

January 2006



Overview of the Installation of a Microsoft Exchange Server Cluster

*By Peter Wang
IBM Systems and Technology Group*

Abstract

This document presents an overview of the process for creating a Microsoft Exchange Server cluster. The intent of this document is not to explain every detailed step involved in creating an Exchange Server cluster, but rather to provide a general guide for setting up Exchange Server in a cluster environment. The information presented is based on our experience in the IBM® eServer™ xSeries® Performance laboratory with setting up and running an Exchange Server cluster on an IBM eServer BladeCenter® configuration. This paper also describes the test environment and presents the results obtained. FOR QUESTIONS RELATED TO THE OPERATION OF MICROSOFT PRODUCTS, PLEASE CONTACT MICROSOFT DIRECTLY.

Contents

Abstract..... 2

Introduction 4

Building the Cluster 4

Test Environment 9

Test Results..... 12

Conclusion..... 12

For More Information 13

Introduction

A cluster is a group of independent servers, called *nodes*, working collectively as a single server to provide high availability, fail over, manageability, and scalability for resources and applications. Server clusters allow client access to applications and resources in the event of failures and planned outages. If one of the servers in the cluster is unavailable because of a failure or maintenance requirements, resources and applications move to other available cluster nodes. When a failed node comes back online, the cluster service can also automatically reassign the workload in a cluster to its predetermined preferred owner. The cluster is managed as a single system; that is, applications are managed as if they were running on a single server. Cluster services can grow to meet increased demand. Additional nodes can be added to a cluster to increase the cluster's capabilities dynamically.

The Exchange cluster described in this paper was created using Microsoft Cluster Service. Every node has a local Exchange installation with a unique configuration for every cluster node. Exchange Virtual Server (EVS) group(s) will create on the cluster with Exchange installed. An EVS is the logical node that will be used for Exchange cluster operations. An EVS contains an IP address, network name, Exchange executable file and data disk(s), and services for Exchange.

An Exchange cluster can be set up as active/active or active/passive. Both Windows Server 2003 Enterprise Edition and Datacenter Edition support a maximum of eight nodes. In an active/passive cluster, if one node in the cluster fails, the active cluster fails over to the passive node, which becomes active. This is called *failover*. When the failed node is back online, a fail back can be manually or automatically configured in the cluster group properties.

Every cluster node must have two network interfaces. One network interface for the cluster communication ("heartbeat") is called *private*. A heartbeat is much like a *ping*, which can be used to test if the node is still available. If the heartbeat fails, the failover process occurs. The other network interface, called *public*, is used for client access and other network traffic.

Building the Cluster

Windows Server 2003 Cluster Installation

This section provides an overview of how to set up a two-node cluster. The basic components required for an Exchange cluster are listed below:

- Microsoft Windows Server 2003 Enterprise Edition (up to eight nodes) and the latest service pack.
- Hardware on cluster HCL.
- Microsoft Cluster Service installed.
- Microsoft Distributed Transaction Coordinator (MSDTC) installed.
- Shared disks that are physically attached to a shared bus and accessible by all nodes in the cluster. All partitions on the shared disk must be configured as basic not dynamic and formatted as NTFS.
- Only physical disks can be used in the cluster and preferably should be configured as RAID-5.

In the tested cluster, shared disks are all RAID-5 on Fibre Channel with Windows Server 2003 cluster as host type.

Windows Server 2003 Enterprise Edition SP1 is installed on three blade servers before the Cluster Service is activated. The first server is then set up as domain controller with domain function as Native Windows Server 2003 with only one of the two network interfaces activated.

The other two servers are set up with two network interfaces, one private and one public, before joining the domain.

