



WAS z/OS Potpourri

IBM Corporation

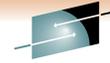
March 2, 2011 - 4:30pm
Session 8380



Slide 1

© 2011, IBM Corporation

WebSphere Application Server Sessions



SHARE

Room	Day	Time	Title	Speaker
208B	Monday	11:00	Lab	Multi
201A	Monday	11:00	The Value of the WebSphere Application Server Job Manager	Loos
205A	Monday	4:30	WebSphere Application Server for z/OS -- I am No Longer a Dummy but...	Loos
205B	Tuesday	9:30	Performance Tuning for WebSphere Application Server for z/OS - Practical Advice	Everett
205A	Wednesday	4:30	WebSphere Application Server for z/OS: Tools and Tricks (Potpourri)	Loos and Co.
205A	Wednesday	6:00	WebSphere Application Server for z/OS: Helping Customers Help Themselves	Stephen
206B	Thursday	8:00	Securing WebSphere Application Server for z/OS	Kearney
206B	Thursday	9:30	Application Improvement and Savings Through Simplification	McCorkle
206B	Thursday	11:00	WebSphere Application Server for z/OS: Batch	Bagwell
206A	Thursday	12:15	WebSphere Application Server 101	Stephen
206B	Thursday	1:30	WebSphere Application Server for z/OS: Availability Considerations	Bagwell
206B	Thursday	3:00	WebSphere Application Server: z/OS Exploitation/Differentiation	Follis
206B	Thursday	4:30	Performance Tuning for WebSphere Application Server for z/OS - WAS and WLM Interactions and Concepts	Follis

Slide 2

© 2011, IBM Corporation

SHARE
in Anaheim
2011

Table of Contents



Dave Follis - follis@us.ibm.com



Mike Loos - mikeloos@us.ibm.com



Mike Everett - meveret@us.ibm.com



Mike Stephen - msteff@us.ibm.com



Don Bagwell and John Hutchinson
dbagwell@us.ibm.com hutchjm@us.ibm.com

[Hiperlink to Sections](#)

Slide 3

© 2011, IBM Corporation

SHARE
In Anaheim
2011



[Back to Table of Contents](#)



SHARE
Technology - Connections - Results

Dave Follis

Slide 4

© 2011, IBM Corporation

SHARE
in Anaheim
2011

Changing the number of servants



- You can define the minimum and maximum number of servants in the admin console
- Prior to V7 you had to recycle the server to change it
- In V7 we introduced two new Modify commands
 - `MODIFY server,WLM_MIN_MAX=(min,max)`
 - `MODIFY server,DISPLAY,WLM`
- This temporarily (until the server restarts) changes the min and max values
 - If the current number of servants is lower than the new MIN value, WLM will **usually** start more to catch up
 - If the current number of servants is greater than the new MIN value, WLM will **probably eventually** shut down the extra servants... **maybe**.
 - If the current number of servants is at the old MAX value and you've increased it, WLM **might** start more
 - If the current number of servants is more than the new MAX value, WLM **may eventually** shut down the extras to get below the new MAX. **Maybe**.



Slide 5

© 2011, IBM Corporation

Change the timeout delay value dynamically



SHARE
Technology • Connection • Results

- When the controller decides to abend a servant for a timeout
- You can specify a delay to let work in-flight finish before the abend occurs
- The dispatch threads in the servant won't ask for more work in this window
 - Although if they are already waiting for work they might get one more
- Now you can change the size of this delay dynamically
 - `MODIFY server,TIMEOUT_DELAY=x`
- Work with affinity to this servant waits in the queue until the delay ends and the servant abends
 - Some requests will be requeued at that point (if they haven't timed out)
 - Difficult to direct to another servant while the the problem servant is still up because it hasn't ended yet and still 'owns' the affinity data (stateful session, HTTPsession etc).

SHARE
in Anaheim
2011

Slide 6

© 2011, IBM Corporation

Default Classification for Internal Work



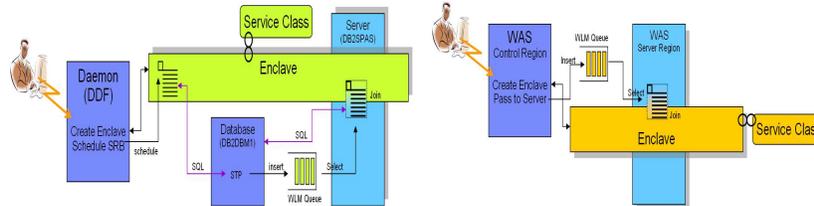
- You can create a classification XML file to tell the controller what transaction class name to provide WLM when a request is classified
- This allows you to separate different work into different WLM service and report classes
- You can specify a transaction class name for 'internal' work
 - When the controller is talking to its servant regions (e.g. JMX Mbeans etc.)
- But what if you want to separate out the internal work into a different report class without bothering with this XML file thing?
 - `default_internal_work_transaction_class`
- Will be ignored if there is a clause in the XML file for internal work

Slide 7

© 2011, IBM Corporation

SHARE
in Anaheim
2011

Overview: WLM Enclaves and Server Management A Short Retrospective



- Enclaves encapsulate Business Units of Work in order to provide performance management based on installation defined goals
 - Enclave management allows to directly apply CPU and I/O dispatch priorities to the units of work (TCBs and SRBs) across multiple z/OS Address Spaces
 - Storage management is applied for the address spaces in which the units of work execute
 - Enclave management provides the base for WLM to dynamically manage (start/stop) the number of server address spaces based on goal fulfillment and need
- But what about the enclave server address spaces?



