

# Deploying Microsoft Exchange Server 2010 with Hyper-V on the Hitachi Virtual Storage Platform

Implementation Guide

*By Jeff Chen and Leo Nguyen*

*November 2010*



## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). Be sure to include the title of this white paper in your email message.

# Table of Contents

<b>Solution Overview .....</b>	<b>4</b>
Hitachi Virtual Storage Platform .....	5
Hitachi Dynamic Provisioning Software .....	6
Exchange Server 2010.....	7
<b>Tested Solution Components.....</b>	<b>8</b>
Hardware Components .....	8
Software Components.....	8
<b>Solution Implementation.....</b>	<b>9</b>
Configure Fabric Switch Zones .....	10
Configure the Virtual Storage Platform .....	10
Install Exchange .....	14
Create a Mailbox Database.....	14
Create a DAG .....	14
Add a Mailbox Server to the DAG .....	15
Configure the DAG .....	15
Replicate Active Databases .....	16

# Deploying Microsoft Exchange Server 2010 with Hyper-V on the Hitachi Virtual Storage Platform

## Implementation Guide

IT administrators need email solutions that provide data protection and simple management in environmentally friendly data centers. Using Microsoft Exchange Server 2010 with SAN storage like the Hitachi Virtual Storage Platform accomplishes those business-critical objectives.

The Hitachi Virtual Storage Platform can help you leverage your information, which is the new currency in today's data-driven economy. Information exists in many forms and must be protected and readily accessible to ensure business survival and success. The Virtual Storage Platform maximizes cost efficiency and return on investment by creating an agile storage infrastructure that reduces costs and increases performance, availability, scalability and reliability.

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage.

Exchange Server 2010 can be deployed for many types and numbers of users in a wide variety of infrastructure topologies. By implementing Exchange Server 2010 with the Hitachi Virtual Storage Platform storage system, you can effectively scale an environment from a few thousand users to hundreds of thousands of users. Hitachi Dynamic Provisioning software also allows you to combine different workloads in a single storage frame for greater flexibility. Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration.

This implementation guide focuses on deploying Exchange Server 2010. It is based on the reference architecture described in the [Deploying Microsoft Exchange Server 2010 with Hyper-V on the Hitachi Virtual Storage Platform Reference Architecture Guide](#) white paper. It is intended for use by IT administrators responsible for Exchange and storage. It assumes some familiarity with Hitachi Storage Navigator software, Microsoft Windows Server 2008 R2 and Exchange Server 2010.

## Solution Overview

This white paper describes how to deploy an Exchange Server 2010 environment that supports 20,000 mailboxes on a Virtual Storage Platform storage system with the Exchange 2010 servers virtualized using Hyper-V. The Exchange databases are protected by using a Database Availability Group (DAG). The documented environment consists of a single Hitachi Virtual Storage Platform, Hitachi Dynamic Provisioning software, three physical servers, and nine Hyper-V virtual machines (one domain controller, four servers that host the client access and hub transport roles and four mailbox servers). The 20,000 mailboxes have a quota of 1GB each and an I/O profile of 0.12 IOPS, which is the result of sending or receiving 100 messages per day. RAID-5 (7D+1P) with 240 300GB, 10,000RPM SAS disks are used. Figure 1 shows the complete Exchange 2010 solution.