Since the processes to set up a cluster are not trivial, it is a good idea to create drive images at this point. You can also find detailed information in the "Guide to Creating and Configuring a Server Cluster under Windows Server 2003" at the Microsoft Web site:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;821833>

Following is a brief procedure for setting up a Windows Server 2003 cluster:

1. From either node, log on as domain administrator. Open *Administrative Tools* and click on *Cluster Administrator*. Choose *Create a new cluster* from the dialog box. Go through the Wizard.
2. Provide a unique NetBIOS name for the cluster (e.g., cluster1).
3. Choose the computer name that will be the first node of the cluster. The current server you are on should be listed. Then click the *Advanced* button. Select *Advanced (minimum) Configuration* in order to manually set up a quorum disk later.
4. The *Analyzing Configuration* window will start. This process may take several minutes or up to hour if you have more Fibre Channel logical drives. You will be asked to type an IP address for the cluster. Enter a unique IP address in the domain; it should be different from the physical IP address. The Wizard will associate this IP address with the public network by default.
5. You will be asked to create a user name and password for the Cluster Service account. Select the correct domain name in the Domain drop-down list.
6. Review the *Summary* page and click the *quorum* button to select the correct quorum disk (on Fibre Channel, a three-disk RAID-5). *Creating the Cluster* will start. It may take several minutes or up to hour if you have more Fibre Channel logical drives. Click *Finish* to complete the process. Bring the newly created cluster group online if it is not yet active. Right click the group and click *Bring on line*.
7. Open *Administrative Tools* and click *Cluster Administrator*; the newly created "Cluster1" should be online by now. Click the *Cluster Groups* in the Groups. Three resources should appear on the right panel: *IP address, Name, and Physical disk (quorum)*. Now you need to add the additional node to cluster1.
8. Add node 2 to this cluster by clicking *File, New, and then Node*. Enter the machine name you want to add to the cluster. Click *Add*. Type the password for the account used to start the cluster service. Review the *Summary* and click to *Next* to finish the process.
9. Start *Cluster Administrator*. Right click the cluster1 group and select *Properties* and add node 2 to the preferred owner list in the *General* tab. Change to *Allow fail back immediately* under the *Failover* tab. Right click on the network name from the right panel and select *Properties*; check *DNS Registration must Succeed* and *Enable Kerbeos authentication* buttons just under the name field from *Parameter* tab.
10. Start *Cluster Administrator*. Click *Networks*, right click *Private* then *Properties*, Check *Internal cluster communications only (private network)*; then click *OK*. Right click *Public* then *Properties*, and check *Client access only (public network)*; then click *OK*.
11. Click *Cluster Group*. Right click for *New Resource* in the right panel. Enter a name as *DTC* and select *Microsoft Distributed Transaction Coordinator* for the type from list. Two nodes should be listed as preferred owner and add IP Address, network name, and physical disk (quorum drive) the group currently listed resources to the dependence list. (I don't understand this sentence.)

Exchange Server Installation

Before installing Exchange, ensure that all the cluster nodes have Cluster Service installed properly. If Cluster Service is not installed and running on each node in the cluster before installation, Exchange Server 2003 Setup cannot install the cluster-aware version of Exchange Server 2003.

Exchange Server must be installed on the same drive across all the cluster nodes, and the drive must be in NTFS format. For example, the operating system is installed on C:\ drive, which is the internal SCSI disk in the blade server used for the testing described in this document, and Exchange Server is installed on D:\exchsrvr, which is a shared disk (three-disk RAID-5).

Following are the basic steps for setting up Exchange Server:

1. Prepare Active Directory for Exchange Server 2003 - ForestPrep

Running ForestPrep is required only if you are installing Exchange Server 2003 for the first time in your organization. If you have already installed Exchange Server 2003 in your organization, you can skip this. To start, you need to log on to the domain in which the schema master of your forest resides, typically the root domain from the first node. The account you use must be a member of the following Windows Server 2003 security groups: Domain Admins, Schema Admins, and Enterprise Admins. You can create a new Organization by entering a new name during this step. This step needs to be done only once on the first node. Do not repeat this step on the rest of nodes in the cluster.

2. Prepare Active Directory for Exchange Server 2003 - DomainPrep

You must run DomainPrep in each Windows domain in which you want to install Exchange Server. Before you can run the DomainPrep process, the replication of the schema updates by the ForestPrep process must finish. To start, you need to log on to the domain from which you want to run DomainPrep on node 1. The account you use must be a member of the Domain Admins Windows security group. This step should be done only once on the first node; do not repeat this step on the rest of nodes in the cluster.

