

This is part 1 of 2.

Day	Time	#	olication Server Ses	Speaker	Rogman
Wednesday	3:00	9483	Using IBM's New Cross-Platform Installer on z/OS	Mierzejewski	Oceanic 5
Thursday	8:00	9482	WAS Version 8 – Overview	Follis	Europe 2
Thursday	9:30	9486	WAS Version 8 – Batch Update	Hutchinson	Europe 2
Thursday	11:00	9485	WAS Version 8 – New z/OS Exploitation/Differentiation Features	Follis	Europe 2
Thursday	1:30	9484	WAS Version 8 – High Availability Enhancements	Follis	Europe 2
Thursday	3:00	9488	WAS - Back to Basics Part 1	Loos	Europe 2
Thursday	4:30	9489	WAS - Back to Basics Part 2	Stephen	Europe 2
Friday	8:00	9490	WAS for z/OS - Level 2 Update	Stephen	Europe 2
Friday	9:30	9487		Follis, Hutchinson, Loos, Mierzejewski, Stephen, etc.	Europe 2





This presentation is the accumulation of experiences installing WebSphere on z/OS at customer locations across the US in combination with those acquired working with the folks from the WSC, with input from WebSphere on z/OS level 2 support.



High Level Overview of WebSphere Application Server

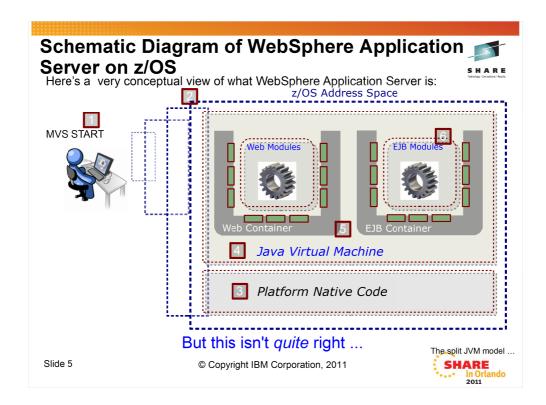
Slide 4

© Copyright IBM Corporation, 2011



This overview is extracted from the WBSR7 wildfire class, the full content of which is available on the techdocs website at the following url:

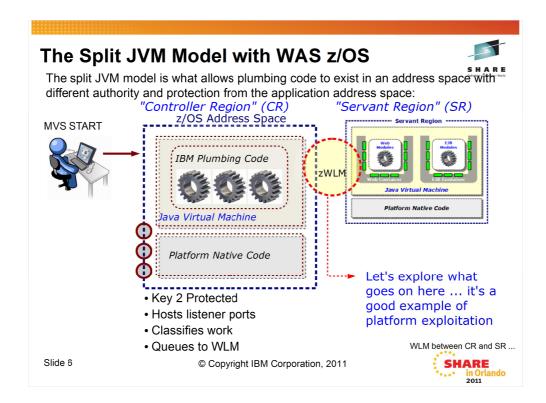
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3422



This chart seems busy on the surface, and it does in fact convey some complex things, but the key message is one of how WebSphere Application Server on z/OS provides the Java runtime environment on the platform. Let's walk through the progression so the picture can painted:

- 1. A WAS z/OS application server is started with a standard MVS START command.
- 2. On z/OS that implies an address space, which is roughly analogous to a UNIX process.
- 3. The initial code to start is not Java at all, but native code (C++) compiled for the z/OS platform. This is what brings up the initial lower-level framework and plumbing code.
- 4. Once the foundation is set, then the Java environment is started up.
- 5. With the JVM in place, WAS can now load the various Java class files that implement the open standard APIs and the function behind it. They are organized into two "containers" -- one for web modules and one for EJB modules.
- 6. Finally your application modules are loaded in and started.

That is the essence of it -- a foundation written in native code that launches the Java environment and then loads in additional Java services and structures to implement the Java EE server specifications. The standards and specifications are common to all vendors -- the way they're *implemented* is up to the vendor.

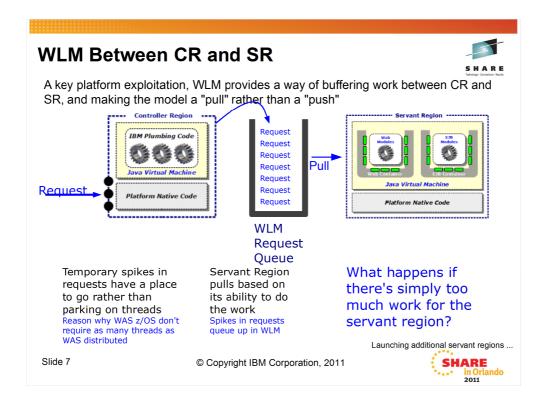


WAS z/OS is different from WAS on other platforms in that each application server operates with a "split JVM" model -- a "controller region" and a "servant region." The CR provides the initial handling of requests and does workload classification. The SR is where the applications run and where the "container" structures are implemented.

Note: for those with some familiarity with WAS z/OS ... yes, multiple SRs is possible. We'll cover that soon

Between the CR and the SR sits the z/OS Workload Manager (zWLM, or just WLM). We put the "z" on the front of that because WAS on the other platforms has something termed "workload manager" but it is not the same thing as z/OS WLM.

WLM is what starts the SR after the CR initializes. WLM is also what maintains work request queues, which is what we'll explore next.

