



# **WebSphere Liberty z/OS**

## **Applications and Application Deployment**



## Objective of this Presentation



- Provide an understanding of the application types supported by Liberty
- Provide a general understanding of the API model of Liberty, particularly as it relates to the API model of WAS traditional
- Provide an understanding of the deployment model of Liberty



For this unit we will explore the topic of applications and application deployment within Liberty. Even though our focus is z/OS, this material relates to Liberty on all platforms, as the programming APIs are the same across all platforms. The deployment model is also the same across all platforms, with the possible exception of the ability to share a file system between LPARs in a Sysplex environment. (You could achieve something similar with distributed systems as well. The point here is that in a Sysplex environment the sharing facility may be present, as it commonly is, where on a distributed system it may not be.)

This unit does not go into programming specifics or examples, as that is outside the scope of what we're trying to accomplish with this collateral.



## Application Types Supported by Liberty



### Web Application

- JSPs and Servlets
- Accessed over the network using a client (browser, or a REST client)
- Packaged in a Web ARchive (WAR) files



### Enterprise Application

- EJBs
- Accessed using network protocols (RMI/IIOP) or message queueing (JMS)
- Packaged in Enterprise ARchive (EAR) files



### OSGi Application

- Java-based architecture that implements a dynamic component model
- Can be installed, started, stopped, updated, and uninstalled without requiring a restart
- Packaged in either Enterprise Bundle Archive (EBA) or Enterprise Subsystem Archive (ESA)



### "Loose" Application

- Not a different kind of application (it's a different way to provide application components to server)
- This allows your application components to be located in different places and run in Liberty
- Particularly good for development: run applications without export and deploy

This chart is intended to provide a brief survey of the types of applications supported by Liberty, and the packaging models for each.

- **Web application** -- this includes JSPs and servlets, and is accessed over a network, typically with HTTP protocol. The packaging as a "WAR" file (Web ARchive). This type of application was supported by WAS traditional as well.
- **Enterprise Application** -- this includes EJBs, and is accessed with protocols such as RMI/IIOP, or message queuing. The packaging is an "EAR" file (Enterprise ARchive). This type of application was supported by WAS traditional as well.
- **OSGi Application** -- the acronym stands for "Open Service Gateway Initiative," and is based on the idea of a dynamic component model. Liberty itself is based on an OSGi design. OSGi applications are packaged in either EBA (Enterprise Bundle Archive) or ESA (Enterprise Service Archive) file formats.

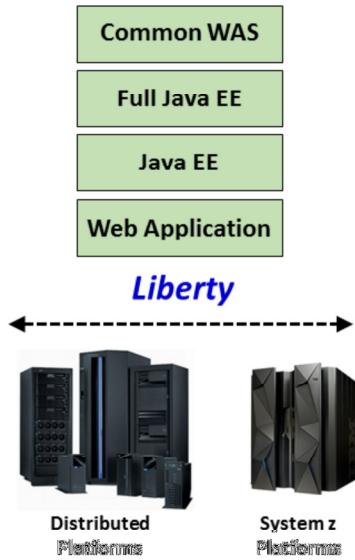
**Note:** as of the time of writing this, WAS traditional supported a few more OSGi APIs than does Liberty, but that gap is closing with each quarterly fixpack.

- **"Loose" Application** -- this is not really a different type of application, rather it is a different mechanism by which to deploy applications to Liberty. This is based on the idea of deploying the artifacts of an application without packaging into an archive file. They can be located at any location, and the location is specified in the server.xml. This is a good format for development environments where a quick change to some file can be done without doing a full repackage / export / deploy. This is probably not the mechanism to be used for production environments, where change control will be focused on the archive file tested and validated prior to deployment to production.

Your Liberty environment may run some or all of these application types.



## Liberty APIs Across the Platforms



### Liberty is supported across many different OS platforms:

- Windows, AIX, HP-UX, Solaris, IBMi, Linux, Linux for System z, z/OS

### Same programming APIs across the platforms:

- When comparing the "Network Deployment" level of Liberty
- Distribute platforms have other "editions" which have different subsets of the full Liberty features (Base, Core, Express, Developers)

### All else equal, applications can move across platforms

- "All else equal" -- same version of Liberty, Liberty features configured the same, data connectivity definitions are in place, security requirements are the same, etc.