3. Install Exchange 2003 Server on Each Node

Before installation, ensure that the Cluster Service is running on all the nodes. And a Distributed Transaction Coordinator resource was created for the cluster. Select Exchange Deployment Tools and then First Exchange Server installation from the Installation Wizard. In the Component Selection dialog box, verify that the action next to Microsoft Exchange Server 2003 is Typical. Change the installation location to the drive for Exchange executables (D:\exchsrvr as tested for this document). When you install Exchange Server 2003 in a cluster node, Setup Wizard prompts you with a dialog box that indicates that the cluster-aware version of Exchange Server 2003 will be installed. In this dialog box, click OK to proceed with Setup. Apply Service Pack 1 and then restart the server. Repeat the installation process on the second node of the cluster by selecting Additional Server Installation from Installation Wizard. For a four-node Exchange Server cluster, you need to install Exchange in the same location on each of the four servers.

4. Create Exchange Virtual Server

The number of EVSs you need to create depends on whether you are creating an active/passive or active/active Exchange Server cluster. If you are setting up a two-node active/passive Exchange Server cluster, you will set up one EVS. You only need to perform the steps below once. If you are setting up a two-node active/active Exchange Server cluster, you will set up two EVSs, one on each node. Each EVS consists of a static IP address, a unique Network Name, at least one shared physical disk, and an Exchange System Attendant

resource. For a high-availability Exchange Server, an active/passive cluster is recommended since active/active cluster does not provide any failover protection, and an active/active cluster generates an extra layer of complexity (cluster environment) for administration.

a) On the first node of the cluster, click Start, point to Programs, point to Administrative Tools, and then click Cluster Administrator. If prompted to specify a cluster, enter the cluster name, or browse and select the cluster in which you want to create an EVS. Right-click the Groups container, click New, and then click Group. The New Group Wizard starts. In the Name box, enter a name for this EVS, and then click Next.

b) In the Preferred Owners dialog box, verify that there is either one or no cluster nodes listed in the Preferred Owners box, and then click Finish. This new EVS (cluster group) is displayed under Groups. For an active/passive setup, ensure that the order of the list in the Preferred Owners dialog box is opposite. For example, if the preferred owners list on the first node lists node1 and then node 2, ensure the second node lists node 2 and then node 1. For active/active setup, only this one node should be listed.

c) For active/passive setup, right click EVS just created and click properties. Check allow fail back immediately under the failover tab or specify how long to wait before the fail back proceed.

5. Create an IP resource

This step is to create an unique static IP address for EVS each individual EVS created and need to repeat on each active EVS.

a) Right-click the EVS group, point to New, and then click Resource. The New Resource Wizard starts.

b) In the New Resource dialog box, enter the IP Address and give a name. In the Resource Type box, choose IP Address, and then click Next. For the active/passive setup, select the node on which this EVS resides and the node to which it will fail over in the cluster in the Possible Owners dialog box, and then click Next. For the active/active setup, remove all other nodes from the list.

c) In the Dependencies dialog box, do not select any resources that appear in the Resource Dependencies box, and then click Next. In the TCP/IP Address Parameters dialog box, in Address, enter the static IP address. Note that it is highly recommended that the Exchange Server cluster have its own dedicated static IP address, separate from all other resources (including the quorum) that are defined in Cluster Administrator. In Subnet mask, verify that the subnet mask for the EVS is correct. In Network, verify that the "Public Network Connection Name" is selected, and then click Finish.

6. Create a network Name Resource

This step is to create a unique NetBIOS name for each EVS created. Do this for each active EVS.

a) Right-click the EVS, point to New, and then click Resource. New Resource Wizard starts.

b) In the New Resource dialog box, enter the Network Name for this EVS in the Name box. In the Resource Type box, click Network Name, and then click Next. For the active/passive setup, enter the node on which this EVS resides and the node to which it will fail over within the cluster in the Possible Owners dialog box, and then click Next. For the active/active setup, remove all other nodes from the list.

c) In the Dependencies dialog box, click the IP Address resource for this EVS in the Available resources box, and then click Add. Click Next. In the Network Name Parameters dialog box, type a NetBIOS name for the EVS in the Name box. This name is the network name that identifies the EVS on your network. Check the buttons of Using Authentication and DNS registration must succeed under this field. Click Finish.

