Extend IBM i HTTP Server high availability to the IPv6 environment

Xu Le
Tian Gang
Xu Meng

August 29, 2012

High availability of the web server in an Internet Protocol version 6 (IPv6) environment can be achieved through the use of the IBM® PowerHA® SystemMirror for i software. This article describes how to extend to an IPv6 highly available web server cluster on the IBM i HTTP Server with a takeover IP model.

Introduction

Today, more and more companies realize the importance of high availability (HA) for their web servers because their customers demand reliable service. At the same time, movement from the IPv4 to the IPv6 standard is occurring at more and more organizations. Thus, a new requirement to support high-availability web service within the IPv6 environment has emerged.

High availability of the web server in the IPv6 environment can be achieved through the use of IBM PowerHA SystemMirror for i software. This article describes an easy method to extend an IPv4 highly available HTTP Server cluster on IBM i to the IPv6 environment.

IBM PowerHA SystemMirror for i

Three web server cluster models are supported for the highly available HTTP servers on the IBM i operating system:

- **Primary or backup with takeover IP model:** The web server runs on the primary and all backup nodes. The backup node or nodes are in an idle state, ready to become the primary web server when the primary web server fails or a switchover takes place. All the client requests are always served by the primary node.
- **Primary or backup with a network dispatcher model:** As with the primary or backup with the takeover IP model, the web server runs on the primary and all backup nodes. The backup nodes are in an idle state and all client requests are served by the primary node. A network dispatcher, such as the IBM WebSphere® Edge Server, sends client requests to the web server.

- **Peer model:** There is no declared primary node. All the nodes are in active state and serve client requests. A network dispatcher, such as IBM WebSphere Edge Server, evenly distributes the requests to different cluster nodes. This guarantees distribution of resources when there is a heavy load. However, linear scalability is not guaranteed beyond a small number of nodes. After some nodes are added, scalability can disappear, and the cluster performance can deteriorate.

This article describes how to extend to an IPv6 highly available web server cluster on the IBM i HTTP Server using the first model. The other two models require an additional network dispatcher, such as IBM WebSphere Edge Server, and are not covered in this article.

The following diagram illustrates the *primary or backup with takeover IP model* in an IPv6 network.

Figure 1.

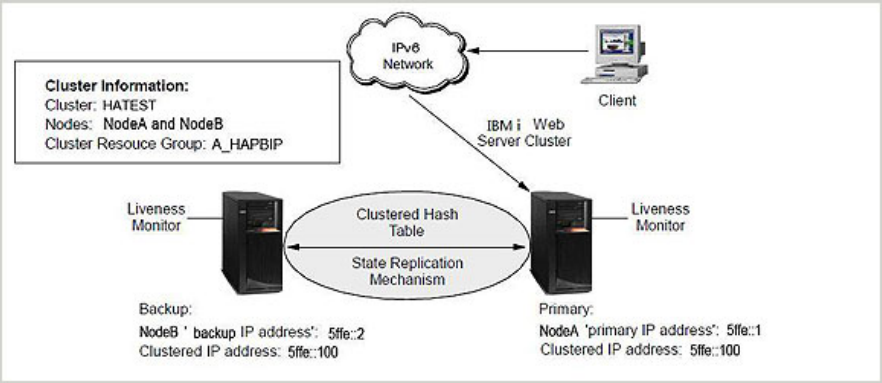

Figure 1. Primary or backup with takeover IP model — solution definition

When the primary node fails or is brought down by the administrator, the failover process begins. The following steps are performed during failover:

1. One of the backup servers becomes the primary node (the first backup in the switchover order).

2. Client requests are redirected to the new primary node. Assuming that this client was not in the process of running a persistent Common Gateway Interface (CGI) application, the fail over is completely transparent.

3. If the new primary node receives a user request that belongs to a long-running-session (a CGI program that is updated to be a highly available CGI program), the server restores the request's state. The new primary node retrieves that highly available CGI program's state information from the clustered hash table. The clustered hash table is part of the state replication mechanism.

4. After the failed node recovers, you can restart the highly available web server instance, which then becomes the backup system. If the system administrator wants the failed node to become primary again, they must perform a manual switchover.

Software requirement

For the *primary or backup with takeover IP* model in IPv6, the following licensed programs and program temporary fixes (PTFs) are required.

License program:

5770SS1 41 HA switchable resources

5770DG1 *BASE IBM HTTP Server for i

5770HAS *BASE IBM PowerHA SystemMirror for i Standard Edition

5770HAS 1 IBM PowerHA SystemMirror for i Enterprise Edition

Required PTFs:

- Current PTF Group for 5770DG1 (minimum SF99368 - level 10)

Although IPv6 has been supported on IBM i since version 5 Release 2 (v5R2), the IPv6 HA cluster support requires the IBM i 7.1 release.