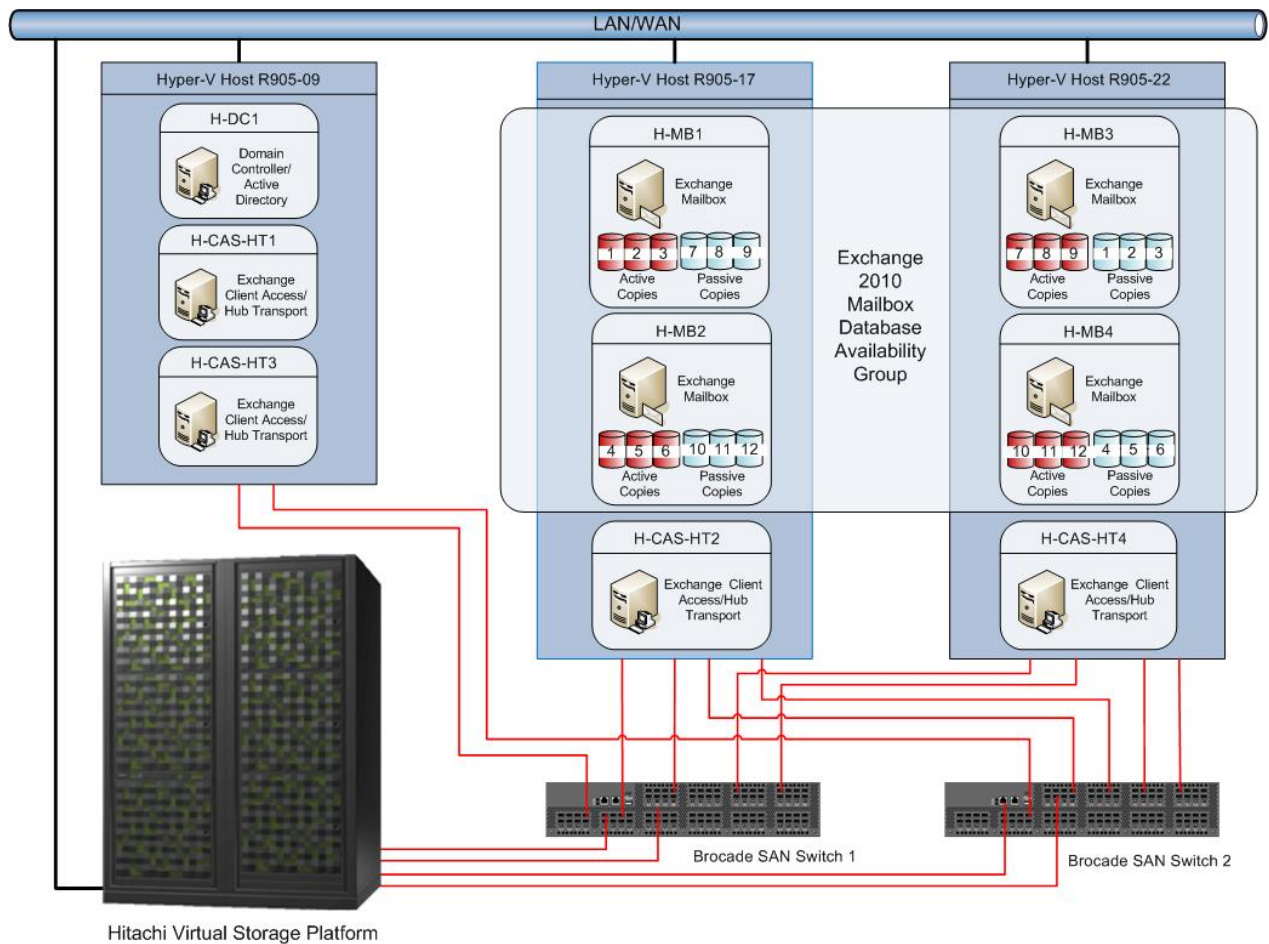
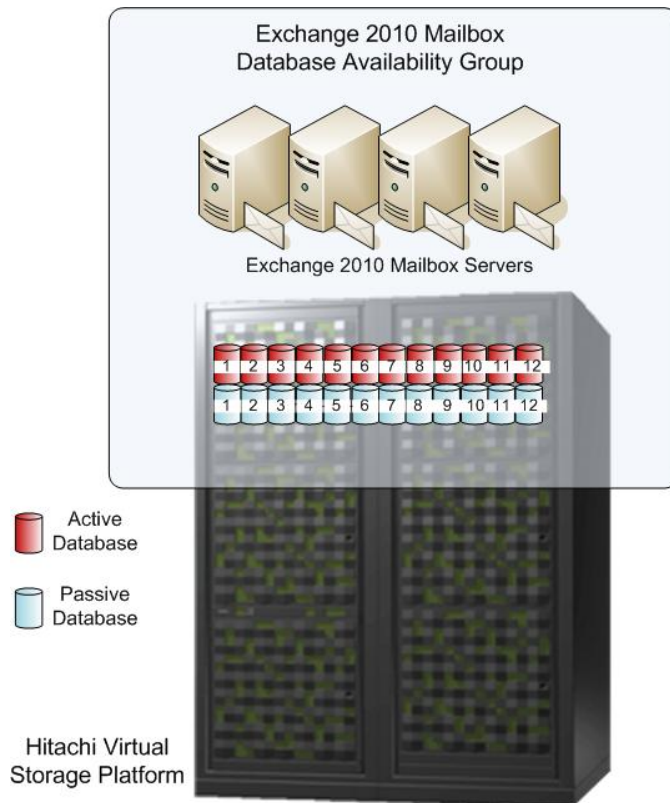


