Information Management

IBM

Data Management

DB2 9.7: Workload Management

Paul Bird
Senior Technical Staff Member, DB2 for LUW
pbird@ca.ibm.com

Agenda

- **Background**
- **What Is New In DB2 9.7 GA?**
- **What Is New In DB2 9.7 FP1?**

IBM

Data Management

Background

Release Objectives for Workload Management

- **Enable simpler implementations**
  - Support a basic “Priority Aging” concept
- **Complete the foundation started in DB2 9.5**
  - Enhancements to workloads, thresholds, and service classes
- **Officially deprecate Query Patroller and Governor**
  - Will be removed in next major release

Monitoring in DB2 9.7

- **The DB2 9.5 Workload Manager introduced both new monitoring and new control capabilities**
- **DB2 9.7 builds on this foundation to introduce dramatic new monitoring capabilities including:**
  - Light-weight SQL access to enhanced in-memory metrics for different aspects of executing work, data objects, and the package cache
  - New “time spent” metrics to better understand where executing work is spending its time
  - And much more!!
- **See a DB2 9.7 Monitoring presentation for more information!**

Tooling

- **The tooling provided for DB2 9.5 has been merged into the Optim (Data Studio) tooling strategy under Curt Cotner**
  - Design Studio and DB2 Performance Expert products are now part of that organization
  - Providing WLM functionality is part of their mandate
- **Tooling development teams will be delivering capabilities on a different and more rapid schedule than the DB2 server**
  - Expect important tooling deliveries in 1H and 2H of 2009 for WLM and monitoring
- **No tooling enhancements discussed as part of this presentation**
  - Website: www.ibm.com/software/data/studio

DB2 WLM Best Practices

- **Published on DB2 Best Practices webpage in October 2008:**
  - **http://www.ibm.com/developerworks/data/bestpractices/workloadmanagement**
- **Reflects IBM (development and field) and customer experiences with new capabilities as of 2008**
  - Assumes you have followed the other DB2 best practices first!
- **Further updates planned this year both in content and to cover any new function or changes**
  - We need your input and feedback to make it more meaningful!
  - Please send me any comments, corrections, or suggestions on:
    - The contents of the current best practices document
    - Your experiences and guidance on any monitoring or control aspects of WLM

Information Management

IBM

Data Management

What Is New In DB2 9.7 GA?

What Is New In DB2 9.7 GA?

- **Enabling Simpler Implementations**
- **Integration with Linux WLM**
- **DB2 Workload Enhancements**
- **DB2 Threshold Enhancements**
- **DB2 Service Class Enhancements**
- **WLM Table Function Enhancements**
- **WLM Event Monitor Enhancements**

Enabling Simpler Implementations

- **Many customers do not need the more complex aspects of DB2 Workload Management but do want to implement a "bias" in the way database work is executed**
  - They don't have explicit workload management conflicts that need to be resolved
  - They want to favour one type of work over another in order to improve the overall processing for that type of work but have no specific performance objectives
- **DB2 9.7 introduces support for the "Priority Aging" concept**
  - Intended for use by those customers with simple environments

