


IBM Workload Deployer

High availability



This presentation covers high availability for IBM Workload Deployer.

Table of contents

- Overview
- Floating IP address
- High availability commands
- Failover
- General information
- Error conditions and troubleshooting
- Summary

This presentation covers all aspect of high availability, also known as HA, in IBM Workload Deployer. You will take a look at the overall design, the use of a floating IP address, various high availability commands and some IBM Workload Deployer functions and how they integrate with the HA function. Finally, you will review some error conditions and how to recover from those situations.

Overview

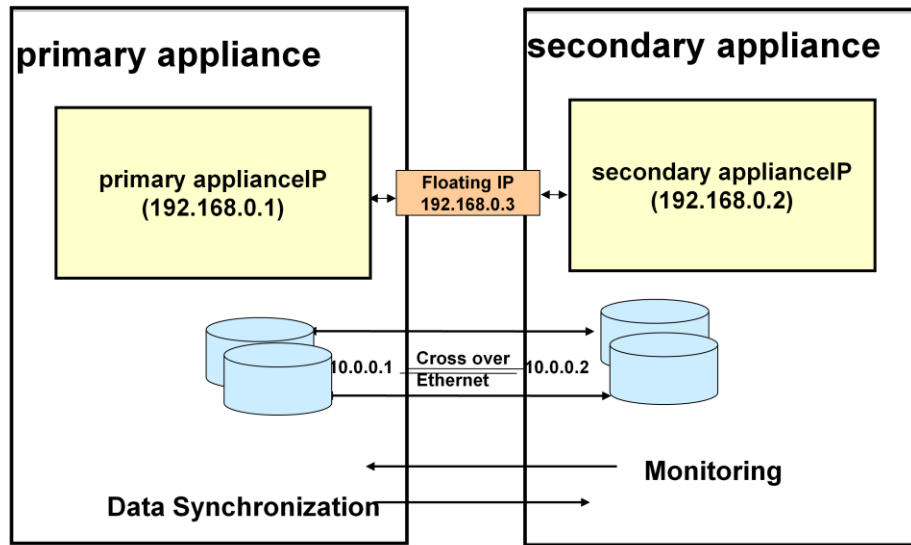
This section will give an overview of high availability for the IBM Workload Deployer appliance.

Overview

- New active/passive failover capability leveraging two IBM Workload Deployer appliances
 - One appliance is configured as the 'primary appliance' with the second appliance in standby as a secondary appliance until needed
 - The secondary appliance will have a replicated database and data so that it is ready to take over when needed
 - The administrator will configure the replication to define the primary appliance/secondary appliance relationship
- Allows customers to seamlessly continue managing their environments in the case of a failure on the primary appliance
- Needs to be initially setup and configured by the administrator
 - On failure of the primary appliance, the secondary appliance takes over control automatically
 - Eventually, the administrator can restart replication between the two appliances where now the secondary appliance has become the 'primary appliance' and a new appliance becomes the secondary appliance

The IBM Workload Deployer appliance can still run in 'standalone' mode in V3.1. If high availability is required you can introduce a second appliance and configure the original one as the 'primary appliance'. This is done by configuring some IP addresses to allow for communication between the two appliances. Once the communication is set up, the administrator can designate the first appliance as the 'primary appliance' which then activates the primary appliance/secondary appliance data replication process. If a failure on the primary appliance occurs, the secondary appliance automatically takes control by detecting that the primary appliance has failed. The secondary appliance will now start serving requests as a new primary appliance. However, the new primary appliance does not have the secondary appliance monitoring the primary appliance device until the administrator designates a new secondary appliance. This will again activate the primary appliance/secondary appliance data replication process but in the opposite direction.

High availability



5

High availability

© 2012 IBM Corporation

The IBM Workload Deployer has three operational modes, stand-alone, primary appliance or secondary appliance. The primary appliance and secondary appliance modes are for failover support. The primary appliance and secondary appliance have a dedicated crossover Ethernet cable that is used to keep the data in sync. When the secondary appliance is connected to the primary appliance for the first time, the entire disk partition is synced to the primary appliance using this dedicated crossover cable. Once this data sync is complete, all future writes to the primary appliances disk partition are written to the secondary appliances partition as well. This will keep user data, files and the database in sync at all times. The secondary appliance will continue to monitor the primary appliance using the ping interval specified on the startHA command. If the primary appliance does not respond within that interval, the secondary appliance assumes the primary appliance is down and takes control of the floating ip address. This transition is transparent to you since the floating IP address is used to access the IBM Workload Deployer console. If the network fails, the primary appliance will continue to update locally and the secondary appliance is re-synced to the primary appliance as soon as connectivity is re-established. In an extreme case where there are numerous updates to the primary appliance, the secondary appliance can be reconfigured and the entire disk is recopied. Both appliances must have the same firmware level. The secondary appliance does not run any processes except monitoring the primary appliance.

