Application Level Resource Monitoring of WebSphere z/OS - DB2 JDBC Workloads

Robert Catterall
David Follis
IBM

August 8, 2014
Session Number 15999

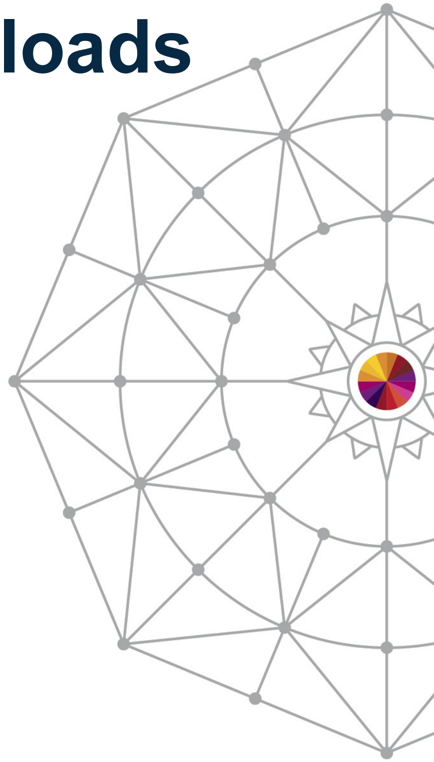

#SHAREorg

Trademarks

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

CICS*
DB2*
GDPS*
Geographically Dispersed Parallel Sysplex
HiperSockets
IBM*
IBM eServer
IBM logo*
IMS
On Demand Business logo

Parallel Sysplex*
RACF*
System z9
WebSphere*
z/OS
zSeries*

* 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.

MIB is a trademark of MIB 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.

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 Overview

#SHAREorg

The CR / SR Structure ... One More Time

It's worth starting with a review of the essential heart of this:

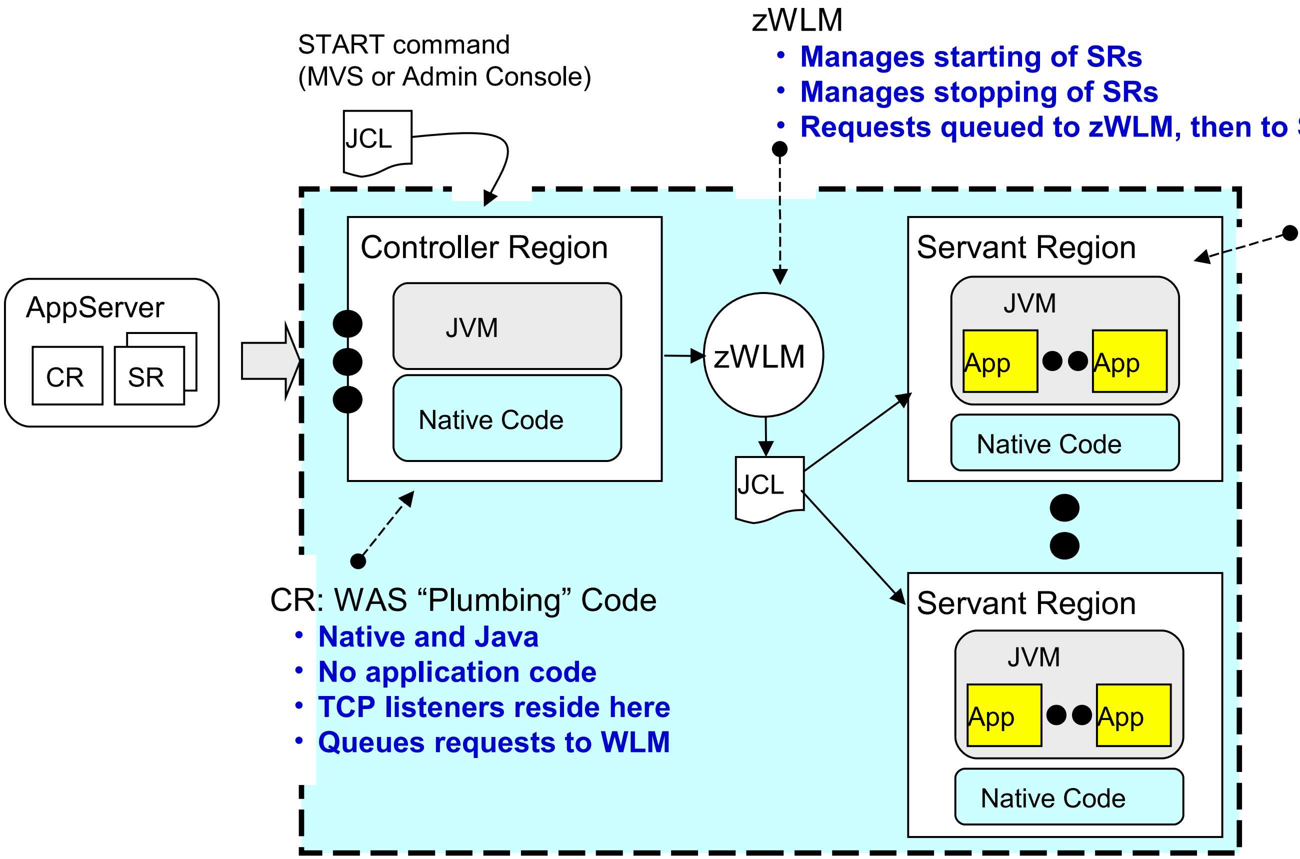

zWLM

- Manages starting of SRs
- Manages stopping of SRs
- Requests queued to zWLM, then to SR

SR: Application Infrastructure

- Maintains app JVM runtime
- May support one or more applications
- Connectivity to data from SR
- Min/Max controllable by admin

CR: WAS “Plumbing” Code

- Native and Java
- No application code
- TCP listeners reside here
- Queues requests to WLM

