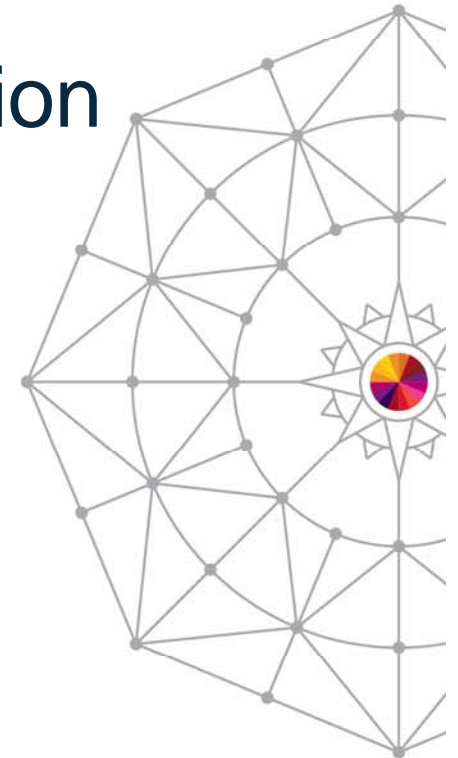# WebSphere Liberty Profile and Traditional WebSphere Application Server – What's New?

David Follis, Gary Picher, Mike Stephen

IBM

August 5, 2014

Session Number 15783

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

CICS*	Parallel Sysplex*
DB2*	RACF*
GDPS*	System z9
Geographically Dispersed Parallel Sysplex	WebSphere*
HiperSockets	z/OS
IBM*	zSeries*
IBM eServer	
IBM logo*	
IMS	
On Demand Business logo	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Oracle.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

HIS is a trademark of HIS Group Inc.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Disclaimer

- The information contained in this documentation is provided for informational purposes only. While efforts were many to verify the completeness and accuracy of the information contained in this document, it is provided "as is" without warranty of any kind, express or implied.

- This information is based on IBM's current product plans and strategy, which are subject to change without notice. IBM will not be responsible for any damages arising out of the use of, or otherwise related to, this documentation or any other documentation.

- Nothing contained in this documentation is intended to, nor shall have the effect of , creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of the IBM software.

- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
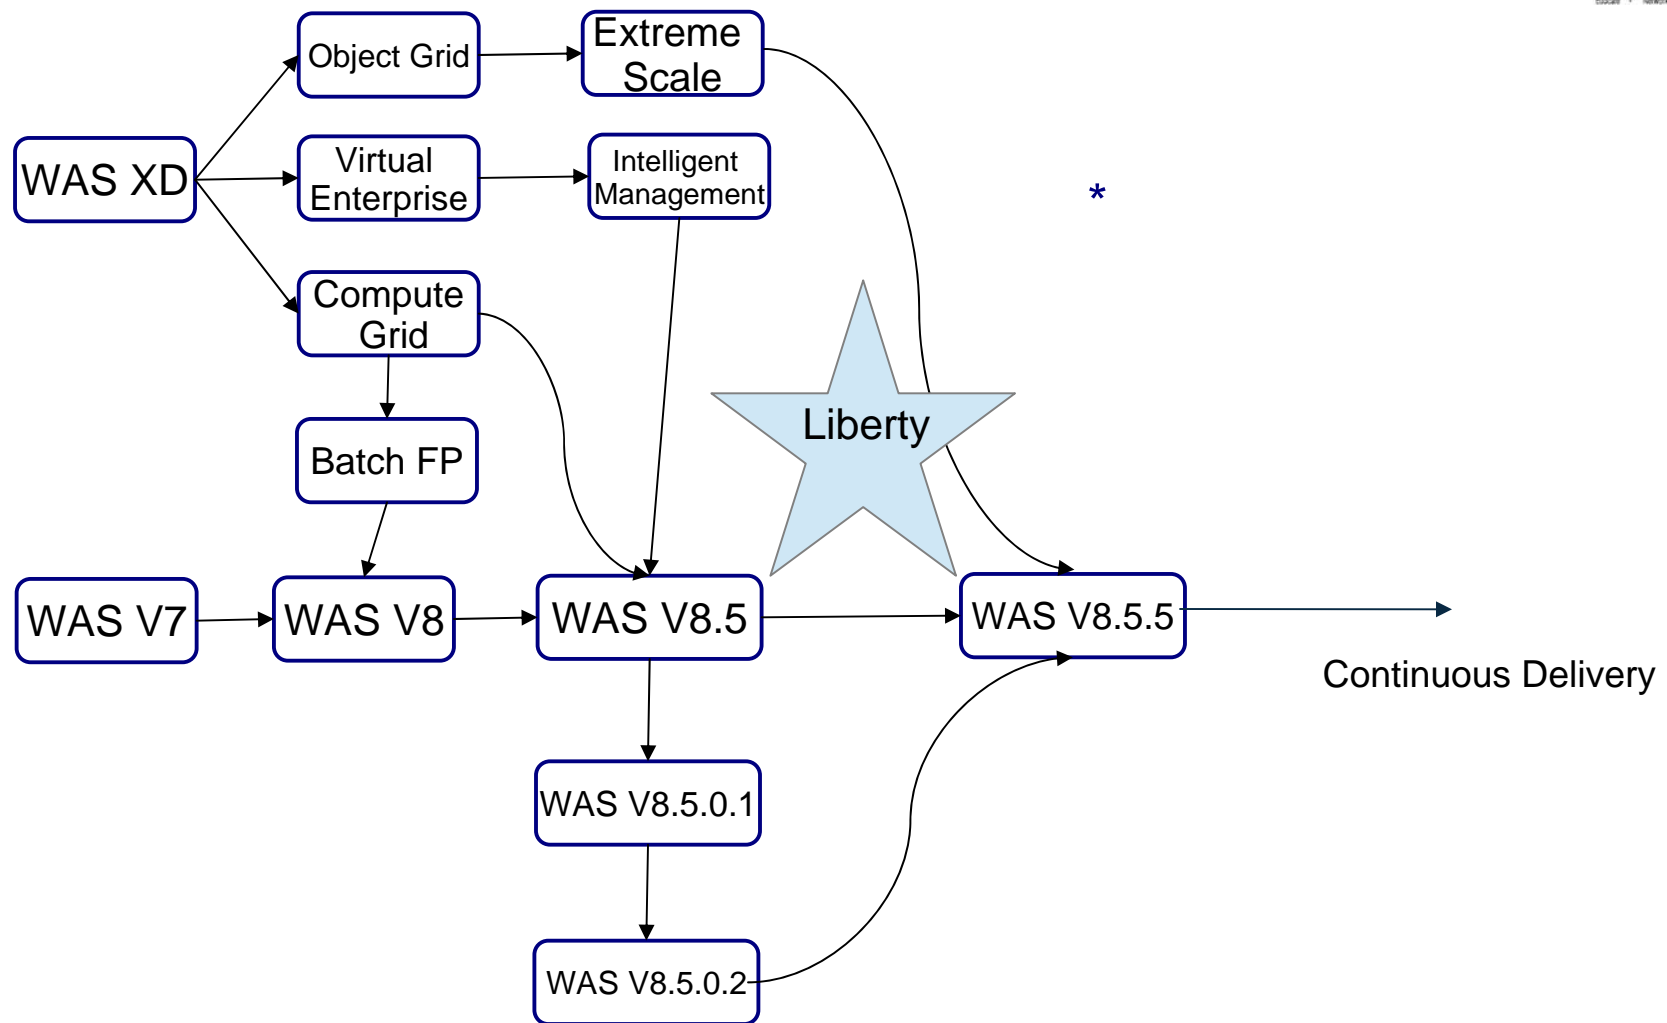
WebSphere Application Server on System Z

| Session	Title	Time	Room	Speaker
15563	Web Apps using Liberty Profile Technology in CICS	Monday 1:30	Grand Ballroom Salon C	Ian Mitchell
15783	WebSphere Liberty Profile and Traditional WebSphere Application Server – What's New?	Tuesday 10:00	Grand Ballroom Salon J	Follis/Picher/Stephen
16278	Feed Your Appetite for Knowledge about z/OS Connect and IBM Mobile Pricing over Breakfast	Wednesday 7:20	Grand Ballroom Salon K	Ian Mitchell
15784	Common Problems and Other Things You Should Know about WAS on z/OS	Wednesday 10:00	Platinum Ballroom Salon 7	Mike Stephen
15609	WebSphere Liberty Profle, Windows and z/OS, Hands-on Lab	Thursday 12:25	Grand Ballroom Salon G	Cast of Characters
15782	z/OS Connect: Opening up z/OS Assets to the Cloud and Mobile Worlds	Thursday 1:30	Grand Ballroom Salon A	David Follis
15785	IBM Installation Manager for z/OS System Programmers: Web-based Installs, Fix Packs, and How iFixes Really Work.	Friday 8:30	Grand Ballroom Salon G	Stephen/Mierzejewski
15999	Application Level Resource Monitoring of WebSphere z/OS – DB2 JDBC Workloads	Friday 10:00	Platinum Ballroom Salon 2	Follis/Catterall

The Big Picture

```
           Object Grid ──────► Extreme
          ▲                     Scale ──┐
         /                              │
        /   Virtual      ──►  Intelligent      *
WAS XD ─────Enterprise         Management
        \                           │
         \  Compute                 │
          ▼  Grid ──────┐           │
                        │        ╭──────╮
          Batch FP      │        │Liberty│
             │          │        ╰──────╯
             ▼          ▼          │
WAS V7 ──► WAS V8 ──► WAS V8.5 ──────► WAS V8.5.5 ──────►
                         │                              Continuous Delivery
                         ▼
                    WAS V8.5.0.1
                         │
                         ▼
                    WAS V8.5.0.2 ──┘
```

* Client Only on z/OS

WebSphere Application Server
Full Profile (T-WAS)
What's New?
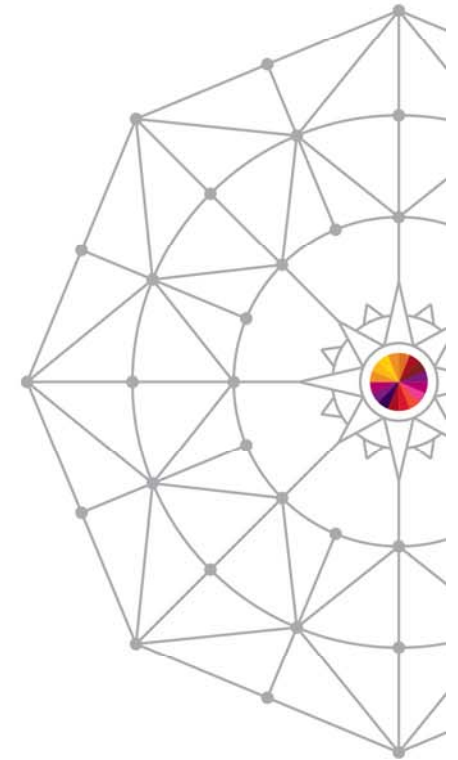
USS APAR OA43650

- z/OS UNIX Services is changing the behavior of several interfaces to restrict the symbolic links that can be used for UNIX files that are loaded or executed from APF-authorized or must-stay-clean (RACF program controlled) address spaces.