Slide 8

© 2011, IBM Corporation

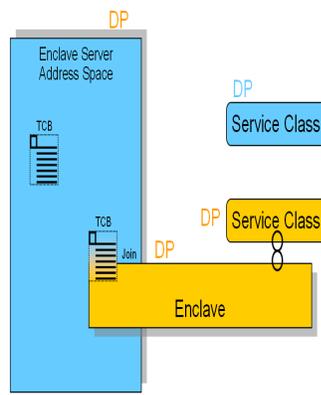
Background: Retrospective on why enclaves exist and how they are managed

The two pictures show typical examples of how enclaves are exploited.

- The picture on the left hand side shows a DDF/DB2 exploitation: A work request is classified by DDF, an enclave is created and an SRB is scheduled to process the work request. The program started for the SRB issues SQL calls to DB2. If these contain a Stored Procedure Call, a work request is inserted to a WLM application environment queue and a DB2 Stored Procedure address space eventually processes the request. For this purpose the task processing the request joins the same enclave.
- The picture on the right hand side a Websphere exploitation: In this a queue insert is always done and a WAS server region processes the request.

Overview: WLM Enclave Server Management

Is There a Possible Problem?



- What if the programming model does not hold true?
 - What happens if there is significant work running in TCBs not associated with enclaves?
 - Example: Garbage collection for a JVM (WAS)
 - Example: Common routines which provide service for the enclave TCBs (for example garbage collection tasks in JVMs)
 - Is it sufficient to manage this work in the same way as the enclaves?
- What happens if no enclaves are running (note: this applies to queue servers only)
 - And the address space is swapped out?
 - A mechanism exists to swap in the address space but this mechanism assumes that the swap in is only for a queue server task which wants to select a unit of work and then joins the enclave
 - And even if the address space stays swapped in?
 - The TCBs running within the address space just stay with the DP and IOP from the last enclave being associated with the address space
 - No CPU or I/O adjustment is performed.



Slide 9

© 2011, IBM Corporation

WLM Enclave and Queue Server Address Spaces are today managed towards the service classes for the enclaves. The programming model these address spaces are based on assume that nearly all work always runs in an enclave. But this programming model no longer holds true. Especially in Websphere environment a substantial amount of work in such address spaces is running outside of enclaves, for example java garbage collection threads. This work is today not managed and the address space inherits management attributes from the enclave service classes.

MPL Problem:

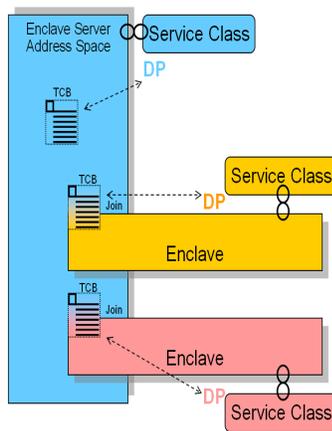
PMR 75009,487,000 showed a problem where no enclaves were running in the address space anymore and the work not associated with enclaves was not managed anymore. Because non enclave work wanted to run it got swapped in with the expectation that the tasks selects a unit of work and joins an enclave. But it was a garbage collection thread. The swap in is only temporarily so the task became swapped out because WLM didn't notice that non-enclave work actually wants to run. When the task was swapped out it became immediately eligible for swap in and this thrashing effect created harm to the z/OS environment

The picture on the left hand side explains how the enclave server address space inherits its Dispatch priority (DP) from the service class of the enclaves. The DP of the original address space service class is not considered. This is the same for all other resource access controls like I/O priority for example.

© 2011, IBM Corporation

© 2011, IBM Corporation

Solution: WLM Non Shell Server Management



- New OPT Parameter
 - ManageNonEnclaveWork = {No|Yes}
 - Default: No (no change to previous releases)
 - Causes everything in the address space which is not associated to an enclave to be managed towards the goals of the external Service Class to which the address space has been classified to.
- Advantages
 - Enclave (Queue) server address spaces in which no enclave is running will be managed as usual address spaces
 - The importance and goal of the service class for the address space now has a meaning
- Attention
 - The importance and goal of the service class for the address space now has a meaning
 - Therefore verify goal settings for server address spaces
 - This is a deviation from the past when the service class for servers was only important for startup, shutdown and recovers



Slide 10

© 2011, IBM Corporation

With non shell server management the inheritance of resource access controls significantly changes. This is shown on the left hand side for the Dispatch priority. Now the tasks which have not joined an enclave inherit the DP of their original service class. Also all data and samples are attributed back to this service class and provide the base for managing the address space.



