

Staged Application Deployment WebSphere V7

Mike Loos

IBM mikeloos@us.ibm.com SHARE 115, Boston, MA Thursday, August 5, 2010 12:15 PM

2010 IBM Corporatio



Overview

The staged deployment of a new version of an application is a component of maintaining high availability in a Web-Sphere Application Server on z/OS environment. Though not difficult to achieve, it is not automatic. This paper will attempt to demonstrate one method of deploying a new version of an application, both manually, and via the scripting interface, in an ND cell, over two members of an application server cluster.



Clustered Server in a Network Deployment Environment.

The most common configuration for a clustered server is to have two application server nodes, with a server (cluster member) in each of them.

The members of the cluster are basically identical. They serve the same applications, have the same configuration, and usually listen on the same ports.

Assumptions

Let's start by assuming that we already have an application deployed in the cell and mapped to a particular server cluster. We will also assume that the new version of the application is built such that running both of the versions (1 and 2 in this case) simultaneously, will not cause any problems. This assumption would not be true in the case of a new version with, for instance, a different and inconsistent backend database structure.

The other basic assumption is that this is something that the user wishes to do. It is certainly "okay" to simply deploy a new application version and have it take effect on all nodes at one time. The procedure we will be documenting here is for that occasion where what is desired is to put the new version of an application into one member (server) of a multimember cluster and run it concurrently with the older version of the application in the remaining member(s) of the cluster. Then, once the new version has been "proven" in the production environment, the rest of the member(s) of the cluster may be switched to the new version either one at a time, or all at once.

Conversely, if the new version of the application doesn't seem to be working as expected, it can be removed from one member that is running it and the older version can be reinstated.



Now when an application is installed, what that really means is that all of its artifacts (files) are copied into the master repository maintained by the deployment manager. Also, deployment information is built depending upon the information that you provide at installation time. Of particular interest here is the locations (which servers) that the modules comprising the application will be run. Once the application is installed into the master configuration, the next step is to synchronize the changes made to the master configuration with the configurations maintained by the separate nodes. The standard or "default" methodology for this synchronization is to simply synchronize the master configuration with all nodes in the cell at one time. When a new version of an application gets to the node configuration, the application is stopped and started to make the changes effective. While this certainly works fine and may be exactly the behavior that we desire, we may want to see the alternative behavior that we described earlier, when we described the assumptions that drove the need for this documentation.

Init	ial Install o	of Application	
Integrated Solutions Console			
Integrated Solutions Console Welcome mileos View: All tasks Welcome Guided Activities Servers			
Applications New Application Application Services	Integrated Solutions Console Welcome mjloos View: All tasks • Welcome	Cell=s1cell, Profile=default New Application	Help L
Resources	 ➡ Guided Activities ➡ Servers 	New Application This page provides links to create new applications of differe	nt types.
Environment	 Applications New Application ∴ Application Types ↔ Services ↔ Resources ↔ Security ↔ Environment ↔ System administration ↔ Users and Groups ☆ Monitoring and Tuning 	Install a New Application New Enterprise Application Image: State of the sta	se application install ation.

The application that we are going to use to demonstrate this is the MyIVT application. This copy of the application has been modified to display a version number on its initial screen. We really don't care at all about the rest of the application. We have two copies of the ear file for the MyIVT application, one for Version 1 (MyIVTV1.ear) and one for Version 2 (MyIVTV2.ear). The **only** difference between the two versions is the text on the index.html page. Everything else is identical, particularly the artifacts that make it unique within the cell, the Application ID (My_IVT_Application) and the context root (/MyIVT).

	Close page		
	?		
epare and install your application	on.		
	repare and install your application	repare and install your application.	repare and install your application.

The first task is to install the application to the master configuration, specifying that it should be run on the cluster of interest (s1sr03 in our examples). We accomplish this by first logging on to the adminconsole application, opening up the Applications section on the left hand pane of the screen, and clicking on New Application

On the next screen we click on New Enterprise Application.

On the next screen we might use the Choose File dialog to select the full path to our ear file, in this case to My-IVTV1.ear for Version 1 of the application. Then we click Next.