- These interfaces will now require that any symbolic link that comprises the pathname to a target executable file have an owning UID of 0.

- For WAS on z/OS, we believe the only affected paths are:
  - startServer.sh shell script provided with WAS
  - node agent start during federation

- Symptoms include:
  - BPXP029I OPEN ERROR FOR FILE PATH asv85config/azcell/azdmnode/DeploymentManager/lib/s390-31/bboorb DEVICE ID 18 INODE 348.
  - ABEND EC6-xxxC04A

- WAS will ship an update in:
  - 7.0.0.35, 8.0.0.10, 8.5.5.4
  - iFixes and ++APARs will be available for other levels

- This update changes how authorized code to issue the START console command is accessed

- If you only start/stop the servers from the admin console or with MVS console START and STOP commands, and manually start node agents after federation then WAS is unaffected.

XML File Extended -- Control Driven to Request Level

As we saw, the XML file identifies requests ... this new function then picks up and drives various WAS behavior controls from server level down to the request level:

```
<Classification schema_version="1.0">
    <InboundClassification type="http" schema_version="1.0"
        default_transaction_class="Z9DEFLT" >
        <http_classification_info
            uri="/SuperSnoopWeb/*" transaction_class="Z9TRANA"
            description="Snoop"                                    />
        <http_classification_info
            uri="/MyIVT/*" transaction_class="Z9TRANB"
            description="MyIVT"                                    />
    </InboundClassification>
</Classification>
```

Granular Control to Request Level

Granular Control to Request Level

Various Timeouts

Stalled Thread Dump Actions

CPU Time Used Limit

DPM Interval and Dump Action

SMF Recording

Tracing

Message Tagging

Timeout Recovery Actions

Topics to Cover In this Section:
- What those functions are and how they work
- How to dynamically reload a new or updated XML file
- How to dynamically revert to previous XML file

Dispatch Progress Monitor (DPM)

Look at the response times of a lot of requests over time....

Most will be fairly short

And collect the configured doc (javacore, etc) at that interval

But occasionally one takes much longer than normal

Set the DPM Interval to a value greater than the normal response time
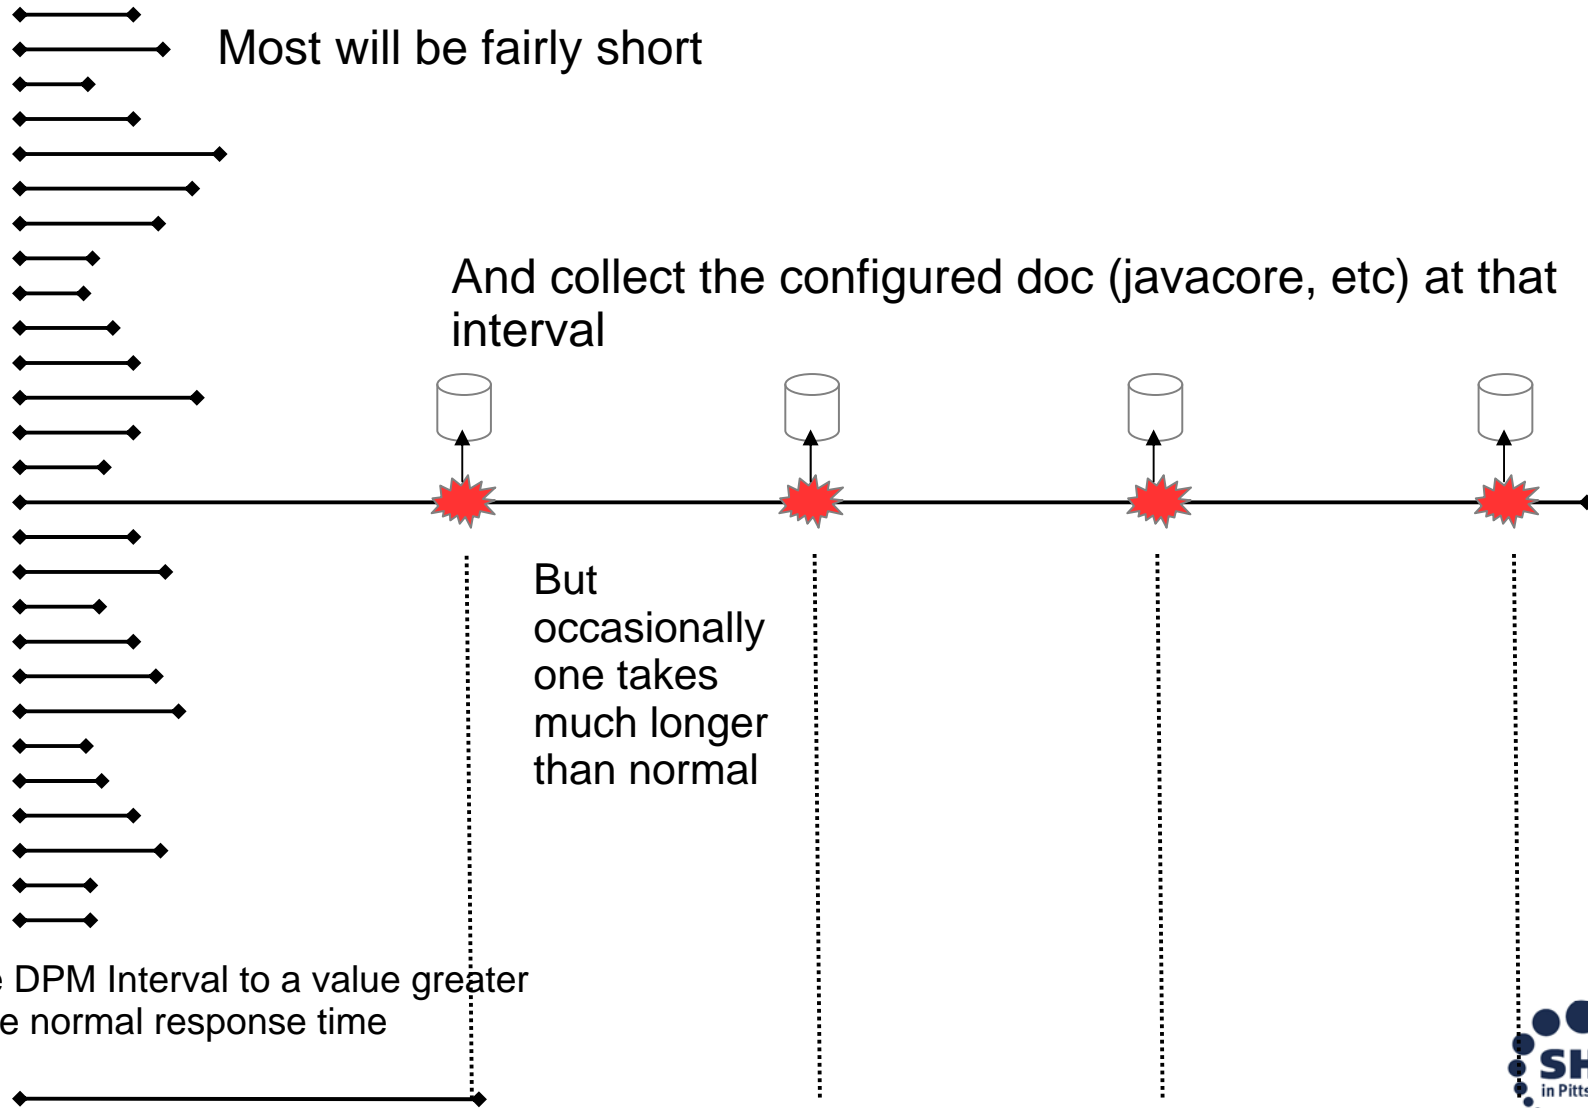
Dispatch Progress Monitor (DPM) Settings

```
dispatch_timeout="_____"
queue_timeout_percent ="_____"
request_timeout="_____"
stalled_thread_dump_action="_____"
cputimeused_limit="_____"
cputimeused_dump_action="_____"
dpm_interval="_____"
dpm_dump_action="_____"
SMF_request_activity_enabled="__"
SMF_request_activity_timestamps="__
"
SMF_request_activity_security="__"
SMF_request_activity_CPU_detail="__
"
classification_only_trace="__"
message_tag="_____"
timeout_recovery="_____">
```

DPM stands for Dispatch Progress Monitor. It is a function that will process a dump action every *n* seconds.
`dpm_interval` is the interval period expressed in seconds
`dpm_dump_action` is the same as we just saw for the other dump action: `svcdump, javacore, heapdump, traceback, javatdump` and `none`

This function has a set of `MODIFY` commands that may be used to clear DPM settings or reset to XML settings See WP102023 for the details on these `MODIFY` actions for DPM

Use of the WLM Health API

Provides a way to ratchet up the WLM server "health" value so Sysplex Distributor flows work gradually to a server that's been restarted, but not yet fully "warmed up"

Prior to this function, a server freshly restarted would be seen as having a WLM health value of 100% ... that told Sysplex Distributor to treat it as equal to other servers in the cluster.
The issue is a freshly restarted server has not yet gone through JIT processing and caching of various objects. It has not "warmed up" to handle work equal to other servers.