Figure 1

To support database mobility and site resiliency, Exchange Server 2010 introduced Database Availability Groups (DAGs). A DAG is an object in Active Directory that can include up to 16 Mailbox servers that host a set of databases; any server within a DAG has the ability to host a copy of a mailbox database from any other server within the DAG. DAGs support mailbox database replication and database and server switchovers and failovers. Setting up a Windows failover cluster is no longer necessary for high availability; however, the prerequisites for setting up a DAG are similar to that of a failover cluster. Hitachi Data Systems recommends using DAGs for high availability and mailbox resiliency. Figure 2 shows 12 active databases and 12 passive databases placed in a single DAG. Microsoft recommends using one DAG if you are deploying fewer than 16 mailbox servers.



**Figure 2**

## Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the new Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. The Virtual Storage Platform provides virtual storage that meets the growing demands of server virtualization.

The trend in server virtualization is to consolidate the I/O workload of many servers onto a single storage system. As more virtual machines are consolidated onto a physical host, storage systems must be able to dynamically add more storage resources to keep up with I/O demand. The 3D scaling capability of the Virtual Storage Platform meets that requirement.

Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. With the Hitachi Virtual Storage Platform, you can increase performance, capacity and connectivity by adding cache, processors, connections and disks to the base system. A virtual server that accesses the storage system can use all these resources, which act as one system managed as a common pool of resources.

Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. By scaling out you can support increased resource needs in virtualized server environments.

Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform's pool of storage resources. Once virtualized, external data can then be migrated, tiered, replicated and managed by the Virtual Storage Platform. In this manner, older data storage systems can gain a longer useful life. You can extend distance replication for business continuity to lower-cost, lower-function storage systems by virtualizing them behind a Virtual Storage Platform.

The switch matrix architecture of the Virtual Storage Platform makes all of this possible. It connects the basic components, front-end directors, back-end directors, global cache modules and virtual storage directors. You can add redundant pairs of directors and cache modules as required without disruption to connected host servers. All these resources are tightly coupled through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

## Hitachi Dynamic Provisioning Software

On the Virtual Storage Platform, Hitachi Dynamic Provisioning software provides wide striping and thin provisioning functionalities. In the most basic sense, Hitachi Dynamic Provisioning software is similar to the use of a host-based logical volume manager (LVM), but with several additional features available within the Hitachi Virtual Storage Platform and without the need to install software on the host or incur host processing overhead. Hitachi Dynamic Provisioning software provides for one or more pools of wide striping across many RAID groups within a Virtual Storage Platform. One or more Dynamic Provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB (with no initial physical space allocated) are created against each pool.

Primarily, you deploy Hitachi Dynamic Provisioning software to avoid the routine issue of hot spots that occur on logical devices (LDEVs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many RAID groups as members of a striped Dynamic Provisioning pool underneath the virtual or logical volumes seen by the hosts, a host workload is distributed across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning software also carries the side benefit of thin provisioning, where physical space is only assigned from the pool to the DP-VOL as needed using 42MB pages, up to the logical size specified for each DP-VOL. A pool can also be dynamically expanded by adding more capacity or reduced by withdrawing pool capacity. Either operation is performed without disruption or requiring downtime. Upon expansion, a pool can be rebalanced so that the data and workload are wide striped evenly across the current and newly added RAID groups that make up the pool.

Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration, as follows:

- Optimizes or "right-sizes" storage performance and capacity based on business or application requirements.
- Supports deferring storage capacity upgrades to align with actual business usage.
- Simplifies the storage administration process.
- Provides performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.
- Eliminates hot spots across the different RAID groups by smoothing the combined workload.
- Significantly improves capacity utilization.

For more information, see the Hitachi Dynamic Provisioning software [datasheet](#).

## Exchange Server 2010

Exchange Server 2010 introduces several architectural changes that need to be considered when planning deployments on a Hitachi Virtual Storage Platform system.

### *Storage Groups*

Exchange Server 2010 no longer uses the concept of storage group objects. This change from previous versions of Exchange facilitates database mobility and means that Exchange data is now protected at the database level instead of at the server level. This results in a simpler and faster failover and recovery process. Limits on the number of databases per server still exist; the maximum for the standard edition of Exchange Server 2010 is five and the maximum for the enterprise edition is 100.

### *Database Availability Groups*

To support database mobility and site resiliency, Exchange Server 2010 introduced Database Availability Groups (DAGs). A DAG is an object in Active Directory that can include up to 16 Mailbox servers that host a set of databases; any server within a DAG has the ability to host a copy of a mailbox database from any other server within the DAG. DAGs support mailbox database replication and database and server switchovers and failovers. Setting up a Windows failover cluster is no longer necessary for high availability; however, the prerequisites for setting up a DAG are similar to that of a failover cluster. Hitachi Data Systems recommends using DAGs for high availability and mailbox resiliency.

### *Databases*

In Exchange Server 2010, the changes to the Extensible Storage Engine (ESE) enable the use of large databases on larger, slower disks while maintaining adequate performance. Exchange 2010 supports databases up to approximately 16TB, but Microsoft recommends using databases of 2TB or less. When using databases that are larger than 1TB, always configure online database scanning to run continuously by enabling the **Enable background database maintenance (24 x 7 ESE scanning)** option for each database.

## Tested Solution Components

The following sections detail the hardware and software components that were deployed in the Hitachi Data Systems lab.

### Hardware Components

Table 1 lists the detailed information about the hardware components used.

**Table 1. Hardware Components**

<i>Hardware</i>	<i>Detail Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Virtual Storage Platform storage system	Single frame 6x 8GB Fibre Channel ports used 128GB cache memory 240x 300GB, 10K RPM, SAS disks	70-00-50-00/00	1
Brocade 5300 switch	SAN switch with 8Gb Fibre Channel ports	FOS 6.4.0b	2
Dell PowerEdge R905 server	4 x Quad-Core AMD Opteron processor 1.9 GHz, 128GB RAM. Equipped with 2 x Emulex LPe12002 8GB HBAs	BIOS firmware 4.2.1	2
Dell PowerEdge R905 server	4 x Quad-Core AMD Opteron processor 1.9 GHz, 64GB RAM. Equipped with 2 x Emulex LPe12002 8GB HBAs	BIOS firmware 4.2.1	1

### Software Components

Table 2 lists the detailed information about the software components used.

**Table 2. Software Components**

<i>Software</i>	<i>Version</i>
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Provisioning	Dependent on microcode version
Windows Server	2008 R2, Datacenter for Hyper-V hosts, Enterprise for VMs
Microsoft Exchange Server	2010

Table 3 lists the detailed information about the virtual machines used.

**Table 3. VM Servers**

<i>VM Server Name</i>	<i>Hyper-V Host Name</i>	<i>Role</i>	<i>Number of vCPUs</i>	<i>Virtual Memory (GB)</i>
H-DC1	R905-09	Active Directory domain controller	1	2
H-CAS-HT1	R905-09	Exchange client access and hub transport	4	8
H-CAS-HT3	R905-09	Exchange client access and hub transport	4	8
H-MB1	R905-17	Exchange mailbox	4	48
H-MB2	R905-17	Exchange mailbox	4	48
H-CAS-HT2	R905-17	Exchange client access and hub transport	4	8
H-MB3	R905-22	Exchange mailbox	4	48
H-MB4	R905-22	Exchange mailbox	4	48
H-CAS-HT4	R905-22	Exchange client access and hub transport	4	8

This Implementation Guide is based on a lab environment that simulated a specific Exchange 2010 environment. Before implementing a solution in a production environment, consider the size and complexity of the environment and design your Exchange 2010 environment using the E2010 Mailbox Server Role Requirements Calculator Exchange.