Our next step is to indicate on which server or cluster we wish to run this application. On this screen follow the callout boxes to select cluster s1sr03.



When presented with the Summary page, click Finish.

When presented with the result screen click on Review.

Initial	Install o	of Application	on
New Application			
New Application > Enterprise Applications > Si	ave	ADMS0203: The automatic synchronization	mode is enabled for node: s1noded.
Save your workspace changes to the master co	onfiguration.	III ADMS0207I: Node Synchronization state for 03:39:21:107 EST result: In Progress Lindate or	node: s1nodec - initiate time: 2010.02 curred
Click Save to update the master repository with begin work again using the master repository co changes.	your changes. Click Di: onfiguration. Click Canc	ADMS0207I: Node Synchronization state for 03:39:21:107 EST complete time: 2010.02.08 at occurred	node: s1nodec - initiate time: 2010.02 03:39:24:231 EST result: Complete Up
		ADMS0205I: The configuration synchroniza	ion completed successfully for node:
\bigcirc		s1nodec.	·
ynchronize changes with Nodes		ADMSU208: The configuration synchroniza	
Save Discard Cancel			
ki internet			
		<u>lok</u>	
		42	
	Internet of Only France One of a		
	Integrated Solutions Console w	elcome mjloos	
	View: All tasks		
	- Welcome		
	Guided Activities Guided Activi		
	Applications	12	
	 New Application 	12	
	Application Types		
	 WebSphere enterprise ap Business-lev hoplicati 	ns	
	- Assets WebSpher	enterprise applica	
	+ Services		

Since this is the First install of Version 1, we will want to synchronize with all nodes, so check the box to Synchronize changes with nodes, and then click on Save.

When the text in the progress box indicates that the configuration synchronization is complete for the cell, you may click OK.

The application is now installed and ready to be started. Open up Applications and Application Types and then click on WebSphere Enterprise Applications on the left hand pane of the adminconsole.