## What about WAS traditional vs. Liberty?

4

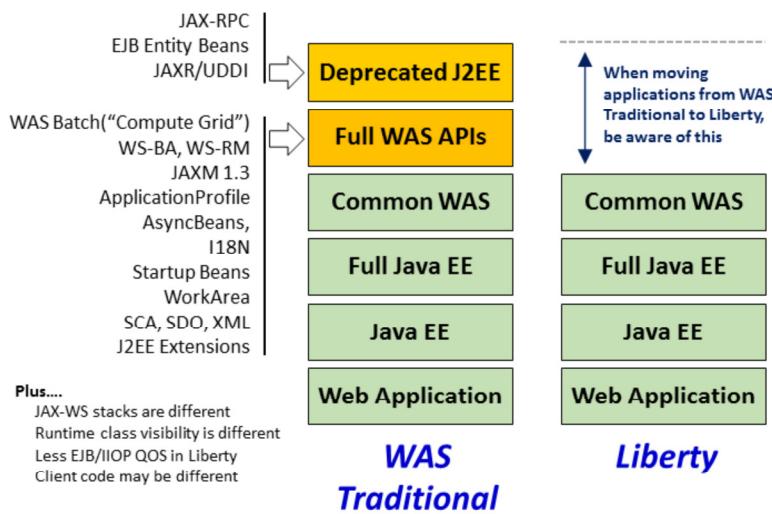
A central design feature of Liberty is that the programming APIs are common across all platforms when comparing like-to-like releases ("editions") of Liberty. The only edition of Liberty supported on z/OS is what's called "network deployment." If you compare that edition on z/OS to that same edition on Windows, AIX, HP-UX, Solaris, IBMi, or Linux, you'll have the same API set (assuming the version and fixpack level is the same).

The value of that is it allows you to develop on one platform and deploy on another without having to worry about programming API incompatibilities. Now, you have to make sure the application has what it needs on the target platform for this to hold true. For example, you can't take an application that runs in Liberty on Windows and move it to a Liberty z/OS server that does not have the same configuration definitions for things like JDBC. Or if the application relied on some security artifact defined to Linux but not present on the target z/OS system. But it holds true that if we go from like-to-like, then the identical programming model provides portability.

That's Liberty-to-Liberty application portability; what about WAS traditional to Liberty?



## WAS traditional APIs Compared to Liberty APIs



**The API set is very similar, but they are not exactly the same**

### Application mobility:

- WAS traditional to Liberty -- be aware of deprecated APIs and any use of "Full WAS" APIs in applications
- Liberty to WAS traditional -- fewer concerns about API differences

**An toolkit to evaluate applications exists ...**

5

When comparing WAS traditional with Liberty (any platform), we see there is an "API gap" with Liberty. That "gap" is comprised of two categories of APIs:

- **Deprecated J2EE** -- these are APIs that are part of the open standard but are marked deprecated. They have been superseded by more current standard specified APIs. Applications use of deprecated APIs is not a good practice, and applications that do make use of those APIs should be investigated to bring up to current API standards. These applications may still work in WAS traditional (because the APIs are deprecated-but-not-yet-removed). But on Liberty the applications will throw an error since the API they are seeking to use is no longer present.
- **"Full WAS" APIs** -- these are APIs IBM supplied over and above the common open standard APIs. There is a set of those "above and beyond" APIs that are not in Liberty. The chart shows some examples of this. An application written to make use of any of those APIs would not work if carried to Liberty.

**Note:** there is an excellent write-up on this topic:

<https://developer.ibm.com/wasdev/docs/was-classic-or-was-liberty-how-to-choose/>

This begs the question: is there a way to evaluate applications to see what issues may exist? The answer to that is yes. We point you to that on the next chart.