Steps of moving IPv4 HA to IPv6 HA

Perform the following tasks to change the IPv4-based HA cluster to an IPv6-based HA cluster.

Step 1: Configure the IPv6 address on both of the IBM i servers

Step 2: Change the node address in the cluster

Step 3: Change the HA clustering directives in both of the httpd.conf configurations

Step 4: Test the primary and backup servers

Step 1: Configure the IPv6 address on both of the IBM i servers

The following two tables are examples of changing the HA IPv4 solution to a HA IPv6 solution.

Old HA solution with IPv4 address：

Table 1.

Node	IP address	Primary/Backup node	Clustered IP address	LOOPBACK IP address
NODEA	192.168.0.1	Primary	192.168.0.100	127.0.0.1
NODEB	192.168.0.2	Backup	192.168.0.100	127.0.0.1
Cluster information: Cluster: HATEST Nodes: NODEA and NODEB Cluster Resource Group：A_HAPBIP				

Note:

1. The primary node NODEA and the backup node NODEB are required to be in the same subnet.
2. Both the nodes must have a routable IP address that is accessible from each other server (NODEA can ping 192.168.0.2 and NODEB can ping 192.168.0.1.)
3. The clustered IP address (also known as the virtual IP address) is routable (that is, if a client on the same subnet wants to communicate to this IP address, it uses the Address Resolution Protocol (ARP) to resolve the Media Access Control [MAC] address on IBM i) but cannot be active on both the systems at the same time.

New HA solution with IPv6 address：

Table 2.

Node	IPv6 address	Primary/Backup site	Clustered IP address	LOOPBACK IP address
NODEA	5ffe::1	Primary	5ffe::100	::1
NODEB	5ffe::2	Backup	5ffe::100	::1

Cluster information:
Cluster: HATEST
Nodes: NODEA and NODEB
Cluster Resource Group：A_HAPBIP

Note:

1. The primary node NODEA and backup node NODEB are required to be in the same subnet.
2. Both the nodes must have a routable IPv6 address that is accessible from each other server (NODEA can ping 5ffe::2 and NODEB can ping 5ffe::1)
3. The clustered IP address (also known as virtual IP address) is routable but cannot be active on both the systems at the same time.
4. The IPv6 network among the clients and cluster must have been already set up and available to access. That means, every client should be able to ping through the clustered IP address when the primary node is active.

1. Create the IPv6 address **5ffe::1** on the primary node and **5ffe:2** on the backup node by using IBM System i® Navigator. As shown in Figure 1, this action can be done from the Network component of the main navigation tree.

Click **Network TCP/IP Configuration IPv6 Interfaces**, then click **New Interfaces Local Area Network** task.

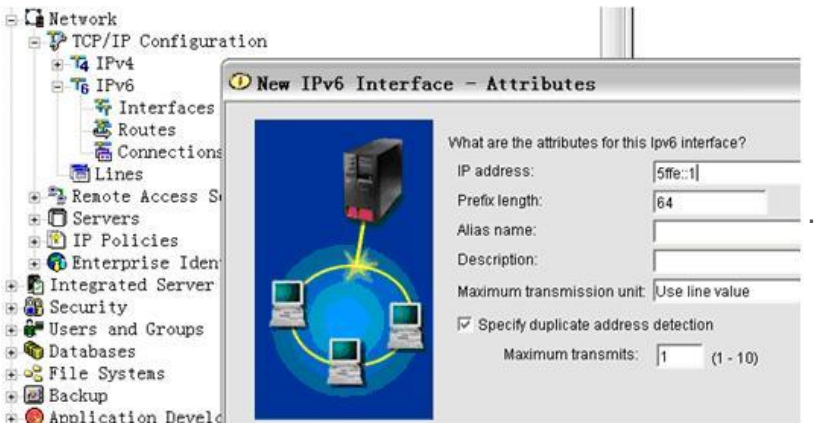

Figure 2. Creating an IPv6 interface

1. Verify that the IPv6 addresses on both: the primary node and the backup node can be pinged by each other.

2. Create the virtual IP address on both the servers by using the virtual IP interface in System i Navigator (Figure 3) or by using the following system command.
   ```
   ADDTCPIFC INTNETADR('5ffe::100') LIND(*VIRTUALIP) AUTOSTART(*NO)
   ```

Note: The virtual IP address is routable in the same subnet but cannot be active on all the nodes at the same time. It is the IP address takeover feature of IBM i HA clustering that automatically allows only one of the IBM i servers to have 5ffe::100 active at one time. To be clear, you never manually make 5ffe::100 active on either of the IBM i servers. HA clustering's IP takeover does this for you. Make sure that a PING request to the clustered IP address **5ffe::100** results in a time out. That is, the IP address is not active on both the servers.