[Back to Table of Contents](#)



Mike Loos



Slide 11

© 2011, IBM Corporation

Topic List



- **zfs setup**
- **smf type 92**
- **BPX.SAFFASTPATH**
- **ISA**
- **Virtual Keyring support**
- **support toolbar**
- **JVM dump mask**

Slide 12

© 2011, IBM Corporation

SHARE
in Anaheim
2011

There are a few things I'd like to talk about, some of which you may know and some of which may be new. This is the list...

zfs setup...

A suggested setup.



```

/wasv7config          (Zfs - approximately 20 to 30 cylinders with secondary
                      extents mounted R/W at the sysplex root)
  |
  |
/wasv7config/xxcell   (mounted R/W)
  |
  | zfs - approximately 50 cylinders, no secondary extents.
/wasv7config/xxcell/xxdmnode (mounted R/W)
  |
  | zfs - approximately 500 cylinders, 100 cylinder secondary
  |   extents.
/wasv7config/xxcell/xxnnode (mounted R/W)
  |
  | zfs - approximately 500 cylinders, 100 cylinder secondary
  |   extents.

```



Slide 13 © 2011, IBM Corporation

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 filesystem 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.

SMF type 92 record.



- File System Activity Record.
- Written for all types of file system activity (mount, quiesce, unmount, and of most interest, file open and close.
- Opening the server.xml file below will result in at least 27 type 92 records.
- **Turn them off!**
- No particularly useful info in the record. May have to turn them on at the request of level 2 if investigating some particular problem.

```
/wasv7config/xxcell/xxnoden/AppServer/profiles/default/config/cells/xxcell/nodes/xxnoden/servers/xxsr01n/server.xml
```



Slide 14

© 2011, IBM Corporation

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.

BPX.SAFFASTPATH SAF profile.



- FACILITY Class profile BPX.SAFFASTPATH.
 - Must be non-generic (no BPX.SAF*)
 - Defined is good enough. No access needs to be granted.
- The existence of this profile will cause the security system to **NOT** be called if z/OS UNIX can quickly determine that file access will be successful (using the permission bits and user/group/other ownership). This will be significantly faster.
- Without this profile, the security system will be called for file access checking, IPC access checking, and process ownership checking.
- Using the example below, at least 14 calls to the security system may be avoided.

```
/wasv7config/xxcell/xxnoden/AppServer/profiles/default/config/cells/xxcell/nodes/xxnoden/servers/xxsr01n/server.xml
```



Slide 15

© 2011, IBM Corporation

This tidbit 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 issuubg the following command:

```
SET OMVS=(XX)
```

where xx represents an empty BPXPRMxx member.

IBM Support Assistant (ISA)

The screenshot shows the IBM Support Assistant Workbench interface. The 'Tools Catalog' pane lists various diagnostic tools. Two tools are highlighted with red circles:

- IBM Pattern Modeling and Analysis Tool for Java Garbage Collector (PMAT)
- IBM Monitoring and Diagnostic Tools for Java™ - Dump Analyzer

Other tools listed include Database Connection Pool Analyzer, HBaseAnalyzer, IBM Thread and Monitor Dump Analyzer for Java (TMDA), IBM Trace and Request Analyzer for WebSphere Application Server, IBM Web Server Plug-in Analyzer for WebSphere Application Server (WSPA), Memory Dump Diagnostic for Java (MDD4J) version 3.0, ThreadAnalyzer (Deprecated), Visual Configuration Explorer, IBM Assist On-site, IBM Monitoring and Diagnostic Tools for Java™ - Health Center, IBM Monitoring and Diagnostic Tools for Java™ - Memory Analyzer [Tech Preview], Log Analyzer, and Memory Dump Diagnostic for Java (MDD4J).

The interface also shows a 'Restrictions' section set to 'None' and an 'Associations' section stating 'Tool is not associated with any products'.

Slide 16

© 2011, IBM Corporation

SHARE
in Anaheim
2011



The IBM Support Assistant (ISA) is a multipurpose tool that you can download for free.

The tools within it are selectable, and once selected, are updated automatically.

I'm going to do a very quick demo of two of the tools, which are the ones I tend to use the most. YMMV.

Both of these tools use the same file as input: A file of verbose garbage collection statistics downloaded from the host. One caveat: Do not include any carriage control characters in the files. They confuse the parsers.

IBM Support Assistant (ISA) - PMAT



IBM Pattern Modeling and Analysis Tool for Java Garbage Collector

File List

Name	First Garbage Collection	Last Garbage Collection	AF/OC
java1_jm.log.dec09	Thu Dec 9 03:59:23 2010	Thu Dec 9 14:40:53 2010	95927

- Number of verbose GC cycles : 1
- Number of Garbage Collections : 927
- Number of Allocation failures : 95
- First Garbage Collection : Thu Dec 9 03:59:23 2010
- Last Garbage Collection : Thu Dec 9 14:40:53 2010
- Number of Java heap exhaustion : 0
- Overall Garbage Collection overhead : 0.44%
- Maximum Garbage Collection overhead : 85% (Thu Dec 9 10:32:53 2010)
- Number of 100% AF overhead : 0
- Total Garbage Collection pause : 169 seconds
- Maximum Tenured Area usage : 1,931,175,448 bytes (Thu Dec 9 10:32:01 2010)
- Average Tenured Area usage : 621,858,748 bytes
- Number of Explicit Garbage Collection : 2
- Maximum Allocation Request : 9,437,208 bytes (Thu Dec 9 10:29:15 2010)
- There is no object request larger than 10 M bytes.