Introduced In

7.0.0.31
8.0.0.8
V8.5.5.1

Provide a way to modify WLM "health" value over time so a server has a chance to "warm up" before being considered equal to other servers in cluster

```
wlm_health_increment=15
wlm_health_interval=15
```
Ratchet up the health value in increments of 15% every 15 seconds

```
11.56.14 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 0
11.56.44 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 15
11.56.59 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 30
11.57.14 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 45
11.57.29 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 60
11.57.44 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 75
11.58.00 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 90
11.58.15 STC00085 BBOO0411I SERVER WLM HEALTH PERCENTAGE IS NOW 100
```
Requires Sysplex Distributor config VIPADISTRIBUTE DISTMETHOD set to SERVERWLM.
Careful with mixed levels – see APAR PI14413 – in 7.0.0.33

Proper handling of CARTs

Modify commands issued with a Command And Response Token (CART) don't have the CART included in WTOs issued in response

Automation that issues console commands via MGCRE that include a CART expect to have that CART value included with WTOs issued in response to the command.

This allows the automation to easily find the response.

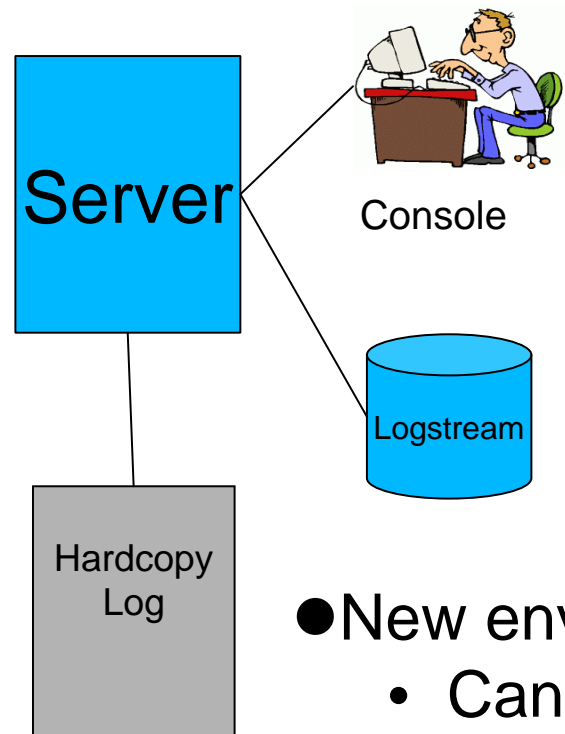WebSphere Modify command handling ignores CARTs

Introduced In

V8.5.5.2

WTOs issued in response to Modify commands issued via MGCRE (or TSO CONSOLE) with a CART will include the CART.

Message routing

Server

Console

Logstream

Hardcopy Log

- ● Messages are written as:
  - ● WTOs to the console
  - ● WTOs to the log
  - ● Writes to SYSOUT or Logstream
- ● The destination for a message is determined by the code that issues it
- ●New environment variables override the code
  - • Can move the message to a new place entirely, or duplicate it elsewhere
- ●Force messages (by ID) to a chosen target
- ●Or 'NONE' to suppress entirely
- ●Update dynamically with MODIFY
- ●Use DISPLAY to see current configuration

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

Taking Action on Java Out Of Memory (OOM)

- The JVMTI (Tool Interface) provides a way to take action when various events occur

- One such event is 'heap resource exhaustion'

- Function added to WAS to register a JVMTI exit and optionally take some action when an Out Of Memory occurs

- Controls:

    - ras_java_oom_action=**NONE** | WTO | SVCDUMP

    - ras_java_oom_interval=**600** | # of seconds

    - MODIFY server,JAVAOOM,INTERVAL= | ACTION= | RESET

    - MODIFY server,DISPLAY,JAVAOOM

- BBOO0404E JVM OUT OF MEMORY IN ASID=nnnnX

- Available in 8.0.0.3 and in 7.0.0.23

PM85194 (Dynamically adjust CR threads)

The number of Control Region worker threads is fixed at startup.

Configurable from Application Servers → <SERVER> → Container Services → ORB Service → Custom Properties → was.controlThreads (default is 25)

Introduced In

V8.0.0.8
V8.5.5.2

Thread pool size is fixed once the server starts.

Under certain conditions, the thread pool can deadlock with all threads synchronously waiting for work to complete, but the work is stuck on the queue.

Although rare, a server restart is required at this point.

Introduced a new environment variable to allow the pool to grow:

```
control_region_thread_pool_maximum_size
```

Default is "0", but under the covers "0" means "determine a sensible maximum size based on the number of Servants and worker threads per Servant".

The thread pool starts at the same size as before, but can now dynamically grow to the maximum when deadlocks are detected.

PM74923 - Better Living Through Server Output Mgmt

Provides a means of routing SYSPRINT and SYSOUT to a UNIX file rather than JES with better file management than simple output routing we had before *Introduced In*

In development and test environments Java developers wish to view server output from UNIX file, not JES. Ability to push server output to UNIX file has existed for some time, but output file management was very limited.

V7.0.0.29
V8.0.0.6
V8.5.0.2

With PM74923 comes two new variables:
`DAEMON_redirect_server_output_dir` (for the Daemon)
`redirect_server_output_dir` (for everything else)

Point to path where output files will be written

The output file names will be uniquely named by WAS
`F <server>,ROLL_LOGS` will close existing and start new
Use whatever UNIX file viewing mechanism you wish

`ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102267`
More detail, including how the IHS Apache webserver can be used to list output files, control access to those files, and browse the files

DISPLAY WORK Support for Message from the CRA

Provides a way to use DISPLAY WORK to show work coming out of the CRA, for either SIBus (internal messaging) or Activation Spec support for MDBs:

Introduced In

Prior to this function, work coming out of the CRA (SIBus and/or ActSpec for MDBs) was not included in DISPLAY,WORK

7.0.0.31
8.8.0.8
V8.5.5.1

`F BBOS001,DISPLAY,WORK,CRA`

```
BBOO0255I TIME OF LAST WORK DISPLAY 2013/04/30
21:31:37.470352
BBOO0256I TOTAL CRA REQUESTS 1000 (DELTA 0)
BBOO0257I CURRENT CRA REQUESTS 100
BBOO0258I CRA REQUESTS IN DISPATCH 100
BBOO0410I HIGHWATER CRA REQUESTS 500
BBOO0267I TOTAL CRA TIMEOUTS 0 (DELTA 0)
BBOO0188I END OF OUTPUT FOR COMMAND
DISPLAY,WORK,CRA
```

100 in flight, all dispatched, so none queued right now

DISPLAY WORK "Highwater" Reporting

Provides a way to know not just how many requests are in the server right now, but also the maximum number seen since the server was started:

7.0.0.31
8.0.0.8
V8.5.5.1

Prior to this function, the DISPLAY,WORK only showed work currently in the server. There was no "high water" indicator.

`F BBOS001,DISPLAY,WORK,SERVLET`

SERVLET is one command operand that displays all HTTP work. Other operands exist: EJB, MDB, etc.

```
BBOO0255I TIME OF LAST WORK DISPLAY 2013/05/13
21:04:27.773915
BBOO0256I TOTAL SERVLET REQUESTS 1000 (DELTA 0)
BBOO0257I CURRENT SERVLET REQUESTS 100
BBOO0410I HIGHWATER SERVLET REQUESTS 500
BBOO0258I SERVLET REQUESTS IN DISPATCH 100
BBOO0267I TOTAL SERVLET TIMEOUTS 0 (DELTA 0)
BBOO0188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,SERVLET
```

Tracing|Logging Large Callstacks and BUFFER OVERFLOW

Provides a way to capture the full Java stack trace into an FFDC record. Still see the buffer overflow in the SYSOUT, but now you have additional doc

Introduced In

V7.0.0.31
V8.0.0.8
V8.5.5.1

Many Java stack traces show up looking something like this:

```
Java stack line
Java stack line
Java stack line
   :
Java stack line
*** BUFFER OVERFLOW ***
Java stack that doesn't show up
Java stack that doesn't show up
```

Original design of WAS z/OS back in the 31-bit JVM day did this to protect against very large output consuming the heap as part of converting the stack to EBCDIC

And all too often the really interesting part of the stack trace would be down here ... not shown because it came after the clip-point for overflow

In maintenance levels shown above the same "BUFFER OVERFLOW" occurs, but the full stack is *also* put out to an FFDC incident record

```
***BUFFER OVERFLOW*** TRACE DATA ROUTED TO FFDC REPORTING A NativeMessageOverflowException***
FFDC1003I: FFDC Incident emitted on <path and file> processEvent NativeMessageOverflow_1
```

The actual FFDC will have the sequence number as part of the title so you can match them up.

Need to trace a single Thread / TCB in WebSphere?

- Why you might need this
  - There is a thread that is accumulating CPU
  - Occasional javacore dumps have not proven beneficial

- You COULD turn on ALL Tracing for a few seconds
  - Get MANY lines of trace output
  - Have to filter to find the one thread you are interested in

- What if you could turn Trace on for the ONE thread

Need to trace a single Thread / TCB in WebSphere?

