CICS, IMS, and DB2 in a WLM World

Steve Grabarits
S/390 Performance
Poughkeepsie, New York
STEVEJG@US.IBM.COM

Session 503

OS/390 & Storage Systems Technical Conference

Trademarks

The following terms used in this publication are trademarks of the IBM Corporation in the United States and/or other countries:

ACF/VTAM
CICS
DB2*
DFSMS/MVS
Enterprise Systems Architecture/390
ES/9000
ESA/390
IBM*
IMS
MVS
MVS/ESA
OS/390
S/390*
System/390*
S/390*
VTAM*

This presentation has not been submitted to any formal IBM test.

It is distributed on an "as is" basis without any warranty either expressed or implied.

Abstract

re you confused about how to classify and set goals for OLTP ansactions in WLM goal mode? Interested in how WLM manages ICS, IMS, and DB2 work?

his session will explore the meaning behind the WLM policy definitions r these subsystems. The externals will be covered, as well as how /LM uses these definitions to manage this activity.

Agenda

- CICS and IMS first
- then DB2

A CICS/IMS Configuration

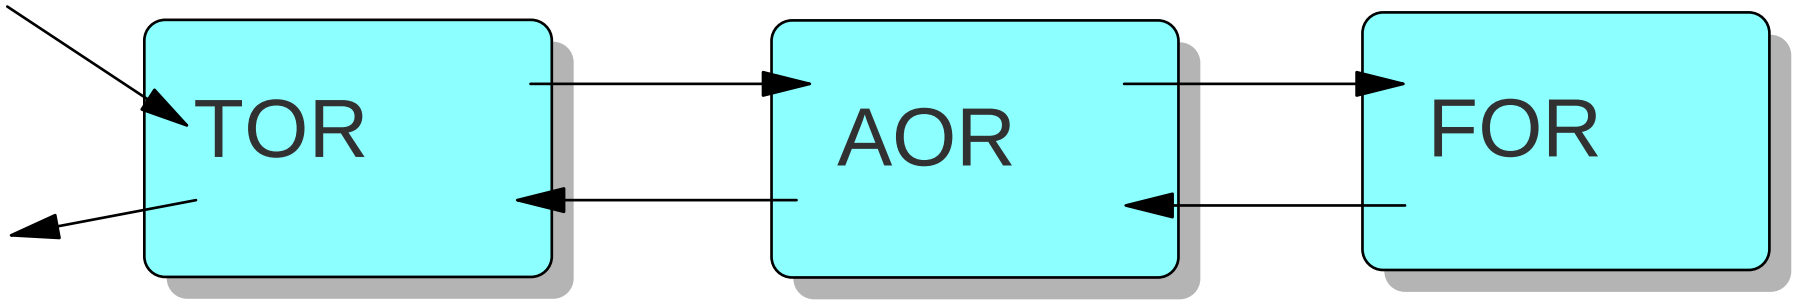

Adjust resources

CPU

Storage

I/O

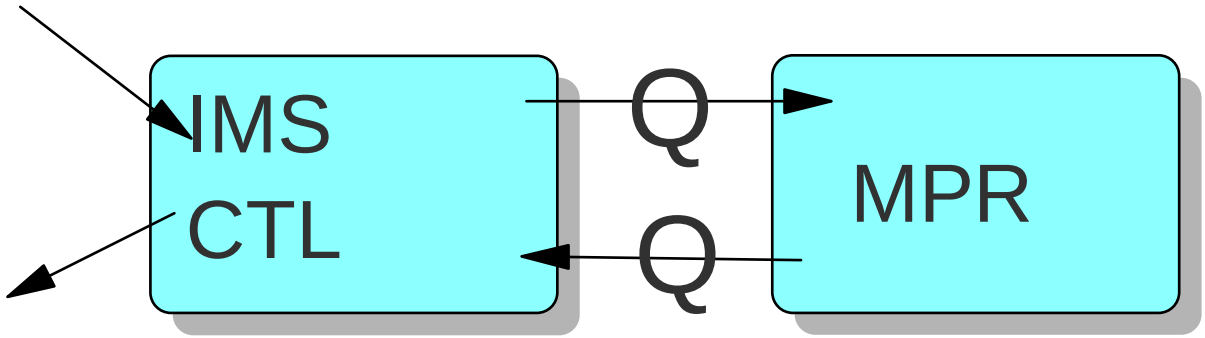

WLM Interfaces

- Work Manager Services
  - Connect/Disconnect to WLM
  - Classify a transaction to a service class
  - Report end of tran and report response time
- Execution Delay Monitoring Services
  - Notify end of execution phase (time in application)
  - Create/Delete a performance block (pb)
  - Initialize a performance block
  - Change state of a performance block
- Exploited by CICS V4 and IMS V5 and higher
  - Issued in goal and compatibility mode
  - Issued regardless of goal types selected

CICS WLM Exploitation

- Transaction classified and reported in TOR (1st CICS region)
- Each transaction associated with a PB in each region
- PBs per region based on CICS MAXTASK