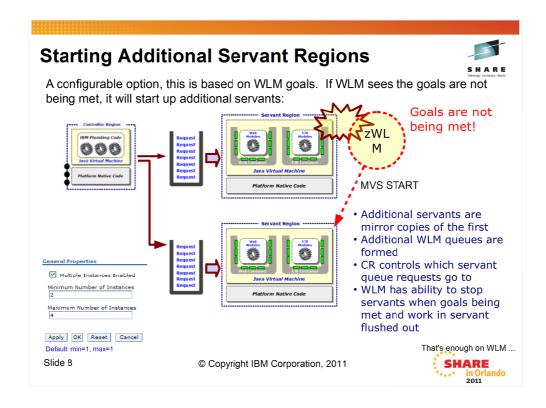


Sitting between the CR and the SR is a WLM work request queue where work is placed prior to the SR taking the work to service the request. This is what provides the ability to take a spike in requests and not overload the system. The CR will take the work in and park it on the WLM queue. These are not execution threads, these are very low overhead memory structures.

The SR takes work as it's able. It won't take more than it can because it's just a matter of a thread freeing up and taking the next unit of work.

Note: there's actually quite a bit more sophistication in the middle of all this. We're simplifying things here to make key points.

What happens if the inbound work is too much for the servant to service? If you have it configured to allow it, the CR can ask WLM to start up additional servant regions. That's next.

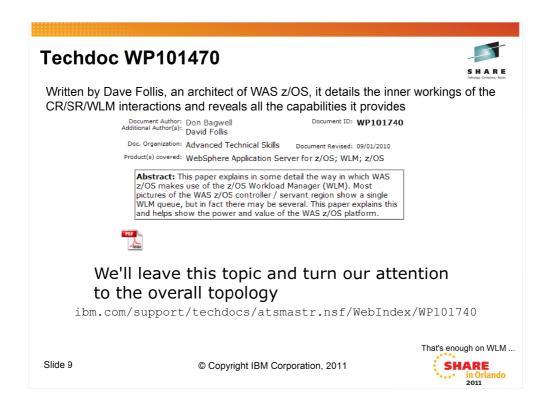


WLM is a very sophisticated system resource monitoring and allocation control mechanism. We can't go into all the details of WLM here, but suffice to say that WLM watches the state of all activity and compares against the work goals you've defined.

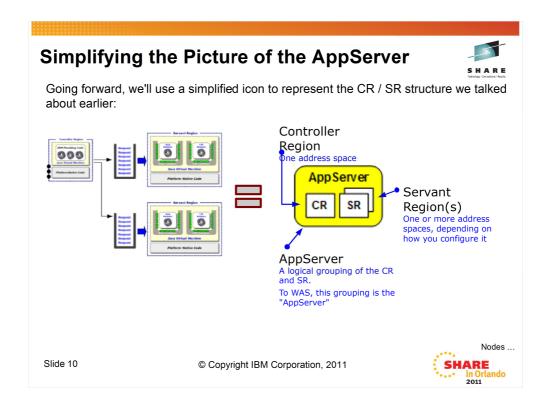
Imagine a servant region is taking work as fast as it can, but WLM still sees that the defined goals are not being met. If you've allowed it (it's a configurable option), WLM will start additional servants. Doing this results in additional WLM work queues being created and the CR placing work requests on the queues of the servants.

Note: there is considerable sophistication in the middle of this ... far more than we can cover on one chart here. We'll point you to a technical document that explains all this in much closer detail.

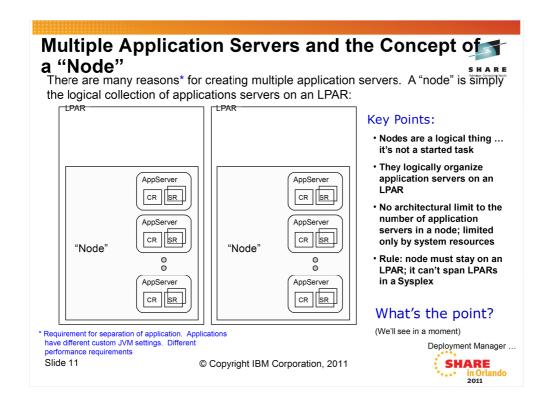
When the workload surge cools off, WLM has the ability to stop allocating work to the excess servants and allow work there to flush out. Then stop that servant.



If the CR / SR interaction with WLM interests you, by all means pull Dave Follis' Techdoc on the subject.



Going forward from here we'll use a little stylized icon to represent the CR / SR structure. The icon we'll use is shown on the chart. The square blocks represent z/OS address spaces; the curved outer box represents the logical collection which WAS itself views as the "appserver."

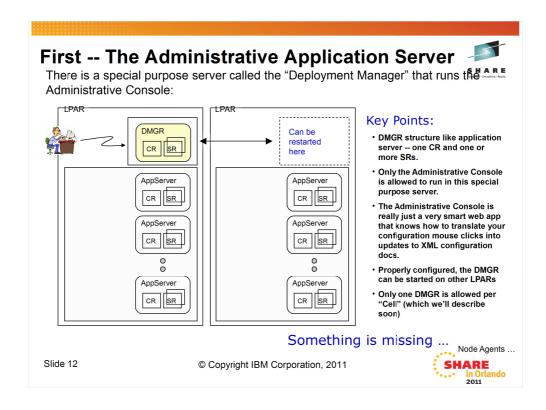


Okay, let's now explore what the other things are in the WebSphere z/OS configuration picture. We'll start with the notion of a "Node." This is actually a cross-platform WebSphere concept. A node is really just a logical collection of servers on a given LPAR. WebSphere maintains the concept of a "node" because that's how it ties servers to a configuration file system, but that's a point we're not quite ready to fully understand. For now, just lock in on the idea of a node being a collection of servers on a given LPAR.