## Solution Implementation

To deploy this Exchange Server 2010 solution, follow these high-level steps:

1. Configure fabric switch zones.
2. Configure the Virtual Storage Platform.
3. Install Exchange.
4. Create a DAG.
5. Add a mailbox server to a DAG.
6. Configure a DAG.
7. Replicate active databases.

These are general tasks that need to be completed for a successful deployment. Your checklist might vary based on your environment. Details about each of these steps are included in the following sections.

For more information about each of these high-level tasks, see the following documentation:

- Hitachi Storage Navigator online help
- Hitachi Virtual Storage Platform Dynamic Provisioning Software user's guide
- Microsoft TechNet article "[Deploying Exchange 2010](#)"

## Configure Fabric Switch Zones

Configure zones on your fabric switches according to the manufacturer's guidelines; in addition, follow these best practices:

- Use World Wide Port Name (WWPN) identification for all zoning configuration.
- Connect a minimum of two HBAs per server for multipath high availability.
- Disable all unused switch ports to increase security.
- Configure ports for point-to-point topology.
- Set ports to a specific speed; do not use the auto negotiate setting.
- Use single initiator zoning.

## Configure the Virtual Storage Platform

This solution uses four Dynamic Provisioning pools, each created from multiple RAID-5 (7D+1P) groups on the Virtual Storage Platform. These procedures assume that RAID groups are already created, which is done when Hitachi installs your Virtual Storage Platform. Two pools are used for databases and two are used for logs with active and passive copies intermixed between them.

Hitachi Dynamic Provisioning software is used to avoid the routine issue of hot spots that occur on logical devices (LDEVs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many RAID groups as members of a striped Dynamic Provisioning pool underneath the virtual or logical volumes seen by the hosts, a host workload is distributed across many RAID groups. This provides a smoothing effect that dramatically reduces hot spots. By reducing hot spots in an Exchange environment, you can reduce the need to move mailboxes as the user load varies.

### *Create a Dynamic Provisioning Pool*

To create a Dynamic Provisioning pool using Hitachi Storage Navigator software, follow these steps:

1. Choose **Actions > Pool > Create Pools**.

The **Create Pools** dialog box displays.

2. From the **Pool Type** drop-down menu, choose **Dynamic Provisioning**.
3. Select the **Disable** radio button for the **Multi-Tier Pool** option.
4. (Optional) Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

These options allow you to filter the available pool volumes.

5. Click the **Select Pool VOLs** button.

The **Select Pool VOLs** dialog box displays.

6. Highlight one or more pool volumes in the **Available Pool Volumes** pane and click the **OK** button.

The **Create Pools** dialog box displays with the **Total Selected Pool Volume** and **Total Selected Capacity** fields populated.

7. Assign a prefix for the pool name in the **Prefix** field.
8. (Optional) Assign an initial number for the pool name in the **Initial Number** field.
9. Expand the **Options** pane.
10. Assign a pool id in the **Pool ID** field.
11. Assign a subscription limit in the **Subscription Limit** field.

This is the percentage of oversubscription for this pool that you allow in your environment.

12. Choose a value from the **User-Defined Threshold** drop-down menu and click the **Add** button.

The **User Defined Threshold** value determines when a pool capacity alert is triggered.

The **Selected Pools** pane is populated.

13. Click the **Finish** button.

The **Create Pools** dialog box displays.

14. Click the **Apply** button.

### *Create a Host Group*

Before you create a host group, you must enable port security. Port security is used to isolate traffic from multiple servers and group them in common configurations. Host group members must run the same operating systems and must be connected to the same port. You can define up to 2,048 LUN paths for one host group. You can assign up to 255 host groups and up to 255 WWPNs to one Fibre Channel port.