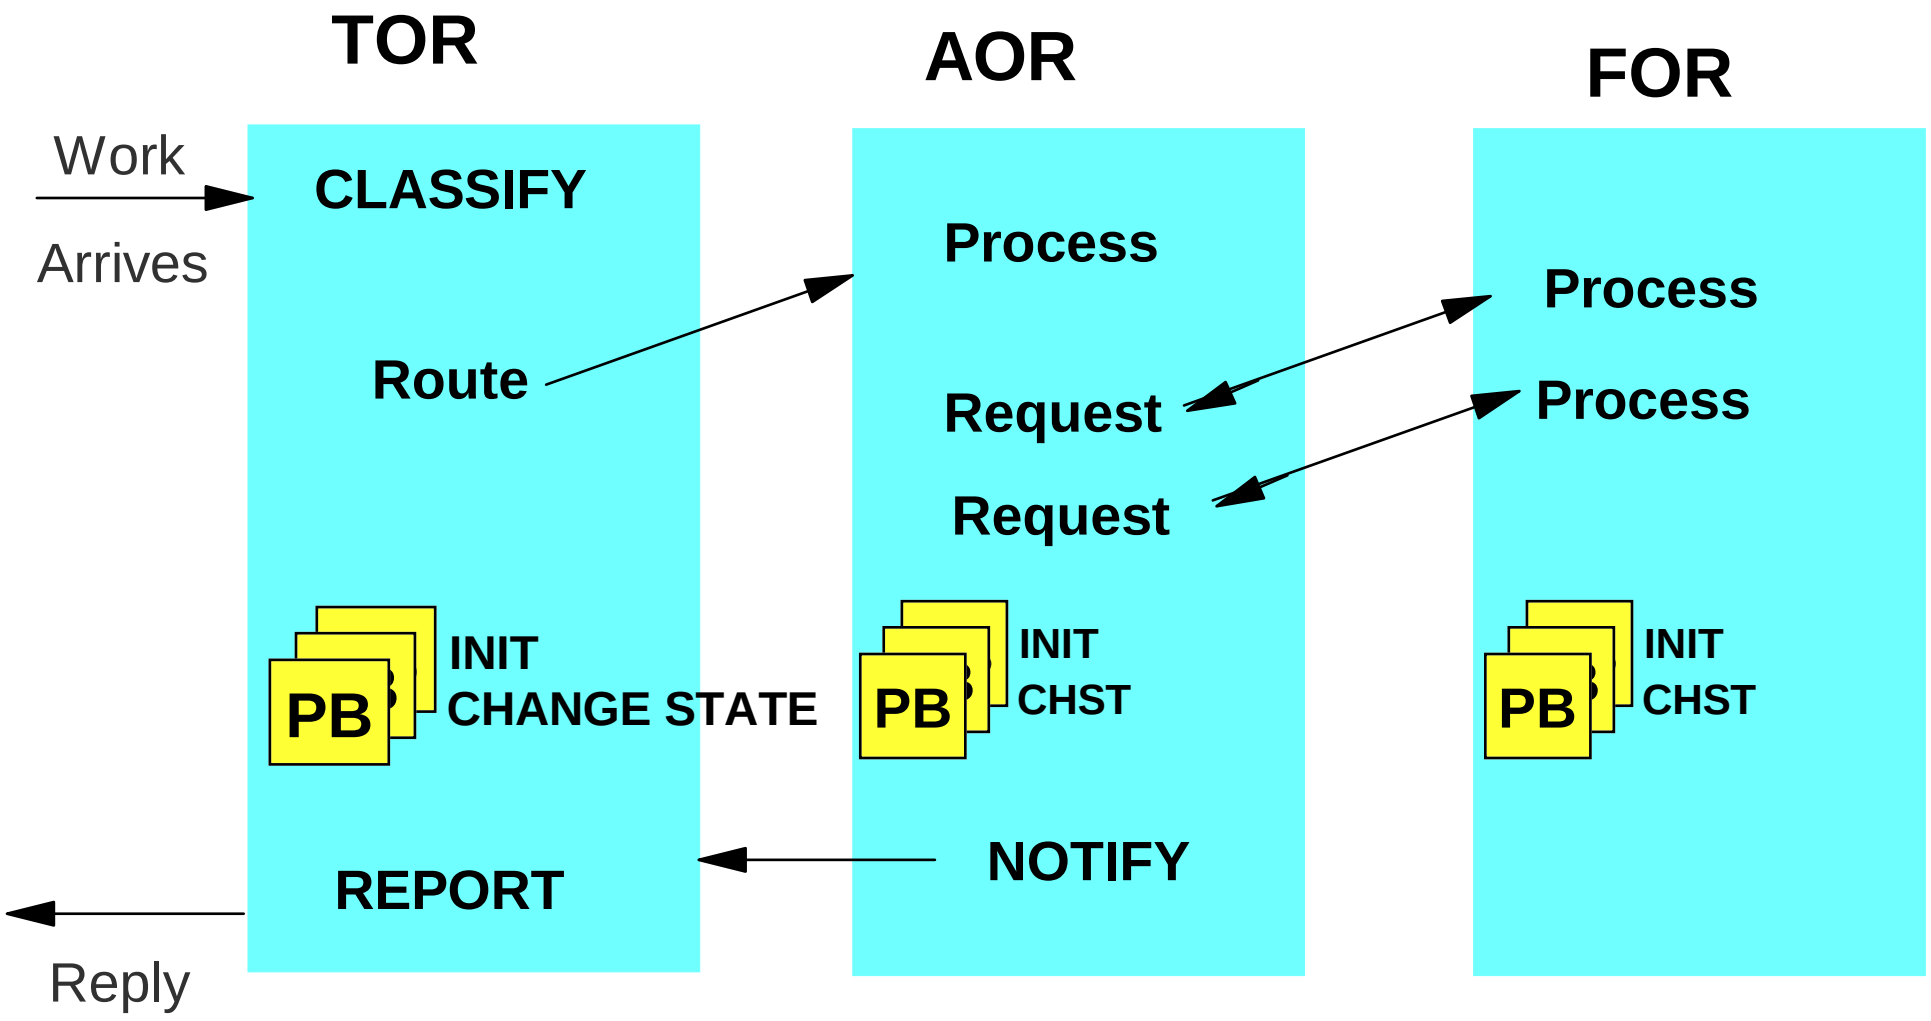

IMS WLM Exploitation

- Transaction classified and reported in Ctl Region
- Each transaction associated with a PB in dependent region
- One PB per dependent region (message driven BMP like MPR)