## Reference: Application Migration Toolkit

Your application binaries



Migration Toolkit



Product Edition	WebSphere Application Server V8.0.5						
	Liberty for Java on IBM Bluemix	Liberty Core	Liberty	WebSphere Application Server	Network	Network Environment	Liberty for iOS
Web app							
Implementing Java API for Application Services	✓		✓				
Java API for RESTful Web Services (JAX-RS)							
Java Server Faces (JSF)	✓	✓	✓				
JavaServer Pages (JSP) 2.2	✓				✓	✓	✓
JavaServer Faces (JSF) 2.0	✓				✓	✓	✓
Enterprise							
Java Database Connectivity (JDBC)							

Summary report of technology used in application and target environments where application can be deployed

File name	Reference details	Match criteria	Line number
PlantsByWebSphere.war/WEB-INF/classes/com/ibm/websphere/lsamples/pbw/war/ShoppingBean.class	Method getShippingCostString java.text.NumberFormat.formatted(double)	✓	155
	Method getTotalCostString java.text.NumberFormat.formatted(double)	✓	175

Detailed report by file name, method name and line number

Main wasDev page:

[https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration\\_Toolkit\\_for\\_Application\\_Binaries](https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration_Toolkit_for_Application_Binaries)

Technical Overview:

<https://developer.ibm.com/wasdev/docs/migration-toolkit-application-binaries-tech/>

Updates page:

<https://developer.ibm.com/wasdev/blog/2015/03/13/announcing-websphere-liberty-migration-tools-updates/>

6

IBM has produced an application migration toolkit to assist you with analyzing applications and assessing the portability to the target platform you specify. This toolkit works on program source as well as program binaries. The chart provides the URLs, which are reproduced here in case you can't use the hyperlink from a chart format of this deck:

[https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration\\_Toolkit\\_for\\_Application\\_Binaries](https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration_Toolkit_for_Application_Binaries)

<https://developer.ibm.com/wasdev/docs/migration-toolkit-application-binaries-tech/>

<https://developer.ibm.com/wasdev/blog/2015/03/13/announcing-websphere-liberty-migration-tools-updates/>



## The WP102110 Techdoc\*

*Located at the bottom of the Techdoc page:*

**Liberty or WAS Traditional ... How to Decide?**

Additional function has been added to Liberty since it was first introduced. It is now Java EE 7 compliant and supports the Java 8 SDK. So the question arises "What should I consider when deciding between using WAS traditional and Liberty z/OS?"

The following guide provides some considerations to take into account when comparing the two.

[WP102110 - Liberty or WAS Traditional - CHARTS.pdf](#)  
[WP102110 - Liberty or WAS Traditional - NOTES.pdf](#)

**Framework of the presentation:**

Hyperlink

- ▶ **Executive Overview**  
A one-chart summary of the usage considerations presented in the document
- ▶ **Setting Context**  
Establishing terminology and providing background on the evolution over time of each runtime models
- ▶ **Application Considerations**  
Exploring the application interface considerations of each runtime model
- ▶ **Operational Considerations**  
Exploring the runtime operational considerations of each runtime model
- ▶ **Performance Considerations**  
Exploring the performance profile of each runtime model
- ▶ **Other Information for Consideration**  
A collection of other information you may find useful when making this decision

**At that Techdoc page there's a section devoted to going through the considerations in a systematic manner**

**No "formula" ... the key considerations are around:**

- The application design and the APIs it uses
- The degree of reliance on the CR/SR structure of WAS traditional on z/OS
- The degree of reliance on automated scripting (WSADMIN)
- Memory and GP (Liberty tends to use less of both)

**For rest of presentation  
we assume Liberty**

7

\* <http://www.ibm.com/support/techdocs/atstrmstr.nsf/WebIndex/WP102110>

To further assist with the broader question of which runtime to use -- WAS traditional z/OS or Liberty z/OS -- we have produced a Techdoc that takes you through a structured discussion of the considerations. The Techdoc comes in both chart format and speaker-note format. It can be found here:

<http://www.ibm.com/support/techdocs/atstrmstr.nsf/WebIndex/WP102110>

For the rest of this unit, we will assume the decision has been made to consider Liberty z/OS, and we will focus on that.