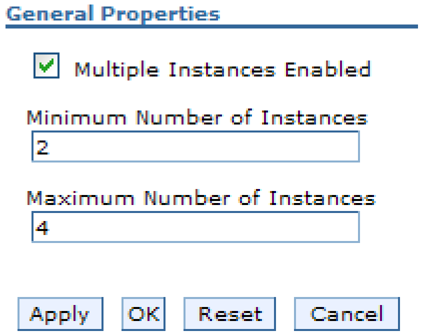

General Properties

☑ Multiple Instances Enabled

Minimum Number of Instances: 2

Maximum Number of Instances: 4

Default: min=1, max=1

Accessing DB2

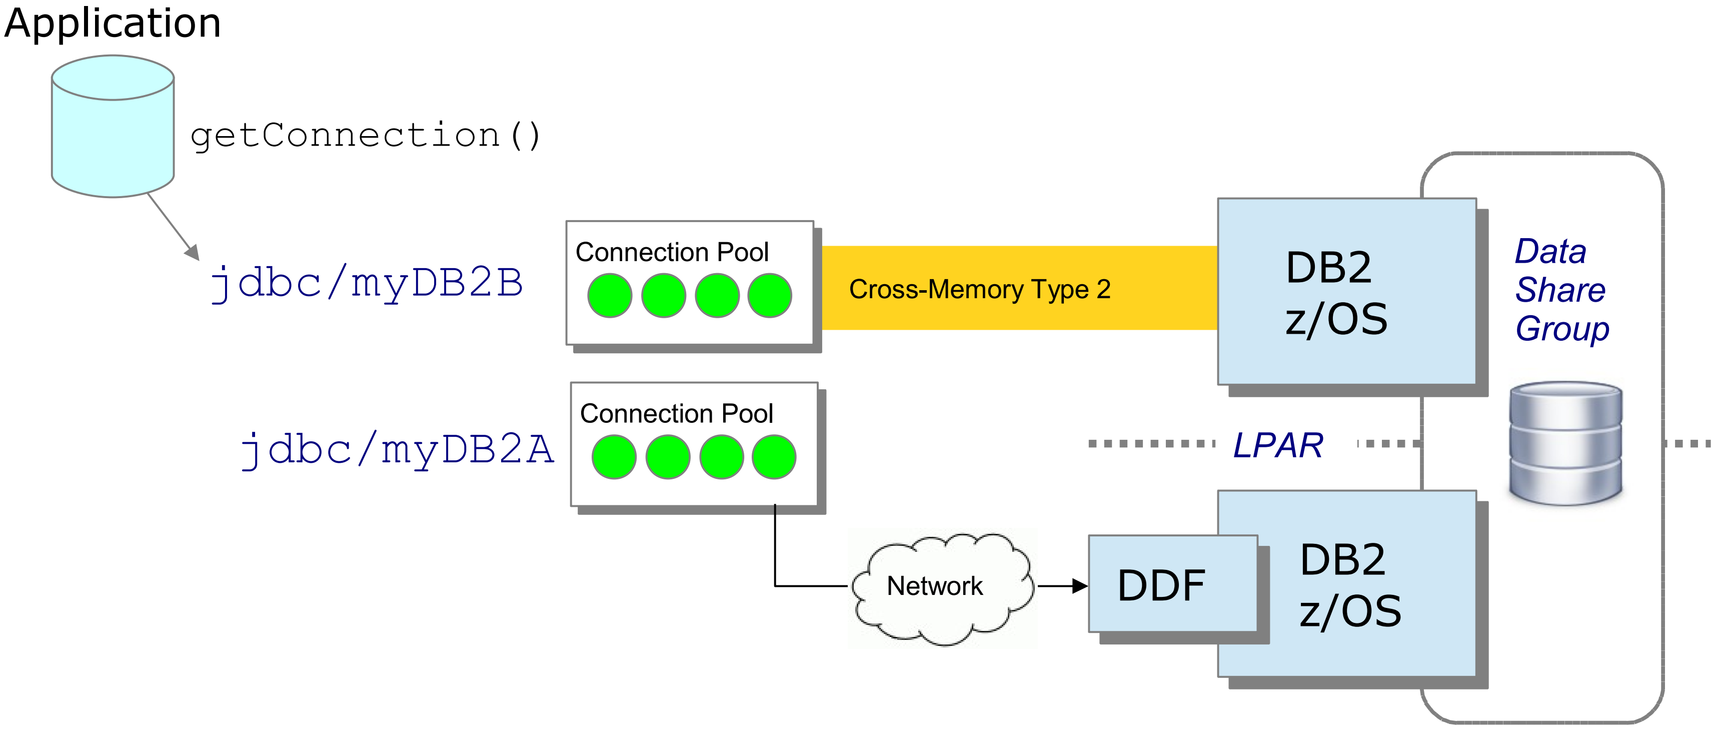

Which JDBC driver for WAS for z/OS?

- When WAS is running on a system other than the one that houses a target DB2 for z/OS subsystem, you have one driver choice for accessing that DB2 server:
  - The type 4 JDBC driver (the one used by Java applications that are network-connected to DB2)
- When WAS and DB2 are running in the same z/OS LPAR, you have two choices for application access to DB2:
  - The type 2 JDBC driver
    - *Local connection to DB2 – uses the DB2 recoverable resources attach facility (aka RRS)*
  - The type 4 JDBC driver
    - *Go into the LPAR's TCP/IP stack, then to DB2 via the Distributed Data Facility (DDF), like any other network-attached application*

How to choose between type 2 and 4 drivers?

- If you have a JDBC choice (as you do when WAS and DB2 are in the same z/OS LPAR), which one should you use?
  - Type 2 driver: more CPU-efficient connection to DB2 (you go from WAS to DB2 – and back – for each SQL statement issued) but zIIP/zAAP eligibility of SQL statement execution is reduced
    - *SQL statements are not Java code, so they are not zAAP-eligible*
    - *SQL statements execute under TCB in WAS address space, so they are not zIIP-eligible*
    - *SQL statement-issuing Java process will "hold on" to specialty engine (zAAP, or zAAP-on-zIIP) for a time after SQL statement execution begins, so you get some specialty engine offload*

Type 2 versus type 4 JDBC driver, continued

- Type 4 driver: more pathlength for DB2 connection (because you have to go through the LPAR’s TCP/IP stack), but SQL statement execution is zIIP-eligible to the tune of about 60%
  - SQL execution is zIIP-eligible because statements execute under enclave SRBs in the DB2 DDF address space
- Keep in mind that the goal is minimization of general-purpose CPU time, not maximization of specialty engine CPU time
  - For applications that issue quick-running SQL statements, type 2 driver might be best choice (more efficient connection to DB2 might trump reduced specialty engine eligibility of SQL)
  - For applications that issue longer-running SQL statements, type 4 driver might be best choice (more zIIP offload for SQL)
  - DB2 monitor accounting report will tell the tale

If you use the type 4 JDBC driver…

- Understand that the application will have to connect to DB2 using an ID and a password
- In the past, people were concerned that someone might use the application's credentials for unauthorized access to DB2
- DB2 roles and trusted contexts (introduced with DB2 9) can be used to greatly mitigate this security exposure
  - Create a DB2 role, and grant privileges needed for successful SQL execution to that role, not to an authorization ID
  - Create a DB2 trusted context to limit role's use
    - *Role's privileges will be usable only by an application that connects to DB2 using a particular ID, <u>and only from a particular IP address</u> (the IP address of the WebSphere Application Server)*

WLM Basics

What is "Workload Management" on z/OS?

It is controlled access to system resources coordinated by a function that keeps watch over all the elements of the system:

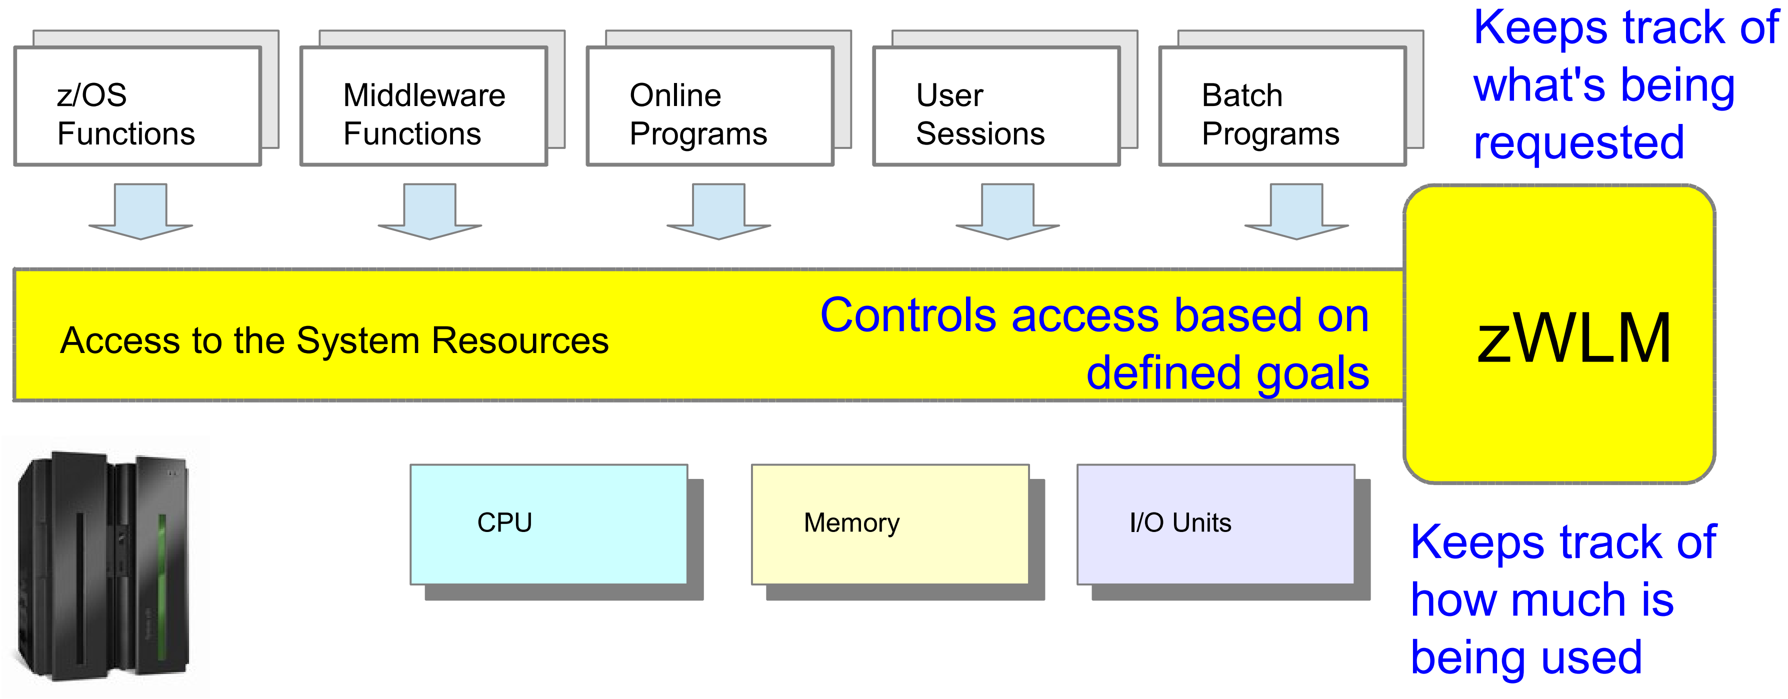

Keeps track of what's being requested

Controls access based on defined goals

Keeps track of how much is being used

There is a tight integration between the System z hardware, the z/OS operating system with WLM having an exclusive view of it all

Key Starting Concepts

To set the stage for the terminology that follows ...

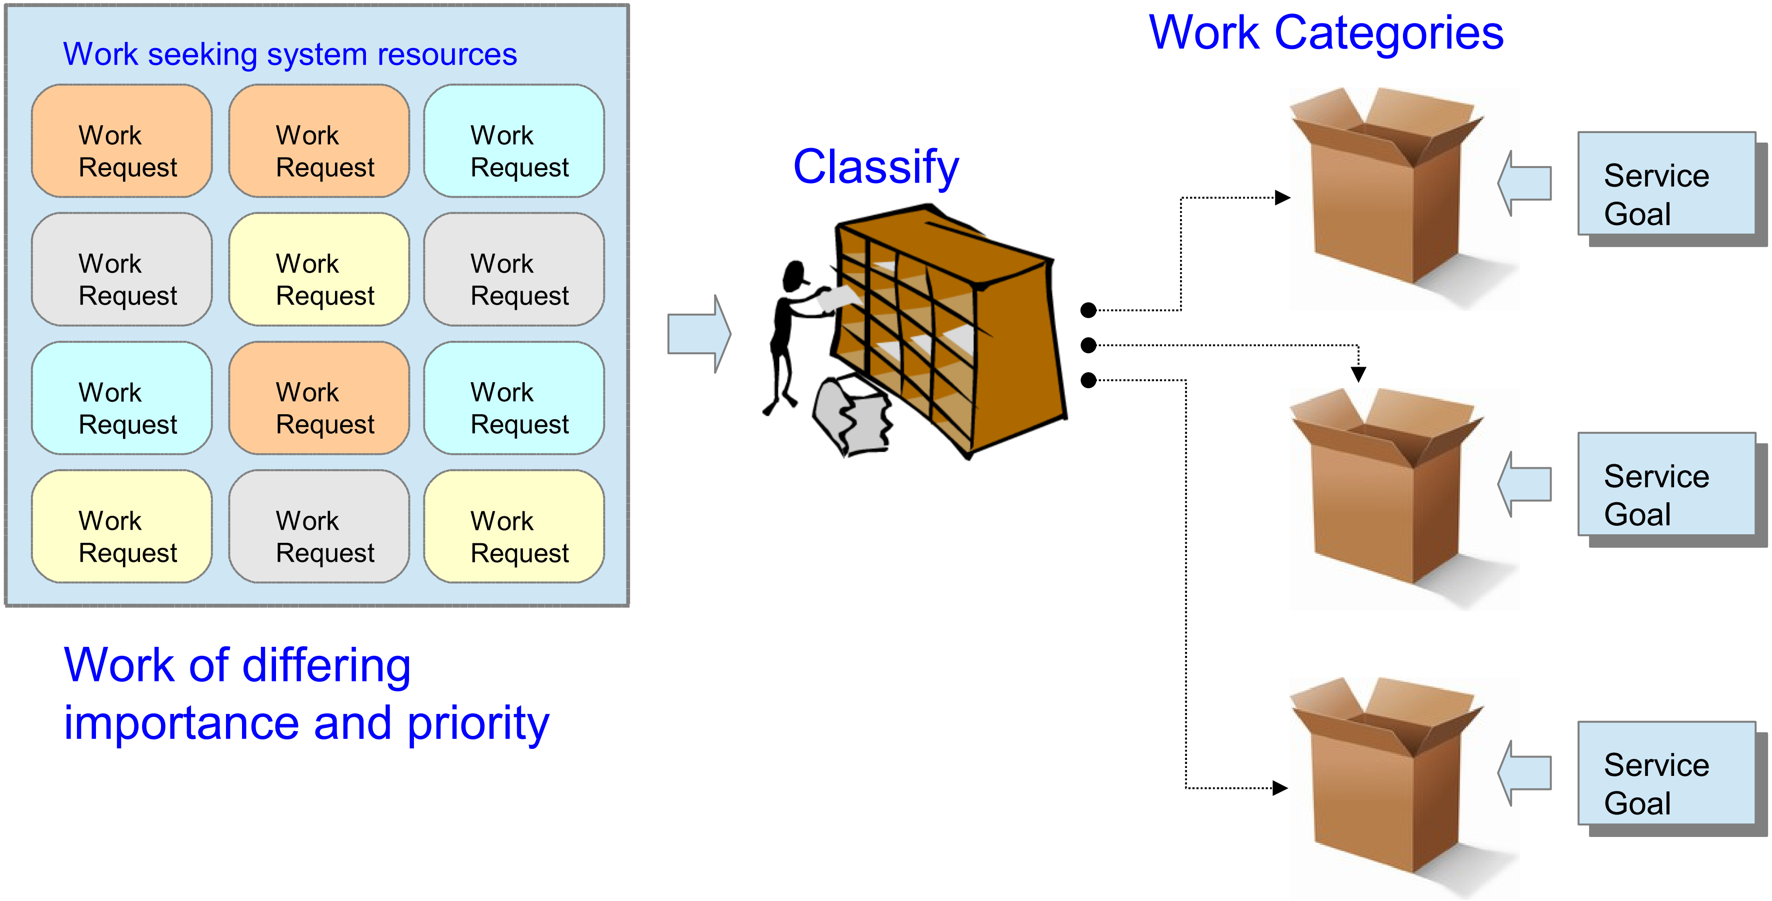

Work seeking system resources

Work Request	Work Request	Work Request
Work Request	Work Request	Work Request
Work Request	Work Request	Work Request
Work Request	Work Request	Work Request