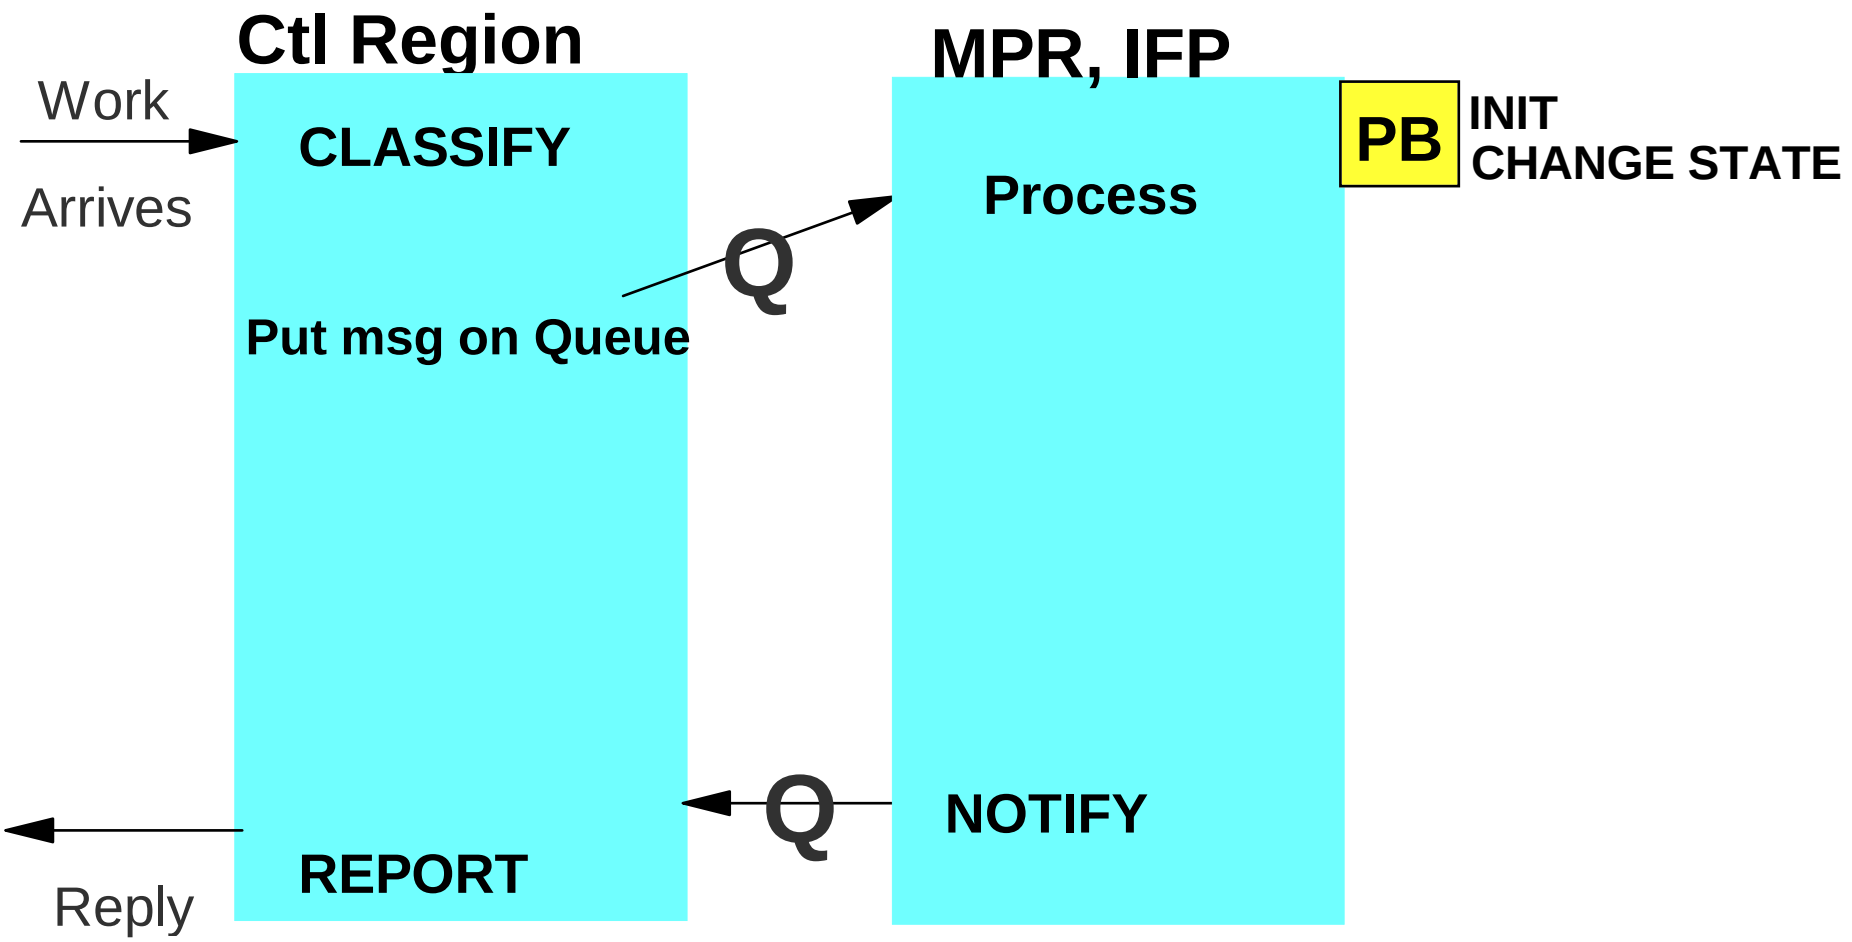

Goals for CICS / IMS

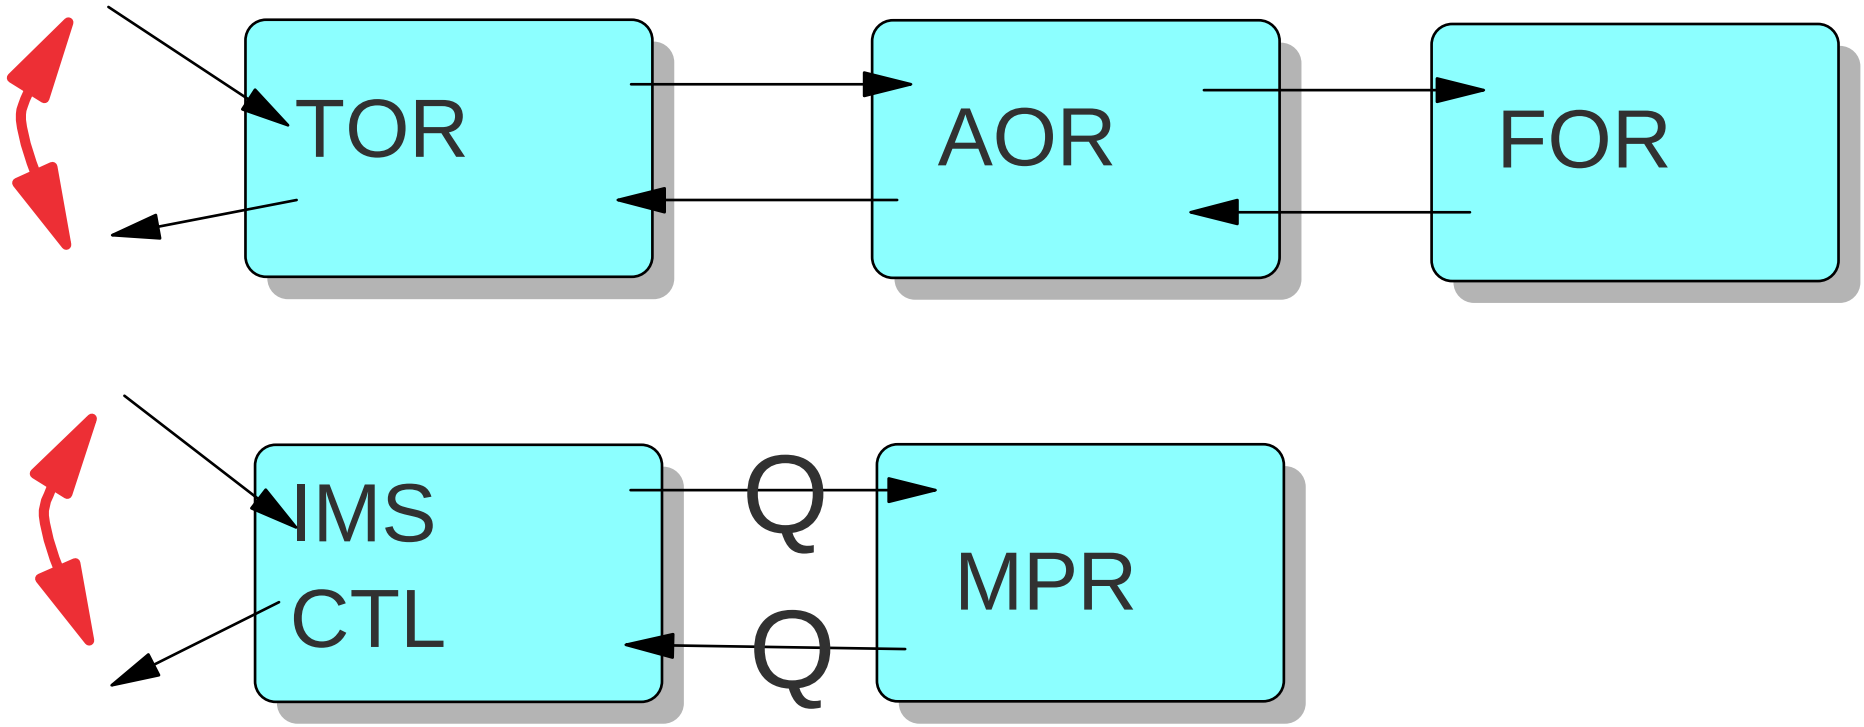

1) Velocity

$$\frac{\text{CPU using+ I/O using}}{\text{CPU using + I/O using +delay}} * 100$$

Delays: CPU, paging, MPL, swap-in, I/O, initiator queue delay

2) Transaction Response times

- average
- percentile

Velocity Goals

- Address Space Level Metric
  - Address Spaces (STC, JES) classified to service class
  - WLM samples address spaces every 1/4 second
  - PBs sampled infrequently (APAR OW32140)
  - WLM management based on class velocity/importance
- Only choice for pre CICS V4 and pre IMS V5
- Velocity goals are very acceptable for production
- Use RMF I Workload Activity or RMF III SYSSUM reports to determine velocity
  - Peak mode data
- Definition of velocity has changed over time
- Requires verification of goals with system changes
- Used during CICS/IMS startup with transaction goals

Response time management

- Velocity goals left in place: then add transaction goals
- WLM detects this condition and sets each participating address space as a transaction server
- Address spaces, not transactions, consume resources
  - yet managed to transaction goals/importance
- Velocity goals / service class ignored for WLM management
  - but used for reporting purposes
- Classify transactions to different service classes than the CICS and IMS address spaces