7. Create Physical Disk Resource

This step is applied to all disk drives that are included in a specific EVS, for example, Exchange executable (D:\), mailboxes and transaction logs.

a) Right-click the EVS, point to New, and then click Resource. The New Resource Wizard starts.

b) In the New Resource dialog box, enter a name for the disk resource in the Name box. In the Resource Type box, click physical disk; the system will pause several minutes depending on how many disks it will have to scan, and then click Next. For the active/passive setup, enter the node on which this EVS resides and the node to which it will fail over within the cluster in the Possible Owners dialog box, and then click Next. For the active/active setup, remove all other nodes from the list.

c) In the Dependencies dialog box, do not select any resources that appear in the Resource Dependencies box, and then click Next to select the correct drive that you want to add as a new resource. This step may take several minutes depending on how many drives you have.

8. Create Exchange System Attendant Resource

The last step is to create Exchange System Attendant Resource for each EVS created. This must be done for each active EVS. Ensure that all resources included in a specific EVS are added and online prior to this step.

a) Right-click the EVS, point to New, and then click Resource. The New Resource Wizard starts.

b) In the New Resource dialog box, enter a name in the Name box. In the Resource Type box, click Microsoft Exchange System Attendant, and then click Next. For the active/passive setup, select the node on which this EVS resides and the node to which it will fail over within the cluster in the Possible Owners dialog box, and then click Next. For the active/active setup, remove all other nodes from the list.

c) In the Dependencies dialog box, select all resources appear in the list for this EVS and then click Add. Click Next.

d) In the Data Directory dialog box, verify the data directory location; it is the directory and drive on which you have installed Exchange Server (in the test, it is D:\exchsrvr, and this location is true for all nodes in the cluster). You must verify that this location points to the shared physical disk resource assigned to this EVS. Exchange will use the drive you select in this step to store the default transaction log files, the default public store files, and the default mailbox store files (pub1.edb, pub1.stm, priv1.edb, and priv1.stm). Click Finish. You can manually point all those mentioned above to the different drives that were included in the Dependencies Resources list (in step d). After this, complete the EVS configuration on the local server.

e) Right-click the EVS, and then click Bring Online. After you successfully create the Exchange System Attendant resource, Exchange System Attendant automatically creates all the other resources for the EVS.

The Exchange System Attendant resource creates the following resources:

- Exchange Information Store Instance
- Exchange Message Transfer Agent Instance
- Exchange Routing Service Instance
- SMTP Virtual Server Instance
- Exchange HTTP Virtual Service Instance
- Exchange MS Search Instance

The following resources are NOT created by default:

- Exchange IMAP4 Virtual Server Instance
- Exchange POP3 Virtual Server Instance

The Message Transfer Agent Instance resource is only created in the first Exchange Virtual Server added to a cluster. All Exchange Virtual Servers in the cluster share the single Message Transfer Agent Instance resource. If you are creating an Active/Active Cluster, repeat to create system attendant resource for each EVS you want to create. For example, if you are creating a two-node active/active cluster, you must repeat this step.

9. Configuration EVS

You can modify all EVS configuration as you would do normally for a stand-alone Exchange Server once the EVS is brought on line. Use Exchange Server Manager on the individual EVS to create and remount mailboxes and transaction log stores. For the active/active setup, you can create an extra public folder, one for each EVS.

a) On each EVS (except the first Exchange virtual Server, which has one public folder store created by default), open Exchange Server Manager (ESM). Right click on the First storage group object, and add a new public folder store. The user must associate this public folder store with the default public folder tree. This will allow the public folder store to be a replica to the free busy folder.

b) In ESM, click "View of system folders" on the public folder object. This will show the Schedule + FreeBusy folder.

c) Open the Properties page of the instance that is in the FreeBusy folder, and add all the public folder stores as replicas.