- NEW MODIFY server command: TRACEBYTHREAD
- Set the ASID / TCB you want to trace: (in hex)
  - → MODIFY server,TRACEBYTHREAD,ASIDX=nnnn,TCBX=nnnnnn
- Enable the tracing you want
  - → MODIFY server,TRACEJAVA='com.ibm.ws.security.*=all;SSL=all'
- IMPORTANT - When you have the tracing you want, turn tracing back off
  - → MODIFY server,TRACEINIT
- AFTER tracing is turned off, disable trace by thread
  - → MODIFY server,TRACEBYTHREAD,RESET

- Introduced in FixPack 8.5.5.2
- For this and other nuggets:
- Whitepaper: **Hidden Gems in the WAS z/OS Maintenance Stream**
  - http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102371

collector.sh on WebSphere App Server z/OS

- collector.sh translates its ouput ASCII > EBCDIC

- causes problem on z/OS

- Issue these 3 USS commands from directory outside of your WebSphere config directory (ie. /tmp)

  export IBM_JAVA_OPTIONS=-Dfile.encoding=ISO8859-1
  WAS_HOME/DeploymentManager/profiles/default/bin/collector.sh
  WAS_HOME/AppServer/profiles/default/bin/collector.sh

- May see 'deprecated' message… don't be alarmed...

collector.sh on WebSphere App Server z/OS

- File created in the directory where the commands were run:

- hostname.cellname.nodename.default-WASenv.jar

- When sending in files prepend PMR# to the filenames:

- eg:

    12251.999.000. boss0181.pok.ibm.com-bcell1-bnode1-default-WASenv.jar

- Feed into ISA Tool:

    'WebSphere Application Server Configuration Visualizer [Report]'

WebSphere Application Server Configuration Visualizer

- Generates an interactive HTML visualization of a WebSphere Application Server configuration. Accepts any combination of archive files containing configuration directories, such as Data Collector output. Configurations from multiple nodes in a cell will be merged into a single visualization, and multiple cells can be displayed in the output.

- It works with any (full profile) WebSphere Application Server based product, such as IBM Business Process Manager Advanced, WebSphere Process Server, WebSphere Enterprise Service Bus, WebSphere Service Registry and Repository.

- Nodes are laid out vertically, with the servers (including nodeagent/dmgr) contained in the node that hosts it.

- The clusters are shown horizontally, spanning the servers in that cluster.

- Default messaging provider (Service Integration Bus) buses are shown vertically on the right, with the messaging engines overlaid on the servers/clusters that host them.

WebSphere Application Server
Configuration Visualizer

- Unique database names are shown at the bottom, from analyzing all the JDBC Data Sources in the cell.

- If you have multiple core groups, the core group membership is shown for each server and a summary of the core groups is shown at the bottom.

- If you click on an item such as a cell, cluster or server, you can see more detail such as the applications deployed, the port numbers, log locations, thread pools etc.

- A list of all the deployed applications is shown if you click on the cell name (top left corner of the outer box).

- The deployment details of each application module are shown under the cluster/server to which it is deployed (highlighted if you click the link in the 'module deployments' on the cell).

- The mapping from resource references to JNDI resources is shown for the application deployments, such as JMS Connection Factories and JNDI Data Sources.

WebSphere Application Server
Configuration Visualizer
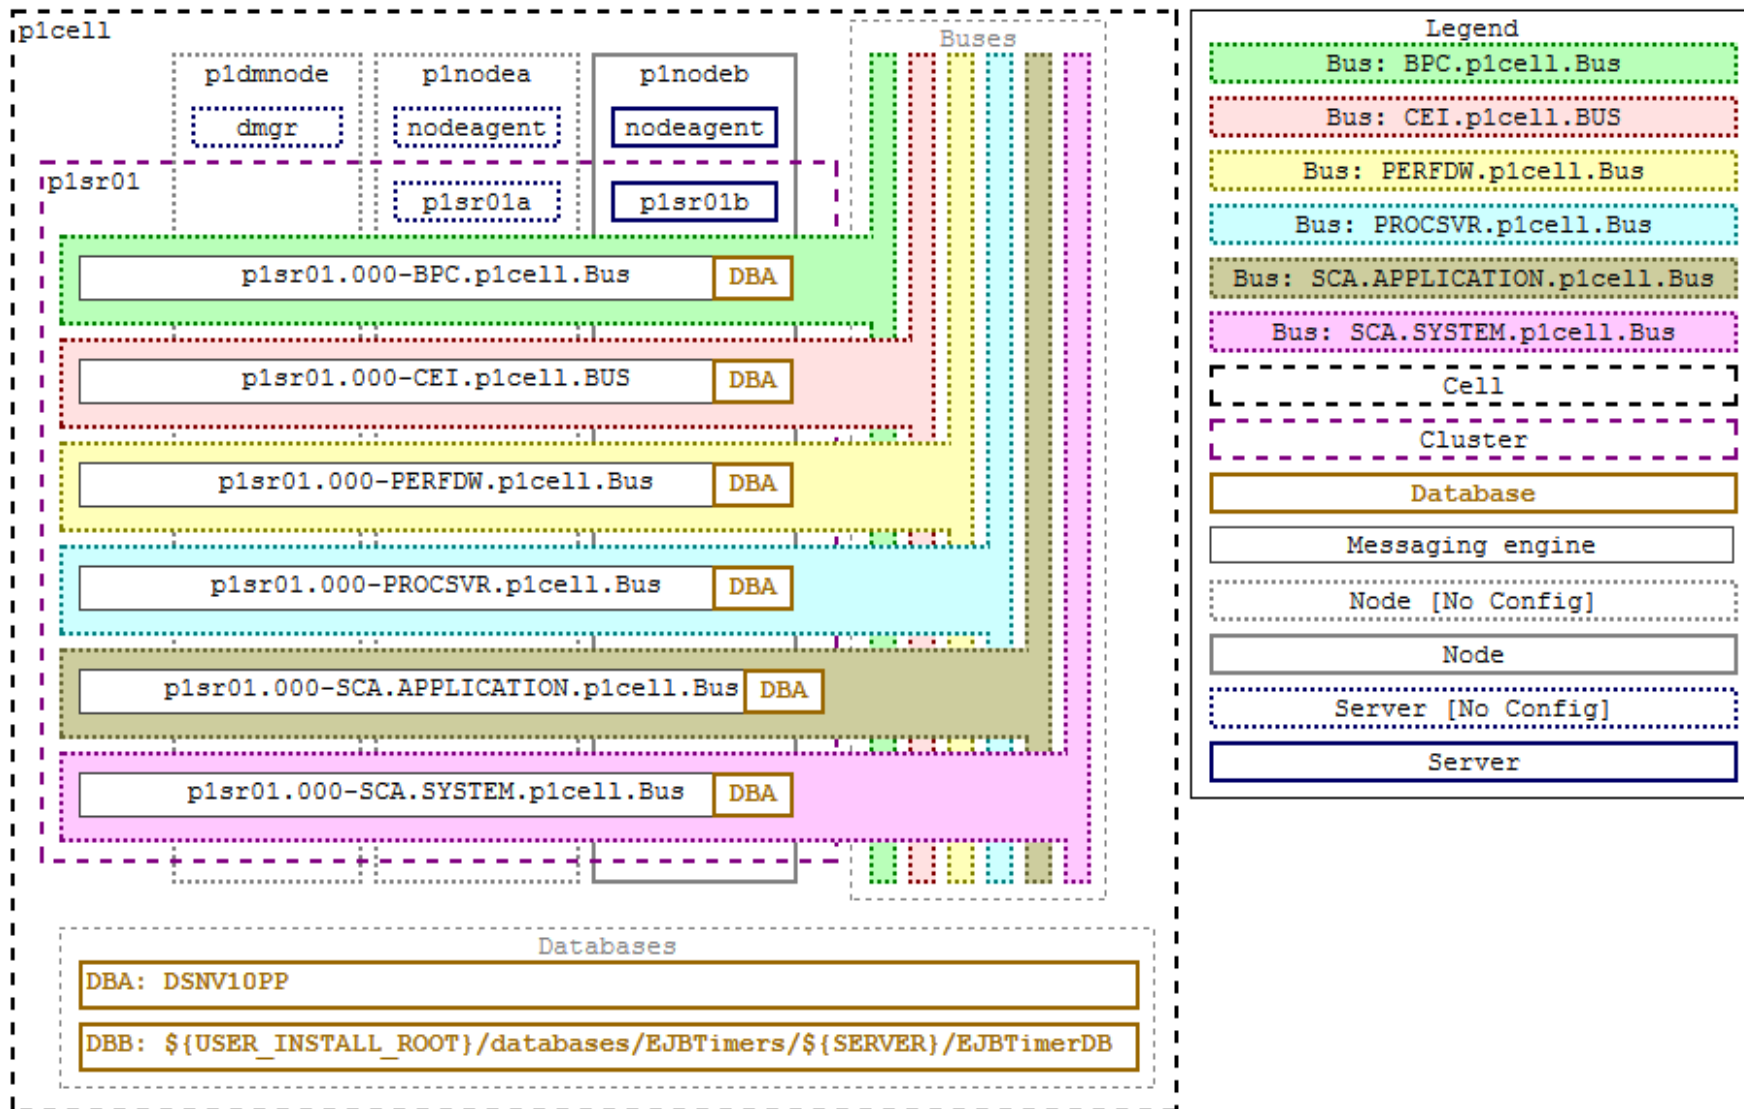
Show all detail for browser search

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

WebSphere Liberty Profile – What's New?
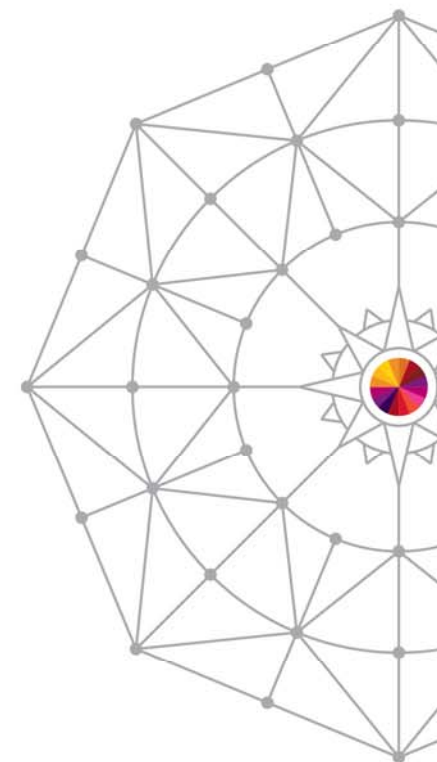
Liberty Administrative Center