Transaction Classification

- CICS
  - LU Name
  - Subsystem Instance :
    APPLID of 1st CICS
  - Transaction Name :
    TRANID
  - Userid
- IMS
  - LU Name
  - Netid (LU 6.2)
  - Subsystem Instance :
    IMSID
  - Transaction Class
  - Transaction Name
  - Userid

CICS / IMS Transaction Goals

- In compat mode use SRVCLASS parm to capture transaction service class data in a RPGN
- If goal mode and velocity goals only, classify to report classes for transaction data
- Transaction class goals
  - Average - used also by CP/SM
  - Percentile - CP/SM uses shortest queue algorithm
- Period durations are meaningless, only first period honored
- No transaction velocity or discretionary goal

Transaction Server Management

- Objective: assign goals to the transactions and let the system determine which regions need the resources to meet these goals.
- WLM identifies what regions are serving which transaction service classes
  - Based On
    - PB sampling : every 1/4 second
    - Notifies and Reports
  - Groups similar regions into same dynamic internal server period ($SRMSxxx)
  - Builds a mapping of this dynamic relationship
- WLM uses address space samples to determine delays for transaction service classes that need "help"

CICS Example

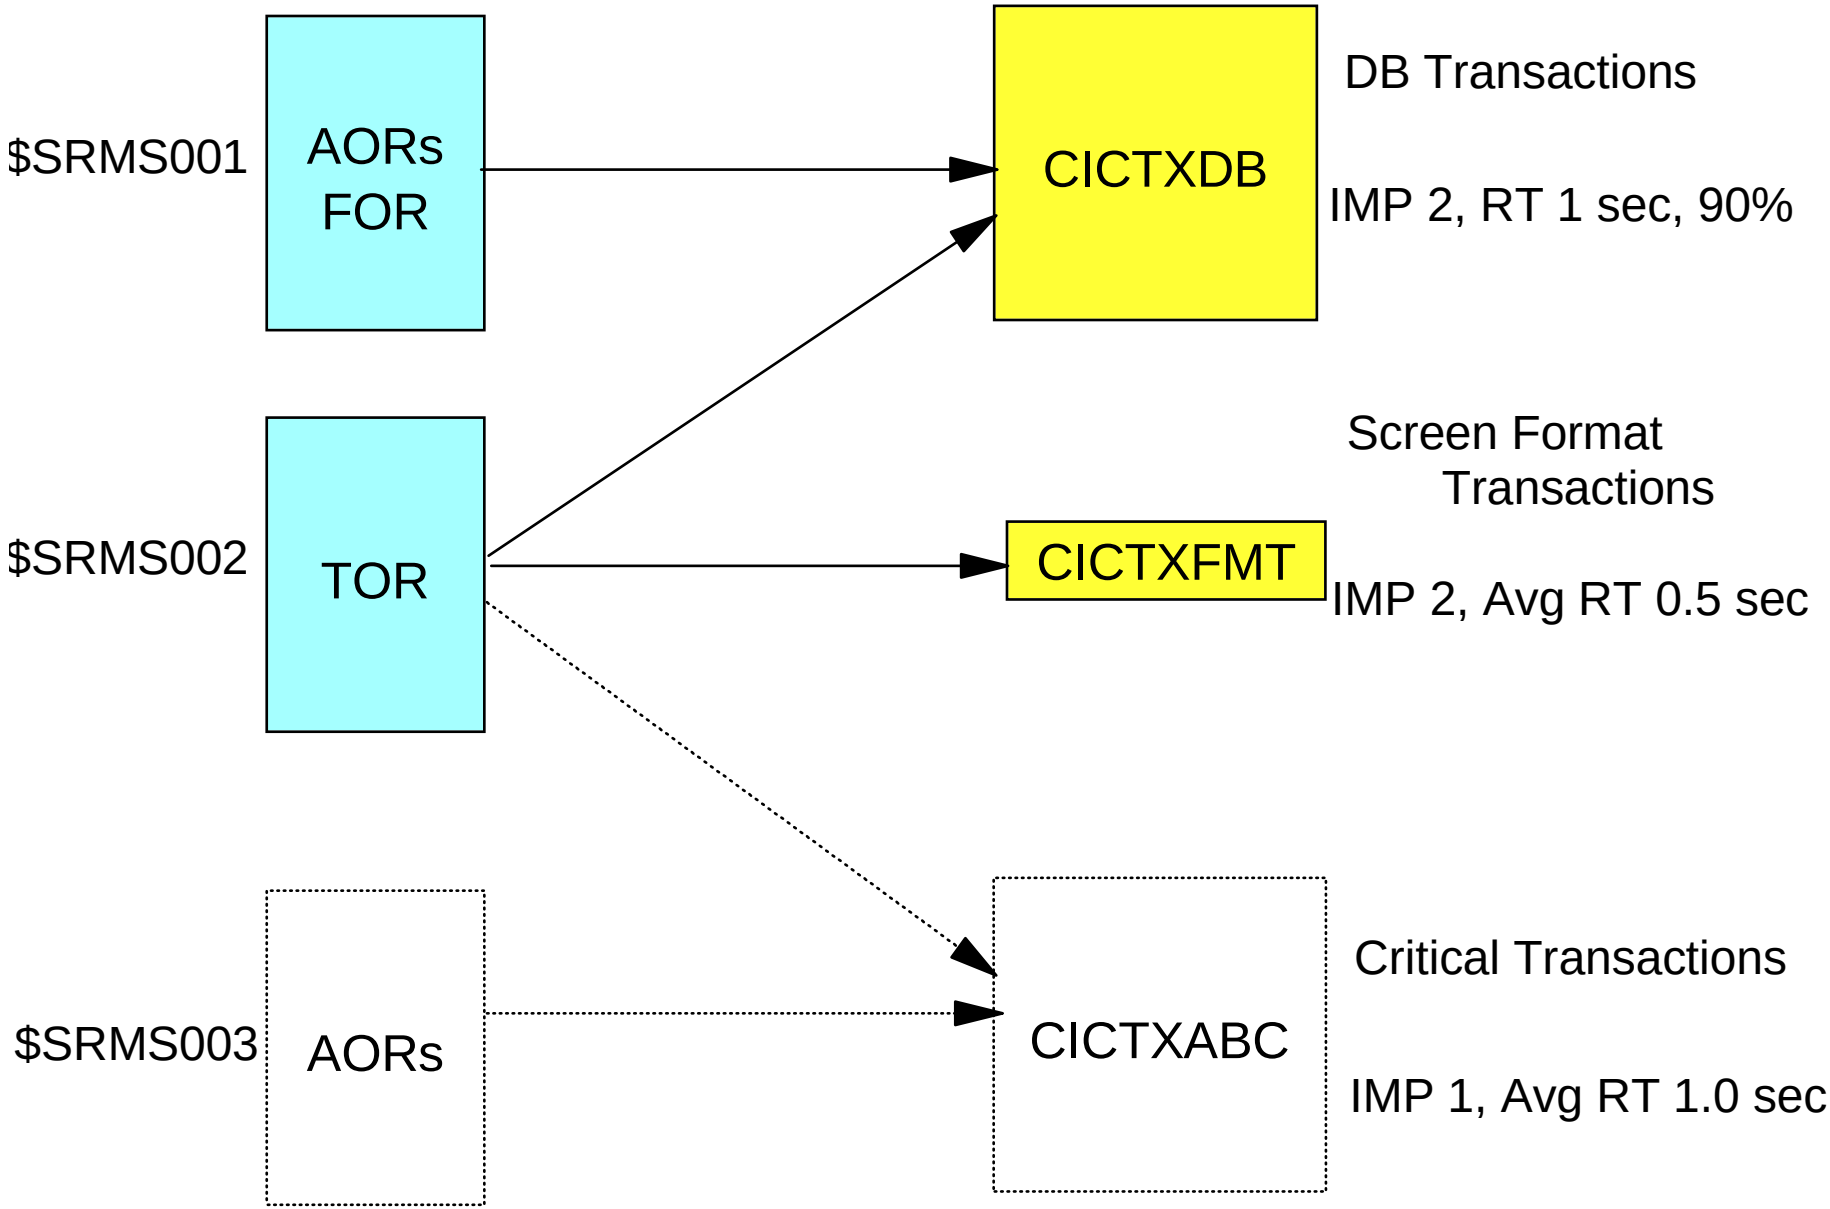

IMS Example

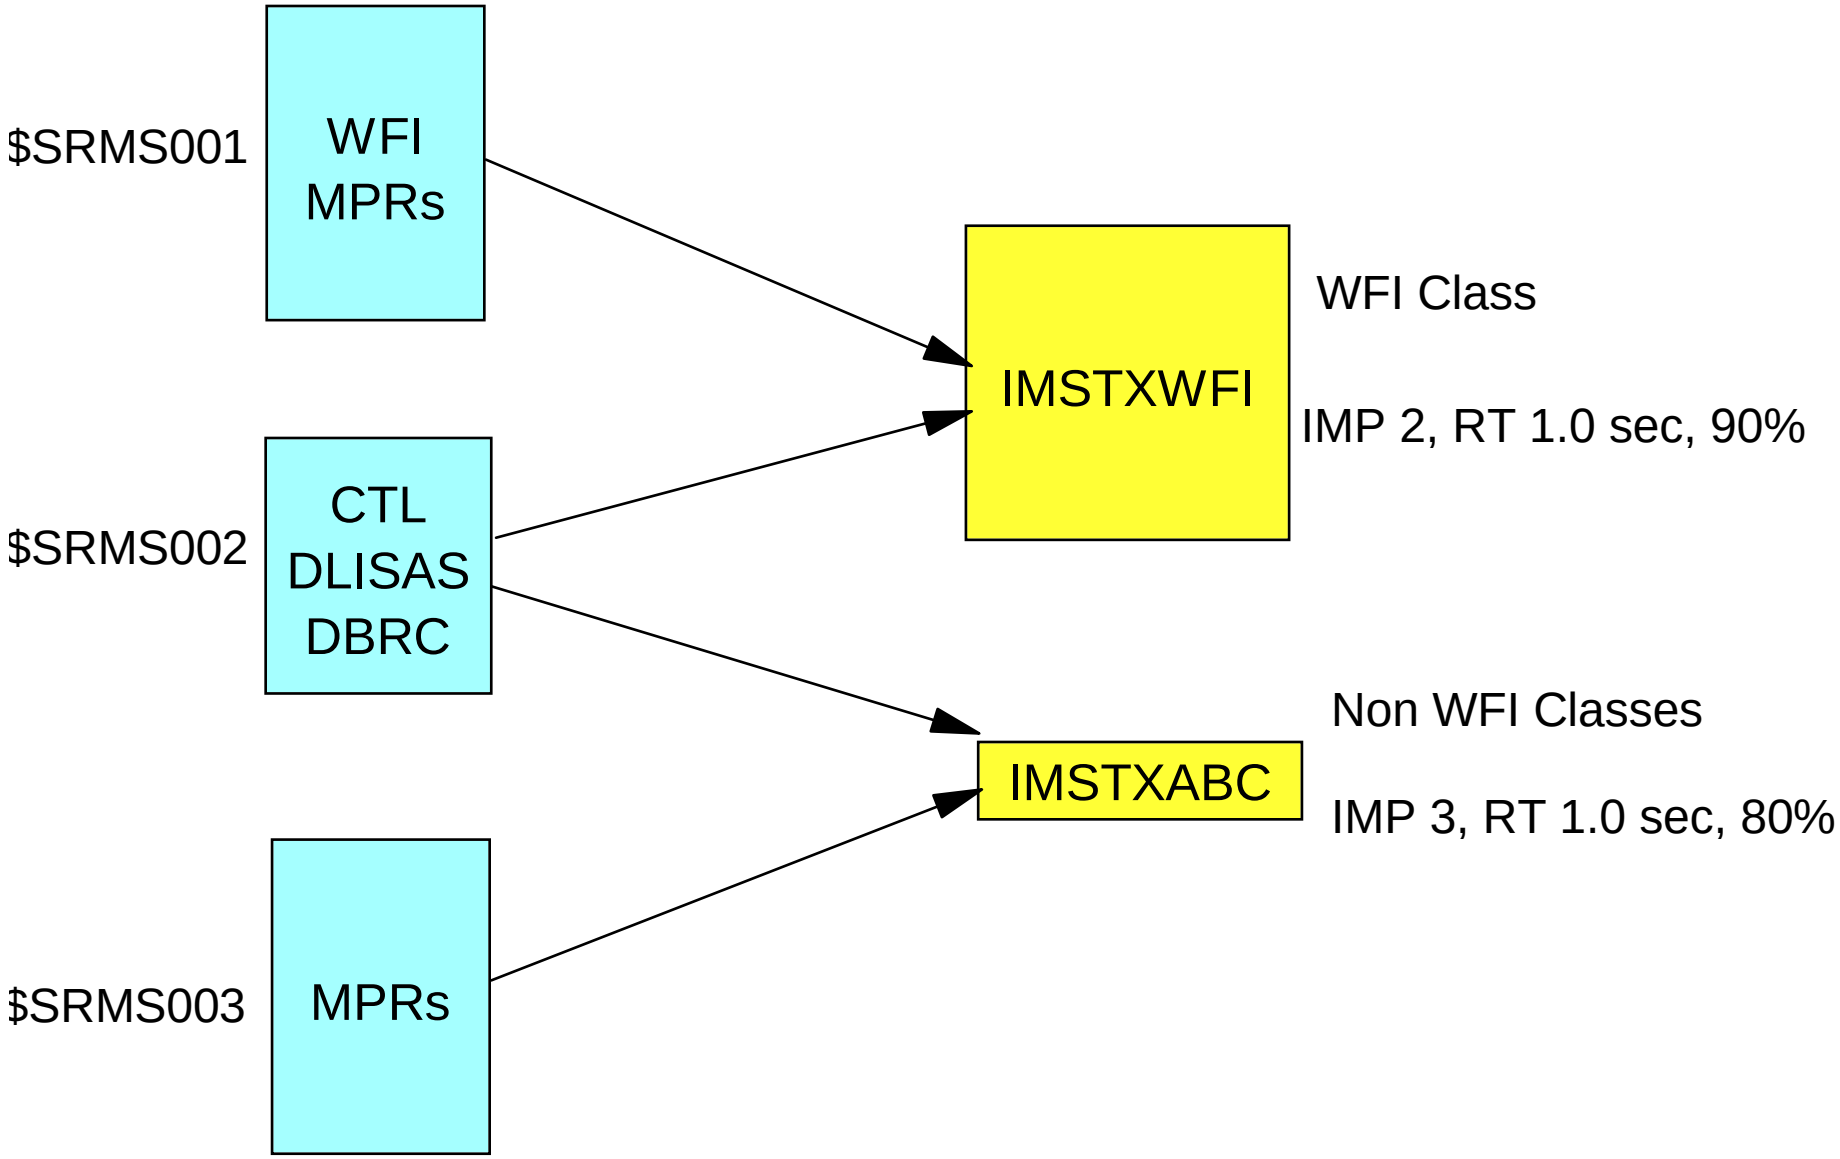

Setting transaction goals

- Start simple, grow to a few transaction service classes
  - Do what makes most business sense
  - few classification rules (TNG if many)
  - Production versus Test
  - Classify most representative and/or critical transactions
    - frequent, predictable, or sensitive to resources
    - default class - non aggressive catcher for rest
  - average of all transactions or percentile? percentile may be most practical
- WLM manages resources at region level only
  - Most benefit when application regions map uniquely to transaction classes
  - Always the case with past tuning