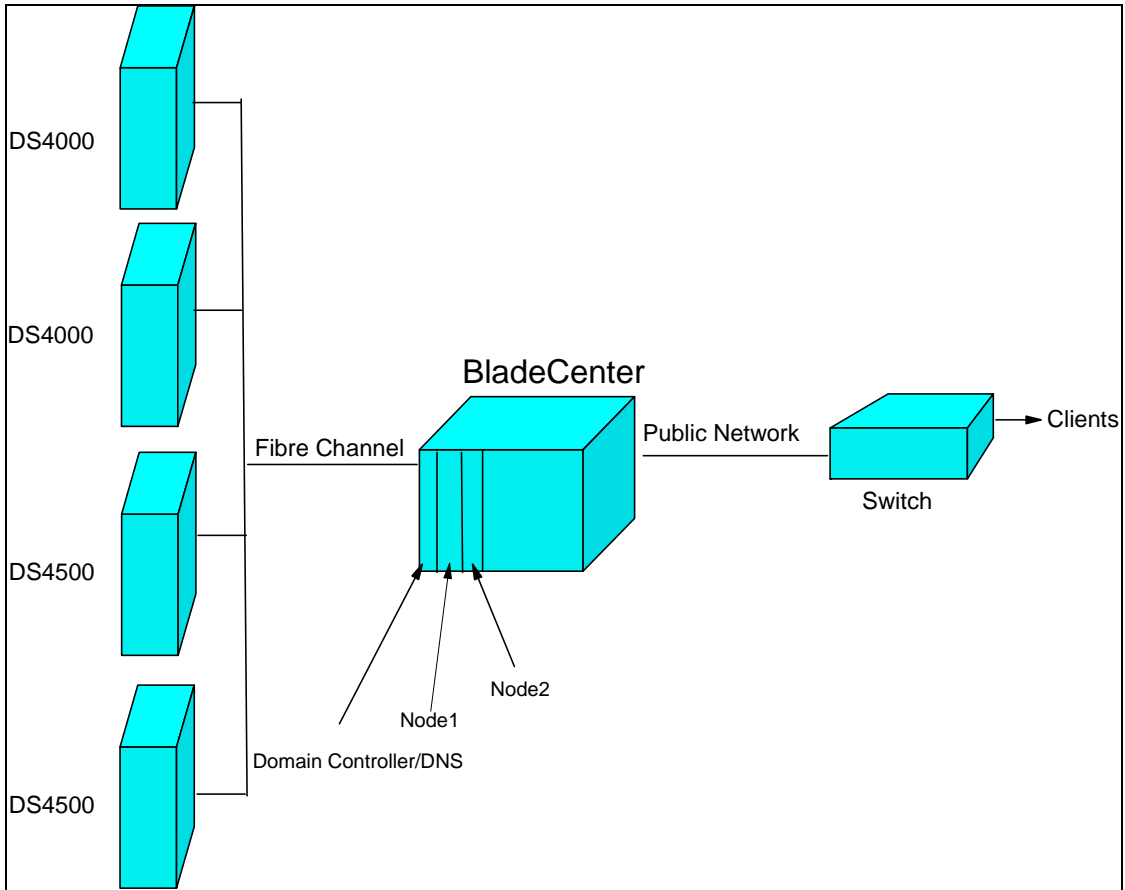
d) Change all mailbox stores to use the public folder store that is on the same EVS; this will help to load balance across all EVSs.