Existing administrative tools for the lightweight Liberty Profile of the WebSphere Application Server include Eclipse developer tools and a command line interface.
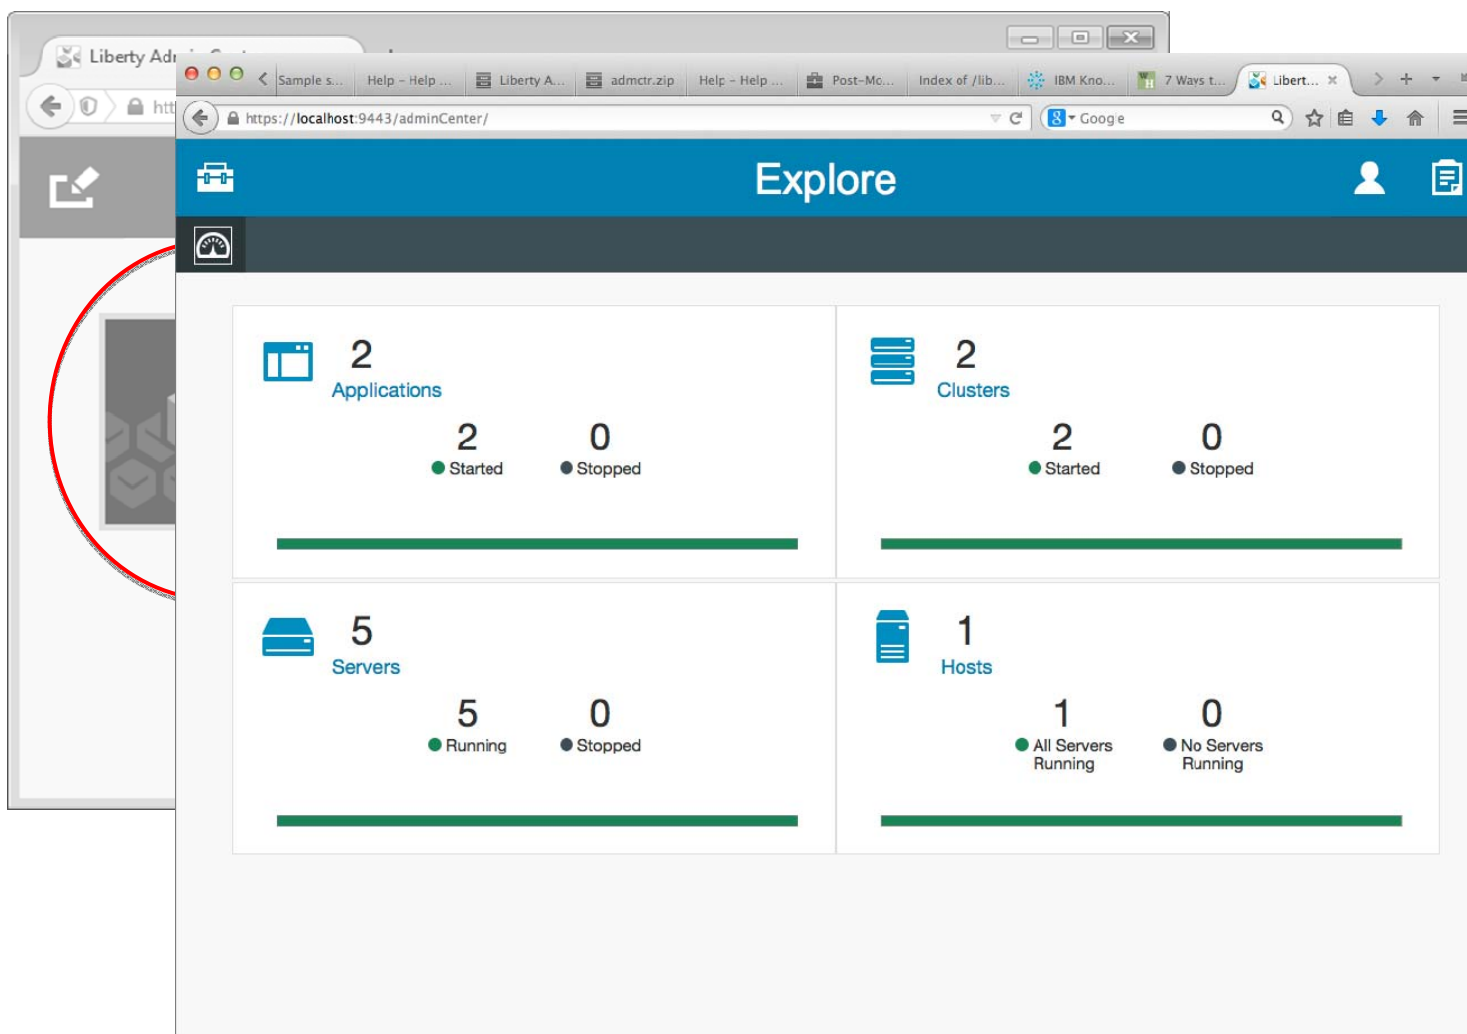
The Liberty admin center extends these capabilities to include a browser-based interface for deploying, configuring, and administering Liberty environments via desktop, tablet and smartphones.
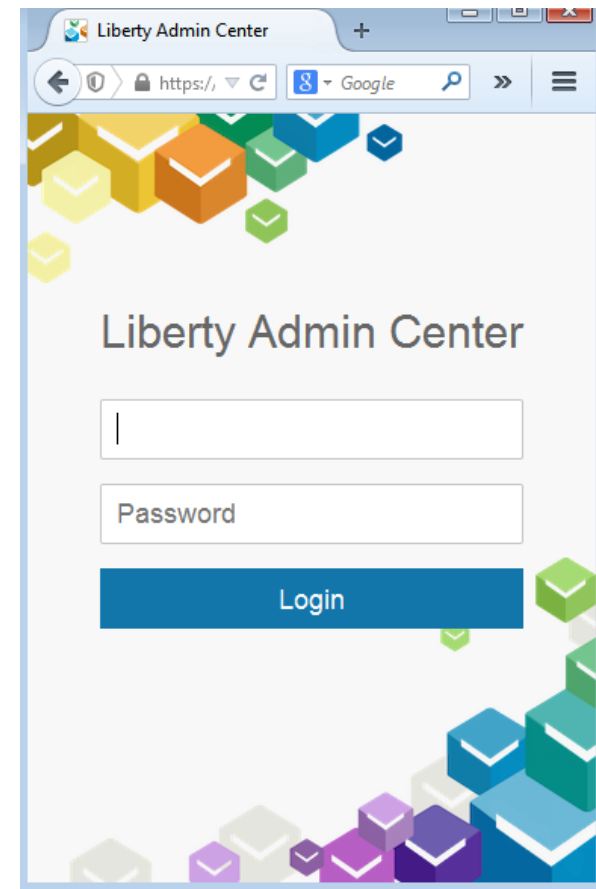
Collective Default Toolbox