Java Heap Activity Analysis and Recommendations report

Garbage collection start / finish	Analysis	Recommendations
#1 Thu Dec 9 03:59:23 2010 Thu Dec 9 14:40:53 2010		
Configuration gcPolicy : Xgcpolicyoptargpau maxHeapSize : 1,342,177,280 bytes initialHeapSize : 536,870,912 bytes compressedRefs : false pageSize : 4,096 requestedPageSize : 4,096	No Java heap exhaustion found	There seems to be a steady increase in Java heap usage (ratio(%): 197.60289 with percentage error(%): 0.53675234)



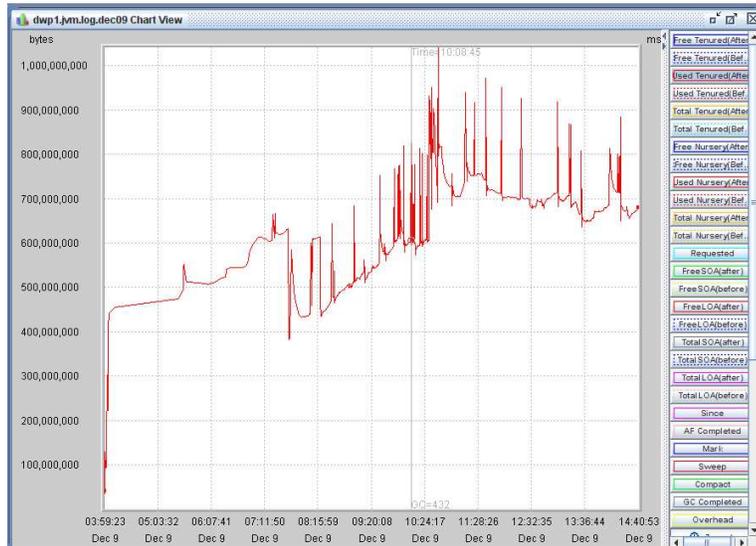
Slide 17

© 2011, IBM Corporation

The first tool is the Pattern Matching and Analysis Tool for java garbage collection.

As you can see, it provides a nice concise analysis of memory usage in the JVM.

IBM Support Assistant (ISA) - PMAT

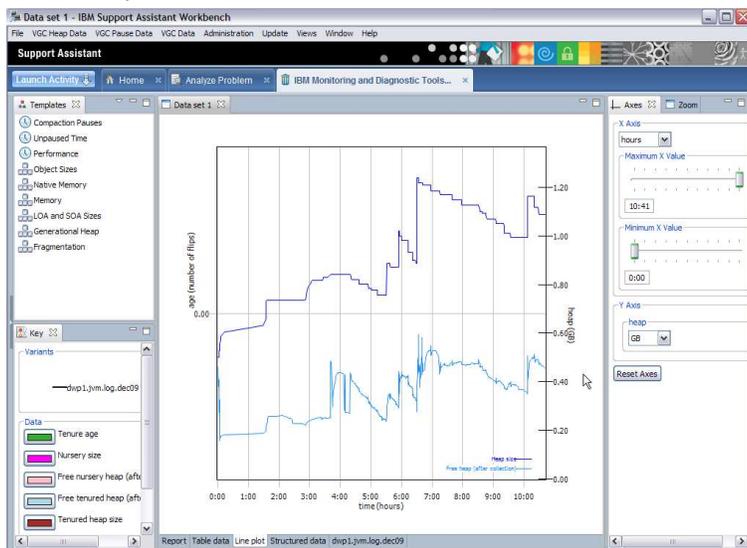


Slide 18

© 2011, IBM Corporation

You can also view all of these statistics graphically...

IBM Support Assistant (ISA) - Garbage Collection and Memory Visualizer



Slide 19

© 2011, IBM Corporation

The Garbage Collection and Memory Visualizer is very similar tool, but it provides different views of the same data.

Virtual Keyring Support



- Primary usage is for wsadmin.sh scripting users.
- Avoid the need for separate keyring for each user.
- Takes advantage of a “virtual keyring” with all CERTAUTH certificates known to RACF connected to it.
- This keyring is called: safkeyring/*AUTH*/*
- To enable the support you must modify the profile's ssl.client.props file.
- Found in: `config root.../Appserver or
/DeploymentManager/profiles/default/properties`
- # com.ibm.ssl.trustStore=safkeyring:///WASKeyring.S1CELL
com.ibm.ssl.trustStore=safkeyring/*AUTH*/*



Slide 20

© 2011, IBM Corporation

Virtual Keyring support was added by RACF at the z/OS 1.9 level.

It is useful for secure FTP clients (which is why I believe it was first introduced) but also can make life much easier for WebSphere wsadmin.sh clients (and their respective RACF admins).