There is no limit to how many servers you can have in a node. It's really a question of resources on the LPAR to support the address spaces.

The rule is this -- a node can't span multiple LPARs. By definition, a node must stay on an LPAR. If you have multiple LPARs in your Sysplex, you would define multiple nodes, such as the picture above illustrates.

Okay ... but why? What's the point of this? We'll see that in a moment. But first we have to introduce the "Deployment Manager," which is a special purpose server that runs the Administrative application.



As mentioned, the Deployment Manager is a special-purpose application server designed to run only one application ... the IBM-supplied administrative application. The DMGR server looks just like any other application server with a CR/SR structure.

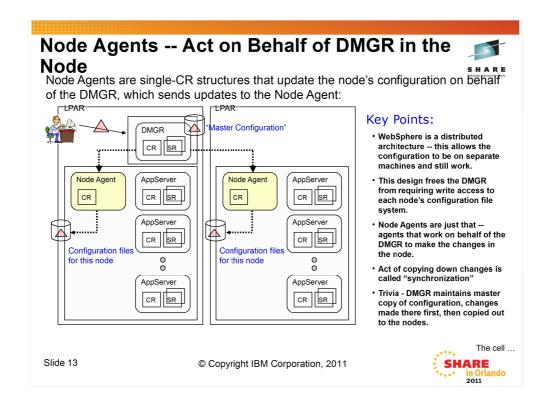
The Administrative Application, or "Console," is really just a very smart web application that knows how to translate your mouse clicks and keypad entry into modifications to the configuration structure, maintained in XML files. You could hand-modify the XML, but what a mess that would be. First, it would take a lot of knowledge of what XML to update and how, and secondly it would mean a typo could keep things from working right. So rather than force you to do that, the Administrative Console does all that updating-of-XML for you. You see a pretty GUI.

The DMGR is capable of being started on another LPAR if you configure things properly. (How that's done is beyond the scope of this presentation, but it involves the use of Sysplex Distributor.) This provides a way to maintain the configuration capabilities of the Administrative Console during periods of planned (or, let's hope note, but it's a possibility -- unplanned) outages of the LPAR.

Note: this is a good time to point out that the DMGR is **not** required for the steady-state operation of the other servers in the configuration, including your applications. The DMGR can be down and your applications, running in application servers, can happily continue on.

You can't have more than one DMGR per administrative "domain," or "cell." This is a definitional restriction of WebSphere -- one DMGR per cell. But again, it's restartable on another LPAR and it's not strictly a critical piece of the application serving role of WebSphere. Just configuration updates.

But there's a piece missing between the DMGR and the other servers. How does the DMGR get configuration changes out to the nodes? That's explained next.



The next piece of this puzzle is the "Node Agent." They're shown in the picture as yellow curved boxes. But wait ... they only have a CR, but no SR. Is that a mistake? No ... Node Agents are pure "plumbing" fixtures -- they only need a CR.

But what do they do? As the name implies, they serve as an "agent" for the node. In particular, they are what receives configuration updates made in the DMGR and apply those changes to the node configuration file system. This is done across the TCP network and involves the exchange of updates files from the DMGR down to the Node Agent. This is known as "synchronization."

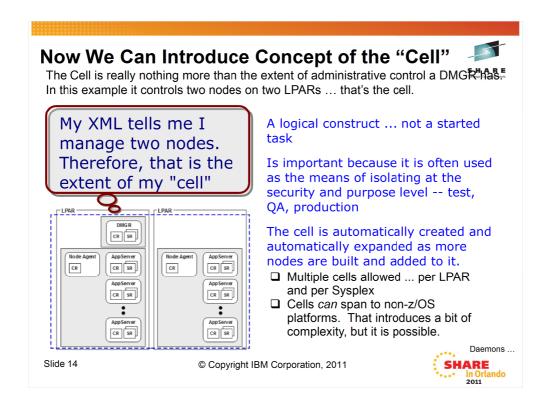
Could the DMGR do that update without the Node Agent? Well ... only if the DMGR had write access to the configuration file system of each node ... but the designers of WebSphere did not want to restrict the construction of the configuration where everything had write access to other things. It's a distributed architecture, which means the file systems don't need to be directly accessible.

Note: yes, cross-Sysplex shared HFS or ZFS is possible. But the design is still distributed, and that's why Node Agents exist. Plus, cross-Sysplex write is not a good performer, so the Node Agent structure is still better.

The process goes like this:

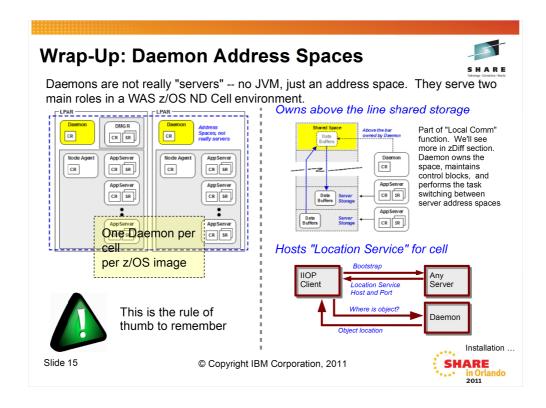
- You make changes to the configuration through the Administrative Console, which runs on the DMGR. The DMGR updates its "master configuration" -- which is maintains for the whole "cell" (which we've not yet explained, but think of it as everything managed by the DMGR).
- The DMGR then taps the Node Agent on the shoulder -- "Hey! You have updates." The updates are copied down to the Node Agent in the act of "synchronization." If the Node Agent is not up, the DMGR simply holds the changes and waits for the Node Agent to come up.
- The Node Agent then applies the changes to the configuration file system for the node.