To create a host group using Storage Navigator software, follow these steps:

1. Choose **Actions > Ports/Host Groups > Create Host Groups**.

The **Create Host Groups** window displays.

2. Assign a name in the **Host Group Name** field.
3. From the **Host Mode** drop-down menu, choose **2C[Windows]**.
4. In the **Available Hosts** pane, highlight one or more hosts.
5. In the **Available Ports** pane, highlight one or more ports.
6. Click the **Add** button.

The **Selected Host Groups** pane is populated.

7. Click the **Finish** button.

The **Create Host Groups** window displays.

8. Click the **Apply** button.

### *Create a LDEV Within a Dynamic Provisioning Pool*

To create a LDEV using Storage Navigator software, follow these steps:

1. Choose **Actions > Logical Device > Create LDEVs**.

The **Create LDEVs** window displays.

2. From the **Provisioning Type** drop-down menu, select **Dynamic Provisioning**.
3. From the **Emulation Type** drop down menu, select **OPEN-V**.
4. (Optional) Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.
5. Click the **Select Pool** button.

The **Select Pool** window displays.

6. Highlight a pool in the **Available Pools** pane and click **OK**.

The **Create LDEVs** window displays with the **Selected Pool Name** and the **Selected Pool Capacity** fields populated.

7. Enter a capacity amount in the **LDEV Capacity** field and choose a unit of measure from the drop-down menu.
8. Enter the number of LDEVs of that size to be created in the **Number of LDEVs** field.
9. In the **LDEV Name** pane, assign a prefix in the **Prefix** field and assign an initial number in the **Initial Number** field.
10. Expand the **Options** pane.
11. Review the value in the **LDKC** field.

Modify the LDKC value if the default of 00 is not appropriate. This is most often the case if the storage will be configured with more than one LDKC.

12. Choose a value from the **CU** drop-down menu.
13. Choose a value from the **DEV** drop-down menu.
14. (Optional) Choose a value from **Interval** drop-down menu.

Leave this value at the default of 0 for sequential numbering of LDEVs. If you want a different numbering sequence, choose a different value.

15. Review the default values in the **Initial SSID** field, the **CLPR** field and **Processor Blade** field.

In most situations, use the default values. Change them only if your environment requires different values.

16. Click the **Add** button.

The **Selected LDEVs** pane is populated.

17. Click the **Finish** button.

The **Create LDEVs** window displays.

18. Click the **Apply** button.

### *Map an LDEV to a Host Group*

To map an LDEV using Hitachi Storage Navigator software, follow these steps:

1. Choose **Actions > Logical Device > Add LUN Paths**.

The **Add LUN Paths** window displays.

2. In the **Available LDEVs** pane, highlight one or more LDEVs.

3. Click the **Add** button.

The **Selected LDEVs** pane is populated.

4. Click the **Next** button.

The **Add LUN Paths** window displays.

5. In the **Available Host Groups** pane, highlight one or more host groups.

6. Click **Add**.

The **Selected Host Groups** pane is populated.

7. Click **Next**.

The **Add LUN Paths** window displays.

8. Click the **Finish** button.

The **Add LUN Paths** window displays.

9. Click the **Apply** button.

## Install Exchange

To install Exchange on the mailbox servers, follow Microsoft's instructions in the "[Deploying Exchange 2010](#)" section of the Exchange 2010 server library on TechNet. For any Exchange management tasks listed in this guide, you can use either the Exchange Management Shell or the Exchange Management Console. This guide uses the Exchange Management Shell.

## Create a Mailbox Database

To create a mailbox database named DB1, run the following cmdlet:

```
New-MailboxDatabase -Name <"Database Name"> -EdbFilePath <full path to database folder> -LogFolderPath <full path to log folder>
```

For the implementation described in this white paper, the following parameters were used:

```
New-MailboxDatabase -Name "DB1" -EdbFilePath E:\DB1\DB1.edb -LogFolderPath K:\Log1
```

This cmdlet mounts the newly created database.