	ell, Profile=default						
erpris	e Applications						
Use th	is page to manage installed applications. A single	application can be o	deployed onto n	nultip			
🕂 Pre	ferences						
Sta	t Stop Install Uninstall Update Roll	lout Update Remo	ve File Export				
Select	Name 🗘	Application	n Status 🖏				
You d	an administer the following resources:						
	DefaultApplication	•					
	ECIDateTimeAD01	8					
	My_IVT_Application	8					
	querv	*					
Total	4						
Total	4		Enterprise A	pplications Messages Application My_IVT_Applicati successfully. Application My_IVT_Applicati s1sr03.	on on server s1 on on server s1 on started succe	sr03c and node sinoc sr03d and node sinoc ssfully on all of the s	dec started ded started ervers in cluster
Total	4		Enterprise A	pplications Messages Application My_IVT_Applicati successfully. Application My_IVT_Applicati successfully. Application My_IVT_Applicati s1sr03.	on on server s1 on on server s1 on started succe	sr03c and node sinoc sr03d and node sinoc assfully on all of the s	dec started ded started ervers in cluster
Total	4		Enterprise A Use this p	pplications Application My_IVT_Applicati successfully. Application My_IVT_Applicati successfully. Application My_IVT_Applicati s1sr03. Applications age to manage installed applications. A snces	on on server s1 on on server s1 on started succe single applicatio	sr03c and node sinoc sr03d and node sinoc sssfully on all of the s on can be deployed or	dec started ded started ervers in cluster nto multiple servers
Total	4		Enterprise A Esterprise Use this p Prefere Start	pplications Messages Application My_IVT_Applicati Application My_IVT_Applicati stocessfully. Application My_IVT_Applicati storess Applications Papplications Papplications A page to manage installed applications. A pages Stop Install Uninstall Update	on on server s1 on on server s1 on started succe single application	sr03c and node sinoc sr03d and node sino ssfully on all of the s on can be deployed or a	dec started ded started ervers in cluster nto multiple servers xportExport DDL
Total	4		Enterprise A Use this p Prefere Start	pplications Messages Application My_IVT_Applicati Application My_IVT_Applicati Cocessfully. Application My_IVT_Applicati s1sr03. Applications Papplications Stop Install Uninstall Update	on on server s1 on on server s1 on started succe single application	sr03c and node since sr03d and node since sssfully on all of the s in can be deployed or can be deployed or	dec started ded started ervers in cluster nto multiple servers xport) Export DDL
Total	4		Enterprise A Enterprise Use this prefere Start Select N	pplications Application My_IVT_Applicati cossfully. Applications Stop Install Uninstall Update advantate the following resource:	on on server \$1 on on server \$1 on started succe single application Rollout Updat	sr03c and node since sr03d and node since issfully on all of the s on can be deployed or a Remove File E Application Status Q	dec started ded started ervers in cluster nto multiple servers xport Export DDL
Total	4		Enterprise Use this prefere Start Select N You can	ppleations Application My_IVT_Applicati Gamma Constraints Application My_IVT_Applicati Gamma Constraints Application My_IVT_Applicati Gamma Constraints Stop Install Uninstall Update Stop Install Uninstall Update Statistication Statistication Application Ap	on on server \$1 on on server \$1 on started succe single application	sr03c and node sinoc sr03d and node sino assfully on all of the s on can be deployed or can be deployed or c	dec started ded started ervers in cluster nto multiple servers xport Export DDL
Total	4		Enterprise Use this prefere Start Vou can C E	pplications Application My_IVT_Applicati Gamma Application My_IVT_Applicati Gamma Application My_IVT_Applicati Gamma Application My_IVT_Applicati Gamma Applications Stop Install Uninstall Update Stop Install Update Stop Install Uninstall Update Stop Install Update	on on server \$1 on on server \$1 on started succe single application	sr03c and node s1noc sr03d and node s1no assfully on all of the s on can be deployed or a Remove File E Application Status Q *	dec started ded started ervers in cluster nto multiple servers xport Export DDL
Total	4		Enterprise A Enterprise Use this p Prefere Salect N You can C E C E Salect N	pyplications Application My_IVT_Applicati Groups Application My_IVT_Applicati Groups Application My_IVT_Applicati Groups Application My_IVT_Applicati Groups Applications Applications Application Stop Install Uninstall Update Cloate TimeAp01 ty_IVT_Application	on on server s1 on on server s1 on started succe single application	sr03c and node s1noc sr03d and node s1noc sssfully on all of the s on can be deployed or b Remove File E Application Status & c c	dec started ded started ervers in cluster nto multiple servers xport Export DDL

Check the box next to My_IVT_Application and then click on the Start button.

The next screen should indicate the successful start of the application on all members of the cluster.



Let's test. First we'll access the member s1sr03c

at url: <u>http://wsc3:15507/MyIVT/index.html</u>



Next let's access the member s1sr03d at url: <u>http://wsc4:15507/MyIVT/index.html</u>

We now have Version 1 successfully installed and running on both members of cluster s1sr03.



The key to controlling the updating of an application in the various members of a cluster is to control the synchronization process.

The default setup for an ND cell is for nodes to run synchronization on a regular timed basis (automatic synchronization) as well as at server startup time. This is fine during "normal" operations, but when we want to control the updating of an application, we need to stop all synchronizations from occurring that we don't explicitly start. That way we can synchronize each node in a controlled fashion and thereby update our application to version 2, one member of the cluster at a time.

 Welcome Guided Activities 	
Gervices	
E Security	
Environment	
🖃 System administration	Node agents
 Deployment manager Nodes Node agents Node group Console Preferences Console Identity 	specific administration functions, such as server process monitoring, configuration synchronization, file transfer, and request routing.
	Select Name 🗘 Node 🌣 Host Name 🗘 Version 🗘 Status
	You can administer the following resources:
	nodeagent s1nodec wsc3.washington.ibm.com ND 7.0.0.8
	nodeagent s1noded wsc4.washington.ibm.com ND 7.0.0.8
	Total 2

Synchronization is controlled by the File Synchronization Service settings which are a part of each nodeagent. In our cell we have two nodes, and thus, two nodeagents. Disabling automatic synchronization involves changing the settings on each nodeagent in the cell. Here is how.

On the left hand pane of the adminconsole click on System administration, and then Node agents.

