCMailboxDatabaseCopyStatus

Use the *Get-EMCMailboxDatabaseCopyStatus* cmdlet to retrieve the status, including the replication state, for a particular Microsoft Exchange Database.

Syntax

Get-EMCActiveMailboxDatabaseCopyStatus –Identity <Database ID Parameter> -DetailedStatus

Parameters

Parameter	Required	Description
<i>Identity</i>	Required	Use the <i>Identity</i> parameter to specify a name that represents the Exchange database from which information will be requested.
<i>DetailedStatus</i>	Optional	Use the <i>DetailedStatus</i> parameter to display EMC specific information of the Exchange database resources.

Example

Get-EMCMailboxDatabaseCopyStatus -Identity DB1 -DetailedStatus

Manually moving active database copies between nodes

Moving database copies from the Exchange console or from the “move-activemmailboxdatabase” shell command will not operate against DAGs where third-party replication is enabled. Any manual database copy movement operations need to be performed from the AutoStart console, as shown in Figure 13 or

from the AutoStart “Move-EMCActiveMailboxDatabase” shell command. Moves are performed at the resource group level and are equivalent to moving an active mailbox database copy between nodes.

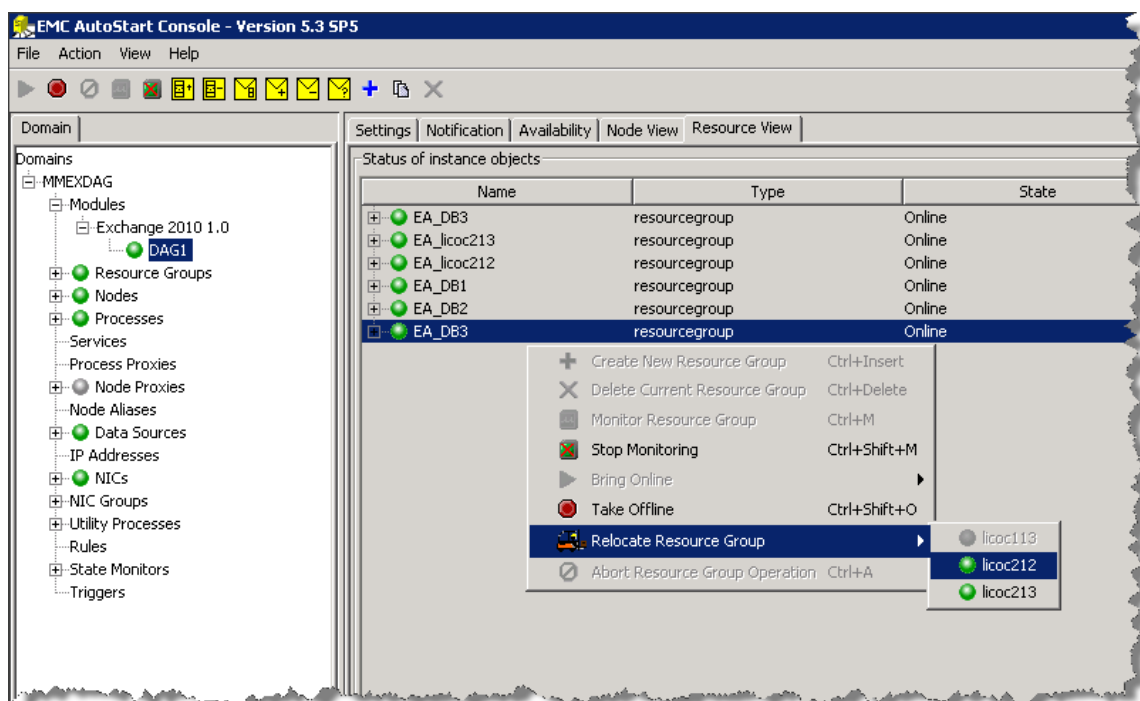


Figure 13. Relocation of an AutoStart resource group representing an active database copy

Controlling automatic database copy activation

Under native DAG replication it is possible to control which DAG members are allowed to automatically take ownership of an active database copy. It should not be expected that any of the native commands to control database activation will take effect where third-party replication is enabled with AutoStart. To give some examples, an individual database can be prevented from moving to a given DAG member with the Suspend-MailboxDatabaseCopy shell command, specifically by adding the “-ActivationOnly” option. A server can also be managed from automatically taking ownership of database copies with the “set-MailboxServer” shell command and specifying the desired “DatabaseCopyAutoActivationPolicy.” Additionally an “activation preference” can be set for each database copy with either the Exchange console or with the “set-mailboxdatabasecopy” shell command. None of these native methods for setting database activation preference should be used when AutoStart and third-party replication is used.