“Priority Aging” Concept

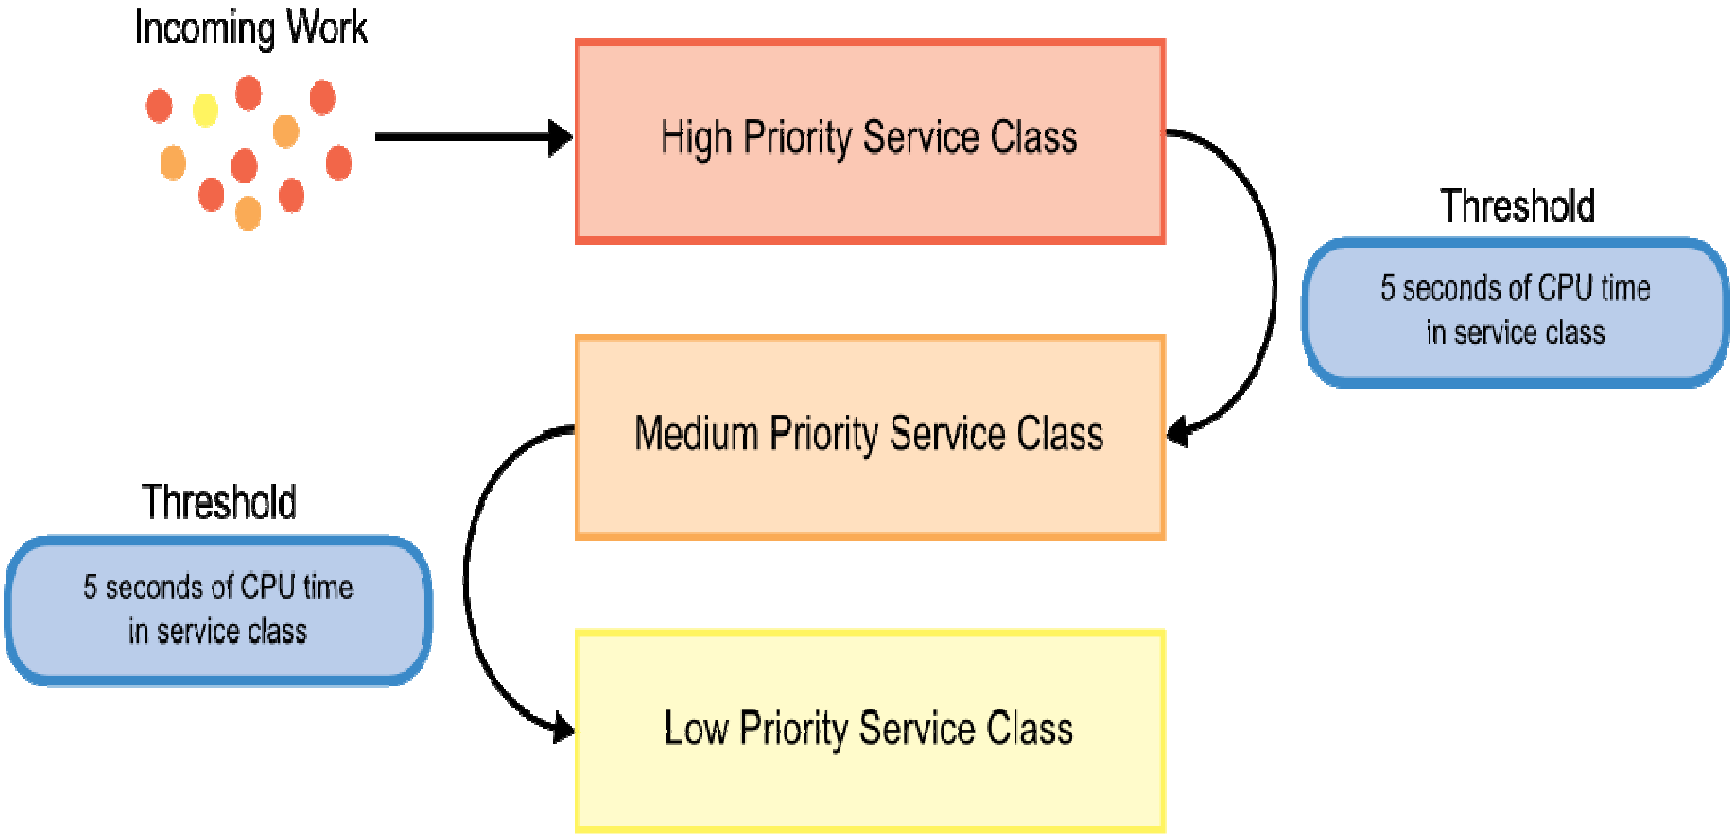

“Priority Aging” in DB2 9.7

- **DB2 9.7 provides an automated way to change the priority given to any SQL statement**
  - DB2 will automatically shift work between service subclasses based on criteria established by the administrator for:
    - CPU usage while in the subclass
    - Rows read while in the subclass
- **Sample scripts provided to allow for easy implementation of two different approaches:**
  - All statements are treated equally at the start
  - All statements are pre-sorted by estimated cost at the start