## Liberty Application Development Good Practices



**Liberty is a Java EE application server, so there's nothing unique about it compared to other Java EE runtimes**

### General good practices Java development ...

- Use good design practices
- Maintain good source control
- For best performance, make heavily used code as efficient as possible
- Profile the application prior to deployment into production
- Use proven change control processes

8

This deck is not intended to focus on Java programming good practices. That topic is fairly well documented elsewhere, and with a platform agnostic language like Java part of the discussion, the good practices are even more broadly-based.

It seems almost cliché to say it, but it's true: if you want a good Java program, then design the Java program to be good. There is a well known set of good coding practices in general, and Java specifically, and it starts with the program objectives and program design. Then: practice good source control, test for performance and where heavily used code is involved, make that efficient as can be accomplished. Profile the application and find out where the hot spots are, and correct areas that are hampering performance and throughput. Finally, exercise good and proven change control processes. A million dollar infrastructure can be made useless by a bad change introduced at the last minute without proper testing and control.

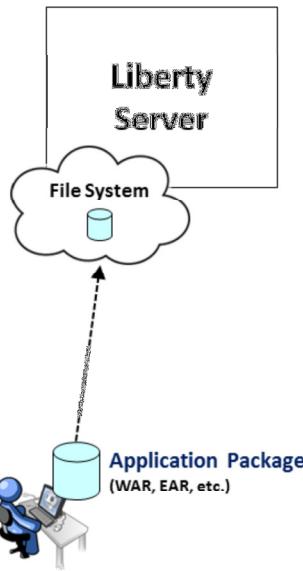


# Application Deployment

Deploying applications into a Liberty z/OS runtime environment



## The Liberty Deployment Model



### Two choices:

1. Drop file in `/dropins` directory and let Liberty detect and dynamically load
2. Place file in another location and configure application into `server.xml`

**Use whatever mechanism you wish to get application package file from your development / source control environment to a file system accessible by the server**

**Then the question is: Dynamic update? Or rely on server restart?**

- In general, we see production environments avoiding dynamic updates.
- If dynamic updates are disabled, then application are loaded at server restart
- If dynamic updates, the trigger mechanism (polled, mBean) can be configured

10

The deployment model for Liberty is quite a bit different than the deployment mode for WAS traditional. WAS traditional required the application be deployed using the administrative interface: either the GUI Admin Console, or the WSADMIN scripting interface. That's because application deployment in WAS traditional involved unzipping the application archive file and updating a number of XML files to let the runtime know about the application.

Liberty is a simpler model -- either (a) have dynamic updates enabled and place the application file in the `/dropins` directory, or (b) place the application file in a some directory and point to it with an XML element in `server.xml`. You can use pretty much any deployment tool to accomplish this.

The next question is whether you wish to rely on the dynamic update capability of Liberty, or whether you wish to control things by mandating a server stop and restart to pick up changes. Let's explore those options over the next few charts.



## Static Definition of Application

### Application Location Pointer from `server.xml`

- 1 `<application location="<app\_pkg\_name>" />`

The server will search both the /apps directory and the /shared/apps directory for the application.
- 2 `<application location="${server.config.directory}/apps/<app\_pkg\_name>" />`

The variable resolves to the server's configuration directory. This definition points to the /apps directory under that.
- 3 `<application location="${shared.app.directory}/<app_pkg_name>" />`

The variable resolves to the /shared/apps directory under WLP\_USER\_DIR.
- 4 `<application location="/<full_path>/<app_pkg_name>" />`

You may provide an explicit pointer to a path and file and load application from any accessible location.

`<config updateTrigger="polled" monitorInterval="500ms"/>`

Controls whether the server.xml change is dynamically detected and acted upon.

11

Let's first focus on the static definition of a file. This is done in the `server.xml` file, or a file that is merged into the `server.xml` file with an `<include>` statement.

**Note:** "static" in this case does *not* rule out the definition being dynamically loaded. It simply means the location and name of the application file is specified in the server configuration. The *configuration* change can still be dynamically loaded, and thus the *application* dynamically loaded based on the new definition in the configuration. The XML shown at the bottom of the chart indicates what controls whether configuration dynamic updates take place, and if so, how often.