To control database activation and server preferences, native AutoStart methods should be utilized. Primarily the preferred node list can be set for each AutoStart resource group that represents an Exchange database and its associated data sources. As previously discussed, by default AutoStart will configure the preferred node list to first activate database copies on nodes with shared disk device access, followed by nodes where database copies are replicated with SRDF/S or MirrorView/S.

The preferred node list can be manually updated from the AutoStart console from within the Resource Groups tree view. Ensure the **Advanced** user mode is selected, as shown in Figure 14 from under the domain level settings, to view all configured resource groups within an Exchange Server 2010 instance.

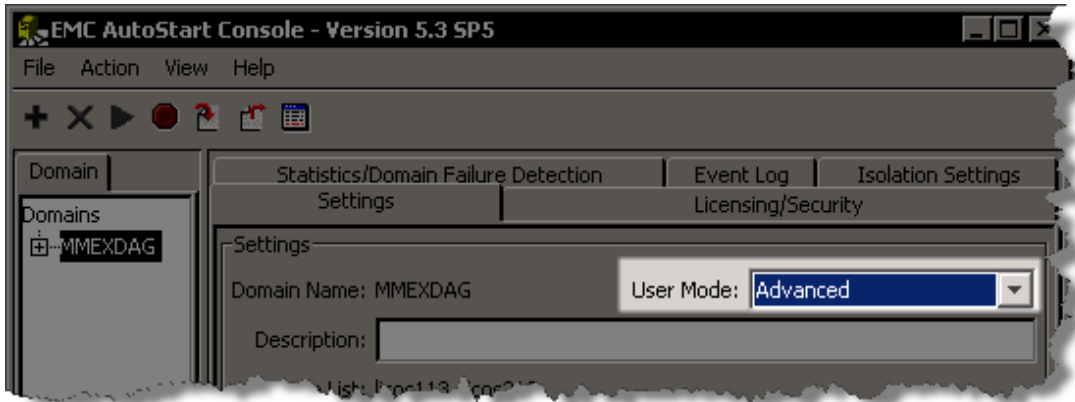


Figure 14. Setting Advanced user mode

Once Advanced user mode is set, the specific resource group can be selected and the preferred node and order can be modified from within the Settings tab, as shown in Figure 15.