All Statements Start in High Priority

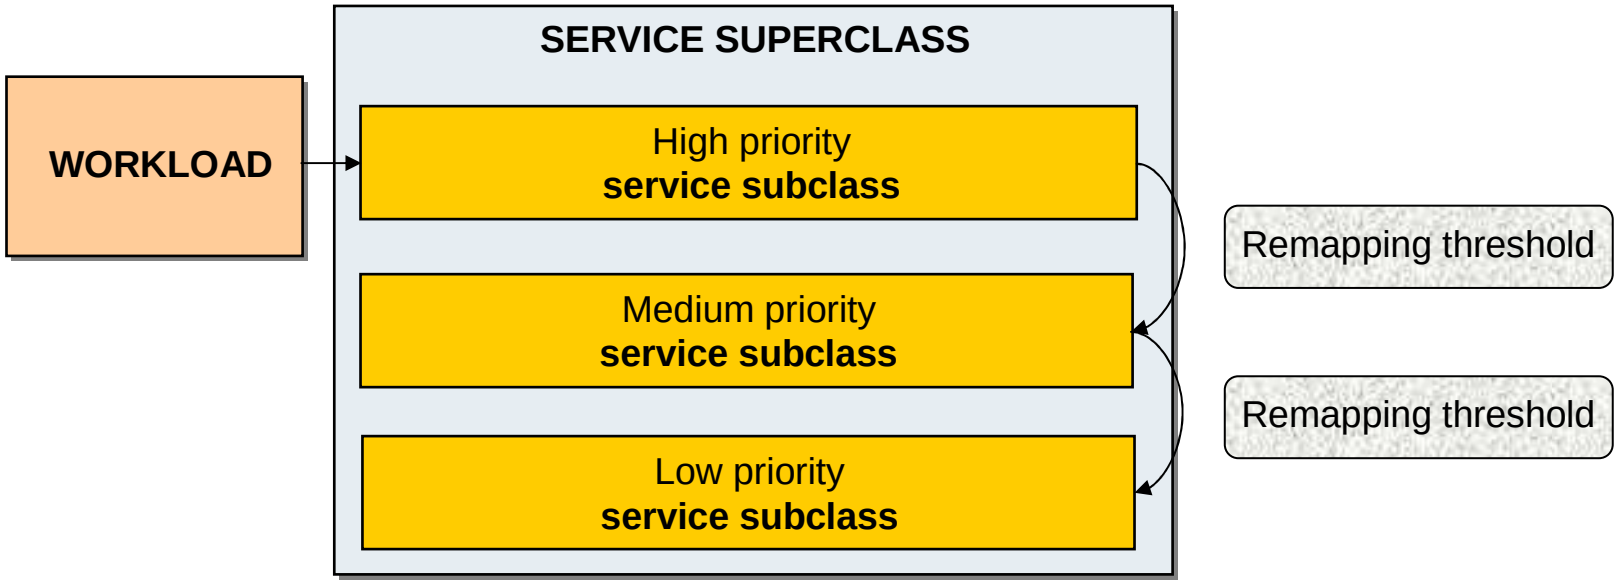

Statements Pre-Sorted by Estimated Cost

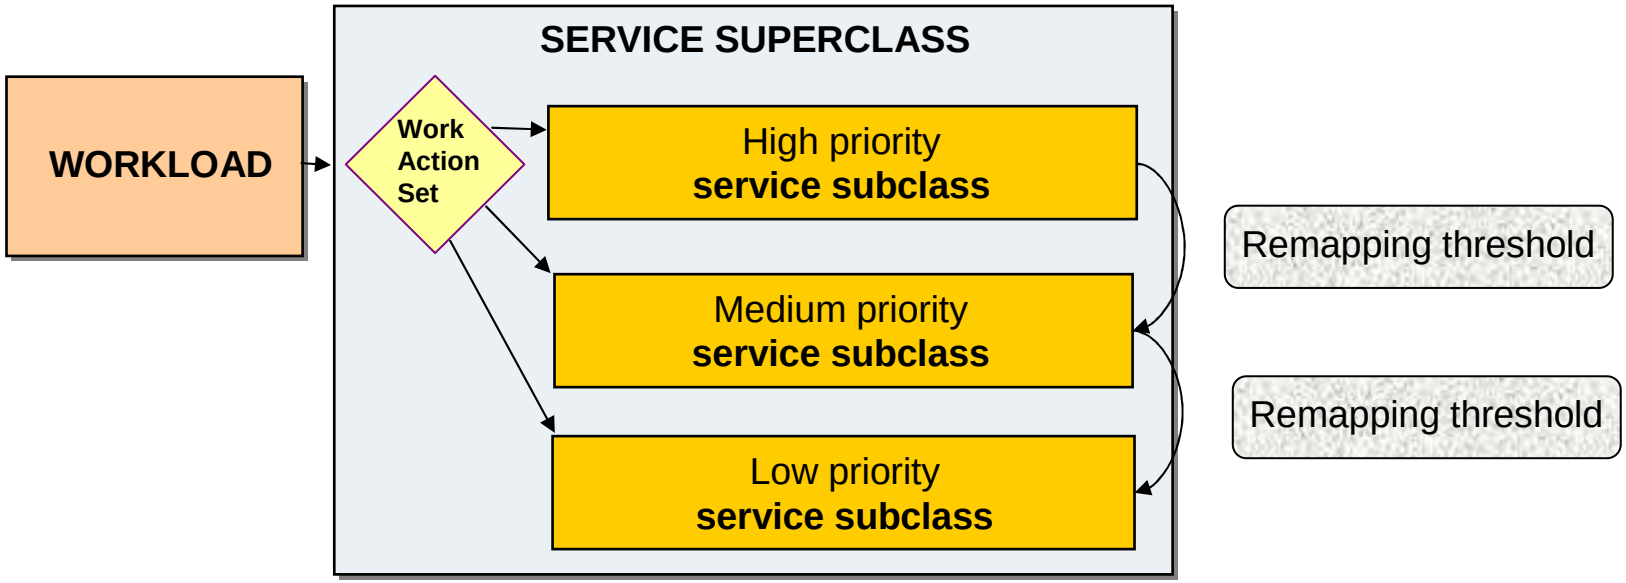

Priority Aging Example

- **Assume a company with 2 main applications:**
  - SALES issues short SQL queries and is a high priority to the business
  - MANUFACTURING issues longer SQL statements and is a lower priority within the business
- **When this workload is run against DB2 9.7 using the default service classes and priority settings, the following results are seen:**

Workload	Average Lifetime (ms)	Standard Deviation (ms)
SALES	**2380**	**4192**
MANUFACTURING	**47318**	**48567**

Modified Configuration

- **Introduce a new DB2 service superclass, WLM_TIERS, with 2 subclasses:**
  - WLM_SHORT
  - WLM_LONG
- **Introduce CPU TIME IN SERVICE CLASS activity threshold to REMAP to lower priority subclass**
  - WLM_SHORT to WLM_LONG after 1CPU seconds
- **Adjust agent and prefetch priorities as follows:**

Service Class Name	Agent Priority	Prefetch Priority
SYSDEFAULTSYSTEMCLASS	**-20**	**HIGH**
WLM_SHORT	**-19**	**HIGH**
WLM_LONG	**19**	**LOW**

Modified Configuration: Results

- **When the workload is run against DB2 9.7 using this modified configuration, the following results are seen:**
  - Average lifetime and standard deviation for SALES queries are greatly reduced

Workload	Average Lifetime (ms)	Standard Deviation (ms)
SALES	**687 (from 2380)**	**1825 (from 4193)**
MANUFACTURING	**96271 (from 47318)**	**81718 (from 48567)**