Like the DMGR, the Node Agent is not required for the steady state operation of the servers. It's a configuration update thing.



The "cell" is another logical thing ... it is the extent of knowledge and management control exercised by a DMGR. The DGMR has configuration knowledge of all the nodes and servers assigned to it, and that comprises the *cell*.

Think of the cell as the boundary of administration. It is the best line of separation for the purposes of administrative isolation. So, for example, if you wanted to isolate "test" from "production" you may think about separating on the cell level. That would mean each cell would have its own DMGR, which you can then lock down with security access policies so testers couldn't touch the production cell, and viceversa.



The Daemon server is something unique to WebSphere z/OS. It consists of a single controller region. It has no JVM and no Java code at all. It's not really a "server" like other WAS servers ... it's really more just an address space. It's not associated with nodes, it's really more a cell-level thing. But those are nuances you can, for now, overlook.

I's purpose in life is two-fold:

- It's what owns shared space above that line that's used for "local comm" -- cross memory data buffer exchanges for inter-server IIOP communications in a cell on the same LPAR. This becomes important because the WOLA function (we'll cover this in the zDiff section) is based on this. That ownership of the shared space is the reason why when a Daemon is stopped all the servers for that cell on that LPAR also stop. Lose the Daemon's shared space control and you lose the local comm.
- It provides the location name service so external clients seeking objects within the WebSphere cell can locate and bind. This is for RMI/IIOP requests coming in from EJB clients outside the WebSphere cell. (Note, a cell comprised of z/OS and distributed boxes -- which is possible -- is still one cell, and in that case an EJB client is operating within the cell. What we're referring to here with the Daemons is the case where you have a server box, unrelated to this cell, with an EJB client looking to connect and bind to an EJB in this cell. Then the location name service hosted in the Daemon servers take over.

Daemons are created at the time the node is created. There are some complex subtleties we'll explore later, but in general their creation is done when you run the jobs to create the cell. And while they can be started manually, in general they are started by WebSphere when the first server for the cell on that z/OS image starts.

The key here really is the yellow box "rule of thumb." Remember that and you'll be fine for this workshop.



Installation and Configuration

- SMPE...
 - PSP Bucket WASAS800
 - SMPE is used to install:
 - The Installation Manager Install Kit
 - The WebSphere on z/OS V8 repository.
- Installation Manager.
 - Install kit is used to build the Install Manager
 - Install Manager is used to create the WebSphere on z/OS V8 binaries.
- A Planning Document.
 - Planning Spreadsheet

Slide 16

© Copyright IBM Corporation, 2011



Before you start it is important to pull the appropriate PSP bucket for your level of WebSphere, and of course, to comply with all of the recommendations that it contains.

The general flow is that a new component on z/OS, the Install Manager is used to create the WebSphere binaries. This starts out by the installation (either SMPE install or download of a zip file and expansion into a file system) of the Installation Manager Install Kit. The Install Kit is then used to create an Install Manager. The Install Manager is then used to create the WAS on z/OS binaries, using a WAS on z/OS repository that is initially installed using SMPE.

Once that much is complete, you may proceed with configuration. Configuration is unchanged from the previous release. The general flow of the construction of an ND (Network Deployment) Cell of WebSphere on z/OS is to build the Deployment Manager cell and server, build an Empty Managed Node or nodes, then add servers, clusters, etc., as necessary.

A good planning document is extremely important. An excellent worksheet may be obtained from the following link:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4686

The sample used here is filled out for a cell which will be referred to as the S8 cell.

The zPMT, is now part of the WebSphere Configuration Toolkit (WCT), is used to build the JCL and data files necessary to create the configuration.



Installation and Configuration

- The zPMT (WCT). (Profile Management Tool, (WebSphere Customization Toolkit)).
 - · Runs on the workstation.
 - Must be installed by Installation Manager on the workstation.
 - Once installed will be used (with the planning spreadsheet) to create...
 - A deployment manager node (within a cell).
 - Any number of Empty managed nodes.
- Where to get this tool, and how to install it will be covered in later foils.
- Servers, server clusters, and any other artifacts can then be created using either the adminconsole, or scripting.

Slide 17

© Copyright IBM Corporation, 2011



The zPMT (WCT) which runs on the workstation must be installed by the Installation Manager on the workstation (which you may also need to install).

Once the WCT has been installed, it will be used (in conjunction with the planning spreadsheet) to create:

- 1. A Deployment Manager (includes the dmgr server, dmgr node, and the cell).
- 2. Any number of Empty Managed Nodes.

Servers, server clusters, and any other artifacts can then be created using either the adminconsole or scripting.

Get Ready...



- Size of the /tmp filesystem.
 - 2 gigabytes seems reasonable.
- IEFUSI Considerations.
 - Many of the jobs and tasks need large region sizes.
 - What does the IEFUSI Do for REGION=0 requests?
 - Does it take OMVS processes into account?
- Paging Space. (AUX storage shortages are not pretty)
 - WAS address spaces are not small (figure approximately 500 Megabytes per...).
 - Minimal ND cell has six address spaces. Do the math...