Work of differing importance and priority

Classify

Work Categories

Service Goal

Service Goal

Service Goal

In order for WLM to manage resources to goals, we must get the work organized into categories based on your goals

The WLM Service Class

The "service class" is at the heart of this ... it's the container into which categorized work is placed

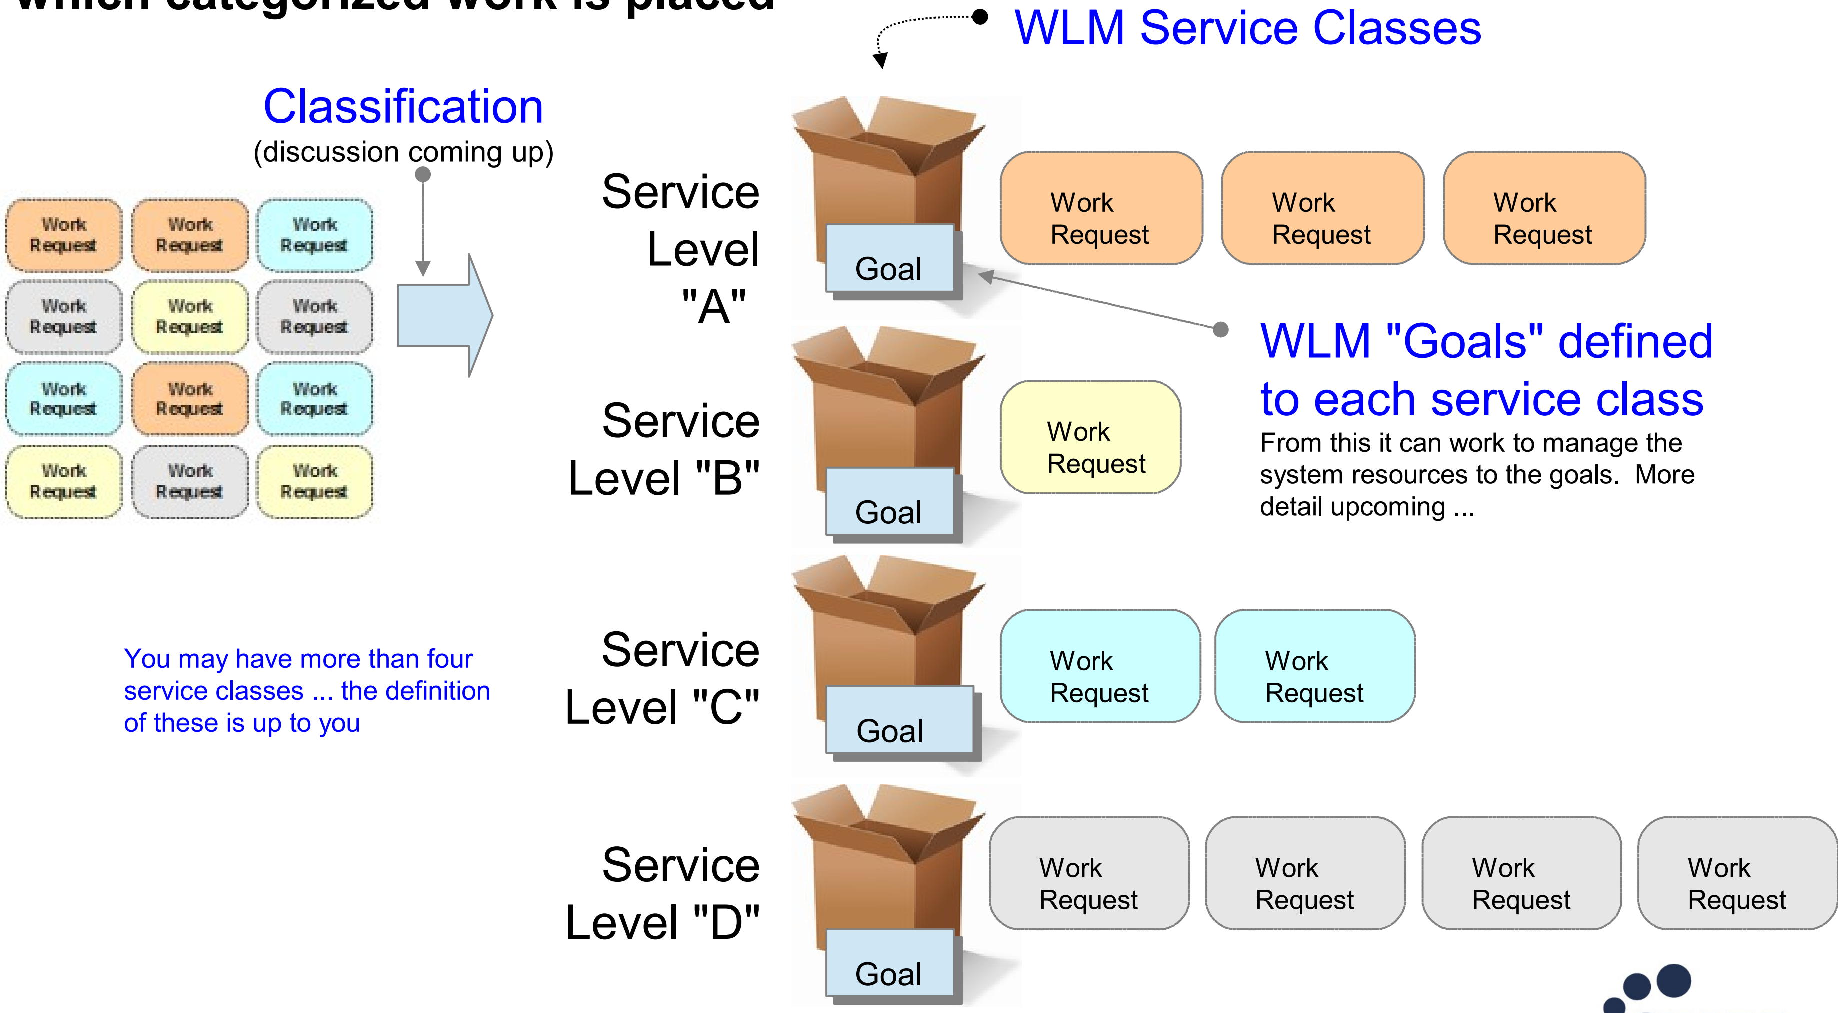

WLM Service Classes

Classification

(discussion coming up)

WLM "Goals" defined to each service class

From this it can work to manage the system resources to the goals. More detail upcoming ...