Integration with Linux WLM

- **Available with Linux kernel 2.6.26 or above**
- **Identical to AIX WLM integration from the DB2 perspective**
  - An optional capability to associate each DB2 service class with a corresponding Linux WLM service class to allow use of Linux WLM to control allocation of system resources
  - DB2 will dynamically associate all agents working in a DB2 service class with the identified Linux WLM service class
- **Example:**

```
ALTER SERVICE CLASS “Marketing”
    OUTBOUND CORRELATOR  “_DB2_Marketing”;
```

Linux WLM definition

Linux WLM Integration

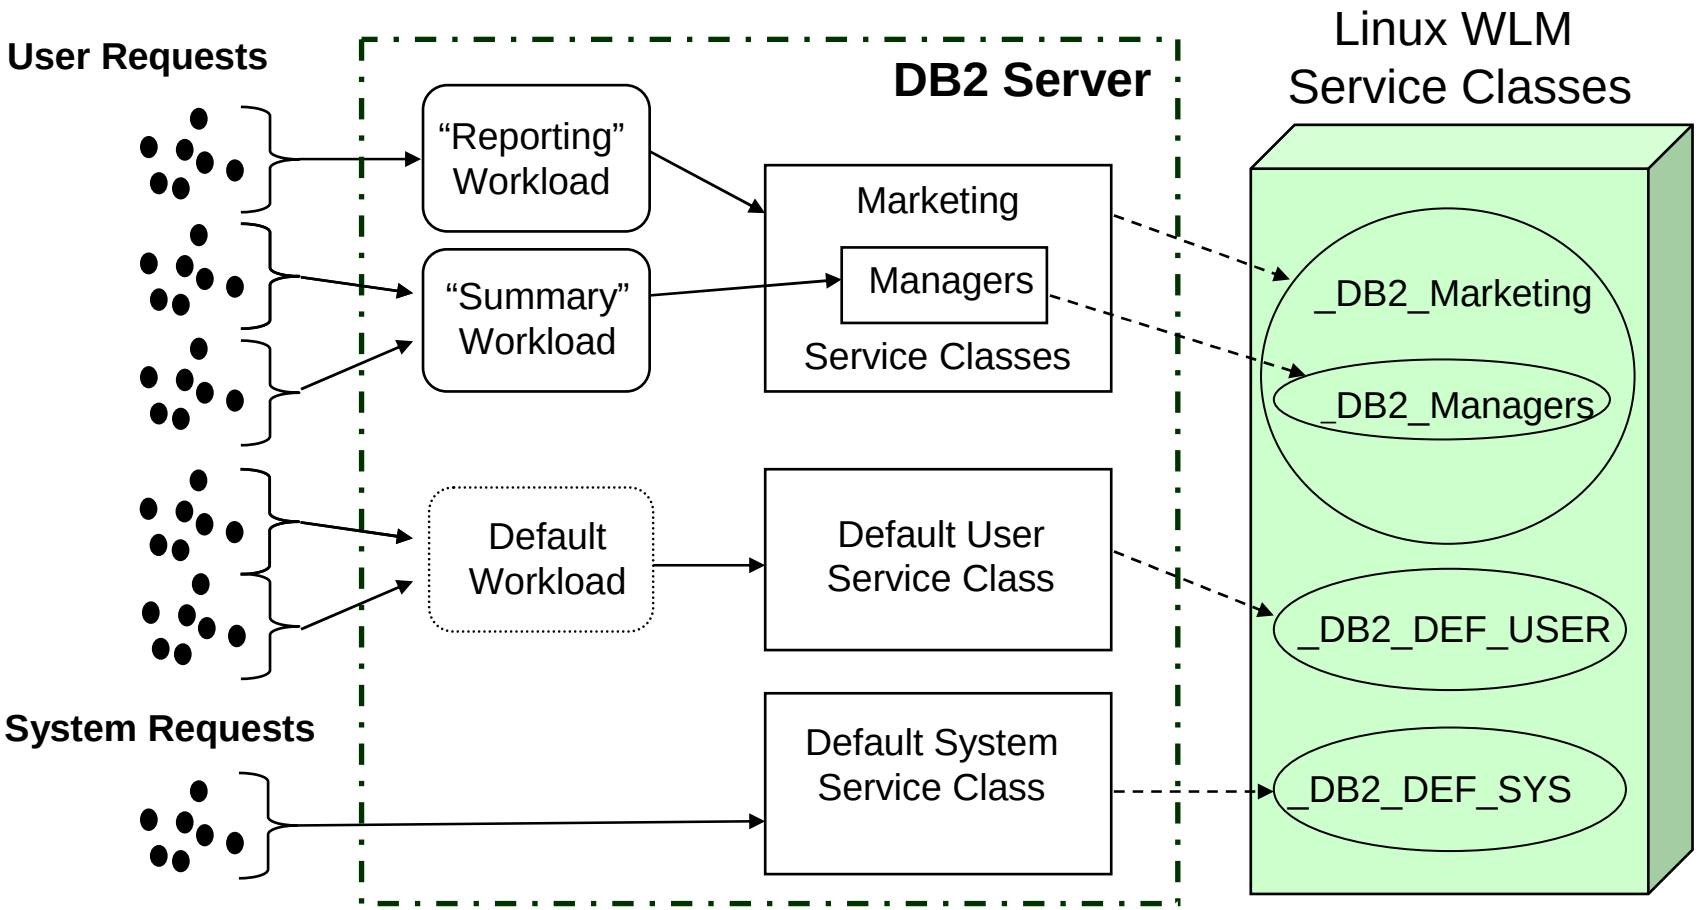

DB2 Workload Enhancements

- **Allow Activity Thresholds to be assigned at the workload level**
- **Support for use of wildcard character in some workload definition attributes**
- **Support for IP address as a workload definition attribute**
- **Support the collection of aggregate activity data at workload level**