To mount the newly created database DB1, run the following cmdlet:

```
Mount-Database -Identity <"Database Name">
```

For the implementation described in this white paper, the following parameters were used:

```
Mount-Database -Identity "DB1"
```

These commands must be performed to create all of the active database copies. After you create all of the active database copies, you can add them to a DAG.

## Create a DAG

To create a DAG, run the following cmdlet:

```
New-DatabaseAvailabilityGroup -Name <DAG Name> -WitnessServer <Witness Server Name> -WitnessDirectory <full path to witness directory> -DatabaseAvailabilityGroupIPAddresses <DAG IP address>
```

For the implementation described in this white paper, the following parameters were used:

```
New-DatabaseAvailabilityGroup -Name DAG1 -WitnessServer H-CAS-HT -WitnessDirectory C:\E2K10DAGFSW -DatabaseAvailabilityGroupIPAddresses 172.17.169.114
```

## Add a Mailbox Server to the DAG

To add a mailbox server to DAG1, run the following cmdlet:

```
Add-DatabaseAvailabilityGroupServer -Identity <DAG Name> -MailboxServer <Mailbox Server Name>
```

For the implementation described in this white paper, the following parameters were used:

```
Add-DatabaseAvailabilityGroupServer -Identity DAG1 -MailboxServer H-MB1
```

The following occurs when the mailbox server H-MB1 is added to DAG1:

- A failover cluster is created using the name DAG1.
- A cluster network object (CNO) is created in default computers container.
- The name and IP address of DAG1 are registered as a Host (A) record in DNS.
- The server is added to the DAG1 object in the Active Directory.
- The cluster database is updated with information on the databases that are mounted on the added server.

## Configure the DAG

After you add the mailbox servers to the DAG, you can configure additional properties, such as DAG networks for MAPI client access and replication.

To create a network for MAPI client access for DAG1, run the following cmdlet:

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup <DAG Name> -Name <DAG Network Name> -subnets <IP address/subnet> -Repl icationEnabled: $True
```

For the implementation described in this white paper, the following parameters were used:

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup DAG1 -Name DAGNetwork02 -subnets 172.17.169.0/24 -Repl icationEnabled: $True
```

To create a network for replication for DAG1, run the following cmdlet:

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup <DAG Name> -Name <DAG Network Name> -subnets <IP address/subnet> -Repl icationEnabled: $True
```

For the implementation described in this white paper, the following parameters were used:

```
New-DatabaseAvailabilityGroupNetwork -DatabaseAvailabilityGroup DAG1 -Name DAGNetwork01 -subnets 10.0.0.0/8 -Repl icationEnabled: $True
```

## Replicate Active Databases

After you configure all the DAG properties, you can replicate the active databases on each server to the other DAG members. To add a database copy to a secondary server named H-MB2, run the following cmdlet:

```
Add-MailboxDatabaseCopy -Identity <Database Name> -MailboxServer <Mailbox Server Name> -ActivationPreference 2
```

For the implementation described in this white paper, the following parameters were used:

```
Add-MailboxDatabaseCopy -Identity DB1 -MailboxServer H-MB2 -ActivationPreference 2
```

The `ActivationPreference` parameter is used to determine which passive copy should be activated in a failover or switchover situation. If the parameters `ReplayLagTime` and `TruncationLagTime` are not specified in the command, such as in the example above, those parameters are set to the default value of 0 seconds. The parameters specify the amount of time the replication service waits before replaying log files that were copied to the database and how long the replication service waits before truncating log files that were replayed, respectively.

The DAG uses a new component in Exchange called Active Manager. Active Manager facilitates failover and failback. In the event of a failure, Exchange 2010 promotes one of the passive copies of the database to active status.

For more information about the cmdlets listed in this guide and high availability and site resilience in Exchange 2010, see the Microsoft TechNet article "[High Availability and Site Resilience](#)."

 **Hitachi Data Systems Corporation**

---

*Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.*

*Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation*

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-064-00 November 2010

**Corporate Headquarters**

750 Central Expressway,  
Santa Clara, California 95050-2627 USA  
[www.hds.com](http://www.hds.com)

**Regional Contact Information**

**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)