It essentially allows the wsadmin client to act a little bit more like a browser. Instead of searching a specific keyring name that must be defined for each userid, for a specific CERTAUTH certificate that must match the signer of the private certificate that the server presents, the virtual keyring may be used by all users and includes all CERTAUTH certs of which RACF knows.

The actual change simply involves changing the keyring name in the ssl.client.props file from the specific keyring name to the name of the virtual keyring, which is a name reserved to RACF.

It is possible to restrict the usage of this keyring with RACF profiles, but if you want more info on that, I'd suggest you take a look at the slides and lab instructions from the WSW07 Wildfire class, available at:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS3961>

IBM Support Toolbar

Found at: <http://www-01.ibm.com/software/support/toolbar/>



IBM Software Support

IBM Software Support Toolbar

Enhance your IBM Software support experience

This is the official notice of the **end of support for the IBM Software Toolbar on December 31, 2010**. IBM customers should migrate to the new award winning IBM Support Portal (<http://ibm.com/supportportal>). It is a consolidated portal that will meet all your technical support needs. The IBM Support Portal is a unified, customizable view of all technical support tools and information for all IBM systems, software, and services. It replaces all legacy IBM technical support sites, in addition to the IBM Software Toolbar. If you have not used the Support Portal yet, try it out.

1. Access the IBM Support Portal: <http://ibm.com/supportportal>
2. Select one or more (up to 10) products of interest to you.
3. View your page
4. Click Sign in and enter your IBM ID and password to customize the view. If you don't have an IBM ID yet, click Register on the Sign in form

Watch our short, introductory videos on YouTube: <http://www.youtube.com/user/IBMelectronicsupport>

Installation
Installing is as easy as clicking "Open" when the download begins.

Note: If prompted with a dialog box after starting the download, click the "Open" button ([example picture](#)).

The IBM Software Support Toolbar is being released under the "Non-Warranted" program. Review the [International License Agreement for Non-Warranted Programs](#) for more information.

Internet Explorer users
[Click here to start download](#)

Firefox users (You must be using Firefox at time of install)
[Click here to start install](#)



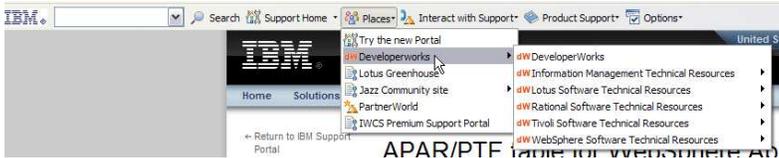
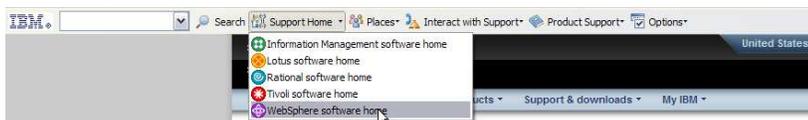
Slide 21

© 2011, IBM Corporation

The IBM Support Toolbar is downloadable, (location shown) and available for both IE and Firefox browsers.

The following slides show some of the function available.
I find it useful, and the price (free) is right.

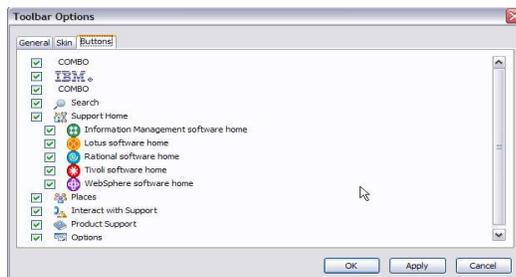
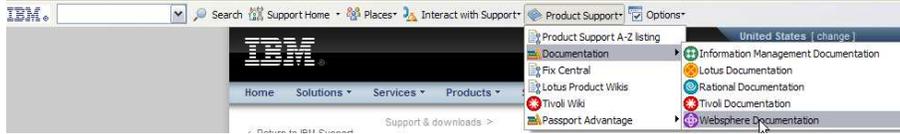
IBM Support Toolbar



APAR/PTF



IBM Support Toolbar





IBM United States [change] Search

Home Solutions Services Products Support & downloads My IBM Welcome Mr. Mike Loos [Not you?] [BM Sign out]

- Return to IBM Support Portal
- Support & downloads**
- Bookmark this page
- View my bookmarks
- Feedback

Support & downloads > APAR/PTF table for WebSphere Application Server V7.0 for z/OS

Product documentation

Abstract

This table includes APAR/PTF information for IBM WebSphere Application Server V7.0 for z/OS.

Content

This table includes committed APAR/PTF information for WebSphere Application Server V7.0 for z/OS.

Notes:

- This table is for reference only. Last updated 24 January 2011
- APARs flagged as Security/Integrity APARs do not have links to the associated APAR text provided.
- WebSphere Application Server for z/OS Fix Pack PTFs are not cumulative service deliveries based on the GA of the product. Fix Pack PTF deliveries will either PREREQ or SUPERCEDE a previous Fix Pack PTF delivery.