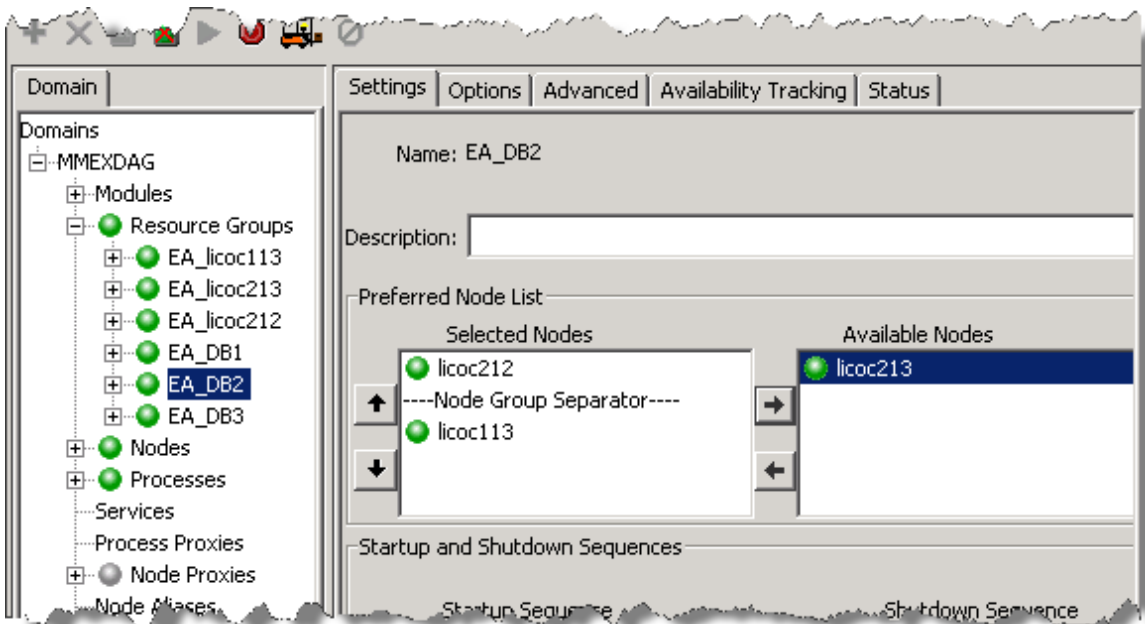


Figure 15. Preferred Node List

Adding or removing nodes for an existing AutoStart Exchange instance

As discussed in the “DAG and database copy creation” section, nodes to be used within an AutoStart instance must first be added to a DAG with native Exchange management commands. Once the nodes are added, AutoStart can be used to configure and manage the node within the instance. This holds true for adding new nodes into an existing DAG and AutoStart instance. Once the nodes are added into the DAG, the **Add Node** option from the AutoStart console, as seen in Figure 16, can be used to discover the newly added DAG member and add it to the instance.

Add node will add the new DAG member to the AutoStart Exchange instance by adding the appropriate node resource group. In order to automatically update the existing database resource groups, the **Update Mailbox Target Nodes** option should be used. Update mailbox target nodes will update the database resource group preferred nodes list based on the configured database copies as defined within the DAG. Whenever database copies are added or removed from members within a DAG, the update mailbox target nodes context menu option should be run.

Removing nodes from an existing AutoStart instance can also be accomplished from the instance context menu using the **Delete Node** option as shown in Figure 16. Prior to removing the node from AutoStart, the server must be removed from the preferred node list on any resources where it is specified. Additionally, the resource group that owns the listener process for that node must be placed offline. Once these steps are done, the node can then be deleted from the instance.

Once the node is removed from the AutoStart instance, it must also be removed from the DAG itself. The removal of the server from the DAG entails first removing any database copies assigned to that node, followed then by removing the node from the DAG. The Exchange Management Console as well as Exchange shell commands like “Remove-MailboxDatabaseCopy” and “Remove-DatabaseAvailabilityGroupServer” can be utilized for these tasks.

Adding or removing databases for an existing AutoStart Exchange instance

Databases can also be added or removed from existing AutoStart instances. The context menu in Figure 16 can be used to accomplish these tasks. When adding a new database, as discussed in the “DAG and database copy creation” section, the active database copy and associated passive copies must already be defined in the DAG environment. Once defined, the database can be added by selecting the **Add Database** option from the context menu. The Add Database option will launch a wizard very similar to the instance creation sequence as outlined in the “AutoStart Exchange Server 2010 Instance Creation Wizard” section. The same processes will be followed as were outlined under the steps from pages 16 through 18. The one difference when adding a database when compared to the create instance wizard is that the DAG will be assumed based on the instance being selected. Only new databases within the previously associated DAG can be added to the selected instance. As previously discussed it is important to understand that the database being added to the AutoStart instance will be dismantled as a part of the resource group creation process.

Removing databases can similarly be performed from the instance context menu. In order to remove a database, the resource group representing the database must first be taken offline. This can be accomplished from the resource group context menu as shown in Figure 13. The act of taking the resource group offline will dismount the associated active mailbox database copy. Once offline, the **Remove Database** option can be selected.

The Remove Database option will delete the appropriate resource group, but will not delete the database object or its associated data files from within Exchange. To delete the database, this must be performed from the Exchange console, or from Exchange shell commands. Also ensure that LUN access to the appropriate devices will be managed once the database resource group is deleted. If the shared disk data source was auto-created as a part of creating the Exchange instance or adding a new database, the data source will be deleted and AutoStart will no longer be managing access to the shared disk resource. If the shared disk data source was manually created it will remain and will not be deleted as a part of the remove database procedure.