DB2

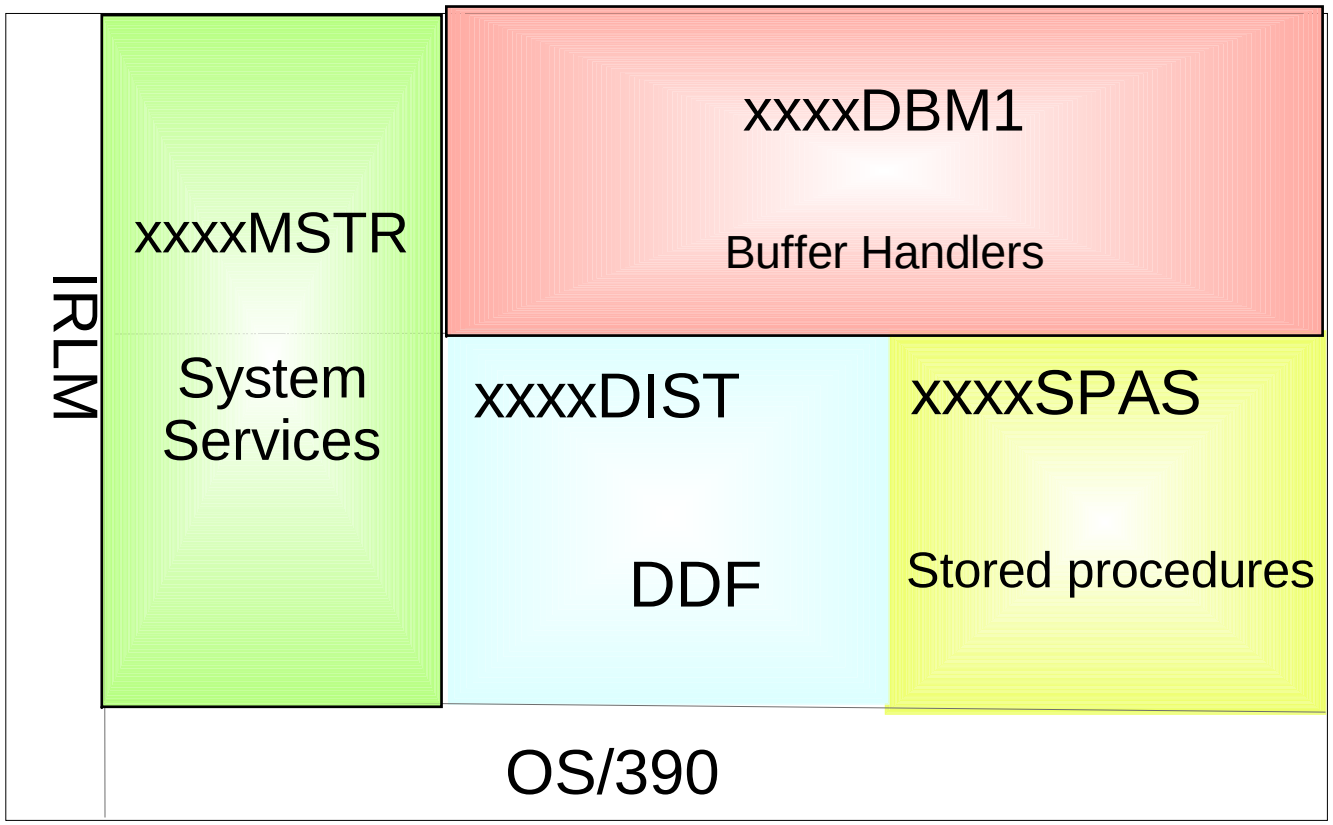

- local
- distributed
- Sysplex Query Parallelism
- WLM stored procedures address spaces