Availability Date	APAR	PTF	Service Level	Build Level	JAVA FOR Z/OS SDK 1.6.0
10/21/2010	PM18125	UK61159	7.0.0.13	cf131039.07	build pmz3460s8fp1-20100924_01 (SR8 FP1) for 64-BIT
10/21/2010	PM23874				pmz3160s8fp1-20100924_01 (SR8 FP1) for 31-BIT



Document information

[WebSphere Application Server for z/OS](#)

Software version: 7.0

Operating system(s): z/OS

Reference #: 7013663

Modified date: 2011-01-24

Translate my page
Select Language

I use it most often to get to here...



JVM Dump Mask



`JAVA_DUMP_TDUMP_PATTERN=string`

Result: The specified string is passed to IEATDUMP to use as the data/set name for the Transaction Dump. The default string is:

`%uid.JVM.TDUMP.%job.D%y%m%d.T%H%M%S`

From:

IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 6 Diagnostics Guide

- The %uid portion of that default string will generally resolve to the userid under which the servant is running.
- You may change this to something which better fits with your systems rules.
- `JAVA_DUMP_TDUMP_PATTERN=WASDUMPS.JVM.TDUMP.%job.D%y%m%d.T%H%M%S`
- Changed in the JVM start arguments of the servant.



Slide 25

© 2011, IBM Corporation

If you allow everything to default, then when a dump is taken by your JVM, it will try to store it in a file of the form:

`USERID.JVM.TDUMP.JOBNAME.DATE.TIME`

You can change that to something that might actually get into a catalog (or better yet, allocated).

JVM Dump Mask



Application servers

Application servers > s1sr01c > Process definition > Servant > Java Virtual Machine > Custom properties > New

Use this page to specify an arbitrary name and value pair. The value that is specified for the name and value pair is a string system configuration properties.

Configuration

General Properties

* Name
JAVA_DUMP_TDUMP_PATTERN

* Value
%job.D%y%em%d.T%H%M%S

Description

Apply OK Reset Cancel





[Back to Table of Contents](#)



Mike Everett



Slide 27

© 2011, IBM Corporation

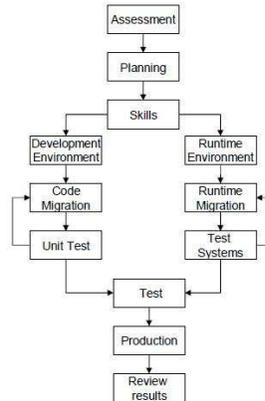
Migration and Interoperability



Where to find the complete history and summary of migration:
<http://www-01.ibm.com/support/docview.wss?uid=swg27008724&aid=6>

Migration Plan Roadmap

- Assessment
- Planning
- Skills
- Development Environment
- Application Code Migration
- Runtime Environment Migration
- Test
- Production
- Review the results



Slide 28

© 2011, IBM Corporation

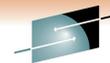
The pdf for Migration is well over 200 pages long. It is comprehensive. No matter what level of WAS you are coming from and going to up to Version 7.0 you will find:

- a roadmap
- the 5 different options for performing a migration
- the differences between the levels
- items found by companies who have migrated

(this presentation is consistently updated as we discover material to add)

This roadmap shows all of the different items that need some planning before moving your infrastructure to the next level.

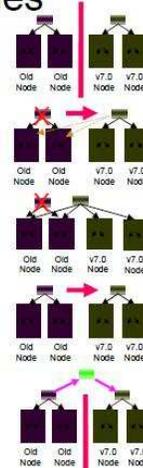
Migration and Interoperability



SHARE
Technology · Connection · Results

Five Network Deployment strategies

1. Manual
 - ✓ Create a new cell and populate with scripts or manually
 - ✓ No runtime migration tools
2. "Copy" and replace the cell
 - ✓ Recreates the **exact** v5.1/v6.1 configuration in v7.0
 - ✓ DMgr and nodes are migrated
3. "Copy" and replace the DMgr
 - ✓ Recreates the **exact** v5.1/v6.x configuration in v7.0
 - ✓ Add new v7 nodes and move incrementally
4. "Copy" and coexist
 - ✓ Recreates the **exact** v5.1/v6.1 configuration in v7.0
 - ✓ Modify the ports in the new cell and coexist
5. Fine Grained
 - Create a new cell and incrementally copy existing configuration
 - Uses an intermediate profile, runtime migration and PBC tools



SHARE
in Anaheim
2011

Slide 29

© 2011, IBM Corporation

I think it is important to start with:

No matter what change in your environment you plan to make, make it a small incremental change.

Companies have to take a look at their needs and decide which of these strategies work for which parts of their environment.

We see all of these strategies used.

- If you have scripting in place already, maybe a manual recreate is best, you have to decide. If you have a checklist of items you would like to do better the next time, maybe a manual recreate is best.
- If you do not have scripting for many of the items in your infrastructure, if you do not have 100's of applications, for test environments maybe you use the migration tooling that comes with WebSphere to migrate the nodes in your cells.
- Migrate the dmgr and create nodes according to your own schedule at the higher levels. Again this works best if you have scripting in place to recreate your resources and deployments. This option is chosen most often to fit into the schedules of developers and also to build an environment and rearrange where the applications live as you go.
- This is the same as number two but instead of not using the old environment, you change the items that collide in the two environments and run both at the same time. I do not see this as often but it is possible.
- Let's say you have a plan to build the newer environment that involves rearranging the various pieces to fit changing business needs. You can use the tooling (profile management tool and property based configuration files to recreate parts of the environment to fit your new plan. The details of this go beyond this small talk but the details are described in the pdf.