Figure 2.

Figure 3. Primary or backup with IP takeover – Creating a virtual IPv6 address

1. Make sure that the loopback IPv6 address ::1 is configured and active on both the servers.

Step 2. Change the node address

Change the primary and backup node address by using the **CHGCLUNODE** (Change Cluster Node Entry) command. For example:

CHGCLUNODE CLUSTER(HATEST) NODE(NODEA) OPTION(*CHGIFC) OLDINTNETA('192.168.0.1') NEWINTNETA('5ffe::1')

CHGCLUNODE CLUSTER(HATEST) NODE(NODEB) OPTION(*CHGIFC) OLDINTNETA('192.168.0.2') NEWINTNETA('5ffe::2')

Note: All the nodes should be active, otherwise the change would be failed.

Step 3. Change HA clustering directives in both of the httpd.conf configurations

Modify the httpd.conf files on both, primary and backup nodes, to support HA IPv6.

1. Update the **Listen** directive to **Listen [5ffe::100]:10102 http**
2. Update the **LmURLCheck** directive to **LmURLCheck http://[5ffe::100]:10108/A_HAPBIP.html**

Figure 3.

```
1  LoadModule ha_module /QSYS.LIB/QHTTPSVR.LIB/QZSRCORE.SRVPGM
2  ServerRoot /www/a_hapbip
3  DocumentRoot /www/a_hapbip/htdocs
4  Options -Includes
5  UseCanonicalName Off          IP takeover address
6  Listen [5ffe::100]:10108 http
7  DirectoryIndex A_HAPBIP.html
8  ## HA Directives ##
9  HAModel PrimaryBackupWithIPTakeover
10 LmURLCheck http://[5ffe::100]:10108/A_HAPBIP.html
11 HACGI On
```

Figure 4. Primary or backup with IP takeover - Updated directives to support HA IPv6

After changing, restart both the primary and backup node.

Step 4. Test the primary and backup servers

1. Use the following DSPCRGINF (Display CRG Information) command, DSPCRGINF CLUSTER(HATEST) CRG(*LIST), to access the Cluster Resource Group information. Figure 5 shows the output of the DSPCRGINF command. You can see that the primary node **NodeA** is active.

Figure 4.

```
                    Display CRG Information

Cluster . . . . . . . . . . . . . . . . . . . . :   HATEST
Cluster resource group . . . . . . . . . . :   *LIST
Consistent information in cluster:     Yes
Number of cluster resource groups:     2

                    Cluster Resource Group List

Cluster Resource Group   CRG Type      Status          Primary Node
     A_HAPBIP              *APP         Active          NodeA
     AXHAU80N5O            *DATA        Inactive        NodeA
```

Figure 5. Primary or backup with IP takeover - CRG A_HAPBIP is active on primary node NodeA

2. Open a browser and go to the following URL: http://[5ffe::100]:10108/A_HAPBIP.html. A page, as shown in Figure 6, gets displayed. Perform this test with a browser such as Mozilla Firefox that can parse IPv6 addresses. It should be noted that the Microsoft® Internet Explorer browser cannot parse IPv6 addresses.

Figure 5.

Figure 6. Primary or backup with IP takeover - Primary server on NodeA is up and running

3. For testing purpose, issue the following CHGCRGPRI (Change CRG Primary) command to switch the primary node over to NodeB.

CHGCRGPRI CLUSTER(HATEST) CRG(A_HAPBIP)

After the CHGCRGPRI command is ran, the IP address **5ffe::100**will be inactive on server NodeA and active on server NodeB. You can also kill the primary node or shut down the primary node's logical partition to simulate a real customer scenario. The result is the same.

4. After the switch, check the CRG information by using the DSPCRGINF CLUSTER(HATEST) CRG(*LIST) command. You can see from Figure 7 that the primary node has been switched to NodeB.

Figure 6.