Floating IP address

This section will discuss the floating IP address used in a high available environment.

Floating IP address

- Floating IP address used to provide redundancy
- Provides transparency to the user
- secondary appliance does not run any processes
 - Maintains an exact data copy of the primary appliance
- If the primary appliance becomes unavailable
 - Processes on the primary appliance are stopped
 - Stand-by appliance will takes over floating IP address and start all processes

A floating ip address is typically used to provide redundancy in a hot standby configuration. One machine is able to take over the IP address of the other and respond to Address Resolution Protocol (ARP) requests. The floating IP address is configured on the primary appliance and is the IP address used to access the IBM Workload Deployer console. It is transparent to you which appliance is controlling the floating IP address.

High availability commands

This section will discuss how to configure high availability on the primary appliance and secondary appliances.

Start high availability

- Support is through command line interface only
- SSH to the appliance using the cbadmin user ID
 - `Console>startHA -mode=master -primary applianceIP=<primary applianceIP address> -secondary applianceIP=<secondary appliancesIP address> -floatingIP=<service IP address> - primary applianceSyncIP=<primary applianceSync IP address> -secondary applianceSyncIP=<secondary applianceSync IP address> - pingInterval=120`
 - Floating IP is the IP used to access the IBM Workload Deployer console
 - primary applianceSyncIP is the IP used by primary appliance to sync data with secondary appliance
 - secondary applianceSyncIP is the IP used by secondary appliance to sync data with primary appliance
 - pingInterval is the amount of time to wait before failover process begins
- For secondary appliance use `-mode=slave`

To setup high availability, using ssh, login with the cbadmin ID to get to the command prompt. From both the secondary appliance and the primary appliance, issue the startHA command with the parameters shown above. Once initial data synchronization occurs, both the local and remote disks are written at the same time ensuring data is not lost between the primary appliance and secondary appliance. The pingInterval specifies the amount of time connectivity must be lost between the secondary appliance and the primary appliance before the secondary appliance takes over.

While the startHA command is running on the primary appliance, the IBM Workload Deployer console is unavailable. Once the startHA command has completed, a success message is displayed and the IBM Workload Deployer console should be accessed using the floating IP address. If either appliance is rebooted the HA setup is lost and will need to be reconfigured.

Re-configure high availability

- To reconfigure the primary appliance or secondary appliance
 - `Console>configureHA -mode=master -primary applianceIP=<primary applianceIP address> -secondary applianceIP=<secondary appliancesIP address> -floatingIP=<service IP address> - -primary applianceSyncIP=<primary applianceSync IP address> -secondary applianceSyncIP=<secondary applianceSync IP address> - pingInterval=120`
- Example: change primary applianceIP address
 - From primary appliance
 - `netif set eth0 address=newIP`
 - `configureHA -primary applianceIP=newIP ... -floatingIP-<floatingIP address> ...`
 - From secondary appliance
 - `configureHA -primary applianceIP=newIP`
 - After issuing `netif set` command, the `floatingIP` address must be set again even though it did not change

To change the parameters from the initial configuration use the `configureHA` command. You can update the primary applianceIP, secondary applianceIP, primary applianceSyncIP, secondary applianceSyncIP and pingInterval.. The same command must be run on both primary appliance and secondary appliances.

After issuing the `netif set` command you must reset the `floatingIP` address even though it has not changed.

Additional commands

- All commands can be issued from either the primary appliance or the secondary appliance

- To stop HA

```
console> stopHA
```

- To determine HA status

```
console> HAstatus
```

```
Console> HAstatus  
Appliance is in Slave mode.  
Console> _
```

At any time the stopHA command can be issued on either the primary appliance, secondary appliance or both appliances. Once the command has run, the appliance is returned to stand-alone mode. If stopHA is run on the primary appliance, the floating IP address is un-configured and you will have to use the primary appliance's IP address to access the IBM Workload Deployer console.

To determine the current status of the appliance issue the HAstatus command. The status command will show the current mode the appliance is running in.

Section

Failover

This section will briefly talk about high availability failover.