Explore the Collective

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval
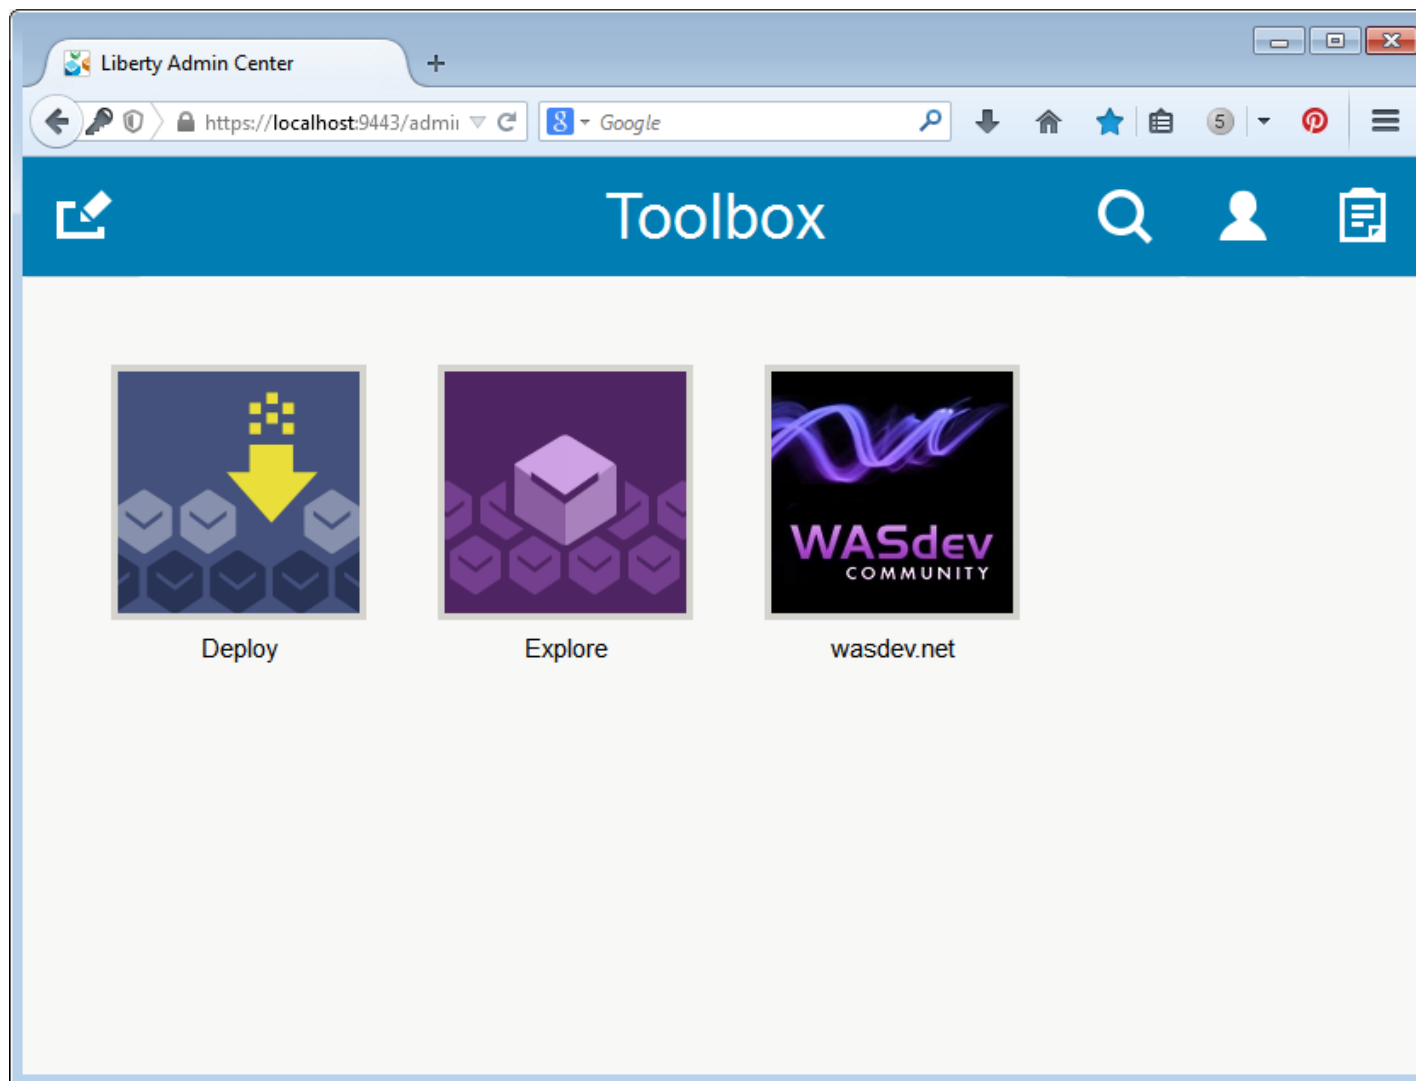
Explore the Collective Servers
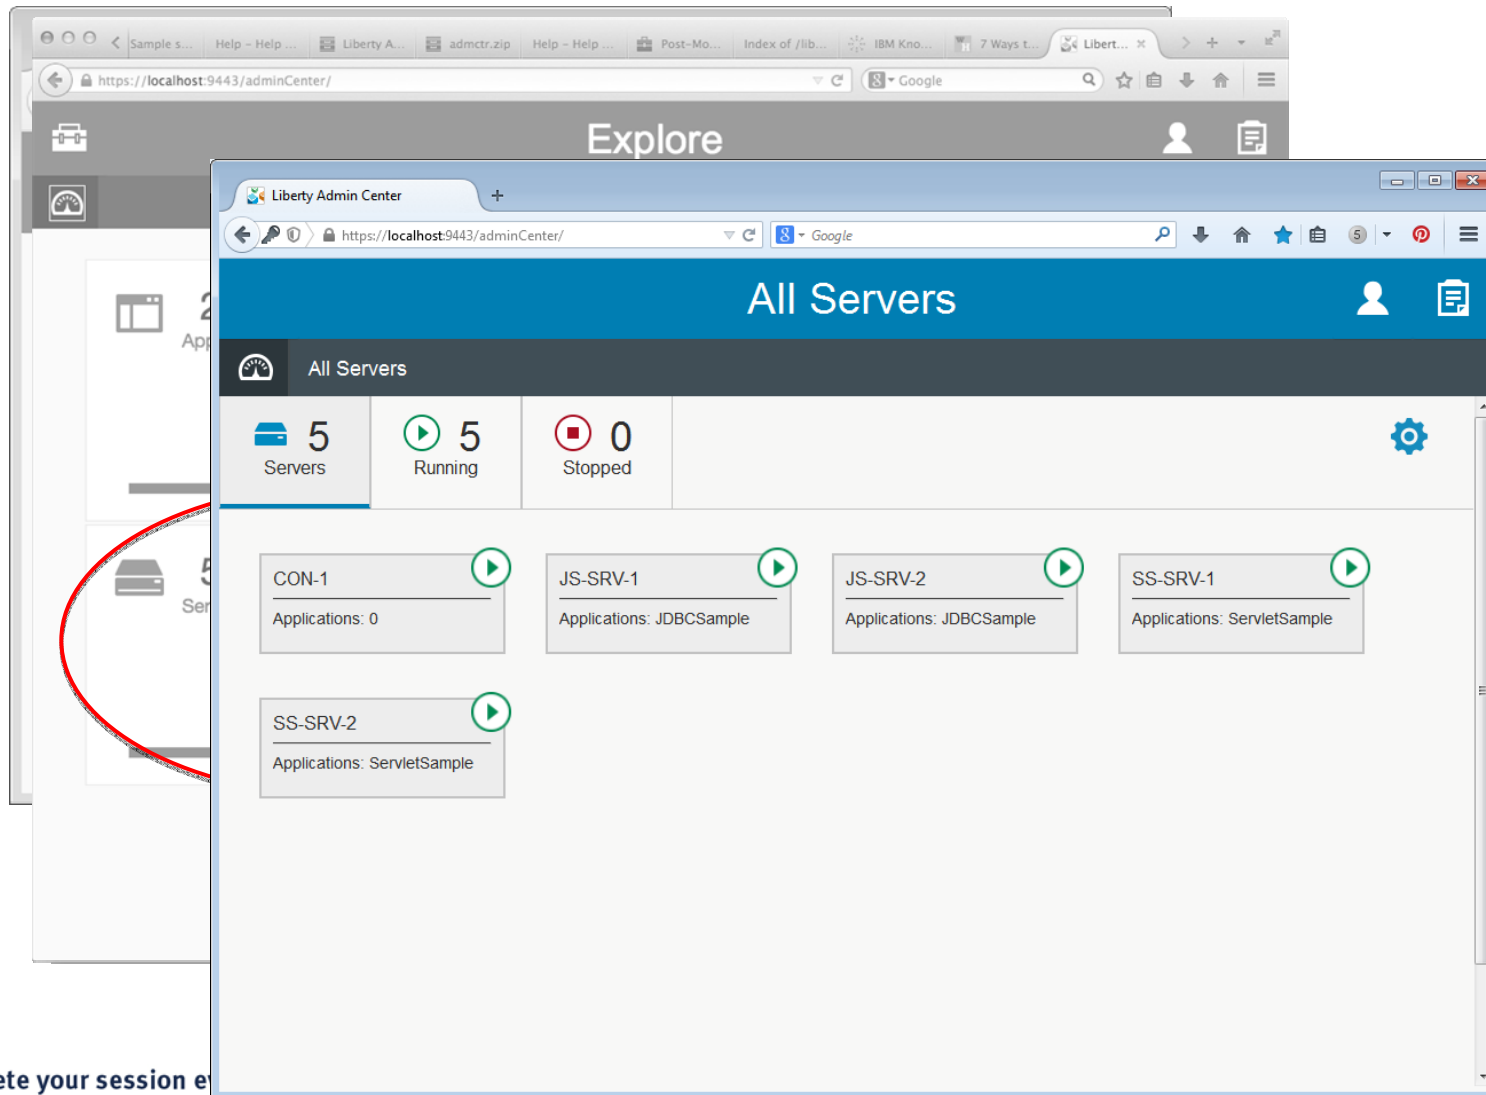
Liberty Administrative Center

Lifecycle Actions

As an Admin…

- I want to start and stop my servers

- I want to start and stop my clusters

- I want to start and stop my applications
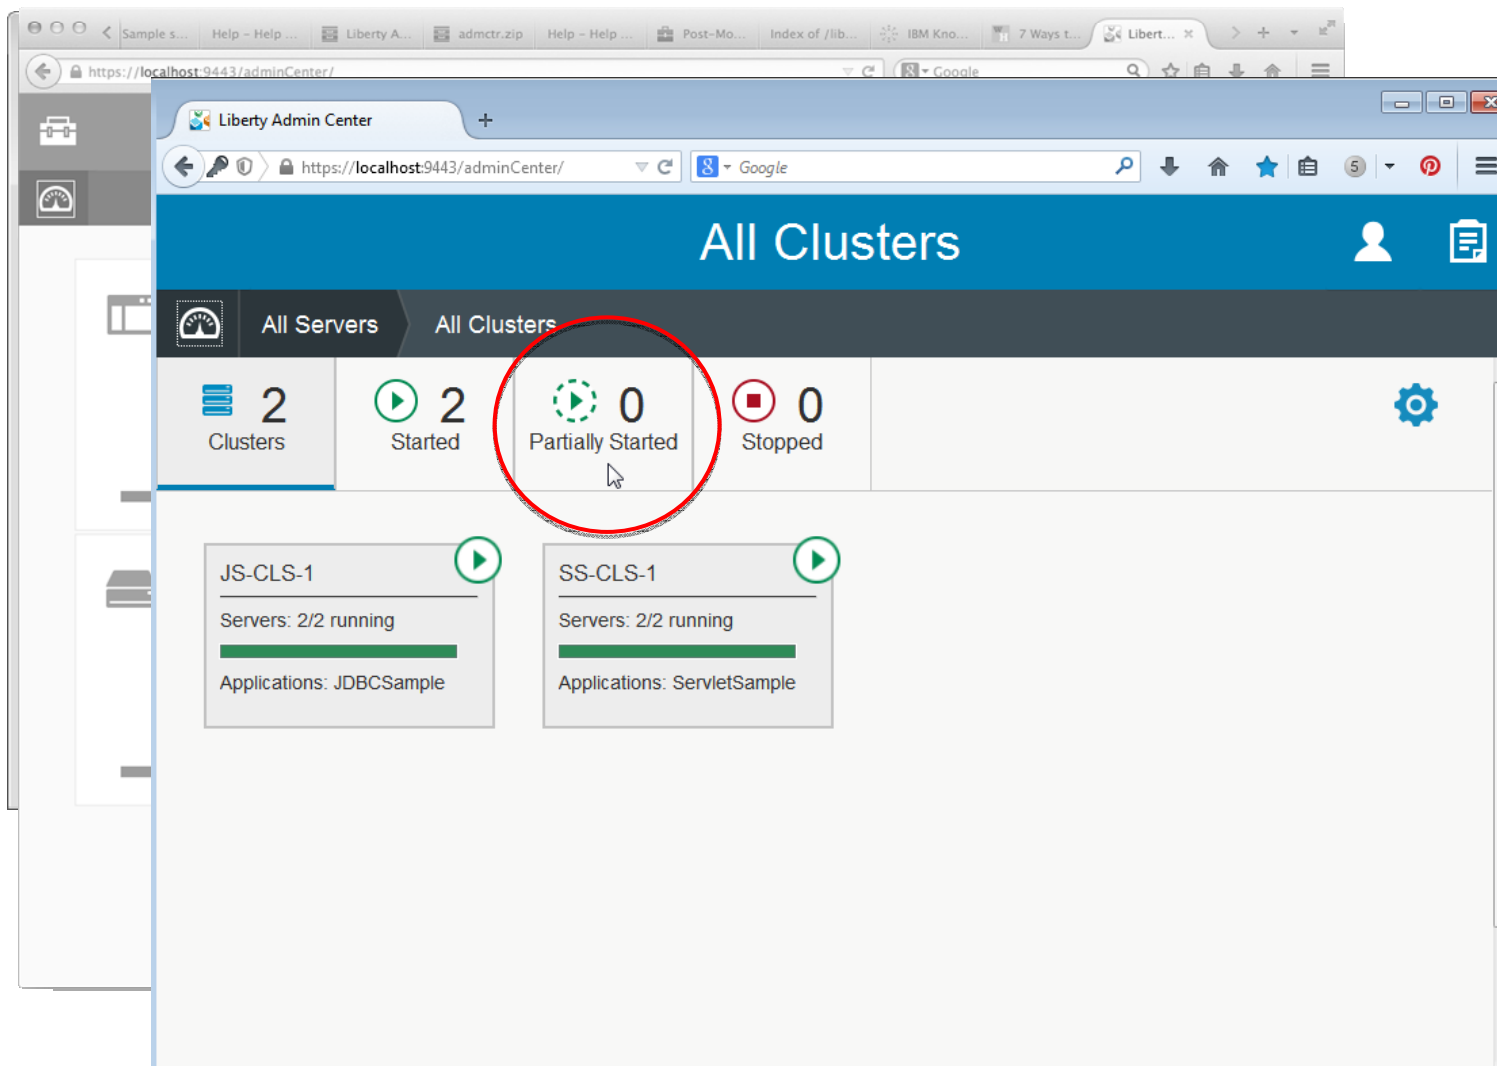
- I want to identify potential problems