You have several options with respect to how you define where the application resides:

1. The server will look in either the /apps directory under the server directory, or it will look in the /shared/apps directory under the WLP\_USER\_DIR location. If the named application is found in either place, it will be loaded.
2. This uses the `${server.config.directory}` variable to resolve the location to the server's directory. The /apps string after that tells Liberty to look in that directory and load the application. This is useful when you are bringing in the application definition using `<include>` processing. Two servers (in a cluster, for example) could merge in this element, and each would resolve the path to *its* server configuration directory.
3. This uses the `${shared.app.directory}` variable to resolve the location to /shared/apps under the WLP\_USER\_DIR. Here again, this would be useful when pulling in common XML where the application to be loaded by different servers was found at that location. A cluster where the two servers are under the same WLP\_USER\_DIR could use this pull in the application from a single location.
4. This is a full path pointer to the application location. This can be anywhere; the only requirement is the server ID have read to the location.

The gray box at the bottom provides an illustration of the XML that will control whether the configuration changes will be picked up dynamically. For example, `<config updateTrigger="disabled" />` means there is no dynamic update, and the application definition will be picked up at next server start. That may be what you want for production. Or you can specify `<config updateTrigger="polled" monitorInterval="60s" />` to indicate every 60 seconds the server will poll for changes and load the changes.



## Application Load from a "Dropins" Directory

### *Dynamic update from "dropins" directory*

The default location is the `/dropins` directory under the server directory

```
<applicationMonitor dropins="${server.config.dir}/myApps" />
```

Sets "dropins" directory location to a directory you specify under the configuration directory

```
<applicationMonitor dropins="${shared.app.directory}/myApps" />
```

Sets "dropins" directory location to a directory you specify under the shared apps directory

```
<applicationMonitor dropins="/<path>/myApps" />
```

Sets "dropins" directory location to a location using an absolute path

```
<applicationMonitor dropinsEnabled="false" />
```

Disables "dropins" monitoring and dynamic loading of applications

12

You can avoid statically defining applications in your server and take advantage of the ability of Liberty to monitor a directory for application files and load them when seen. By default that will be the `/dropins` directory under the server's directory, and you can use that if you'd like. But if you'd like to have the "dropins" directory somewhere else you can do that. The `<applicationMonitor dropins=" " />` value is what defines the location.

The first two examples are showing the use of some built-in variables. The  `${server.config.dir}` will resolve to that server's configuration directory (which is where `server.xml` resides). The  `${shared.app.directory}` will resolve to the `/shared/apps` location under the `WLP_USER_DIR` location.

**Note:** those are not the only two built-in variables. There are many more.

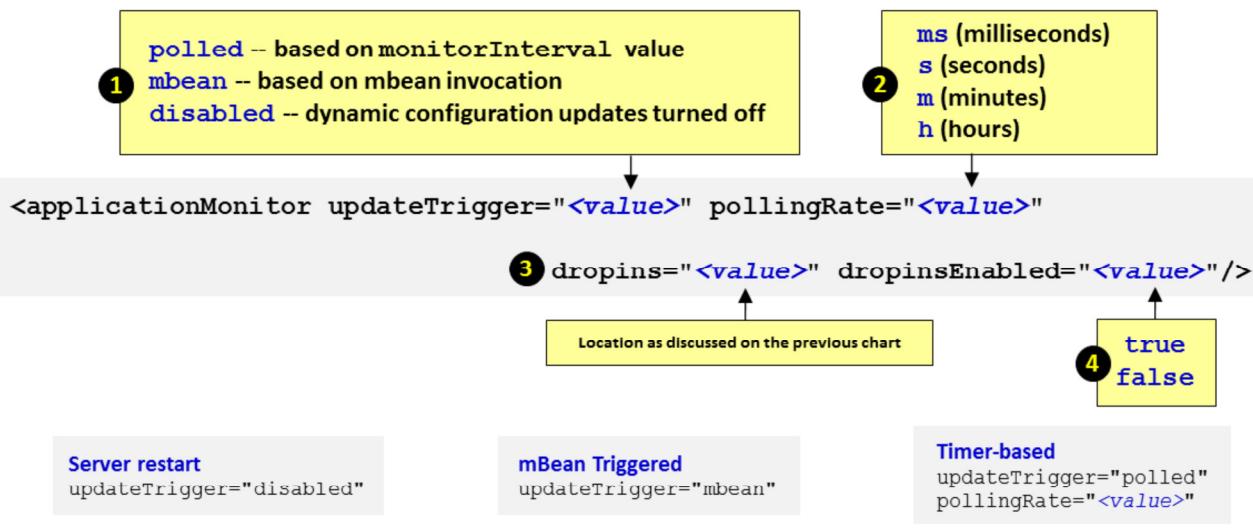
The third example shows an explicit path to a location. This can be anywhere. The only requirement is the server ID must have read access.