Then on the nodeagent list panel, click on the first nodeagent (in our case, the one for node s1nodec).

On the nodeagent configuration panel, click on File synchronization service.

		*	IBM.
Staged D	Deployment		
Runtime		More informatio this page	
General Properties Name nodeagent Node \$lnodec \$short Name \$Short Name \$StaGNTC Unique ID C3BCA7A7C0EE964600000110000000309521847 Run in 64 bit JVM mode Start components as needed Apply OK Reset Cancel	Server Infrastructure Java and Process Management Administration Additional Properties Core group service Ports Elle synchronization service File transfer Unice Performal Specifies the service logic that keeps t Infrastructhis node Administration services Custom services Custom services CoRB service Change Log Detail Levels	the files synchronized for	
	18		
18 IBM Advanced Technica	l Skills – Washington Systems Center	© 2010 I	BM Corporation

Uncheck both the Automatic synchronization and the Startup synchronization boxes, then click on the OK button.

	EM
Staged Deployment	
Node agents ? –	
<u>Node agents</u> > <u>nodeagent</u> > File synchronization service Use this page to configure the file synchronization service. The file synchronization service runs in the deployment manager and node agent. It ensures that configuration changes made to the cell repository are propagated to the appropriate node repositories.	
Configuration	
General Properties Additional Properties Enable service at server startup Custom properties	
Synchronization interval In minutes Device synchronization Device synchronization	
Apply OK Reset Cancel	
IBM Advanced Technical Skills – Washington Systems Center @ 2	010 IBM Corporat

Repeat this for the next node (in our case s1noded). When you have completed this for each node that contains a member of the cluster, you need to save the configuration, and synchronize the changes with all nodes. At this point, the only synchronization that should occur will be one that is manually initiated. To make this take effect, you'll need to **recycle the nodeagents**.



Uninstall V1 Integrated Solutions Console Welcome mjloos View: All tasks • Welcome Guided Activities	
Integrated Solutions Console Welcome miloos View: All tasks • Welcome • Guided Activities	
View: All tasks	
 Welcome Guided Activities 	
Guided Activities	
+ Servers	
Applications	
 New Application 	
Application Types	
 WebSphere enterprise applications Business-level Whilestions 	
- Assets WebSphere enterprise app	
÷ Services 21	

We can now install version 2 of the application. We'll approach this by first uninstalling (to remove all traces) of the application, then installing the version 2 copy of the ear file. Remember that all synchronization is disabled, so any changes we make to the master configuration have no effect on the servers which are members of the cluster.

Uninstall is simple. Just go into the adminconsole and in the left hand pane, click on Applications, Application Types, and WebSphere enterprise applications.

		IBM.
	Staged Deployment	
Uninsta	all V1	
	Interprise Applications Save your workspace changes to the master of Click Save to update the master repository with begin work Make sure that this box is unchecked! 27 Synchronize changes with Nodes Save Discard Cancel	

Now in the application list panel, put a check mark in the box next to the My_IVT_Application and click on the Uninstall button. Click OK on the confirmation prompt. Then go through the Review and Save dialog, remembering to **Uncheck** the synchronization box. This will save the master configuration without the My_IVT_Application.

If you wish, you could now verify that the application (version 1) is still operating successfully in both members of the cluster..

Staged De	ployment
Install V2 Enterprise Applications Preparing for the application installation	
Specify the EAR, WAR, JAR, or SAR module to upload and install. Path to the new application O Local file system Full path Choose File MyIVTV2.ear	• •
Full path Browse Next Cancel	Cell=s1cell, Profile=default Enterprise Applications Enterprise Applications > Enterprise Applica Save your workspace changes to the master Click Save to update the master repository w
	2 begin work changes. B Total of box is box is Unchecked. Discard Cancel Cancel
	νς

We can now begin the installation of Version 2 of the application ear file into the master repository. This is the exact same process as we previously did with the exception of the use of the new ear file, and the elimination of the synchronization step during synchronize. Here are the two steps which differ:

If you were to check, both members of the cluster are still serving Version 1 of the application. I did and they were, but to save repetitive pictures, I'll skip displaying them here

	ie node
System administration	
Cell	
 Save changes to master repository Deployment manager 	
Nodes Nodes	
No <mark>Nodes</mark> ps	
- Console Identity	Nodes ?
	Nodes
	Use this page to manage nodes in the application server environment. A node corresponds to a physical computer system with a distinct IP host address. The following table lists the managed and unmanaged nodes in this cell. The first node is the deployment manager. Add new nodes to the cell and to this list by clicking Add Node.
	Add Node Remove Node Force Delete Synchronize Full Resynchronize Stop
	Note that
	Select Name 🗘 Host Name 🗘 Version 🗘 Di nodes show 🗘 Status 👲
	You can administer the following resources: up as being
	sinodec wsc3.washington.ibm.com ND 7.0.0.8 TCP
	sinoded wsc4.washington.ibm.com ND 7.0.0.8 TCP

The next thing to do is to get Version 2 working in one of the two cluster members. We'll use the member on s1noded. We do this by synchronizing that node, and that node only, with the master configuration. Start by clicking on System administration and then Nodes on the left hand pane.

Now check the box next to node s1noded and click on the Synchronize button.



Checking the member on s1noded we see that we are now running Version 2 of the application.



Once the node (s1noded) is synchronized, you can test.

Checking the member on s1nodec we see that we are still running Version 1 of the application.



Next we can synchronize s1nodec and verify that the member there is running Version 2



And sure enough, after the synchronization it is.

Restore synchroniza	ation settings
 Welcome Guided Activities Servers Applications Services Resources Security Environment System administration Cell Save changes to master repository Deployment manager Node agents Node group Console Preferences Console Identity 	Node agents > nodeagent > File synchronization service. Use this page to configure the file synchronization service. The file synchron deployment manager and node agent. It ensures that configuration change propagated to the appropriate node repositories. Configuration Configuration Recheck both boxes. Synchronization interval Synchronization Image: Synchronization Synchronization Synchronization Image: Synchronization Synchronization Synchronization Image: Synchronization I

All that remains is to return the nodeagent's file synchronization service settings to their normal state and we are done. To do this, simply reverse the changes we make earlier. Here are the screenshots.



Repeat for all nodes, then save and synchronize...

To return synchronization to normal operation, you'll need to recycle the nodeagents. This completes the update cycle.



The install of the application using the scripting interface is done in exactly the same way as with the adminconsole. The same steps are performed, in the same order, with the same result. The only difference is that you avoid all of the "point and click" work, and avoiding that manual work, you have the potential for a greater degree of accuracy.

The script fragments used in this demonstration have all been tested and are provided as part of the techdoc on this subject. Some of them depend upon the setting of a variable which is not set within the script. During testing, the variable was set interactively, followed by including the script file "fragment". This allows the reader to later use these files as-is as part of a larger "controller" script.

Staged Deployment Initial Application Install. appversion = import sys # Sample invocation: ./wsadmin.sh -lang jython -javaoption -Dscript.encoding=IBM-1047 -f /u/mjloos/tb/wp101641/myivtinst.py # Uninstall the app if it is there applist = AdminApp.list().splitlines() for ap in applist: if ap == "My_IVT_Application": print "Found and uninstalling ", print ap AdminApp.uninstall(ap) continue # Installing app # print "Installing application My_IVT_Application" earfile = "/u/mjloos/tb/wpl01641/MyIVT" + appversion + ".ear" print earfile appopts = "[-appname " appopts = appopts + "My_IVT_Application " appopts = appopts + "-reloadEnabled -reloadInterval 0 " appopts = appopts + "-preCompileJSPs " appopts = appopts + "-precompliedSPS " appopts = appopts + "-MapModulesToServers [['.*' '.*' " appopts = appopts + " WebSphere:cell=slcell,cluster=slsr03]]]" # print "appopts = " + appopts 32 AdminApp.install(earfile, appopts) # Save the config AdminConfig.save() print "All Done!" # All Done!

As with the adminconsole install, the first task is to install the application to the master configuration, specifying that it should be run on the cluster of interest (s1sr03 in our examples). We accomplish this by first logging on to either telnet or starting OMVS from TSO. This could also very well be done in batch using either BPXBATCH directly or as a command under the terminal monitor program (IKJEFT01). For this demonstration, we'll be running in a telnet session.