Liberty Administrative Center

Start/Stop/Restart a Server
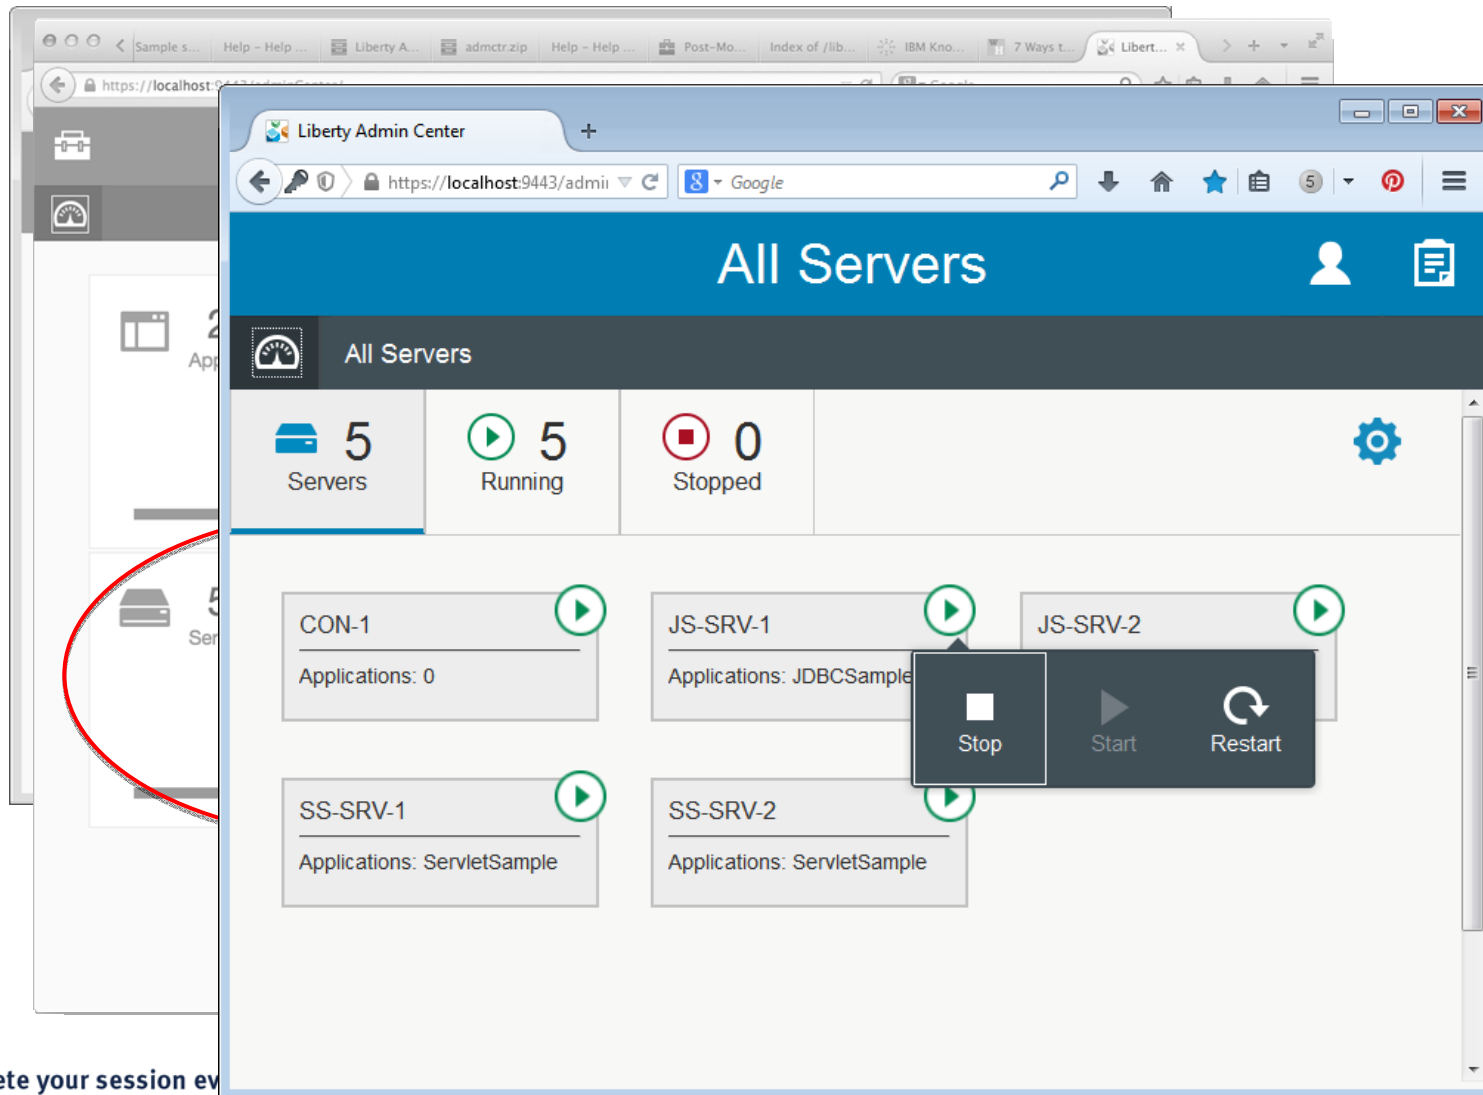
Liberty Administrative Center

View Potential Problems

Liberty Administrative Center

Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval

More Info…

- **The 'GO-TO' place for Liberty**
  - https://developer.ibm.com/wasdev
  - https://developer.ibm.com/wasdev/downloads
    - ALL the features.. Including the Admin Center

- **Setting up Admin Center**
  - https://www.youtube.com/watch?v=NkNXxMp5fAc

- **Touring the Liberty profile Admin Center**
  - https://www.youtube.com/watch?v=Gn4wiYNWkNE

Third Party Extensibility

The Liberty Extensions System Programming Interface (SPI) provides the ability to extend the Liberty profile with custom features, including full life cycle management

- Supports third party extension of the runtime
  - ▸ Provide extension features enabled from server.xml just like base Liberty features
  - ▸ Receive configuration information from server.xml
  - ▸ Integrate with runtime by accessing Liberty SPIs
  - ▸ Expose APIs that applications can access

- WDT project type for feature development

- A product extension is a directory on disk structured like the wlp dir

- All content for a feature is relative to the extension location the feature is installed to