You may have more than four service classes ... the definition of these is up to you

The WLM Report Class

The "report class" is a variation on the "service class" ... WLM uses it to *report* on activity, but *not to manage* resources

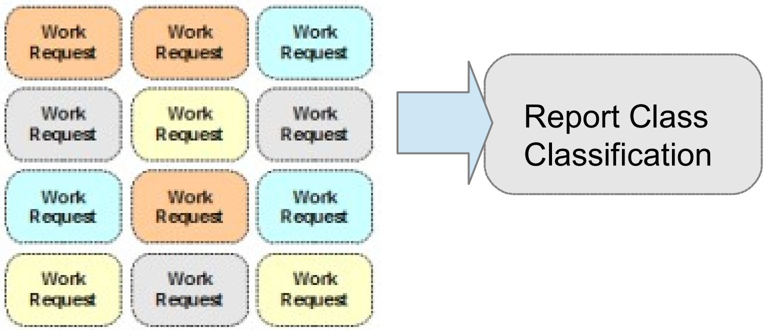

Report Class Classification

Report Class

Ex: "Work related to WAS servers in cell ABCell"

Report Class

Ex: "Work related to CICS region XYZ"

●
●
●

Report Class

Ex: "Work related to transaction DEF"

Provides useful detail on things like CPU usage, zAAP usage and many other system statistics

Generally speaking -- you'll have a handful of service classes and a lot more reporting classes ... based on your needs:

Service Classes -- enough to reasonably categorize work priorities
Reporting Classes -- based on the granularity of your reporting needs

Goals and Importance -- Defined in Service Class

Goals tell WLM what to strive for in terms of service; Importance is used to determine relative importance when resources tight

Goals

Velocity	How fast work should be done without being delayed Number 1 to 99	Started tasks and batch programs
Response Time	Percentage of work completed within a specified period of time Example: 95% within 1 second	Online transactional work
Discretionary	WLM services when other priorities not competing for resources	Work that's okay to push aside if resources are needed

Importance

1 = Most important
2
3
4
5 = Least important

Importance indicates how important it is to you that the service goal be met.

Importance applies only if the service goal is not being met.

Request Flow

#SHAREorg

The WLM "Enclave"

An "enclave" is a way to identify and manage individual pieces of work *within* the many parts of a running z/OS system

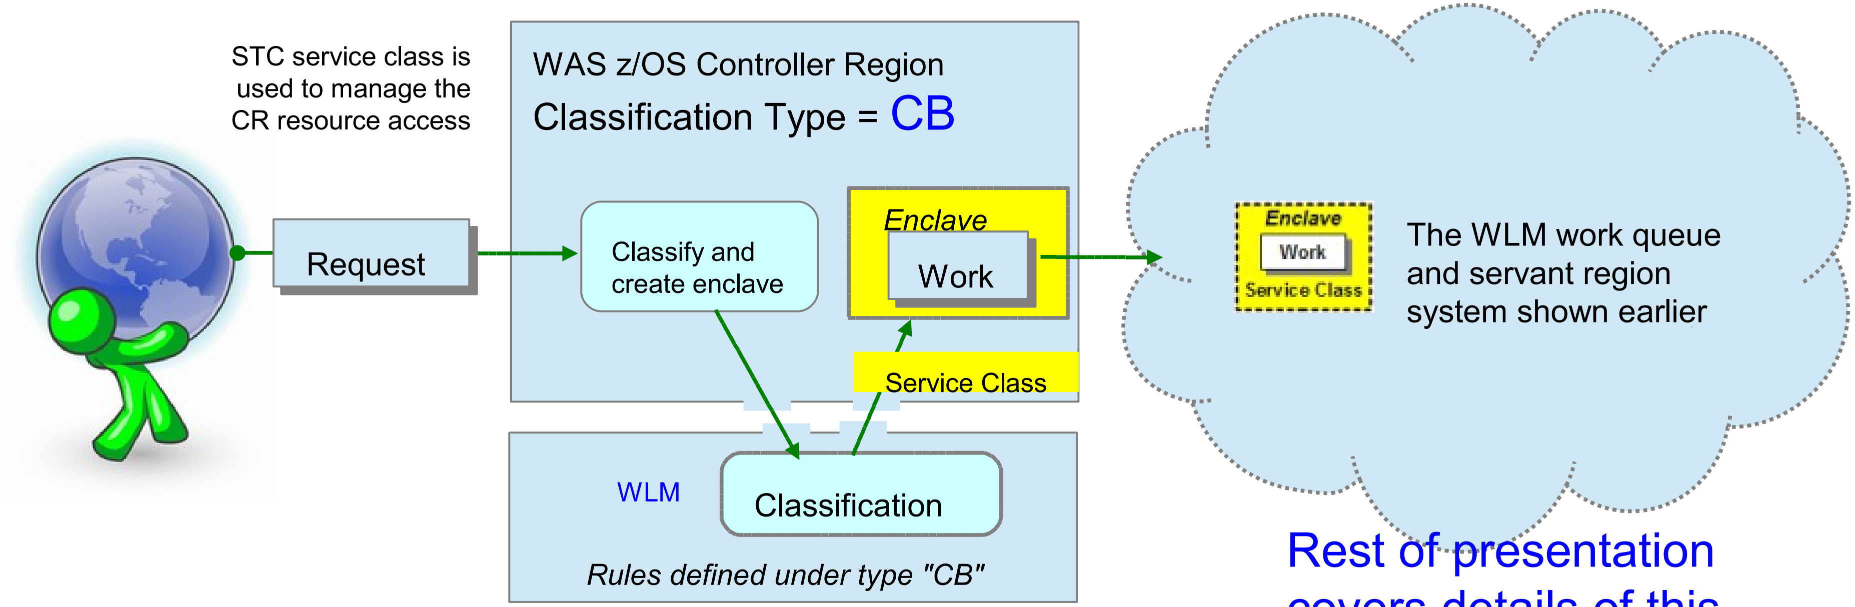
STC service class is used to manage the CR resource access
Request
WAS z/OS Controller Region
Classification Type = CB
Classify and create enclave
Enclave
Work
Service Class
WLM
Classification
Rules defined under type "CB"
Enclave
Work
Service Class
The WLM work queue and servant region system shown earlier

Rest of presentation covers details of this

Key points from this chart

- An "enclave" is simply a way for WLM to understand priorities at a work unit level
- WAS does this automatically ... if you do no other configuration it'll still do this with default values

Assigning a Service Class to the Enclave

This is for the work request ...