The first script used depends on a variable (appversion) being set to the proper version number for the application. The rest of the scripting statements can be found in file **myivtinst.py**.

Clearly, many of the parameters specified could be passed in as arguments to the script, but are "hard coded" here for clarity.



Again, since this is the first or initial install of Version 1, we will want to synchronize with all nodes. The script that we will use can be found in the file **nodesyncall.py.**

This script simply gets a list of all possible nodeagents in the cell and then invokes synchronization on each of them in turn.



The application is now installed and ready to start. To do this, yet another script fragment is used. This one can be found in **myivtstart.py**.

This script first gets the object id of the cluster, then gets a list of all members in the cluster. Each member has as an attribute the servername and nodename of the member. This is used to get and invoke the mbean for starting the application.

Once the application has been started, we could (and should) test the application. This would be done exactly the same way as in the manual install, so we'll skip showing it here..



Next we'll begin a staged update using scripting. Again, it works exactly the same as in the manual, adminconsole, point and click version. The first thing we need to accomplish is to get automatic synchronization disabled. The script file we'll be using is **autosyncoff.py**.

This script gets a list of all of the synchronization services in the cell, and for each one sets the two variables (synchOnServerStartup and autoSynchEnabled) to false. Next we need to synchronize all nodes again, and the same script we used previously for that will work here as well (**nodesyncall.py**).



We now need to restart the nodeagents to make the previous change take effect. The script we'll use this time is **restartnodeagents.py**.

This script gets a list of nodeagents in the cell, and for each nodeagent it extracts the node name. The node name is used in a parameter used to invoke the query-Names service. The result of that service is used to invoke a nodeagent restart service. The two "false" parameters are specified to prevent a synchronization from being run before the nodeagent is restarted, and to prevent all running application servers on the node from being restarted when the nodeagent restarts



We can now install version 2 of the application. Just as we did in the adminconsole method, we'll approach this by first uninstalling (to remove all traces) of the application, then installing the version 2 copy of the ear file. Again, remember that all synchronization is disabled, so any changes we make to the master configuration have no effect on the servers which are members of the cluster.

We can use the same script as we did previously with the minor change of setting the **appversion** variable to **V2** and using the script, **myivtinst.py**.

If you were to check at this point, you could verify that both servers are still running V1 of the application, since we have not synchronized with either node.



The next step is to synchronize the master configuration with just one node (we'll use s1noded), just as we did previously. The script we'll use is **nodesyncone.py**. To indicate to the script which node we wish it to operate upon, we need to set the variable, **nodetosync**.

This script operates similarly to the nodesyncall.py script, with the simple change of conditional statement to check to see if we are operating on the correct node.

We can now test the application on each of the two nodes, and as before, verify that we are now running version 1 on s1nodec and version 2 on s1noded.



Next we can synchronize s1nodec and verify that the member there is running Version 2.

This simply involves synchronizing the master configuration with the other node (s1nodec). Again we'll use the same script, **nodesyncone.py**. This time we'll set the variable, **nodetosync** to s1nodec.

You may now test to verify that Version 2 is running on both nodes.



All that remains is to return the nodeagent's file synchronization service settings to their normal state and we are done. To do this we run a script named **autosyncon.py**.



Next we need to synchronize all nodes again, and the same script we used previously for that will work here as well (**nodesynchall.py**).

We now need to restart the nodeagents to make the previous change take effect. The same script we used previously will work again (**restartnodeagents.py**). This completes the use of the scripting interface to accomplish a staged application update.





The following script uses the provided scripts to install the application the first time, synchronize the nodes, and start the application. It uses the provided scripts by using the execfile command to run them. This script is provided as file **initialinst.py**.



This script uses the provided scripts to turn off automatic synchronization, synchronize the nodes, restart the nodeagents, install V2 of the application, and synchronize one node. This script is provided as file **stagesetup1.py**.



This script uses the provided scripts to synchronize the other node, turn on automatic synchronization, synchronize the nodes, and restart the nodeagents. This script is provided as file **finishup.py**.