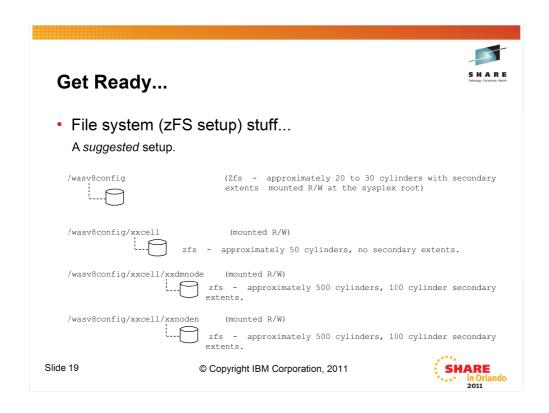
Slide 18

© Copyright IBM Corporation, 2011



Some things to think about before you start doing anything with WebSphere on z/OS, or for that matter the Installation Manager, are:

- 1. Check and correct if necessary the size of the /tmp filesystem. The default is about 4 megabytes and is way too small. I'd suggest that about two gigabytes is a reasonable starting point.
- 2.If you have an IEFUSI exit active (other than the default one delivered with z/OS), you'll need to make sure that it accounts for your need for large region sizes. Also it should take OMVS processes into account by effectively ignoring them.
- 3. You may need to alter the amount of paging space available. Even if you are not paging, there needs to be enough auxiliary storage available to support the region sizes you are requesting. The WAS address spaces are large, roughly half a gigabyte per address space. A minimal ND cell has six address spaces, so that means an additional 3390-3 worth of paging space. You DO NOT want to run out!

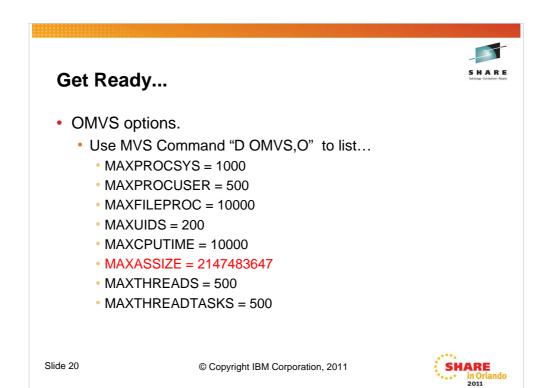


One of the first things that should be considered when creating a new WebSphere on z/OS configuration is the underlying file system(s).

The first file system to consider is what we'll refer to as the WebSphere root. This is basically a filesystem to hold other mountpoints so as to keep us out of the root, always a good thing. A good starting point is to make this file system 20 to 30 cylinders with secondary extents allowed and mount it read/write, usually in the sysplex root.

Next up is what I'll refer to as the "cell" root. There should be one of these for each cell, and it should be about 50 cylinders with no secondary extents, mounted read/write. The configuration file systems are mounted within this filesystem, as well as all of the userid "home" directories for the cell. An advantage to this is that, by default, some java dumps end up defaulting their location to the userid's home directory. Having them within this filesystem with no secondary extents should allow you to capture a couple of them without filling up a lot of space, and hopefully correcting the problem.

Last are the actual node configuration file systems.



Make sure that your OMVS options in BPXPRMxx are realistic for actually using USS for something other than TCPIP. The above values are basically minimums.

Of particular importance, is the value for MAXASSIZE with is basically set to 2 gigabytes minus 1. This essentially removes memory restrictions on OMVS processes, assuming you aren't doing something else with an exit (IEFUSI).



Get Ready...

- Consider turning off SMF Type 92 Records.
 - No particularly useful info in the record and they are written for every type of filesystem and socket activity.
- Create a SAF FACILITY class Profile.
 - BPX.SAFFASTPATH
 - Greatly reduces the number of calls to SAF for file, directory, and IPC access checking that UNIX can already determine (from mode bits and ownership fields) will be successful.

Slide 21

© Copyright IBM Corporation, 2011



SMF type 92 records are basically written for any and all file system activity, as well as socket activity. If you have them turned on they may quickly become the prevalent type of record in your SMF files.

There is very seldom a need for the information they provide, so to save yourself the performance hit of collecting them and throwing them away (or storing them and never looking at them), just turn them off.

The BPX.SAFFASTPATH is very similar to the SMF 92 record hint. Instead of the type 92 record, this can cause an excess of RACF audit records.

All you have to do to implement this is to define the RACF profile:

RDEFINE FACILITY BPX.SAFFASTPATH UACC(NONE) OWNER(SYS1)

and either IPL, restart OMVS, or cause it to refresh its self by issuing the following command:

SET OMVS=(XX)

where xx represents an empty BPXPRMxx member.



Some Setup stuff...

- Install the WebSphere Customization Toolbox on the workstation.
 - · Search on "WCT" in the WAS V8 Infocenter for instructions.
 - Download the Installation Manager for the workstation.
 - · Unzip and install the Installation Manager GUI.
 - Start IM.
 - Go to file >> preferences >> repositories and click "add repository".
 - Add the url that you get from the infocenter.
 - · Click Apply...click ok. Exit the installaltion manager.
 - · Start the installation manager.
 - · Click on "Install". Follow prompts.
 - · Select "WebSphere Customization Toolbox".
 - Follow prompts....click Next a lot.
 - If you do this correctly, you should be able to shut down the Installation Manager and start the WCT.