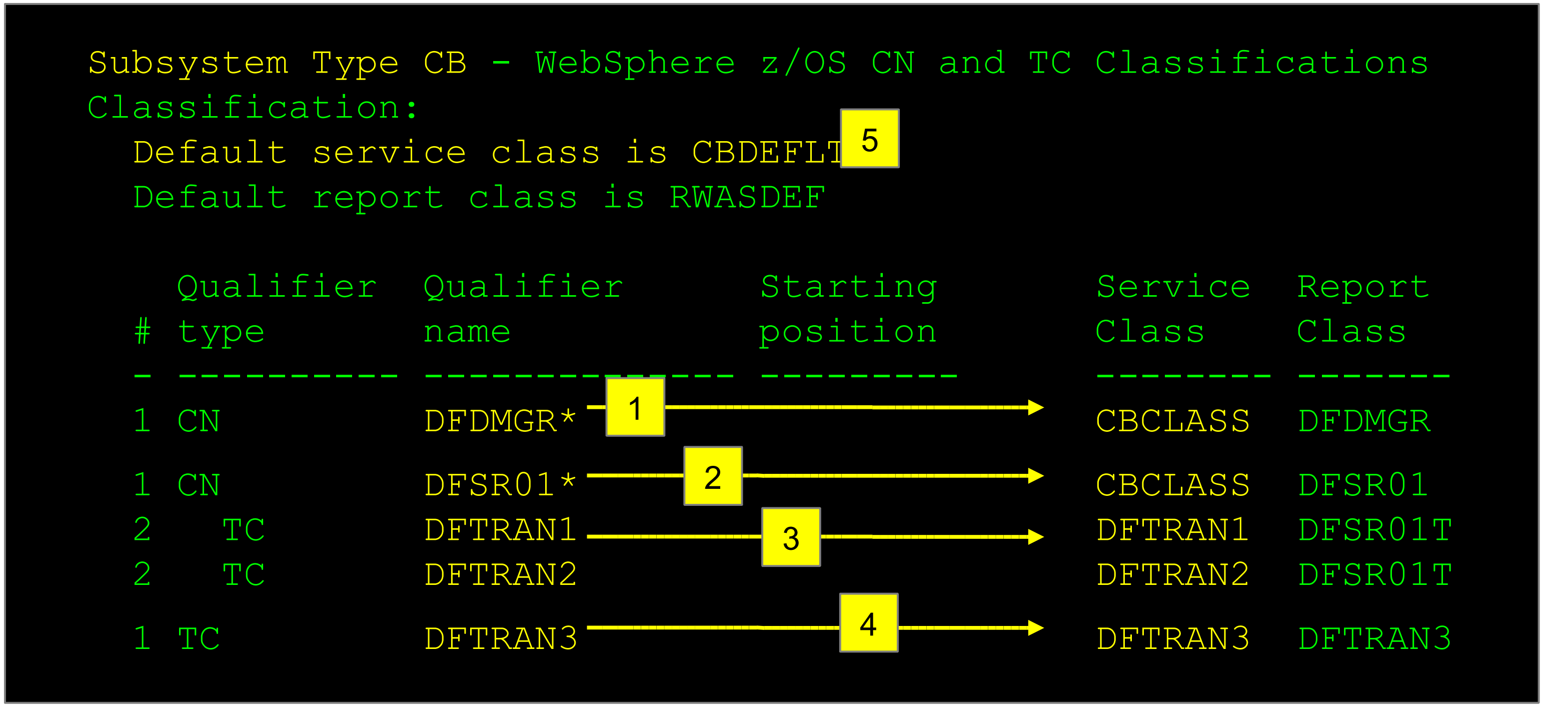

```
Subsystem Type CB - WebSphere z/OS CN and TC Classifications
Classification:
  Default service class is CBDEFLT
  Default report class is RWASDEF

    Qualifier  Qualifier      Starting       Service  Report
  # type       name           position       Class    Class
  - ---------- -------------- ---------      -------- -------
  1 CN         DFDMGR*                       CBCLASS  DFDMGR

  1 CN         DFSR01*                       CBCLASS  DFSR01
  2   TC       DFTRAN1                       DFTRAN1  DFSR01T
  2   TC       DFTRAN2                       DFTRAN2  DFSR01T

  1 TC         DFTRAN3                       DFTRAN3  DFTRAN3
```

Enclaves created in WAS CR are classified by rules in CB subsystem type:

1. CN of `DFDMGR*` matches the Deployment Manager. Work there goes to `CBCLASS`.
2. Work in `DFSR01*` cluster *without* a transaction classification gets `CBCLASS` as well.
3. Work in `DFSR01*` cluster *with* TC of `DFTRAN1` or `DFTRAN2` get service classes as shown
4. Work that matches the TC of `DFTRAN3` *regardless of WAS CN* gets service class `DFTRAN3`
5. Anything that doesn't match any specific rules gets the default service class of `CBDEFLT`

RMF Report with WLM – Example

```
  REPORT BY: POLICY=STANDARD   WORKLOAD=NEWWORK    SERVICE CLASS=WASCLASS   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=2
                                                   CRITICAL     =NONE

  -TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE---   SERVICE TIME  ---APPL %---  --PROMOTED--  ----STORAGE----
  AVG      17.70  ACTUAL                 6  SSCHRT   0.0  IOC         0   CPU 1022.793  CP    341.49  BLK    0.000  AVG        0.00
  MPL      17.70  EXECUTION              6  RESP     0.0  CPU    57020K   SRB    0.000  AAPCP   0.00  ENQ    0.000  TOTAL      0.00
  ENDED   832303  QUEUED                 0  CONN     0.0  MSO         0   RCT    0.000  IIPCP   0.00  CRM    0.000  SHARED     0.00
  END/S  2778.86  R/S AFFIN              0  DISC     0.0  SRB         0   IIT    0.000                LCK    0.000
  #SWAPS       0  INELIGIBLE             0  Q+PEND   0.0  TOT    57020K   HST    0.000  AAP      N/A                -PAGE-IN RATES-
  EXCTD        0  CONVERSION             0  IOSQ     0.0  /SEC   190376   AAP      N/A  IIP      N/A                SINGLE      0.0
  AVG ENC  17.70  STD DEV               21                                IIP      N/A                              BLOCK       0.0
  REM ENC   0.00                                          ABSRPTN    11K                                            SHARED      0.0
  MS ENC    0.00                                          TRX SERV   11K                                            HSP         0.0

GOAL: RESPONSE TIME 000.00.00.250 FOR  80%

           RESPONSE TIME EX   PERF  AVG   --EXEC USING%--  -------------- EXEC DELAYS % -----------  -USING%-  --- DELAY % ---     %
  SYSTEM      ACTUAL%    VEL% INDX ADRSP  CPU AAP IIP I/O  TOT CPU                                    CRY CNT  UNK IDL CRY CNT  QUI

  SP5            100     55.1  0.5  17.8   14 N/A N/A 0.0   11  11                                    0.0 0.0   75 0.0 0.0 0.0  0.0

                                         ----------RESPONSE TIME DISTRIBUTION----------
     ----TIME----    --NUMBER OF TRANSACTIONS--    -------PERCENT-------   0    10   20   30   40   50   60   70   80   90   100
     HH.MM.SS.TTT    CUM TOTAL        IN BUCKET    CUM TOTAL   IN BUCKET   |....|....|....|....|....|....|....|....|....|....|
  <  00.00.00.125           830K             830K          100         100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
  <= 00.00.00.150          830K              560          100         0.1  >
  <= 00.00.00.175          831K              377          100         0.0  >
  <= 00.00.00.200          831K              311          100         0.0  >
  <= 00.00.00.225          831K              223          100         0.0  >
  <= 00.00.00.250          831K              198          100         0.0  >
  <= 00.00.00.275          831K              162          100         0.0  >
  <= 00.00.00.300          832K              113          100         0.0  >
  <= 00.00.00.325          832K              108          100         0.0  >
  <= 00.00.00.350          832K               85          100         0.0  >
  <= 00.00.00.375          832K               76          100         0.0  >
  <= 00.00.00.500          832K              210          100         0.0  >
  <= 00.00.01.000          832K              215          100         0.0  >
  >  00.00.01.000          832K               46          100         0.0  >
```