Test Environment

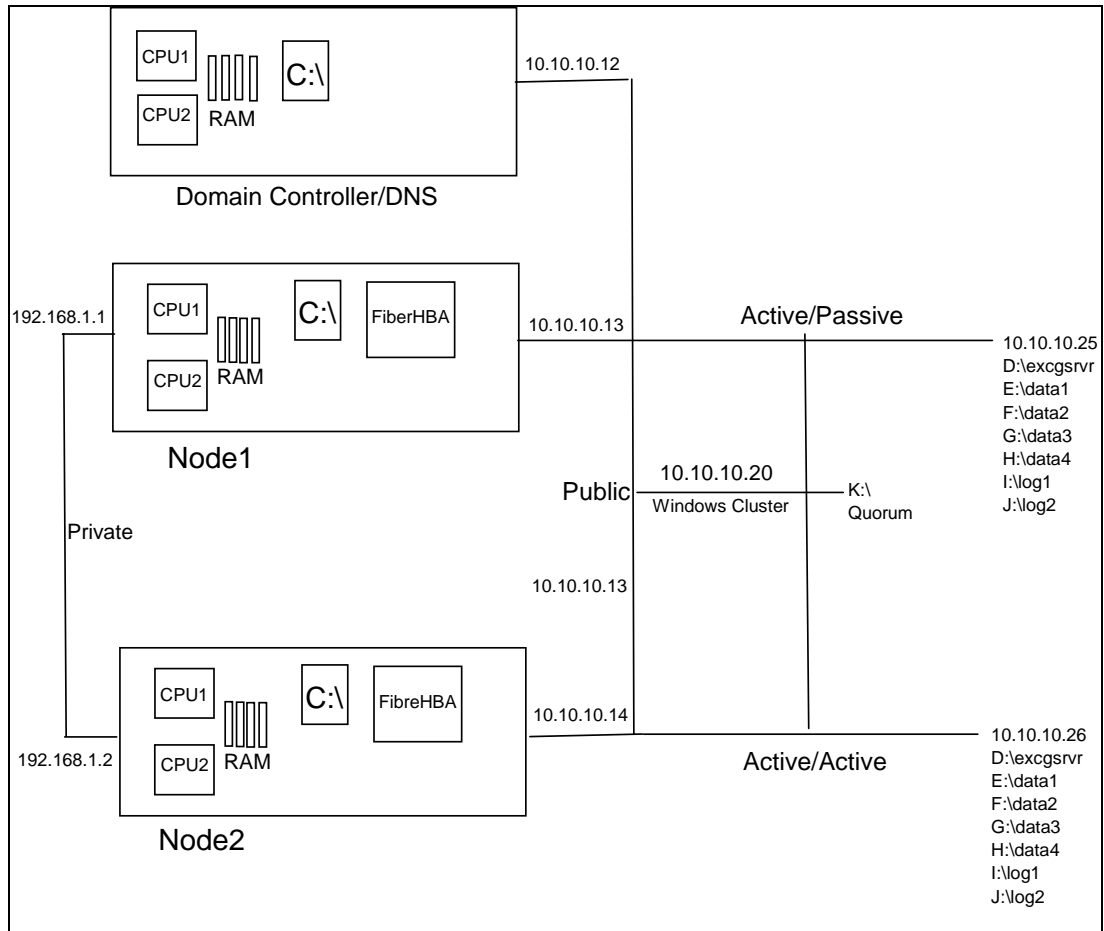
This document is based on a test conducted with an IBM eServer BladeCenter (equipped with two Ethernet switches, one two-port Fibre Channel switch, and one six-port Fibre Channel switch).

The Exchange Server cluster test was conducted with one HS20 blade (dual 3.6GHz Intel Xeon processors, a 36GB SCSI hard drive, four 1GB PC3200 DIMMs, and dual integrated gigabit Ethernet adapters) acting as the DNS/Domain Controller. Two other HS20 blades (dual 3.6GHz Intel Xeon processors, a 36GB SCSI hard drive, four 1GB PC3200 DIMMs, dual integrated gigabit Ethernet adapters, and a Fibre Channel daughter card) were used for the Exchange Virtual Server (EVS). Two DS4000 Host Bus Adapters and two DS4500 Fibre Channel disk controllers were used for disk management.

Following is a diagram of the test bed:



Following is a simple diagram below that shows the individual blade servers.