Slide 22

© Copyright IBM Corporation, 2011



As we indicated earlier, the first thing you need to begin configuring WebSphere is to install the WCT on the workstation.

Before you can install the WCT, you need to have the Installation Manager installed on your workstation at an appropriate level. Once it is installed, unless you do something to prevent it, it will update itself to the most current level whenever it starts.

To install the Installation Manager, you have to download it (it is a zip



Some Setup stuff...

- Install Installation Manager Install Kit on z/OS.
 - Direct download of Install Kit from IBM.
 - Extract into /usr/lpp/InstallationManager/V1R4
 - Run the ./set-ext-attr.sh from the above directory to set the extended attributes properly.
 - Shipped in SMPE installable format with products that require its use.
 - Follow the instructions in the Program Directory.
 - Create an Install Manager from the Install Kit.
 - Follow directions in the InfoCenter.

or...

· Follow the instructions in the Program Directory.

Slide 23

© Copyright IBM Corporation, 2011



You will also nee to install the Installation Manager on z/OS. The first step in this process is to install the Installation Manager Install Kit. The Install Kit may be either:

- Download the install kit (zip) directly from IBM.
- Extract the zip into (usually)
 /usr/lpp/InstallationManager/V1R
 4
- Run the provided ./sset-extattr.sh script to properly set the extended attributes.

And Now, WebSphere...



- WebSphere V8 for z/OS is shipped as a local IM repository in SMPE format with associated products (WAS V8 DMZ secure proxy server, Web Server plugins for WAS V8, and IBM HTTP Server V8).
- JCL to do almost everything is contained in the .F1 relfile and should be copied to an installation specific dataset and modified.
- Installation choice as to whether to use a separate SMPE zone or an existing one.
- The program directory has very good directions and the defaults that it suggests are good.

Slide 24

© Copyright IBM Corporation, 2011



To actually install the WebSphere binaries, you start with a local IM repository that is installed with SMPE. It includes the WAS V8 for z/OS base code, the WAS V8 for z/OS DMZ Secure Proxy Server, the Web Server plugins for WAS V8 for z/OS, and the IBM HTTP Server V8 (apache based).

The JCL needed to use Install Manager to create all of the appropriate file systems and load them is included in the .F1 relfile and should be copied to your own dataset and modified to fit your installation standards. Basically the only thing you should have to modify is the jobcard and the names for filesystems, mountpoints, etc. The sizes for the artifacts are correct.

You have the choice of using an existing SMPE zone, or creating a new one for this repository. Up to you...

The program directory has very good directions and the defaults it suggests are reasonable.

WebSphere V8 on z/OS



At this point the IM repository for WebSphere V8 on z/OS is installed and you can
move on the actual installation of the delivered code into their respective file systems.

These are the jobs.

• BBO1CFS Create and moutn the File System for WAS V8.

BBO1INST Install the code.

BBO2CFS Create and mount the File System for Secure Proxy Server.

BBO2INST Install the code.

BBO3CFS Create and mount the File System for WAS V8 Plugins.

BBO3INST Install the code.

BBO4CFS Create and mount the File System for IBM HTTP Server V8.

• BBO4INST Install the code.

· Change the mount attribute on the four recently created file systems to read only.

- · chmount -r /usr/lpp/.....
- · We're DONE! (With the code installation).
- At this point the code is installed and you can move on to configuration, which looks the same as it did in V7.

Slide 25

© Copyright IBM Corporation, 2011



Once the repository is installed (SMPE work completed), you can then use the Install Manager that you created previously to actually install the code into the appropriate file systems. There is a pair of jobs for each of the components in the repository.

Of course, while you are creating these, the file system must be mounted read/write, so once you are done, you should change the mount attribute to read only, as the example command shows.

At that point you are ready to move on to configuration, which is basically unchanged from the previous release.

Configuration



- Configuration is the same as it was for V7 (and V6.1 if you used the WCT).
 - Fill out the spreadsheet and create the needed response files.
 - Use the WCT to create the configuration jobs.
 - · Use the WCT to upload the jobs.
 - Run the jobs.
- Start with a naming convention. (The spreadsheet uses a good one, and enforces it).

Slide 26

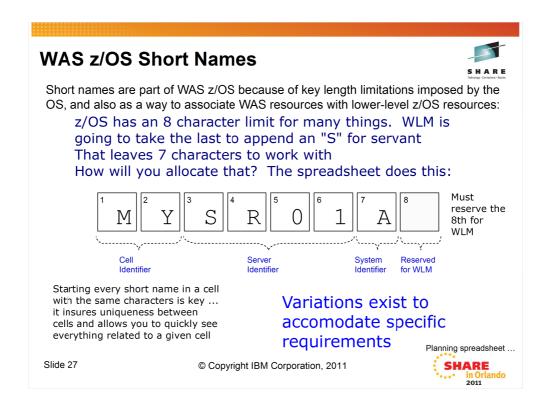
© Copyright IBM Corporation, 2011



The sequence of events is:

- 1. Fill out the spreadsheet.
- 2. Use the spreadsheet to create the necessary response files.
- 3. Use the WCT (PMT) to create the configuration jobs using the response files as input.
- 4. Use the WCT to upload the jobs to z/OS.
- 5. Run the jobs.

It is imperative to start with a good naming convention which the spreadsheet uses and enforces. This is a STRONG recommendation, but unfortunately not a requirement.