DB2 Server

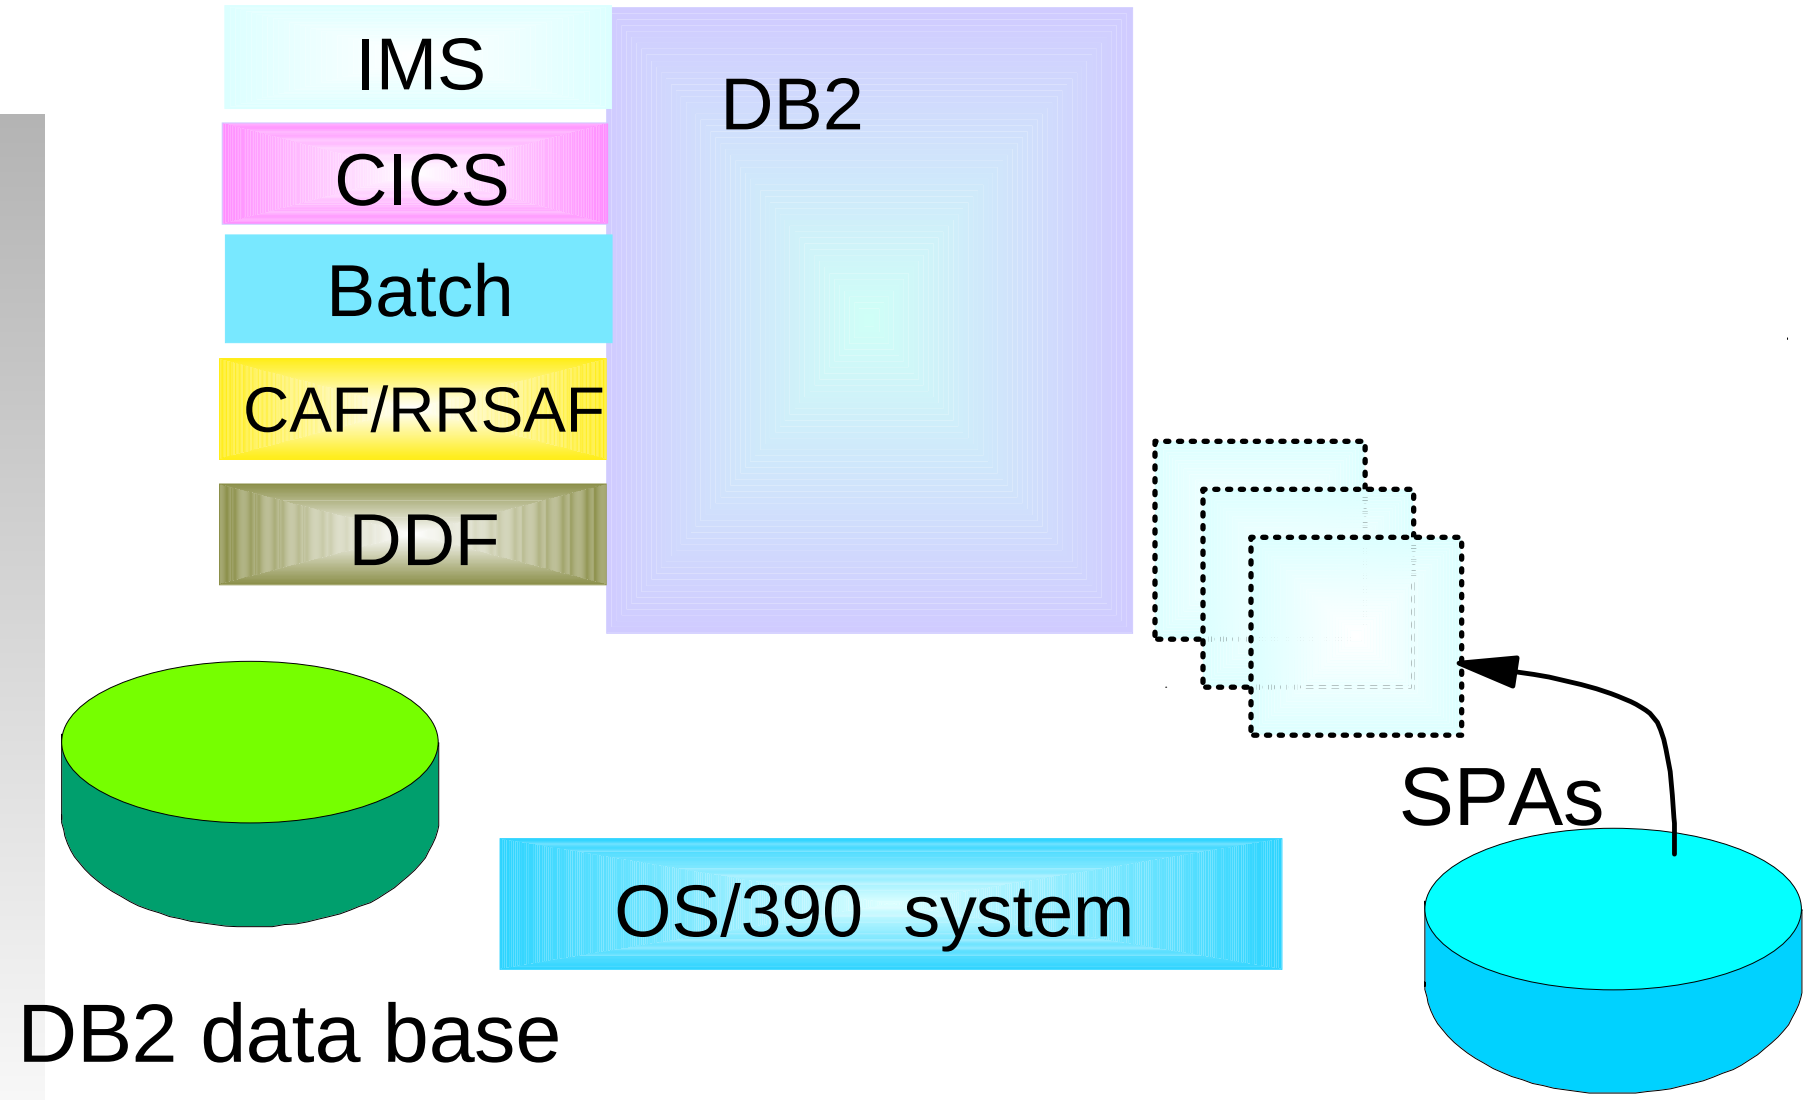

Local Attachments

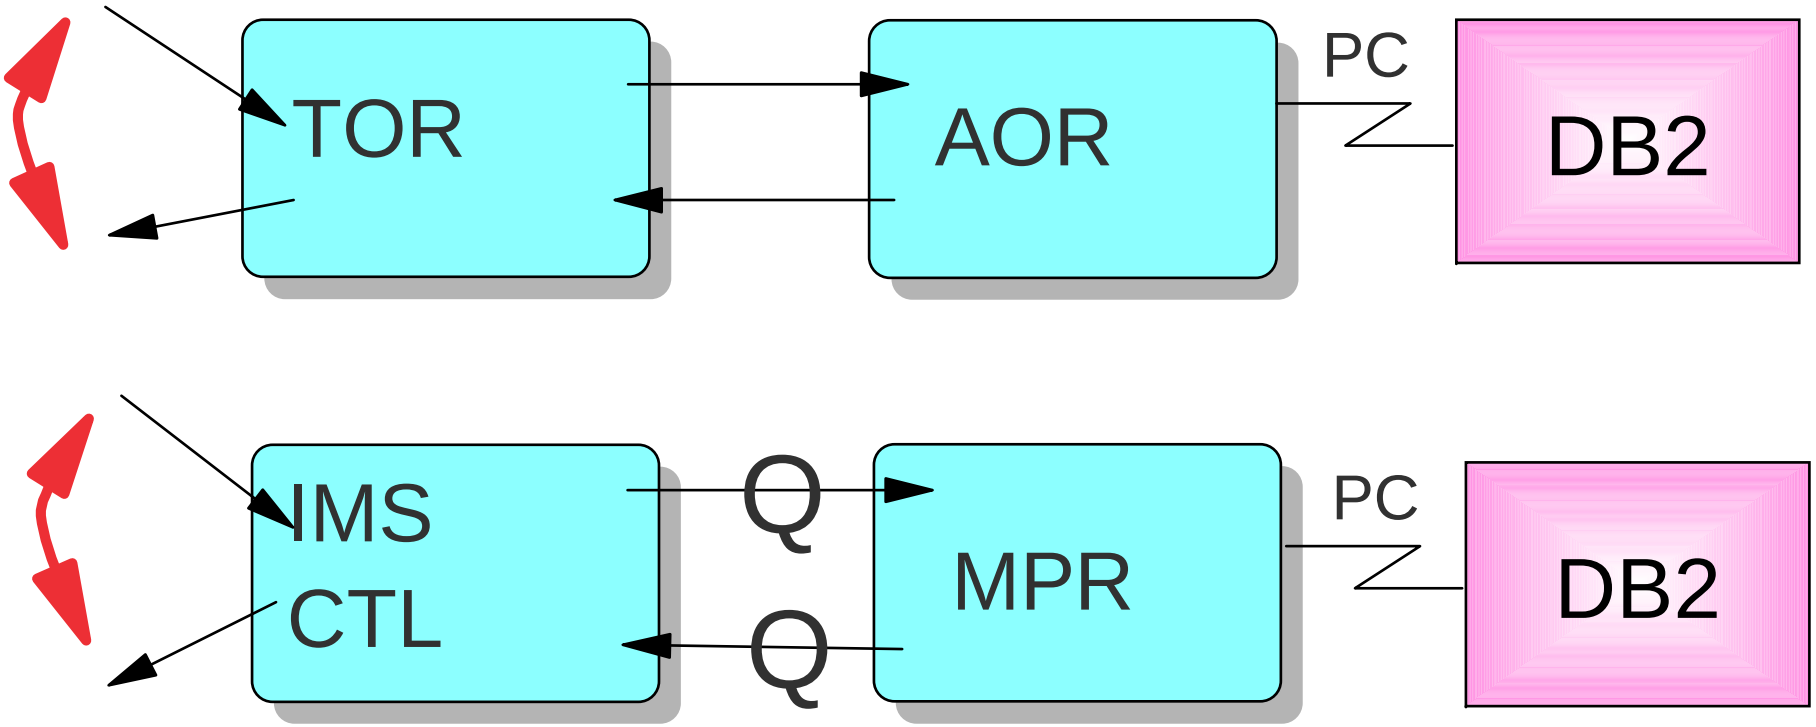

DB2 SQL activity runs under dispatchable unit of invoker

- IMS,CICS, TSO, Batch, etc.
- Inherited classification class of invoker
- Priority and management of home unit
- Service attributed back to invoker