Migration and Interoperability



WebSphere Application Server Migration Toolkit

- ✓ Assists Application Developers migrating between WebSphere Application Server versions!
- ✓ Supports moving to v7.0 from
 - V5.1.x, V6.0.x, and V6.1.x
- ✓ Assists by automatically identifying and change assist of code impacts in all the key areas
 - Changes to the Java™ Runtime environment
 - Changes to the level of supported Java Enterprise Edition version
 - Removal of previously deprecated features
 - Behavior changes in the product
- Appendix in this deck will note changes addressed by the Toolkit ->
- ✓ Free download
- ✓ Eclipse plug-in
 - Support for pure Eclipse (3.4.2 and 3.5) as well as RAD/RSA
- ✓ See http://www.ibm.com/developerworks/websphere/downloads/migration_toolkit.html



Tools that save me time

WinSCP – get or edit files with secure FTP

Putty – with some basic UNIX commands
Now that most files are in UNIX, cat, tail, ls -la, are not bad to know. Just don't tell anyone you might like UNIX.

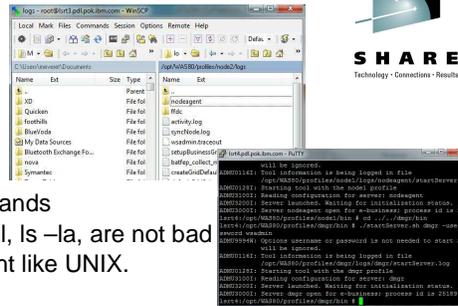
Image Capture(as in sametime) – A picture says a thousand words. If you are remote from a co-worker show them instantly what you are talking about by sending an image.

Free Version of WebSphere on your personal machine.

The file structure is the same for the Express edition.

The Open Source Community addition is different.

<http://www-01.ibm.com/support/docview.wss?uid=swg27013254>





[Back to Table of Contents](#)



Mike Stephen



Slide 32

© 2011, IBM Corporation

What's Slowing WebSphere Down ??



- RACF AUDIT was active for the following classes:
 - DIRACC, DIRSRCH, FSOBJ, FSSEC - AUDIT ALL.
 - None of these classes were RACLISTed.
 - Issued command SETR LOGOPTIONS(NEVER(DIRACC))
 - for all above classes to turn off auditing
 - Following the change
 - Portal restarted in 4 minutes compared to 30 minutes
 - F ZFS, QUERY, ALL showed avg access time 0.003 instead of 1.6
 - CPU usage returned to normal which means that the zAAPs were being used instead of the GCP.
 - When running a load the GCP% is now close to zero.
- "The total response times are now excellent"**



Slide 33

© 2011, IBM Corporation

Common Migration Problem

- One of the hardest to diagnose as well
- Out-Of-Memory problem
 - Not always surfaced
- Prevent sharing of Address Spaces
 - `_BPX_SHAREAS=NO`
- Increase JVM Heap Min and Max
 - `-Xms256M -Xmx768M`

```
BPXBATCH SH + export _BPX_SHAREAS=NO; +  
export IBM_JAVA_OPTIONS="-Xms256M -Xmx768M"; +  
/wit/bigtmp/bbomigrt2.sh WASPreUpgrade +  
/wit/bigtmp/24173105/_ +  
1>> /wit/bigtmp/24173105/BBOWMG3D.out +  
2>> /wit/bigtmp/24173105/BBOWMG3D.err;
```

- InfoCenter has good article:

- search for ***tmig_troubleshoot***

Slide 34

© 2011, IBM Corporation



Security Certificates Expiring / Expired



- Seems Certificates are reaching (or have reached) the end of their 'Shelf Life'
- Renewing expiring RACF certificates for WebSphere
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PR9>

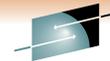
Slide 35

© 2011, IBM Corporation

SHARE
in Anaheim
2011



Back to Table of Contents



SHARE
Technology • Connections • Results

Don Bagwell

John Hutchinson

SHARE
in Anaheim
2011

Slide 36

© 2011, IBM Corporation

Topic List



- **WOLA Update**
 - **IMS support in 7.0.0.12**
 - **2PC WAS ⇒ CICS TS 4.1**
- **Java Batch in WAS batch containers**
 - **Batch FEP**
 - **WebSphere Compute Grid**
- **Running the DB2 command line processor**

Slide 37

© 2011, IBM Corporation

SHARE
In Anaheim
2011

Using the DB2 command line processor



SHARE
Technology • Connections • Results

- What if the DDL is not well-formed for SPUFI or DSNTEP2 ?
 - Like the **CreateMailerTablesDB2.ddl** in WebSphere XD which contains 1000 insert statements with line lengths over 250 chars!

Set-up: Export the following variables (use .profile)

- JAVA_HOME
- PATH
- CLASSPATH
- CLPHOME
- CLPPROPERTIESFILE

Create an alias 'db2' in the .setup file to invoke the clp.
Create clp.properties file to specify ; as the term char.