Finally, the last example on the chart illustrates how to turn off dynamic loading from whatever "dropins" directory you specify. With `dropinsEnabled="false"` set, the server will not look for or load applications from either the default `dropins` directory or a defined one. It won't load from there even on a server restart. So `dropinsEnabled="false"` is how you can force the use of statically-defined applications.



## If Dynamic Update, then Controlling When Dynamic Update Takes Place

**Note:** dynamic update and "dropins" are related, but are not the same thing. You can have a statically-defined application and replace the package with a new file. Liberty can detect change and reload dynamically if you choose.



13

Here is the `<applicationMonitor>` element one more time. Let's walk through this and review the different values and what they do:

1. **updateTrigger** = this determines *how* updates will be accomplished. The default is "polled."

**Note:** this applies to applications deployed with a static definition or deployed via a "dropins" directory.

- If "polled," and "dropinsEnabled" is either unspecified or set to true, then the "pollingRate" value is used to determine how often the deployed are checked for updates.
- If "mbean," then it will check for changes to the deployed applications only when the management bean is invoked. This provides a way to perform application updates without polling and without a server restart.
- If "disabled," then the applications are not checked; a server restart is required to pick up changes.

2. **pollingRate** = if "updateTrigger=polled," then this determines how often the polling takes place. Your options are expressed in milliseconds, seconds, minutes, or hours. The default is 500ms, or one-half second.

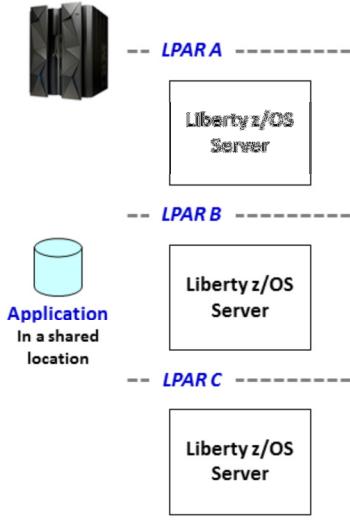
3. **dropins** = this defines the directory path location where the server will look for applications to deploy. This defaults to the /dropins directory. This is ignored if "dropinsEnabled" is set to "false." This is how you would control the ability of others to deploy an application simply by copying in a file.

4. **dropinsEnabled** = this determines if processing of a "dropins" directory is done at all. If "true" (which is the default), then either the default location is checked (/dropins) or the value seen on "dropins=". If set to "false," the server does no processing of applications out of any dropin directory; that means application deployment must be accomplished with a static definition of the application path and file name.

We anticipate for tightly controlled production environments applications will be statically defined and loaded with a server restart. In that case code "updateTrigger=disabled" and "dropinsEnabled=false".



## Rolling an Application Update Across LPARs



### Framework of the approach:

- Assume application is in a shared location accessible to all servers  
If each server has its own unique copy of the application, then it's a matter of updating each unique copy.
- Dynamic update is either disabled or mBean invoked (that is, not polled)
- Update shared copy of application
- Update in each server -- either server restart or mBean-invoked update

### Potential complicating factors:

- When session affinity is required due to application design  
Then you need to insure session persistence is enabled so affinity can be re-established
- Stopping the flow of work to a server in which the application is to be updated  
This is a function of the front-end routing mechanism you're using to route across LPARs
- Cases where the application update implies simultaneous mixed-levels can't be tolerated; for example: a significant change to the backend data model  
In this case you may need to schedule a Sysplex-wide update during a maintenance window

14

The question often comes up how one would achieve the "rolling" of an application through a set of like-configured servers (which could be a Liberty collective "cluster" but it does not have to be). The purpose of "rolling" an application is to maintain at least one instance of the application up and servicing clients while other instances of the application are refreshed (which implies a brief time when it's not available).