The following modifications were made to enhance Exchange Server operation:

- Added /3GB and /USERVA=3030 switches to boot.ini file for Exchange Server booting string
- For the best operation of the server, applied the following registry value:
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager
HeapDeCommitFreeBlockThreshold = REG_DWORD **0x00040000**
- The following value can be added for system monitoring purpose:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ESE\Performance
- Set parameter Show Advanced Counters = REG_DWORD **0x1**
- Stopped all non-essential services such as DHCP server/client
- Spread transaction log and mailbox storage across different physical disks to balance the load
- Increased Exchange Server's buffer maximum

The *msExchESEParamCacheSizeMax* parameter controls the ESE buffer size. Its value is expressed as a page count and must be set to an actual multiple of 8192 for best function. Use

the Install Windows Support Tools from the Windows Server 2003 CD to open *adsiedit.msc* from \Support Tools. Then click through following tree:

*Configuration->Services->Microsoft Exchange->OrganizationName->Administrativ
Groups->Your Administrative Group->Servers->Your Server*

Right click on *InformationStore*; then click *Properties*. Scroll down the list and click *msExchESEParamCacheSizeMax*. In the *Edit* field, enter the value. The recommended value is **303104** (37 * 8192). Click OK.

Test Results

Active/Passive Test

An Exchange Server cluster was created with three BladeCenter HS20 blade servers as mentioned in Test Environment section. The first part of the test is the failover function. The active EVS has two Storage Groups containing one transaction log and two mailbox storages. LoadSim was used to create a 6,000 MMB3 users test case. A 1,500-user mail database was initiated for each mailbox. After all mailboxes were properly initiated, MMB3 users were logged on from 15 client systems (xSeries 330 with 933MHz Pentium III and 1GB of memory with 500 MMB3 users on each client).

Two hours after all users were logged on, the node where EVS resides was shut down by pressing the power button. On client systems, LoadSim reports errors and Node 1 went off line while Node 2 took over EVS. This fail-over process takes a minute or two. Once the fail-over process finished, the client systems resumed normal activities.

Node 1 was started again one hour after fail-over finished. Because fail back was immediately configured for the EVS, the fail-back process began after Node 1 was back online. There is a transition period and LoadSim reports errors on client systems. The fail back can be configured as No, Immediately or Wait for a predetermined period as Node 1 comes back online.

The 95% average response time for this test is 333 msec with 43% CPU utilization through out the entire 8-hour Active/Active Test.

Active/Passive Test

In this test two EVSs were created for both Node 1 and Node 2. Both EVSs had the same configuration as in the first test. Both EVSs ran without interruption with 6,000 MMB3 users each. Fail-over protection is not feasible for a two-node active/active Exchange Server cluster.

Conclusion

Server clusters provide high availability, fail back, scalability, and manageability for resources and applications. An active/passive Exchange Server cluster will ensure uninterrupted mail service for your organization. An active/active Exchange Server cluster cannot provide any protection for system failure. This type of cluster creates an extra layer of administration, and is, therefore, not recommended.

With all essential resources (e.g., network interfaces) integrated and with the various options that can be installed in the on the chassis, the IBM eServer BladeCenter is the logical choice to set up a server cluster. For a 7U rack space, a maximum 7+1 Exchange cluster can be set up. Not only does the BladeCenter reduce the space needed, but it also minimizes the administrative complexity. Although the cluster test in this document was done under benchmarking conditions, the results provide an excellent indication of the function customers can expect from an IBM eServer BladeCenter HS20 blade server configuration running their production applications.

For More Information

To read more about the Microsoft Exchange Server, visit the Microsoft Web site at

<http://www.microsoft.com/technet/prodtechnol/exchange/default.mspx>

You can also find useful information from the Web site:

<http://www.msexchange.org/> that dedicate to Exchange Server



© IBM Corporation 2006

IBM Systems and Technology Group
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
01-06
All rights reserved

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

IBM, the IBM logo, eServer, xSeries and BladeCenter are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml.

Microsoft and Windows are registered trademarks of Microsoft Corporation.

All other products may be trademarks or registered trademarks of their respective companies.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Performance is based on measurements using industry standard or IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve performance levels equivalent to those stated here.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.