- Registered in wlp/etc/extensions
  - One file per extension
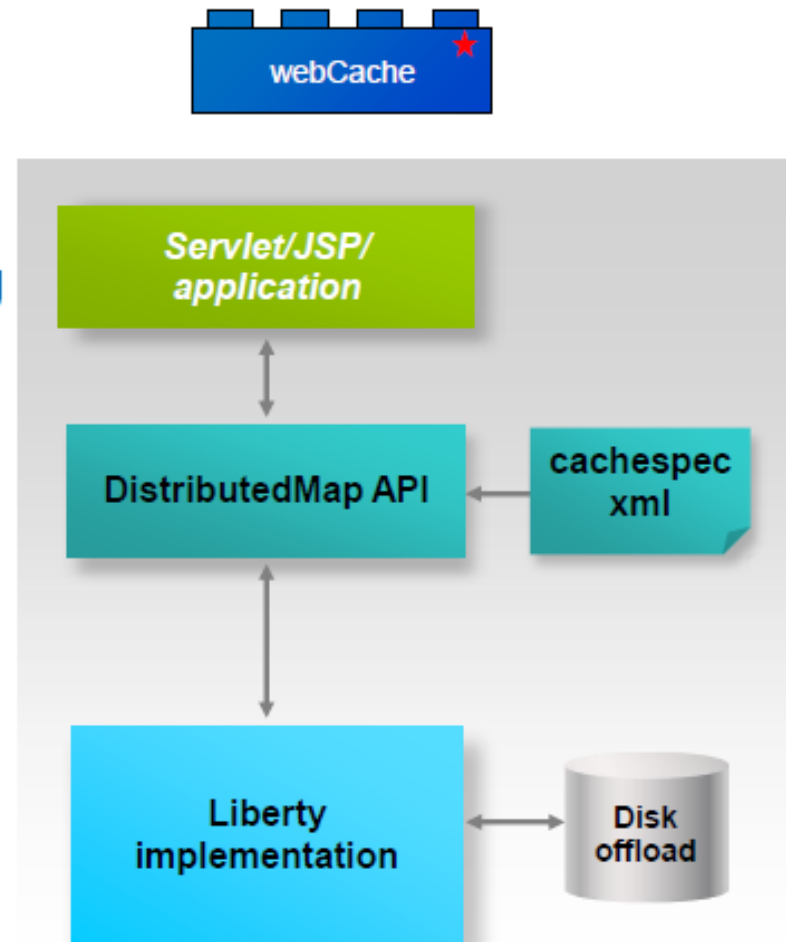
  - File is named <extension name>.properties

  ```
  com.ibm.websphere.productId=<your product id>
  com.ibm.websphere.productInstall=<absolute, or relative file path>
  ```
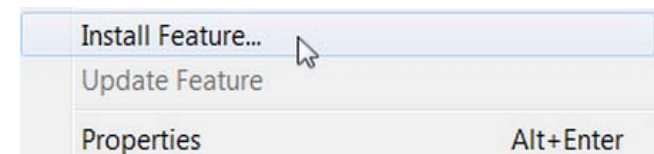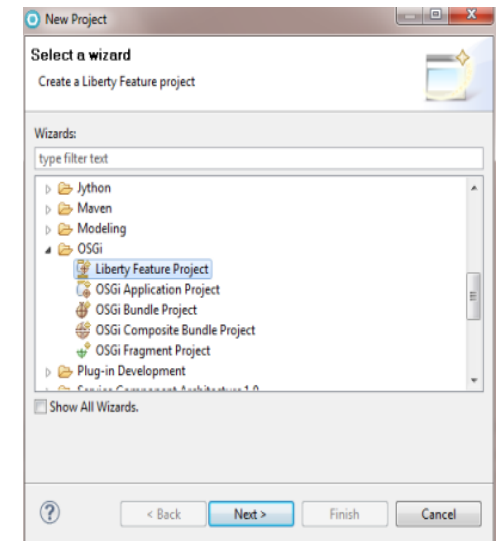
Liberty – Support for WebSphere Web Cache

> **WAS v8.5.5 provides support for WebSphere Web Cache, a.k.a. DynaCache, enabling a local caching service for simplified development of applications**

- Local server cache for dynamic web content

- Enables developers to easily develop / test applications for local server caching solutions

- Same API as full WAS

- Optional disk offload

- LRU replacement policy

- Run as local server cache in production or move applications unchanged to distributed caching via WebSphere eXtreme Scale

webCache

Servlet/JSP/ application

DistributedMap API ← cachespec xml

Liberty implementation ↔ Disk offload

# Liberty "Collective" Administration

> **WAS v8.5.5 introduces a new administrative model enabling the grouping of Liberty servers into collectives for common management and static clustering**

- New Liberty collective management infrastructure
  - Lightweight centralized management of Liberty profile servers
  - Agent-less administration
  - Deployment of topology of WAS Liberty profile servers
  - Static cluster management
  - Application and server deployment/update
  - JMX API through Java, Jython, and Jconsole clients

- Configuration:
  - Controller (requires WAS ND or WAS z/OS license):
    - *<feature>collectiveController-1.0</feature>*
  - *Highly available*
  - Managed server:
    - *<feature>collectiveMember-1.0</feature>* *(All Editions)*
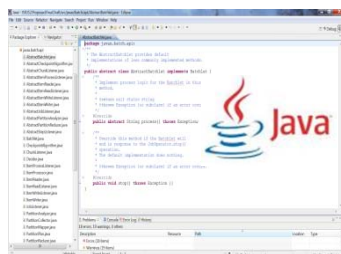   - *<feature>clusterMember-1.0</feature>* *(ND+ only)*
- Scales to 1000s of JVMs
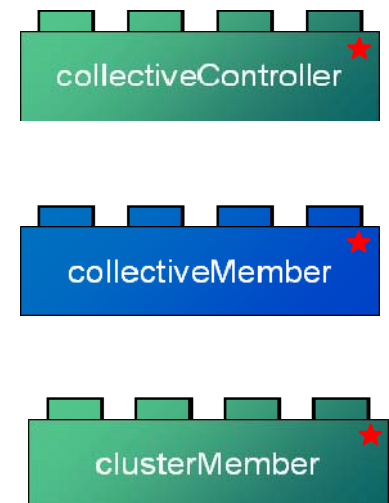
Liberty Collectives (WAS ND)
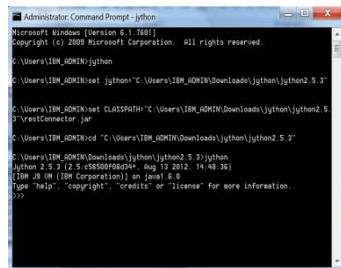
jython

jconsole

java

Collective
Controller

<collectiveMember controllerHost=…/>

Liberty Profile
(Collective Member)

Collective Member

<collectiveMember controllerHost=…/>
<clusterMember name=…/>

Liberty Profile
(Collective Member)
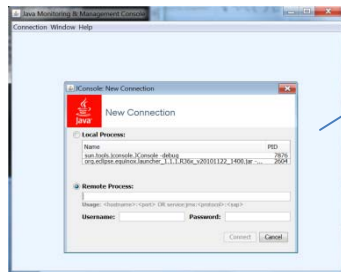
App Cluster

- Provides Operational Registry
- Access all members through controller
- Highly Available
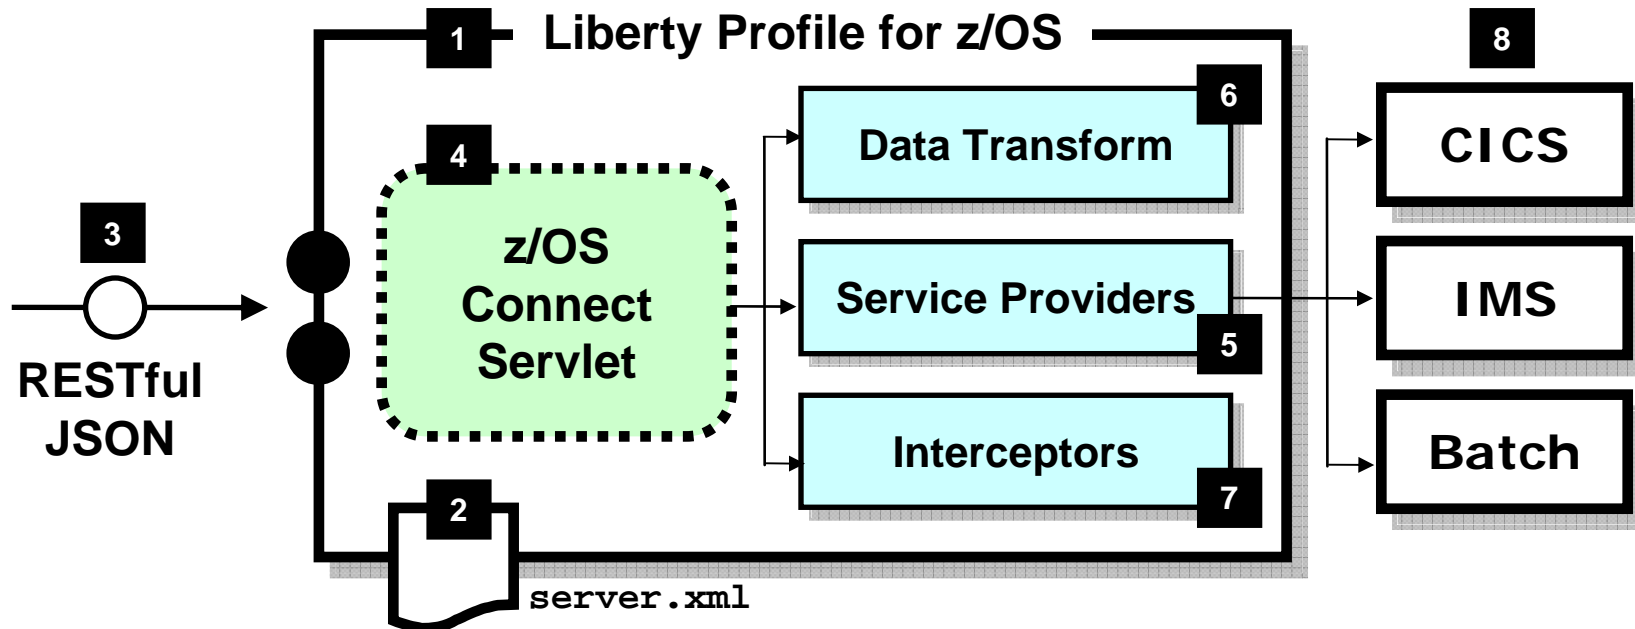- Secure, Scalable
- Agentless
- No central config

- Members "join" collective
- Profiles self-assign to clusters
- Simple config updates
- Easily reversed, reconfigured

z/OS Connect – a brief summary

Liberty Profile for z/OS

1 z/OS Connect is software function that runs in Liberty Profile for z/OS.

2 z/OS Connect is described and configured in the Liberty `server.xml` file

3 z/OS Connect is designed to accept RESTful URIs with JSON data payloads

4 One part of z/OS Connect is a servlet that runs in Liberty Profile z/OS.

5 A 'Service Provider' is software that provides the connectivity to the backend system

6 z/OS Connect provides the ability to transform JSON to the layout required by backend

7 'Interceptors' are callout points where software can be invoked to do things such as SAF authorization and SMF activity recording

8 Initially the backend systems supported will be CICS, IMS and Batch

MOBILE MAINFRAME
APPTHROWDOWN

Will you be our mobile champ?

2014

CICS | IMS | WAS | DB2

Open to existing System z clients

The challenge: Build a proof-of-concept demonstrating mobile enablement of your existing mainframe apps.

Get IBM help to build your mobile PoC

Call us 'Coach': We provide getting started guides and access to IBM zMobile Experts for questions and queries.

Win a week with IBM experts & more

Make it real: Win help from IBM to bring your mobile app to life.

ibm.biz/mmathrowdown

No submission of code required, only screenshots.
Entries must be complete and submitted by **17 Sept 2014**.