Activity Thresholds on DB2 Workloads

- **Able to specify unique activity thresholds for each workload definition**
  - Thresholds defined at the workload level take precedence over service class thresholds
- **You can now have different limits for different applications even within the same service class!**
  - No need to isolate applications in separate service classes just to apply different thresholds
- **Affects:**
  - ACTIVITYTOTALTIME
  - ESTIMATEDSQLCOST
  - SQLROWSRETURNED
  - SQLTEMPSPACE
  - CPUTIME
  - SQLROWSREAD

Wildcards in DB2 Workload Definitions

- **Makes defining a workload for a large number of similar values much easier**
  - Supports a “Best Practice”: use a standardized format for client information values!
- **The following attributes can use wildcards:**
  - APPLNAME
  - CLIENT_USERID
  - CLIENT_APPLNAME
  - CLIENT_WRKSTNNAME
  - CLIENT_ACCTNG

```
CREATE WORKLOAD “work1” CLIENT_USERID(‘db2user*’);
```

New Workload Definition Attribute: Address

- **Allows you to specify the connection IP address or domain information as an attribute of a DB2 Workload definition**

Secure domain syntax

```
CREATE WORKLOAD “IPADDR1” ADDRESS(‘mydomain.ibm.com’)
```

IPv4 syntax

```
CREATE WORKLOAD “IPADDR2” ADDRESS(‘9.26.53.111’)
```

IPv6 long syntax

```
CREATE WORKLOAD “IPADDR3” ADDRESS(‘2002:91a:519:13:204:acff:fe57:6135’)
```

IPv6 short syntax

```
CREATE WORKLOAD “IPADDR4” ADDRESS(‘fe80::202:55ff:fe9a:6eee’)
```

Aggregate Activity Data for a DB2 Workload

- **Aggregate activity data can now be accumulated at the workload level**
  - This is in addition to the existing capability in service sub-classes and work action sets
  - Reflects aggregates across all connections mapping to that workload definition
- **You can now track the distribution of key information over activities for an individual workload definition**
  - No need to create a separate service sub-class to do this!

DB2 Threshold Enhancements

- **Reduced the granularity for time thresholds down to 1 minute**
  - ACTIVITYTOTALTIME activity threshold
  - CONNECTIDLETIME threshold
- **Refinement to the CONCURRENTDBCOORDACTIVITIES threshold**
- **New thresholds**
  - CPUTIME
  - SQLROWSREAD
  - AGGSQLTEMPSPACE

CONCURRENTDBCOORDACTIVITIES Refinement

- **The CONCURRENTDBCOORDACTIVITIES threshold is a very useful and powerful tool**
  - Establishes and maintains a concurrency limit for any activities within its domain
  - Concurrency thresholds have proven their value as a control mechanism many times over in Query Patroller
- **To allow this threshold to be used in even more scenarios, we have made the following changes in DB2 9.7:**
  1. This threshold will ignore any CALL statement and only will control SQL issued from within the invoked routine
  2. This threshold will still control DML statements but will ignore any SQL issued by routines invoked by that DML statement (e.g. a UDF)
- **Best Practice:**
  - Always have an Activity Total Time threshold in place for any activity subject to a concurrency threshold!

New Activity Thresholds

- **CPUTIME**
  - Controls the maximum amount of processor time that an activity can use on a particular database partition
  - Unlike CPUTIMEINSC, tracks CPU usage for entire life of activity across service classes
- **SQLROWSREAD**
  - Controls the maximum number of rows that an activity can read on a particular database partition
  - Unlike SQLROWSREADINSC tracks number of rows for entire life of activity across service classes
- **More weapons in the hunt for rogue queries!**

New Aggregate Threshold

- **AGGSQLTEMPSPACE**
  - Controls the maximum amount of system temporary table space that can be used by all activities in a service subclass