Short names are used only by WAS z/OS. They came about because z/OS has key length limitations that need to be taken into account. In general the magic number is 8 ... many names and values associated with z/OS are limited to 8 characters. These include the JCL start procedure names, most SAF values, and the z/OS JOBNAME used for the started task.

We have determined the best place to start the planning process for short names is the controller JOBNAME value.

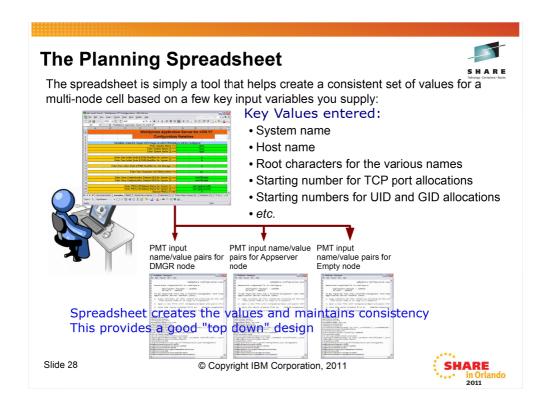
The key to this is understanding that WLM starts the servant region. And by doing so, it provides a JOBNAME for the servant. What it does is add an "S" to the controller JOBNAME. Therefore, the JOBNAME for the controller should be limited to 7 characters so WLM has the space to add the S for the servant.

Further, we have found that it's best to start all controller JOBNAMES with the same characters. What those characters are isn't so important, other than they adhere to z/OS standards such as starting alpha and staying within alpha, number and national characters.

The rest of the controller JOBNAME is allocated to some sort of identifier of the server in question. Again, the specific characters doesn't really matter, provided they're meaningful to you and there's some consistency between servers.

The last thing is an LPAR identifier. That occupies character 7.

You could plan this out by hand. Or you could use the planning spreadsheet. That's what we'll cover next.

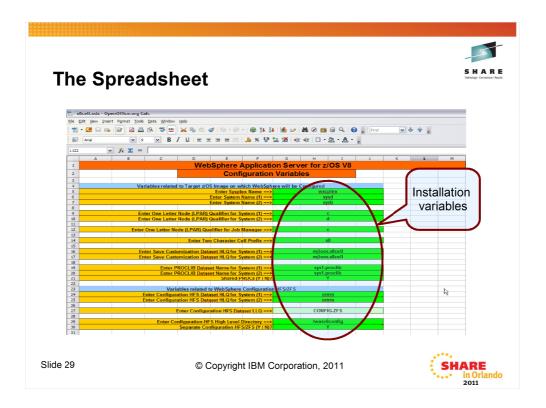


The "Planning Spreadsheet" is simply a spreadsheet (both Excel and OpenOffice) that takes a few key variables from you on one worksheet and *generates a consistent set of values for various PMT node options on other sheets.*

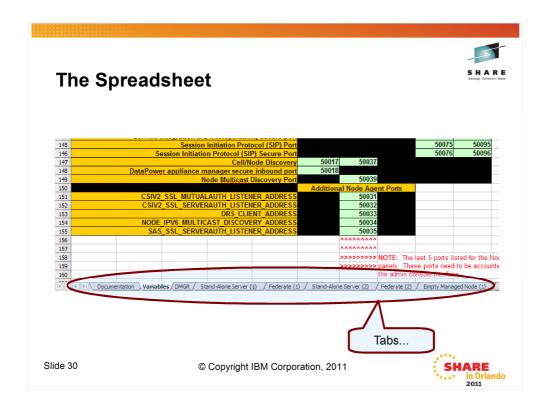
The produced output is a set of name/value pairs that serves as input to the PMT. It's simply a matter of copying the values from the spreadsheet and paste them into a text file. From there you import it into the PMT. The PMT then has *all the values it requires to build the node*.

Again, the key is that this produces a set of consistent values across nodes for a cell. There's no magic to it. It's simply programming within the spreadsheet that enforces the consistency.

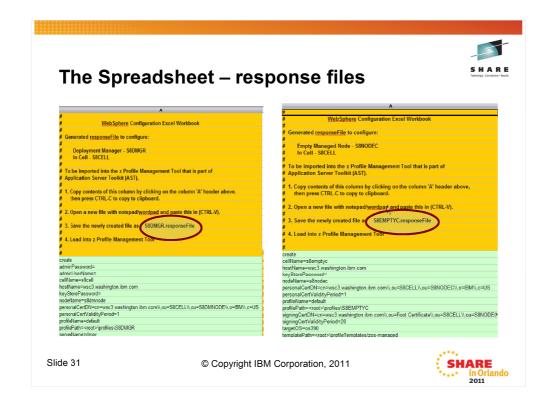
This is what we call a "top down" design. By that we mean a design that takes into account the end objective of the cell design, and keeps consistent those things that require it.



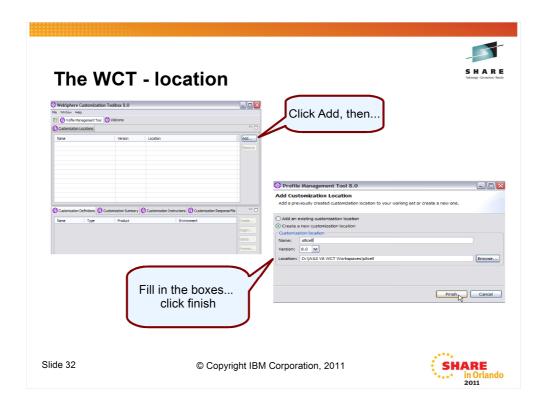
This is the variables tab of the spreadsheet that was used for the cell in this presentation.