Reporting CPU Usage

Where CPU is reported depends on whether or not it's an enclave thread

CPU for enclaves attributed to the Controller -- it created the enclave. This true despite fact the enclave is dispatched and run on a *servant thread*

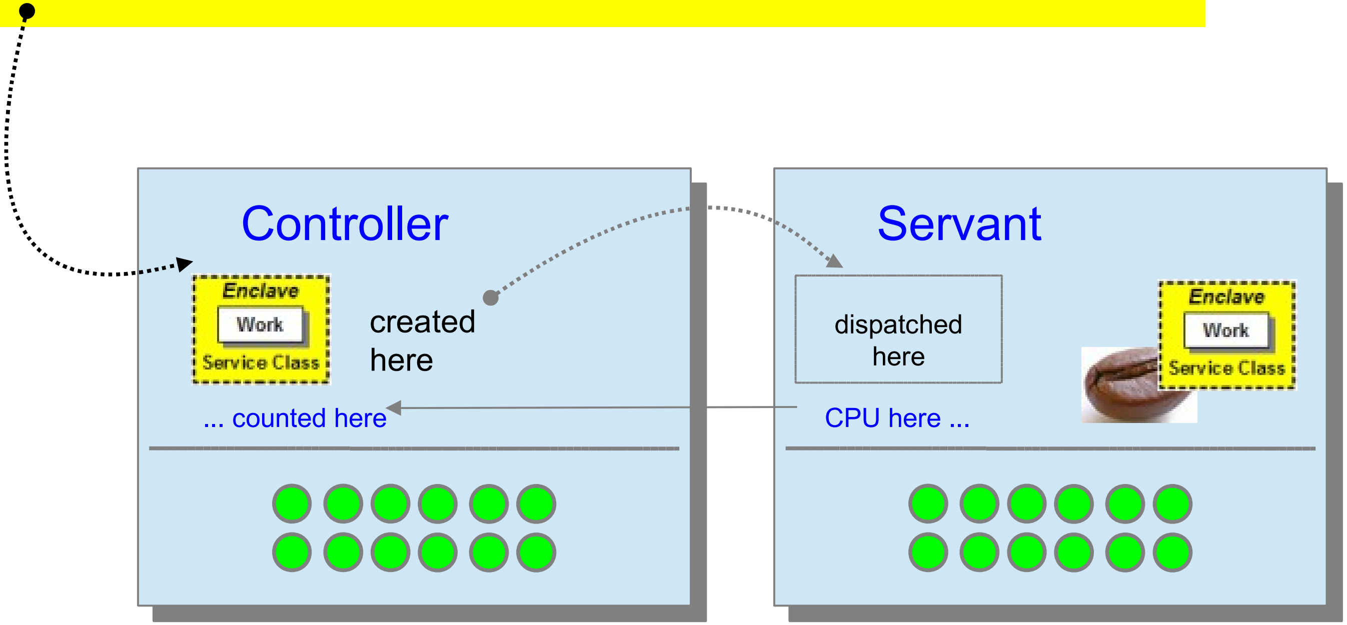

CPU for non-enclave threads used by CR is attributed to the CR region

CPU for non-enclave threads used by SR is attributed to the SR region

What’s in the SMF 120-9?

- One record written for everything dispatched from the controller to the servant

- Includes HTTP, IIOP, MDBs (even from the CRA), Mbeans, and other stuff

- Includes:

  - Who ran it (security section)
  - What it was (URI, EJB AMC names)
  - When it ran (various timestamps)
  - Where it ran (server names etc)
  - Why? User data section
  - How much resource it used (CPU, zAAP..)

- Other 120 subtypes exist, ignore 1-8, the 120-10 is interesting

- WP101342 – Overview of SMF 120-9

An Example: Response Time Summary

- Summarize by each URI we saw:
  - Number of requests
  - Average Response time
  - Average time on the queue
  - Average time in dispatch
  - Average CPU time
  - And average bytes received and sent in response

Requests	AvgResponse	AvgQueue	AvgDisp	AvgCPU	AvgBytesRcvd	AvgBytesSent	URI
10627	352	11	330	84	675	187248	/some/http/request
5484	354	8	341	84	663	210074	/some/other/http/request
4019	268	0	266	63	711	62883	/yet/another/http/request
3989	96	12	78	17	582	86293	/really/just/one/more/http/request
…	…	…	…	…	…	…	…
60578	311	7	296	70	679	152760	Overall

And now we call DB2

#SHAREorg

Monitoring WAS application's DB2 activity

- Primary tool: DB2 monitor
  - Specifically, an accounting long (or detail) report, or an online display of DB2 accounting (i.e., thread-related) information
- From a DB2 monitoring perspective, how do people distinguish one WAS-based application from another?
  - Usually done by way of an application's DB2 authorization ID (other application identifiers will be discussed later in this session)
    - *Often, different WAS-based applications have different DB2 authorization IDs (e.g., master data management application has ID ABCD1234, ERP application has ID WXYZ6789)*
    - *In generating an accounting long report, you can tell your DB2 monitor to aggregate data at the DB2 primary authorization ID level – that will give you detailed data for each different auth ID included in the report*
    - *Similarly, in viewing online display of DB2 thread information, look for the application auth IDs of interest*

More on DB2 monitoring for WAS applications

- You probably want to have DB2 accounting trace classes 1, 2, and 3 active at all times
  - Very low overhead, extremely useful information
  - Records usually written to SMF for historical reporting purposes
  - DB2 can optionally compress its SMF trace records (60-80% compression with less than 1% CPU overhead)

Key DB2 performance metrics for WAS apps

- In DB2 monitor accounting long reports (or online displays):
  - Average class 2 (i.e., in-DB2) CPU time
    - *This is typically average per transaction*
    - *It's the CPU cost of SQL statement execution, broken out in 2 fields: general-purpose CPU time and specialty engine CPU time (latter can be particularly substantial when type 4 JDBC driver is used)*
  - Average class 3 synchronous read wait time
    - *That's the average time that transactions spend in wait mode, waiting for on-demand, single-page reads from disk subsystem*
    - *It's usually your #1 leverage point for reducing per-tran elapsed time*
    - *Best way to bring this down: bigger DB2 buffer pools (more and more organizations have buffer pool configurations that are tens of GB in size for a production DB2 subsystem)*
  - Percentage of dynamic SQL statements found in cache
    - *Want this to be > 90% (if not, enlarge statement cache)*

More key DB2 performance metrics…

- Also from DB2 monitor accounting long reports (or online display):
  - Class 1 CPU time (total CPU time from first SQL statement to end of transaction) for work that was zIIP-eligible but ran instead on general-purpose engines
    - *Want this to be really low, especially if using type 4 JDBC driver*
- In DB2 monitor statistics long reports (or online displays):
  - (If you use the type 4 JDBC driver) High-water mark for active database access threads (DBATs) used for packages bound with RELEASE(DEALLOCATE)
    - *These threads are also known as high-performance DBATs – introduced with DB2 10 for z/OS*
    - *By driving thread re-use, they provide a CPU efficiency boost for applications that are network-connected to DB2*

More Applications….

#SHAREorg

Classification XML File

The XML file identifies requests ...

```
<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>
```

InfoCenter `rrun_wlm_tclass_dtd`

Assigning a Service Class to the Enclave

This is for the work request ...