Automatic failover

- primary appliance becomes unavailable
 - secondary appliance becomes the primary appliance
 - Old primary appliance restarts as stand-alone appliance
- secondary appliance becomes unavailable
 - primary appliance continues processing, replication to secondary appliance is stopped
- Failover is not immediate; processes need to be restarted on the secondary appliance
 - pureScaleApp
 - Kernel Services
 - Storehouse
 - IBM Workload Deployer management processes, etc
- Two networks are pinged to ensure there is a connectivity loss
 - Dedicated data sync network
 - Network that primary appliance and secondary appliance are connected to other hypervisors

If the primary appliance becomes unavailable, it then restarts the secondary appliance and takes over to configure the floating IP address. When the primary appliance comes back up, it will not have the floating IP configured and therefore, it comes back as stand-alone appliance. You must issue the startHA command on the old primary appliance and resume it's processing as the new secondary appliance. If the secondary appliance becomes unavailable and restarts, the primary appliance continues to process, however, replication to the secondary appliance is stopped. The startHA command must be run on the secondary appliance after the appliance resumes. The secondary appliance does not run any processes except monitoring the primary appliance.

Failover is not immediate. Two networks are pinged to provide some redundancy checking to avoid the secondary appliance assuming control pre-maturely. Both the dedicated data network and the network that the appliances are connected to are checked. Once failover has occurred processes must be restarted on the secondary appliance. Existing deployments or virtual machines are not effected and can be managed from the IBM Workload Deployer console.

Forced failover from primary appliance to secondary appliance

- To issue a forced failover
 - first from the primary appliance issue :
`console> failover`
 - once complete from the secondary appliance issue :
`console> failover`

```
Failover Success. Go to slave appliance and invoke failover.  
Console>
```

You can force a failover from the primary appliance to the secondary appliance by invoking the failover command from the console.

Section

Other topics

This section discusses how high availability integrates with existing IBM Workload Deployer functions.

Firmware upgrade

- Issue failover command on primary appliance
`console > failover`
- When complete, issue failover command on secondary appliance
`console > failover`
- primary appliance is now idle; secondary appliance becomes active primary appliance
- Upgrade both appliances
`console > firmware upgrade newimage`
- After appliances reboot restart high availability on both appliances
 - Original secondary appliance becomes new primary appliance
 - Original primary appliance becomes current secondary appliance

Since a firmware upgrade requires a reboot of the appliance, upgrading the primary appliance will cause the secondary appliance to takeover as the primary appliance. Therefore, before upgrading the firmware on the primary appliance, the secondary appliance should be forced to takeover processing while the primary appliance is being upgraded. Once the reboot is complete the appliances are in stand-alone mode. You must restart HA by issuing the startHA command on both the primary appliance and the secondary appliance. The original secondary appliance will serve as the new primary appliance and the original primary appliance assumes the secondary appliance role.

Backup and restore

- Backup can be done anytime
- High availability must be turned off for restore

Appliance A : primary appliance; Appliance B : secondary appliance

Appliance A:

```
console > stopHA
```

Restore on Appliance A

Appliance B:

```
console > stopHA
```

Restore on Appliance B

Appliance A:

```
console > startHA -mode=master ...
```

Appliance B

```
console > startHA -mode=slave ...
```

The backup process can run anytime, however high availability must be stopped to run the restore process. You must stop HA on both appliances and then run the restore process. Once complete restart HA first on the primary appliance, then on the secondary appliance.

Migration

- High availability must be turned off for migration

Appliance A : primary appliance; Appliance B : secondary appliance

Appliance A:

```
console > stopHA
```

Migration on Appliance A

Appliance B:

```
console > stopHA
```

Migration on Appliance B

Appliance A:

```
console > startHA -mode=master ...
```

Appliance B

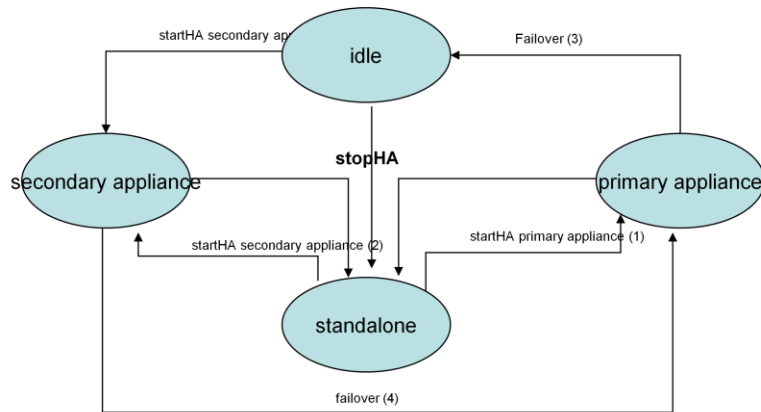
```
console > startHA -mode=slave ...
```

Migration must be run with HA turned off. Migration is similar to the restore process where you issue the stopHA command and then run migration. Once complete on both appliances restart HA.

Error conditions and troubleshooting

This section will discuss various error conditions and troubleshooting techniques.