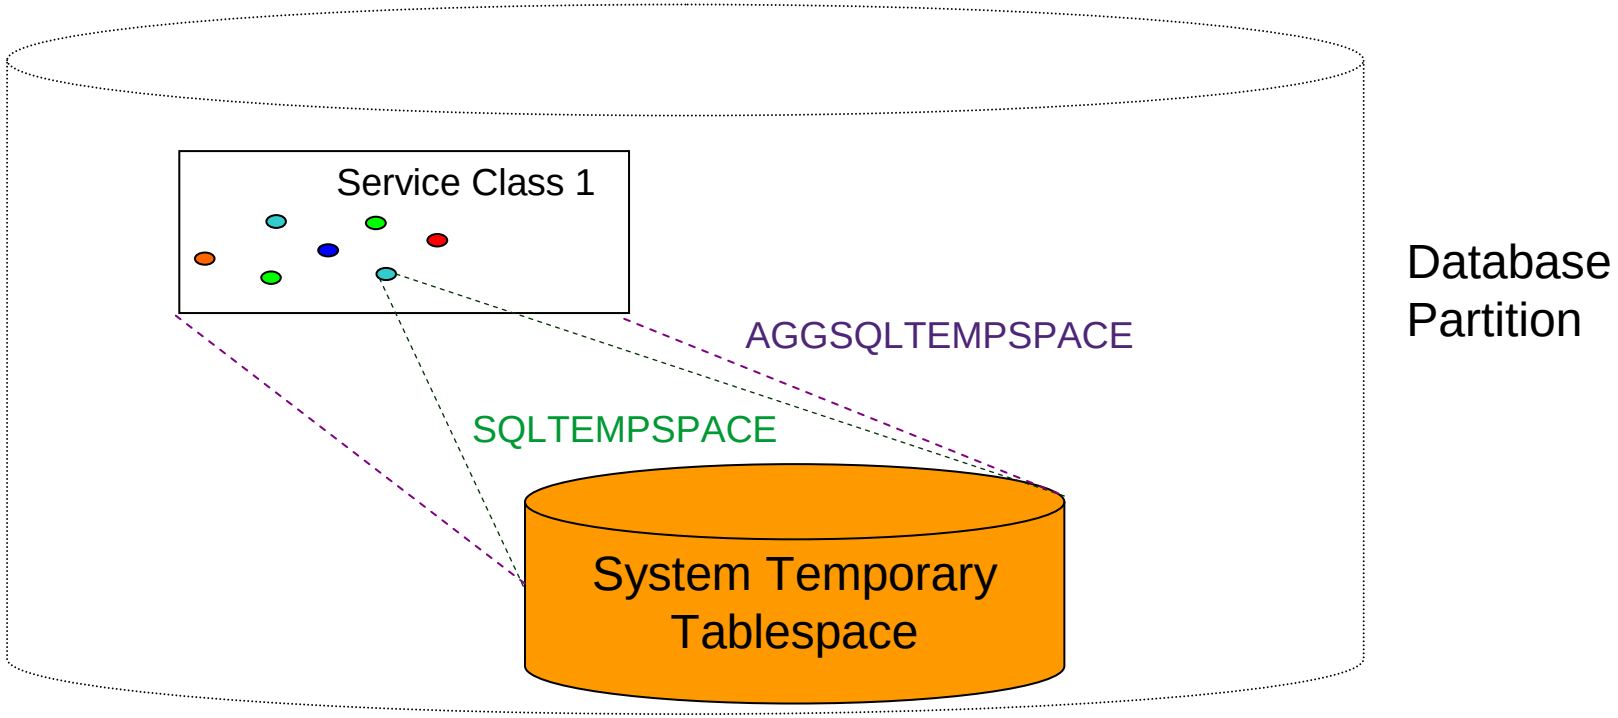

DB2 Service Class Enhancements

- **New Service Class attribute: Buffer Pool I/O priority**
  - Possible values: HIGH, MEDIUM, LOW
- **Bias victim selection in Buffer Pool by assigning priority to pages visited by activities executing in a service class**
- **Reduces likelihood of high priority pages being selected as victim by low priority work**
  - Useful in specific scenarios where competing workloads of different priorities cannot be separated into different bufferpools

WLM Table Function Enhancements

- **A number of the WLM table functions introduced in DB2 9.5 have been enhanced in DB2 9.7 to contain more information**
  - See documentation for more information
  - Some fields are marked as reserved and will be filled in with FP1
- **Changed functions:**
  - WLM_GET_WORKLOAD_STATS_V97
  - WLM_GET_SERVICE_CLASS_WORKLOAD_OCCURRENCES_V97
  - WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97
  - WLM_GET_SERVICE_CLASS_AGENTS_V97
  - WLM_GET_SERVICE_SUBCLASS_STATS

WLM Event Monitor Enhancements

- **Can now have more than one WLM event monitor active at any one time**
  - Affects activity, statistics, and threshold event monitors
  - Allows for easier maintenance by establishing new event monitor and disabling old one for maintenance
- **WLM Activity event monitor move to new fast event monitor infrastructure**
  - Greatly reduces impact on processing agent overhead
  - Improves system throughput to reduce system impact under high volumes of capture
  - Still uses relational table output

Information Management

IBM

Data Management

What Is New In DB2 9.7 FP1?

Disclaimer

It is IBM's current intent to make these technologies available in the next release. However this intent is subject to change and withdrawal, and does not represent a commitment or obligation on the part of IBM.

What Is New In DB2 9.7 FP1?

- **Change to WLM License Coverage**
- **Migration Script for Query Patroller to WLM**
- **Synchronizing WLM Statistics Collection**
- **DB2 Workload Enhancements**
- **DB2 Threshold Enhancements**

Change to WLM License Coverage

- **CREATE WORKLOAD will no longer require the WLM license**
  - Will enable more granular application identification and monitoring
- **Still under WLM License:**
  - CREATE SERVICE CLASS
  - CREATE THRESHOLD
  - CREATE WORK ACTION SET

Migration Script for Query Patroller to WLM

- **New script will be provided to aid in the conversion from Query Patroller to DB2 Workload Management**
  - Generates the DDL statements needed to establish control mechanisms that closely emulate those implemented in a current Query Patroller installation
  - The objective is to migrate from a Query Patroller configuration over to DB2 WLM configuration using the same (relatively) known, stable configuration which can then be modified at leisure
- **Provided for customers who want to move to the new capabilities but who:**
  - Don’t want to migrate in a piece-meal fashion
  - Want to convert quickly but don’t want to risk losing control of the system while they determine how best to use the new capabilities