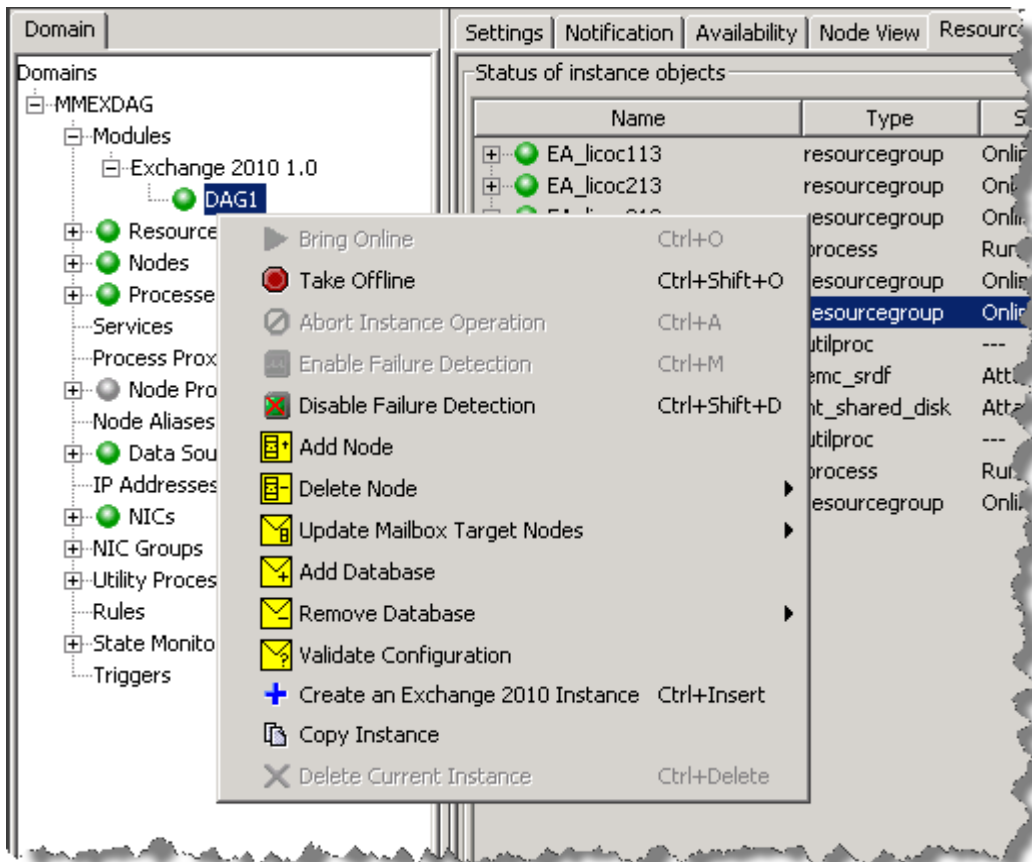


Figure 16. Instance context menu

Querying or controlling replication

Native Exchange management functions and commands to control or query replication will not function as expected when executed against DAGs where third-party replication is enabled. Specifically commands to get the status of replication, either through the Exchange console or with the “Get-MailboxDatabaseCopyStatus” shell command, will only return that “NonExchangeReplication” is being used.

Additionally, suspending replication via the Exchange console or with the “Suspend-MailboxDatabaseCopy” command will not have any effect and will lead to an exit with an appropriate error stating that third-party replication is enabled.

To query replication, native SRDF/S, MirrorView/S or AutoStart-provided commands should be used. To use the AutoStart console to query a data source, highlight the datasource and select the **Status** tab. From within the Status tab the **Get information from the current Data Source** button can be selected. Figure 17 highlights the areas within the AutoStart console where the replication state for an SRDF/S data source can be queried.

Additionally the AutoStart-provided shell command “Get-EMCMailboxDatabaseCopyStatus” can be used to query the state of replication in the environment.