And these are the "tabs" which allow you to select the various response files.



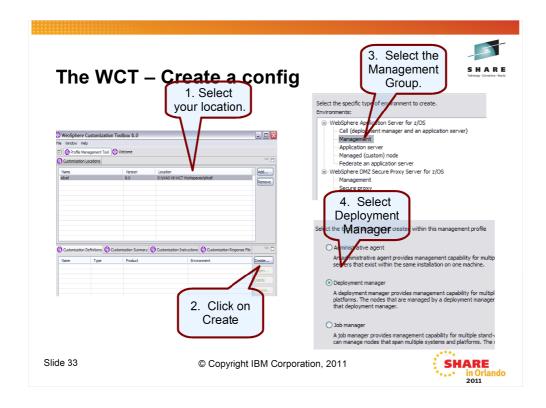
Looking at two of the tabs, we can see the deployment manager response file on the left, and the first empty managed node response file on the right.



Once we've started the WCT, the first thing we'll need to do is add a location. A location is just what it sounds like. A named area on workstation disk that is used to store all of the elements of a configuration.

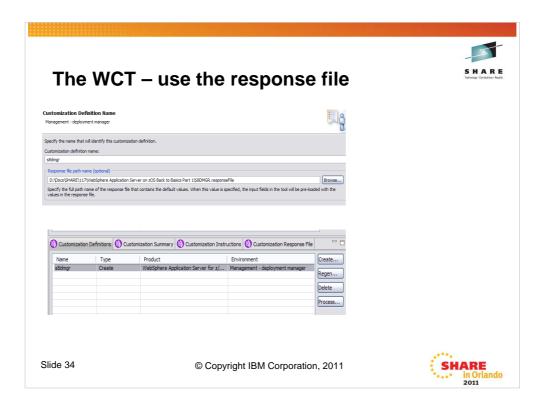
To create a new one, you click on Add, then fill in the blanks on the next panel.

Name for the location in this case is the cell name, you select the version in the drop down box, and then indicate where on disk you want the location to reside.



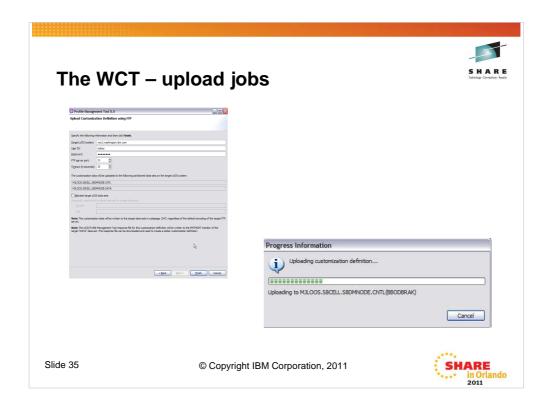
Once you have a location, you can start to create configurations.

In this case we are going to create a deployment manager, so you first make sure you have the correct location highlighted, then click on Create, then select the management suite, then select deployment manager on the next panel.



On the next panel you'll be given the opportunity to "name" your configuration. This isn't anything that will "carry through" but is how this configuration is named within the WCT. It is on this panel that you are allowed to point at a pre-existing response file. In our case, we'll point at the one we save for the deployment manager out of the spreadsheet.

Now if you accept what the spreadsheet has done on your behalf, you'll simply click next a lot



Next you will select the option to Process the configuration and essentially be guided through uploading the jobs that are the result of the the create process to the z/OS system where the configuration will reside.



This is a member list view of the CNTL dataset which was uploaded for the deployment manager build process, with the members containing the jobs needed to build the configuration highlighted.



Run the jobs

- · Start with the security jobs.
 - Either run the standard BBOSBRAK and BBODBRAK, or
 - Run a custom job. See techdocs WP101427.
- · Run one at a time, in order:
 - BBOSBRAM
 - BBODCFS
 - BBODHFSA
 - BBOWWPFD
 - BBODPROC
- Start the Deployment Manager

Slide 37

© Copyright IBM Corporation, 2011



To actually build the configuration, you will run the jobs, one at a time, in the order the instructions (from either the WCT or the BBOxxINS member of the CNTL dataset) specify.

First you'll be running either the BBOSBRAK and BBODBRAK jobs to define all of the necessary SAF profiles for the configuration.

Alternatively, you have the option of running a custom job. There is a techdoc that describes in detail an alternative that builds a set of generic profiles that will cover the



Run the jobs

- Repeat the procedure for any empty managed nodes.
 - Use the WCT with the response file from the spreadsheet.
 - · Generate and upload the jobs.
 - · Run the jobs.
- Start the nodeagent for each managed node.
- You now have the shell of a cell.
 - · Add servers, resources, clusters, etc.
 - · All of which are the subjects of the "Part 2" session.

Slide 38

© Copyright IBM Corporation, 2011



Once the deployment manager is up and running (you must leave it running to complete the federation jobs in the empty managed node configuration jobs), you can repeat the process for each of the managed nodes that you have decided you need.

After each set of jobs for the empty managed nodes are completed, if you didn't allow it to happen automatically, you'll need to start the nodeagent for the node.

At this point, you have a cell with a



If you have any questions, now is the time...