Command flow



20

High availability

© 2012 IBM Corporation

Initially, the appliance is in a stand-alone state. High availability is started on the primary appliance first, and after completion, started on the secondary appliance. Once high availability is configured on both appliances, you can force a failover, first on the primary appliance and then on the secondary appliance. The primary appliance will go into an idle state and does not run any processes. While in an idle state, HA metadata information is retained enabling the appliance to quickly become the secondary appliance of the new primary appliance without synching the entire disk. The stopHA command will put the appliance in a stand-alone state, regardless of its starting state.

Error conditions

- secondary appliance becomes unavailable; currently no notification available
- Connection lost during data sync
 - Inconsistent state
 - startHA must be re-issued
- Command failure – failover; startHA
 - Repeat command
 - If still unsuccessful stopHA
- If stopHA fails reboot appliance
- Use configureHA to correct any not valid parameters specified during startHA
- Failure during deployment will require manual clean up

Since the secondary appliance does not run any IBM Workload Deployer processes, it is assumed that the secondary appliance is always available to monitor the primary appliance. However, if the secondary appliance terminates for some reason the primary appliance will continue to process requests. There is no notification to indicate that the secondary appliance has terminated. You must manually determine that the secondary appliance has failed by pinging the secondary appliance's IP address. After a failure you must go directly to the appliance, restore the appliance and restart HA on the appliance redefining it as a secondary appliance to the primary appliance.

It is possible that the connection between the primary appliance and secondary appliance is lost during a data synch. If this event occurs, or if the secondary appliance becomes unavailable during the sync process, the secondary appliance is considered to be in an inconsistent state. When the connection is restored the secondary appliance must be reconfigured using the startHA command.

To recover from other command errors, such as failOver or startHA, in either the primary appliance or secondary appliance reissue the command passing the correct command parameters as indicated by the command. configureHA can be used to correct some configuration errors. If there is still a failure, the appliance is considered to be in a bad state and the stopHA command should be issued. If the stopHA command fails, the appliance needs to be rebooted.

There is no error checking to ensure the secondary appliance's IP address is correct when issuing the startHA command from the primary appliance. This is because the primary appliance does not need the secondary appliance to function. If this occurs use the configureHA command to correct any not valid parameters.

If a failure happens during deployment, the deployment is considered to be partial, and must be cleaned up manually.

Existing deployments

- IP address must be maintained if there are existing deployments

Appliance A : 172.16.1.1; Appliance B : 172.16.1.2: Float IP: 172.16.1.3

Appliance A: contains existing virtual machines – needs to be primary appliance

```
console > netif set eth0 IPAddress=172.16.1.3/21 enabled=true
console > startHA -mode=master -primary applianceIP=172.16.1.3 -secondary
applianceIP=172.16.1.2 -floatingIP=172.16.1.1
```

Appliance A is the primary appliance; original IP is now defined as the floatingIP

```
console > startHA -mode=slave -primary applianceIP=172.16.1.3 -secondary
applianceIP=172.16.1.2 -floatingIP=172.16.1.1
```

Special consideration has to be given to an appliance that has been running in stand-alone mode before setting up HA. If the appliance has existing deployments, the actual appliance ID has been used to communicate with the appliance. In order to ensure existing virtual machines continue to operate as expected, you must use the appliance's IP address as the floating IP address and change the appliance IP address with another IP address on the network.

Shown here are the steps to configure HA to preserve already deployed virtual machines on a stand-alone appliance. Take note that the floating IP address becomes the primary appliance console IP address, and the primary appliance IP address assumes the role of the floating IP.

Troubleshooting

- Three available log files
 - ha.log – logs all activities
 - ha.out – standard output from the last HA command
 - ha.result – result of the last HA command
- After a failover situation Trace.zip contains all three logs from both primary appliance and secondary appliance
 - Log name is prefixed with IP address of the appliance
 - IP1_ha.log IP2_ha.log
 - IP1_ha.out IP2_ha.out
 - IP1_ha.result IP2_ha.result

There are three log files available to help debug HA failures. The first is the ha.log, which logs all activities; the second is the ha.out, which is the standard output for the last HA command; and the third is the ha.result, which contains the result of the last HA command that was issued. After a failover has occurred, the trace.zip file will contain the three logs from both the primary appliance and secondary appliance. The appliance IP address will prefix the log file names.

Summary

This section is a summary of high availability for the IBM Workload Deployer appliance.

Summary

- Overview
- Floating IP address
- High availability commands
- Failover
- General information
- Error conditions and troubleshooting

This presentation discussed the high availability function in IBM Workload Deployer. You looked at the overall design, the use of a floating IP address, various high availability commands and some IBM Workload Deployer functions and how they integrate with the HA function.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2014. All rights reserved.