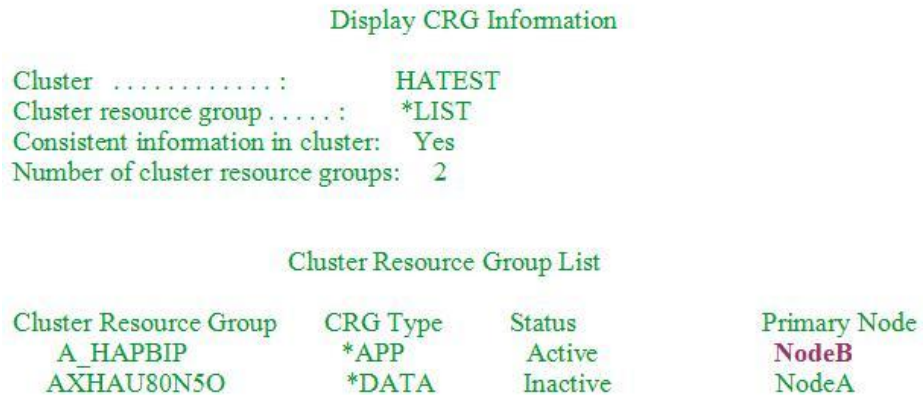

Figure 7. Primary or backup with IP takeover - CRG A_HAPBIP is active on the primary node, NodeB

5. Access the following web page, http://[5ffe::100]:10108/A_HAPBIP.html, again and now, a page, similar to the one shown in Figure 8 is displayed.

Figure 7.

Figure 8. Primary or backup with IP takeover - Backup server on NodeB is up and running

Now that you have successfully extended your web server highly available to the IPv6 environment with the primary or backup with takeover IP model, your organization can enjoy the benefits of a more reliable web service.

Trouble shooting

1. If the HTTP Server cannot be started, check if the primary and backup nodes are all active.

Run the WRKCLU (Work with Cluster) system command, and then select option 6. If the nodes are not active, select option 8 to start the node.

If the nodes cannot be started, perform the following steps.

a. Use the DSPNETA (Display Network Attributes) command to check the value of **Allow add to cluster**. This value should be set to ***ANY** on both the servers. If the value is not ***ANY**, change it by using the following CHGNETA (Change Network Attributes) command, CHGNETA ALWADDCLU(*ANY).

b. Make sure that the Internet Daemon (INETD) server is running on both the servers. This can be done from the System i Navigator client by clicking **Servers->TCP/IP->INETED** (as shown in Figure 9) or by using the following WRKACTJOB (Work with Active Jobs) command, WRKACTJOB JOB(QTOGINTD) to verify if the QTOGINTD job is active.

Figure 8.

Figure 9. Primary or backup with IP takeover - Start INETD server

If the INETD server is not active, start the INETD server with the following command: STRTCPSVR *INETD

c. Make sure that Liveness Monitor can run unimpeded in the QBATCH subsystem on both the servers.

1) To check the QBATCH subsystem setting, use the following DSPSBSD (Display Subsystem Description) command, DSPSBSD SBSD(QBATCH), and select option 6 to display the description shown in Figure 10.

Figure 9.

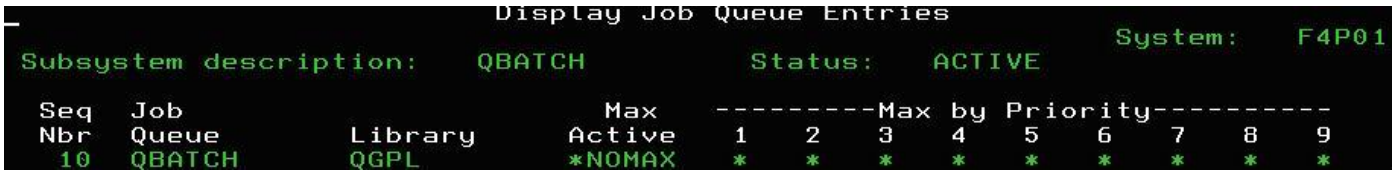

Figure 10. Primary or backup with IP takeover - Max Active for QBATCH should be *NOMAX

2) If the maximum active jobs allowed in QBATCH is not *NOMAX, change the job queue using the following command to change it to *NOMAX

CHGJOBQE SBSD(QBATCH) JOBQ(QBATCH) MAXACT(*NOMAX)

3. If the CHGCRGPRI command failed to change the primary node and fails with the following error message: *Primary node of cluster resource group A_HAPBIP not changed.*, then this means that the primary node cannot find the backup node. Try restarting all of the HTTP servers to correct this error.

Resources

- IBM HTTP Server (powered by Apache) An Integrated Solution for IBM iSeries Servers
- IBM i 7.1 information center
- You cannot use IPv6 addresses in Internet Explorer
- IPv6 specification

- Implementing PowerHA for IBM i