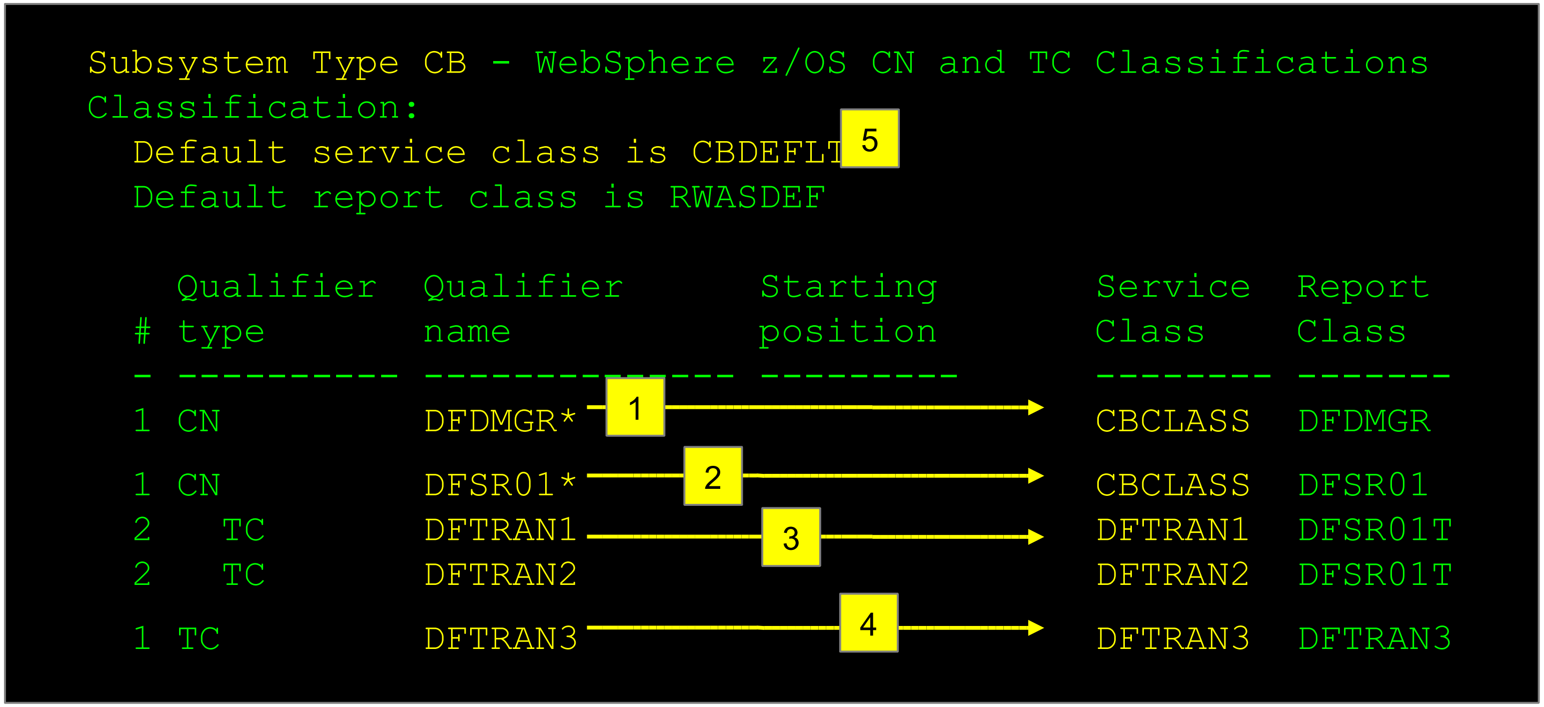

```
Subsystem Type CB - WebSphere z/OS CN and TC Classifications
Classification:
  Default service class is CBDEFLT
  Default report class is RWASDEF

    Qualifier   Qualifier      Starting       Service   Report
  # type        name           position       Class     Class
  - ----------  -------------- ---------      --------- --------
  1 CN          DFDMGR*                       CBCLASS   DFDMGR

  1 CN          DFSR01*                       CBCLASS   DFSR01
  2   TC        DFTRAN1                       DFTRAN1   DFSR01T
  2   TC        DFTRAN2                       DFTRAN2   DFSR01T

  1 TC          DFTRAN3                       DFTRAN3   DFTRAN3
```

Enclaves created in WAS CR are classified by rules in CB subsystem type:

1. CN of DFDMGR* matches the Deployment Manager. Work there goes to CBCLASS.
2. Work in DFSR01* cluster *without* a transaction classification gets CBCLASS as well.
3. Work in DFSR01* cluster *with* TC of DFTRAN1 or DFTRAN2 get service classes as shown
4. Work that matches the TC of DFTRAN3 *regardless of WAS CN* gets service class DFTRAN3
5. Anything that doesn't match any specific rules gets the default service class of CBDEFLT

More Applications….calling DB2

#SHAREorg

If WAS app's DB2 ID isn't granular enough…

- It's not unusual for several WAS-based applications to use the same authorization ID to connect to a DB2 subsystem
  - Some organizations go this route in order to simplify security management
- When that's the case, separating DB2 accounting trace information by primary authorization ID won't give you an application-specific view of DB2 activity
- What about alternative identifiers that are included in DB2 accounting trace records?
  - DB2 plan name probably won't help
    - *Will be DISTSERV for all applications using type 4 JDBC driver*
    - *Will default to ?RRSAF for all applications using type 2 JDBC driver (though different plan names could be set for such applications)*

More on identifiers in DB2 accounting records

- Requesting location won’t distinguish between multiple applications running in one instance of WAS
  - Running multiple applications in one WAS instance is particularly commonplace in WAS for z/OS environments – here’s why:
    - *z/OS LPARs often have a very large amount of processing capacity*
    - *The sophisticated workload management capabilities of z/OS allow multiple, non-homogeneous applications to coexist in one LPAR*
- Main DB2 package probably not granular enough
  - WAS-based applications typically issue SQL statements in the form of JDBC calls – in that case, the main DB2 package for all applications will be one associated with the IBM JDBC driver

More DB2 accounting trace identifiers

- Transaction name and end user ID will typically be too granular

So, what should you do?

A popular choice: workstation name

- Easily set for an application, through several means:
  - Via the WAS administration console GUI (as an extended property of an application's data source)
  - Via the IBM Data Server Driver for JDBC (the driver provides a JAR file that contains the DB2Connection class, which supports the Java API setDB2ClientWorkstation)
  - Via application code, for JDBC 4.0 and later (you'd use the Java API setClientInfo)
    - *This approach is recommended over the JDBC driver-implemented approach, because the latter was deprecated with JDBC 4.0*
- Once workstation name is set for an application, you can direct your DB2 monitor to generate an accounting report with data aggregated at the workstation name level

And there's your application-specific view of DB2 activity

Some additional items

- Lots of information on setting workstation name and other client information can be found in this IBM redbook:
  - *DB2 for z/OS and WebSphere Integration for Enterprise Java Applications*

    http://www.redbooks.ibm.com/abstracts/sg248074.html?Open
  - In particular, see sections 5.5 and 8.2
- Identifiers provided by Java client information APIs can also be very useful for workload classification in a WLM policy

Robert Catterall
rfcatter@us.ibm.com

David Follis
follis@us.ibm.com

MOBILE MAINFRAME APP THROWDOWN

Will you be our mobile champ?

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**.