Synchronizing WLM Statistics Collection

- **DB2 9.5 introduced the ability to automatically collect workload management statistics to event monitor**
  - Collection interval controlled using the wlm_collect_int database configuration parameter
- **Automatic WLM statistics collection will now be synchronized to a fixed base time**
  - New implementation has intervals starting from a fixed base of Sunday 00:00
    - E.g. Interval is 1 hour, catalog partition starts at 3:35 PM
      - Old implementation would collect next interval at 4:35 PM
      - New implementation would collect next interval at 4:00 PM

DB2 Workload Enhancements

- **Introduce the ability to assign a work action set to a workload definition**
  - Will apply to all connections mapping to that workload definition
  - Can be used in same way as a database level work action set:
    - Apply DB2 Thresholds
    - COUNT ACTIVITY
    - PREVENT EXECUTION
    - COLLECT ACTIVITY DATA
  - Can also specify COLLECT AGGREGATE ACTIVITY DATA to gather distribution information for a work class at the workload level

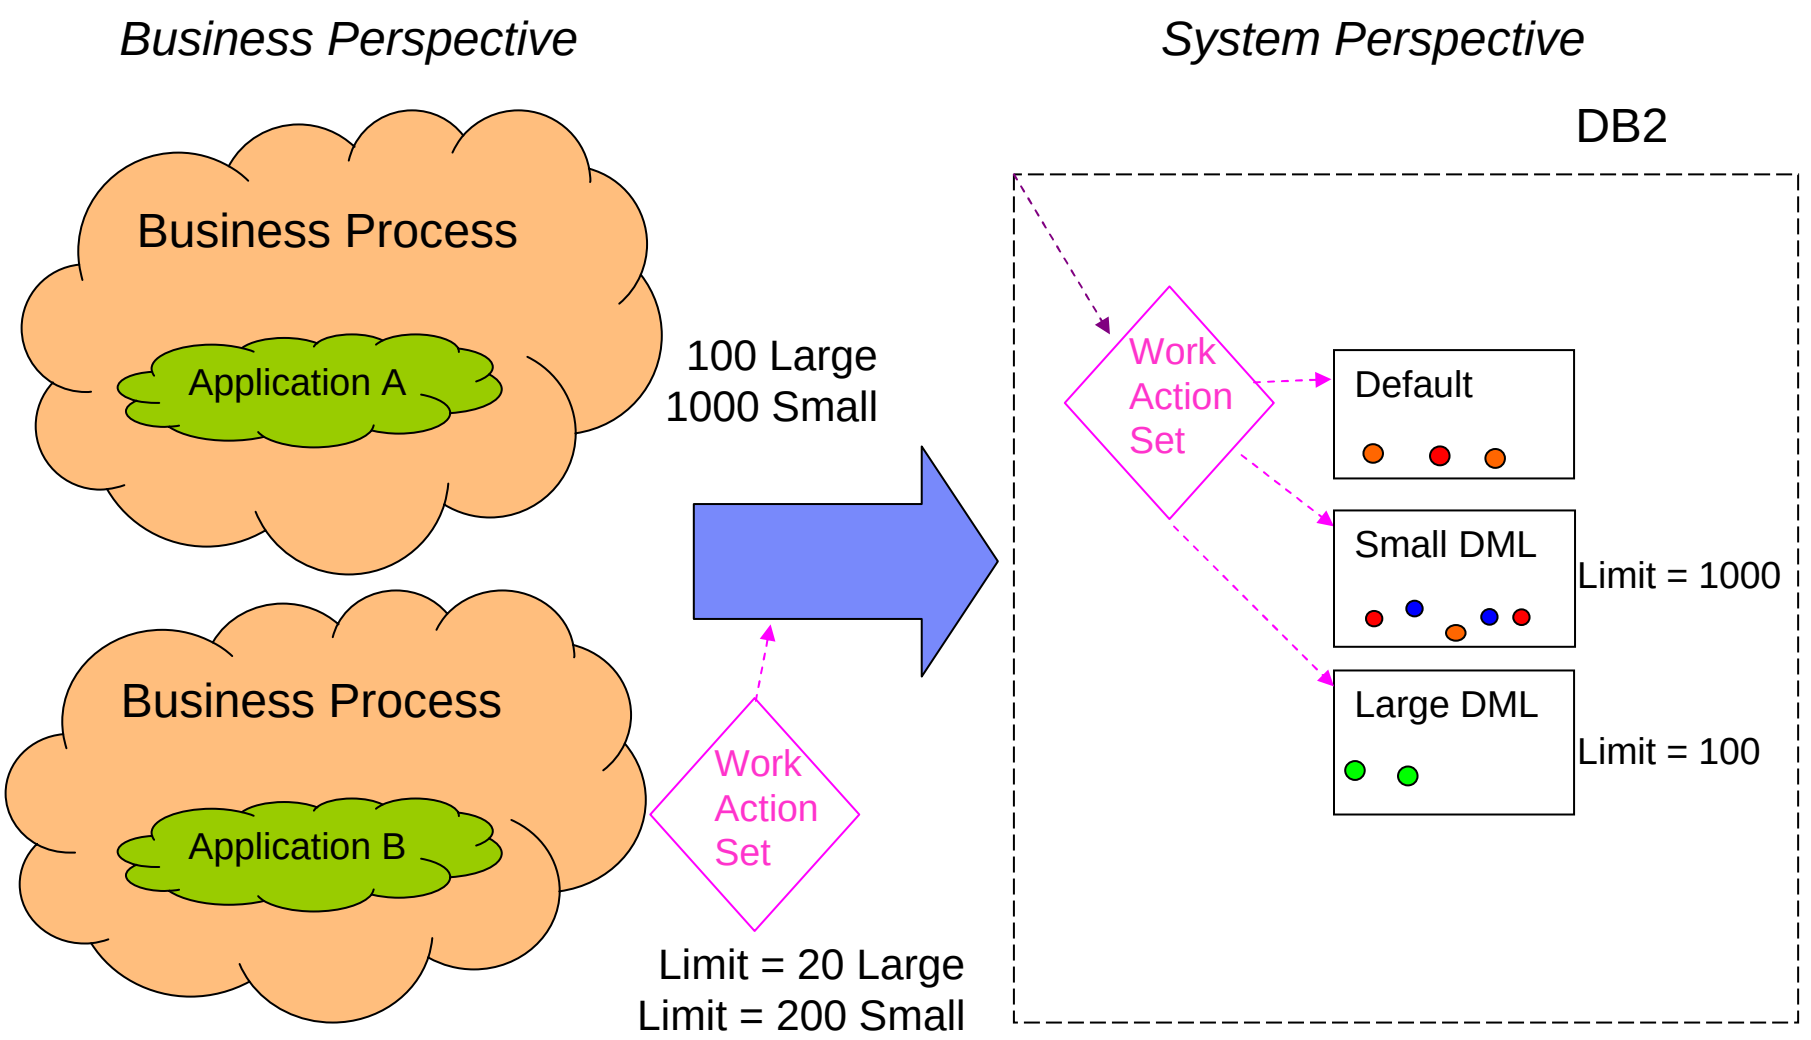

DB2 Threshold Enhancements

- **New UOW Time Threshold**
  - Provides a unit of work time threshold that will control the amount of time a unit of work can run before the threshold is exceeded.
  - When time exceeded, the offending UOW can be rolled back or the connection can be forced
  - Will allow existing DB2 Governor scripts that use the UOW rule to be migrated

IBM

Data Management

Summary

Highlights

- **Simpler WLM implementations are now possible with "Priority Aging" support**
- **Linux WLM integration allows for leveraging of native OS WLM capabilities on AIX or Linux**
- **Major enhancements to workload capabilities:**
  - Provides the ability to vary monitoring and control at the application level which is a more natural focus point
  - Allows for simpler implementations of DB2 service classes
- **Functionality added to allow for official deprecation of Query Patroller and DB2 Governor**
- **WLM license change enables very granular identification and monitoring capabilities for all customers**

Data Management

Questions?

Compare Query Patroller and DB2 WLM

Query Patroller	DB2 Workload Manager
acts as a 'gate keeper': once work is admitted, it is free to execute as it desires	acts as a 'hall monitor': it ensures that work goes to its correct place and follows the relevant rules during execution