Set environment variables to locate the JDBC & SQLJ:

- Set the location of the Java & jcc home, and DB2 OS data set high-level qualifier.
- Add jcc-related jar files to the CLASSPATH.
- Add jars from JCC_HOME/classes for jcc.
- Add current directory to pick up local classes.
- Add jcc-related directories to the PATH & LIBPATH.
- Add java directory to the PATH, if needed.

Sample setup scripts on the following foil...

Running the clp:

Add the following statement to the front of your sql files to connect to the DB2 subsystem:

- `CONNECT TO D9FG;`

Run the command line processor by typing this on the command line of a telnet session:

- `db2 -f <path>/xxx.sql`

Slide 38

© 2011, IBM Corporation

SHARE
In Anaheim
2011

Setting up to run the clp (Back-up foils)



SHARE
Technology - Connected - Results

Sample .profile for DB2:

```
export DB2_PATH=/usr/lpp/db2/d9fg/db2910_jdbc
STEPLIB=DB9FU.SDSNLOAD:$STEPLIB
STEPLIB=DB9FU.SDSNLOAD:$STEPLIB
STEPLIB=DB9FU.SDSNLOAD:$STEPLIB
PATH=$DB2_PATH/bin:$DB2_PATH/bin:$PATH
LIBPATH=$DB2_PATH/lib:$DB2_PATH/lib:$LIBPATH
CLASSPATH=$DB2_PATH/classes/sqlj.zip:$CLASSPATH
CLASSPATH=$DB2_PATH/classes/db2jcc_javax.jar:$CLASSPATH
CLASSPATH=$DB2_PATH/classes/db2jcc_license_cisuz.jar:$CLASSPATH
CLASSPATH=$DB2_PATH/classes/db2jcc.jar:$CLASSPATH
CLASSPATH=/etc/d9fg/DB2JccConfiguration.properties:$CLASSPATH
export PATH STEPLIB LIBPATH CLASSPATH

export JAVA_HOME=/usr/lpp/java/J5.0/bin/java
export PATH=$JAVA_HOME/bin:$PATH
export CLPHOME=/usr/lpp/db2/d9fg/db2910_base
export CLASSPATH=$CLASSPATH:$CLPHOME/lib/clp.jar
export CLPPROPERTIESFILE=$HOME/clp.properties
. /jcc/jcc3_env.sh

alias db2="java com.ibm.db2.clp.db2"
```

Sample clp.properties file.

```
#Specify the value as ON/OFF or leave them blank
DisplaySQLCA=ON
AutoCommit=ON
StopOnError=
TerminationChar=;

#SERVER1=<URL>,<username>,<password>
DSNA=wsc3.washington.ibm.com:8446/DSNA,sysadm1,sysadm1
Slide 39
```

© 2011, IBM Corporation

Sample .jcc3 script for DB2:

```
#!/bin/sh
## set location of the java & jcc & local homes
JAVA5_HOME=/usr/lpp/java/J5.0
JCC_HOME=/usr/lpp/db2/d9fg/db2910_jdbc
BASE_DIR='pwd'
# Add jcc related jar files to the CLASSPATH
JCP=
# Add jars from JCC_HOME/classes for jcc
JCP=$JCC_HOME/classes/db2jcc.jar
JCP=$JCC_HOME/classes/sqlj.zip:$JCP
JCP=$JCC_HOME/classes/db2jcc_javax.jar:$JCP
JCP=$JCC_HOME/classes/db2jcc_license_cisuz.jar:$JCP
# Add current directory to pick up classes anchored here.
export CLASSPATH="$BASE_DIR":$JCP:$CLASSPATH
# Add jcc related directories to the PATH and LIBPATH
export PATH="$JCC_HOME"/bin:$PATH
export LIBPATH="$JCC_HOME"/lib:$LIBPATH
# Add java directory to the PATH if needed
whence java | grep -q -e "J5.0" -e "J6.0"
rc=$?
if [ rc -eq 1 ]; then
export PATH="$JAVA5_HOME"/bin:$PATH
fi
```

SHARE
In Anaheim
2011

Controlling the timestamps in SYSPRINT and SYSOUT



Setting time stamps to show local time in all the Server logs (SYSOUT, SYSPRINT):

```
ras_time_local = 1 and  
DAEMON_ras_time_local = 1
```

Setting time stamps in Application logs (i.e., joblogs in WAS XD Compute Grid for Java Batch jobs and basically anything written from your java application programs):

```
TZ = EST5EDT
```

CUT0GDT	Coordinated Universal Time	CUT
EST5EDT	Eastern United States, Colombia	CUT -5
CST6CDT	Central United States, Honduras	CUT -6
MST7MDT	Mountain United States	CUT -7
PST8PDT	Pacific United States, Yukon	CUT -8
AST9ADT	Alaska	CUT -9
HST10HDT	Hawaii, Aleutian Islands	CUT -10



Slide 40

© 2011, IBM Corporation



Questions? (for anyone?)

Slide 41

© 2011, IBM Corporation



Q and A?