Doing this implies the servers are *not* operating with a polled dynamic update.

**Note:** polled dynamic update would work if the polling interval was sufficiently long to prevent all the applications being out of service simultaneously, and also that the polling intervals don't "pop" at the same time; that is, they are staggered. That's a fair amount of conditions. It's easier to assume either an mBean update or a server restart update.

It's assumed a shared application is used between the servers -- after all, if each server had its own copy the application you'd simply deploy the application into each server in turn and roll through the servers manually. So with those assumptions in place (no dynamic update, and a shared application), then it's just a matter of updating the shared application and then triggering the update, either with an mBean or a server restart.

Some things that complicate this scenario:

- If the applications are maintain session *state*, then that implies affinity, which implies the need to make sure any state information is available in the other servers if an established client lands elsewhere during the rolling of the applications. This is possible using session persistence (copying state information out to a database, then retrieving it back to another server). So while this makes rolling updates a bit more complicated, it's not impossible to engineer for this.
- When an application is being refreshed, there will be a brief period of time when it is not available. Clients routed to that server will get an error in that brief period of time. This can be avoided if the front-end work distribution function is capable of "turning off" work to a given endpoint. Work can then "drain" from that endpoint, the application restarted, and the router out front "turned back on" to allow work to flow to the updated server.
- "Rolling" an application implies a period of time when you have mixed levels of the application available to clients. That may be okay, or it may not be okay if the nature of the change being introduced is so significant you can't have mixed levels. If that's the case, then you will have to schedule an outage and refresh all instances of the application at the same time.



## Application Monitoring?

**Operations, Performance and Problem Determination**

**Monitoring**

*Application Monitoring*

**Monitoring your applications is an important subject**

**It's part of the broader "monitoring" topic**

**Which is part of the even broader topic on operations, performance and PD**

**We have an entire unit dedicated to those topics.**

15

What about application monitoring? It's an important topic ... not just application monitoring, but the broader topic of monitoring in general. We have a unit dedicated to monitoring, as well as operations and problem determination.



## Summary

**Application deployment is fairly simple -- upload application package and make it available to the Liberty z/OS server. No "deploy through Admin Console" needed.**

**The key questions are:**

- Will the /dropins mechanism be used? Or <application> tag reference to file?
- How will updates be handled -- manually, mBean invoked, or timer-based polling?

16

And we're to the summary page. Liberty application deployment is relatively simple because it is really just a matter of uploading the file and letting the server know about it. There are two things we discussed: (1) using "static" application definitions (an XML pointer to the application path and file name) vs. using the "dropins directory" location; and (2) whether application changes are dynamic, or are more manually controlled by you. We anticipate the dropins/dynamic model will be used for development and test, while static/manual will likely be the model for production settings. But you may choose differently based on your own set of criteria. In either case, Liberty z/OS can be configured to support what you want it to do.



## Reference

### WebSphere Liberty 16.0.0.x Knowledge Center

[http://www.ibm.com/support/knowledgecenter/en/SS7K4U\\_liberty/as\\_ditamaps/was900\\_welcome\\_liberty\\_zos.html](http://www.ibm.com/support/knowledgecenter/en/SS7K4U_liberty/as_ditamaps/was900_welcome_liberty_zos.html)

### WebSphere Knowledge Center collection on the topic of migration

<http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg27008724>

### Migration Toolkit for Application Binaries

[https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration\\_Toolkit\\_for\\_Application\\_Binaries](https://developer.ibm.com/wasdev/downloads/#asset/tools-Migration_Toolkit_for_Application_Binaries)

<https://developer.ibm.com/wasdev/docs/migration-toolkit-application-binaries-tech/>

<https://developer.ibm.com/wasdev/blog/2015/03/13/announcing-websphere-liberty-migration-tools-updates/>

17

Here is a reference page that lists a few resources on the topic of Liberty z/OS and migration.

End of Unit