can show current state, SQL and ApplHandle, but no in-flight information about work until execution has completed	can show the current state of the application, SQL, and even agents at any time during execution
is aware only of the coordinator perspective for submitted work	is aware of the work across all database partitions
does not provide any mechanism to explicitly control the resources used for execution	provides mechanisms to control and influence resources used during execution
details on all managed activities written to control tables (disk)	nothing gets written to disk unless requested by a user-created event monitor
monitoring is fine-grained only (individual activities are collected)	monitoring can be fined-grained (individual activities) or coarse-grained (aggregate statistics)

IBM

Data Management

Supplemental Information

DB2 WLM Resources

- **DB2 9.5 documentation:**
  - http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp
- **DB2 Best Practices:**
  - http://www.ibm.com/developerworks/data/bestpractices
- **WLM FAQ in DB2 9.5 documentation:**
  - http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/topic/com.ibm.db2.luw.admin.wlm.doc/doc/c0052604.html
- **WLM Hands On Tutorial in DB2 9.5 documentation (partially) and on IBM Developer Works (for download):**
  - http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/topic/com.ibm.db2.luw.admin.wlm.doc/doc/c0053139.html
  - http://www-128.ibm.com/developerworks/forums/servlet/JiveServlet/download/1116-179878-14005115-301960/wlmiodlab.zip

DB2 WLM Resources (cont.)

- **Article on DB2 Workload Management Histograms (3 Parts) in the Smart Data Administration e-Kit**
  - http://www.ibm.com/developerworks/data/kits/dbakit/index.html
- **White paper: Workload Management with MicroStrategy Software and IBM DB2 9.5**
  - http://www-01.ibm.com/software/sw-library/en_US/detail/G407381L49488H62.html
- **Exploitation of DB2's Workload Management in an SAP Environment**
  - https://www.sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/d046f3f5-13c5-2b10-179d-80b6ae7b9657
- **IBM Redbook**
  - http://www.redbooks.ibm.com/redpieces/abstracts/sg247524.html
- **IOD 2007 Session 1212: What's New in Workload Management**
  - Covers contents of DB2 Workload Manager in DB2 9.5
- **IDUG 2008 Session E04: DB2 9.5 Workload Manager: A Handyman's Tour**
  - Explores less well-known aspects of DB2 Workload Manager in DB2 9.5