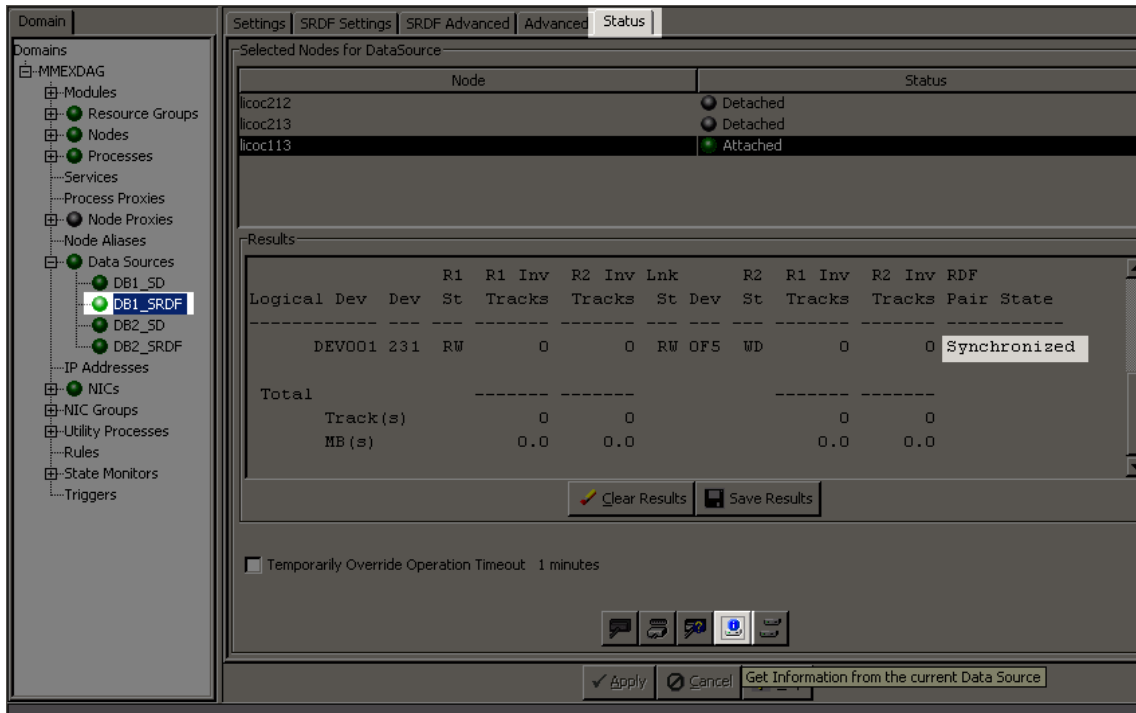


Figure 17. SRDF/S example for querying replication from the AutoStart console

Extensive functionality and fine-grained control of a given environment can be managed from the AutoStart console. The full extent to which AutoStart can be configured is outside of the scope of this document. For additional and more detailed information for configuring AutoStart, please see the *EMC AutoStart Version 5.3 SP5 Administrator's Guide* available on Powerlink.

Unplanned failover when node majority is lost

As discussed in the “Exchange Server 2010 database availability group overview” section, at its core a DAG is a node majority quorum based cluster. If majority is ever lost for a given set of nodes within the cluster, those nodes will be forced offline in order to prevent what is generally known as a “split-brain,” where the same database resources may become active on multiple nodes. If the DAG environment is configured in such a way that a remote data center with active nodes in the DAG cannot maintain majority, or cannot access the file share witness when half the nodes are lost, it should be expected for all nodes in the remote data center to go offline.

In order to recover from the loss of node majority, it must be determined that the primary data center, or previous majority of nodes, are indeed offline and not functioning, and the problem is not just a network outage. Once it is determined that it is appropriate to bring a minority of nodes within the cluster online, the /forcequorum switch must be used to start the cluster service on an appropriate node within the DAG. In order to create a new majority, the nodes no longer functioning within the DAG must be manually evicted and the DAG needs to be reconfigured to remove those members that are no longer functioning. Once failed servers are evicted, the remaining nodes can be brought online to form the new node majority in the cluster. Once the new majority within the cluster is formed, AutoStart can then be used to bring the appropriate database mailbox copies online to the appropriate DAG nodes.

Additional information regarding this scenario can be found under the Exchange Server 2010 section on Microsoft TechNet. Specifically the topic titled “DataCenter Switchovers” (<http://technet.microsoft.com/en-us/library/dd351049.aspx>) should be referenced. Please keep in mind that all steps in the DataCenter Switchover scenario as outlined on TechNet do not need to be executed. Specifically, the steps to manually perform mailbox database switchover using native Exchange Server 2010 commands need not be performed. As discussed above, once node majority is reconfigured via the

eviction and forcequorum process, AutoStart should be used to activate and switch over the specific database copies on the appropriate servers remaining in the DAG. This is accomplished by bringing the database resource group online by using the resource group context menu (Figure 13) within the AutoStart console.

Conclusion

The use of EMC AutoStart in combination with SRDF/S or MirrorView/S replication technologies helps to greatly enhance the usability and process associated with high availability and disaster recovery/business continuity operations. By integrating with the Microsoft Exchange Server 2010 third-party replication API, AutoStart further extends the benefits of the database availability group framework by allowing for shared LUN access and synchronous replication across nodes in both virtual and physical server environments. Integrating directly with Microsoft-supplied interfaces helps to enhance integration and provide for a more reliable experience for administrators utilizing both Microsoft Exchange and EMC AutoStart and storage technologies.

References

White paper

- *EMC Symmetrix with Microsoft Hyper-V Virtualization*

Product documentation

- *EMC AutoStart Administrator's Guide 5.3 SP5*
- *EMC AutoStart Release Notes 5.3 SP5*

Compatibility guide

- *EMC AutoStart Compatibility Guide*