Previous DDF Implementation

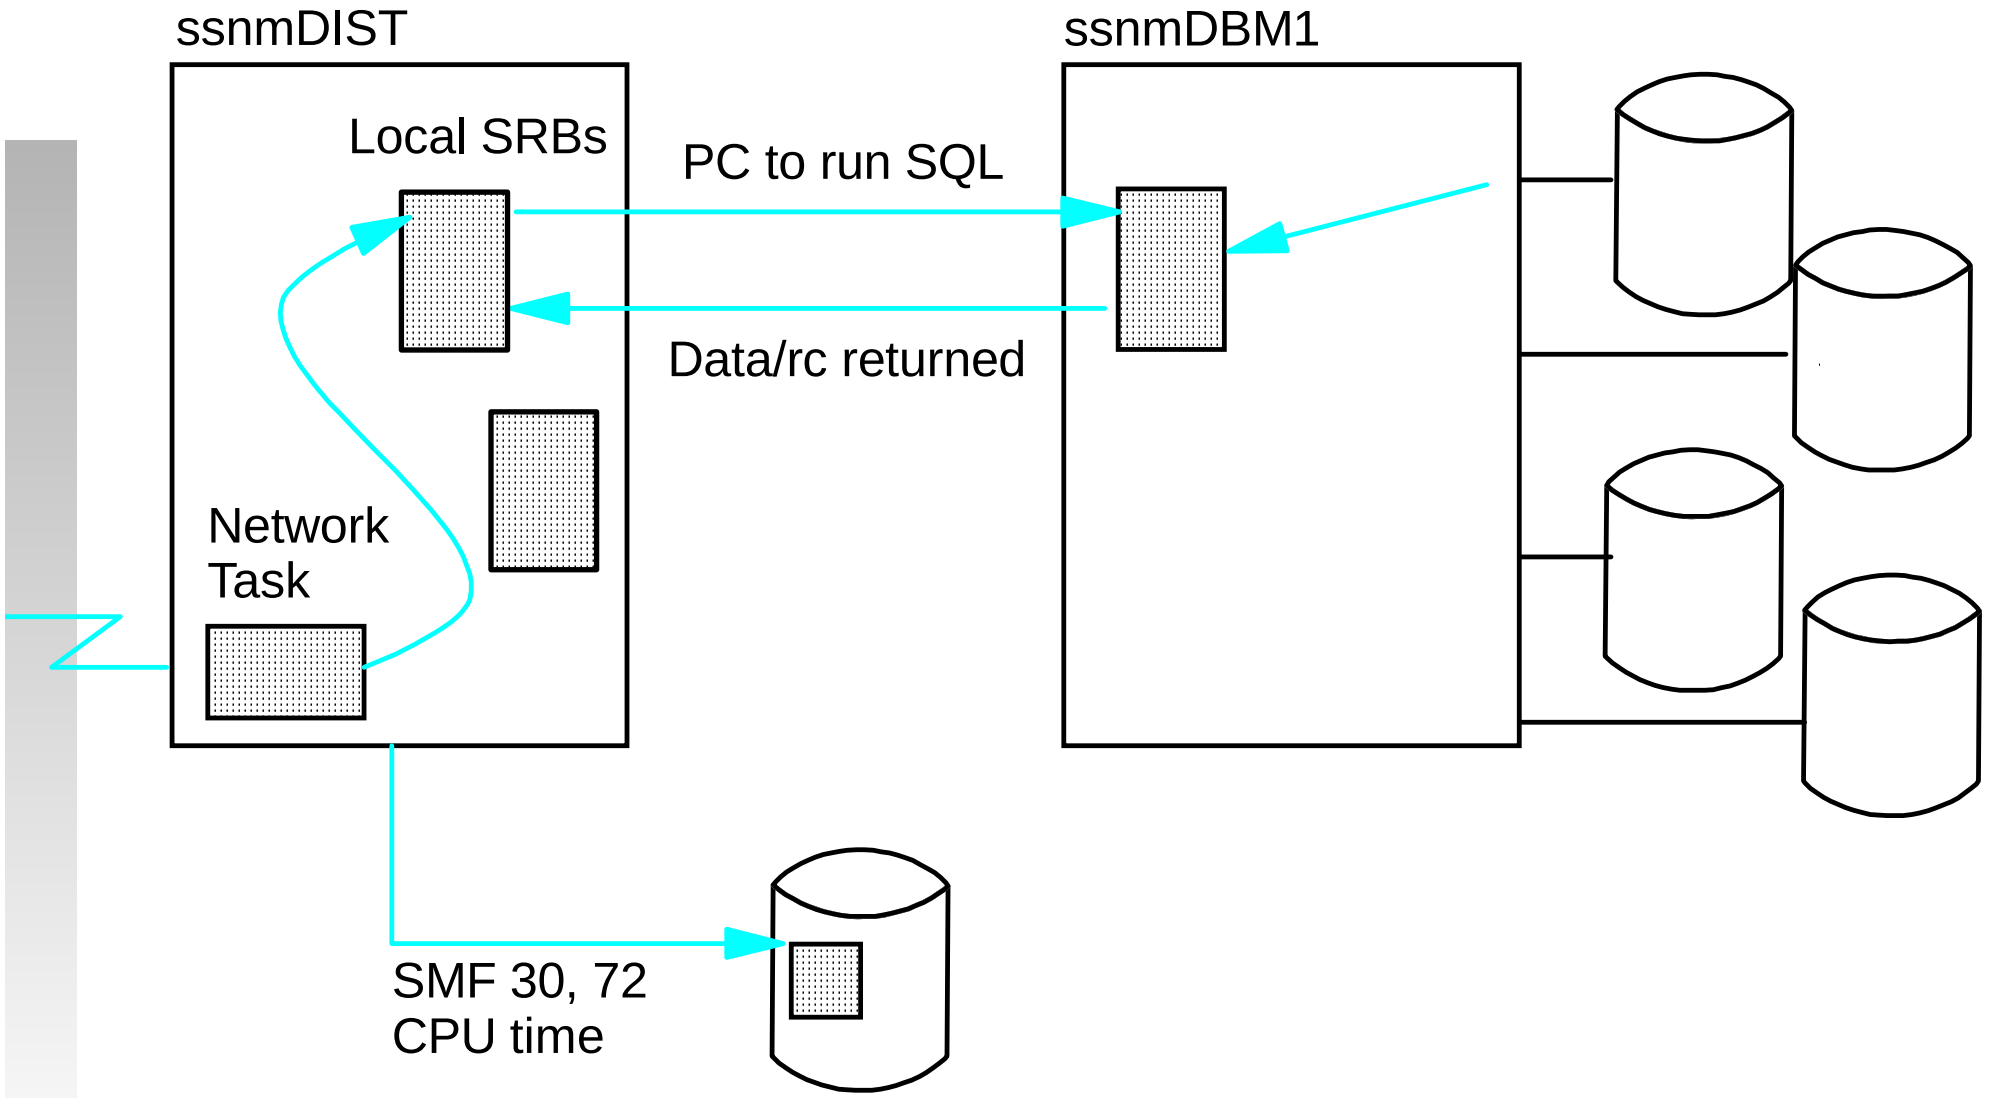

Enclaves

- Represents a "business unit of work"
  - and an individual SRM transaction
- Is managed separately from the address space
- Can include multiple SRBs/TCBs
  - SRBs are preemptible type
- Can span multiple address spaces
- Can have many enclaves in a single address space
- Exists in both goal and compatability mode
  - assigned to service class or PGN
  - managed separately from address space
- DB2 V4 exploits the WLM Enclave Interfaces
  - create/delete the enclave, join an enclave, etc.

Managing DDF Work in DB2 V4

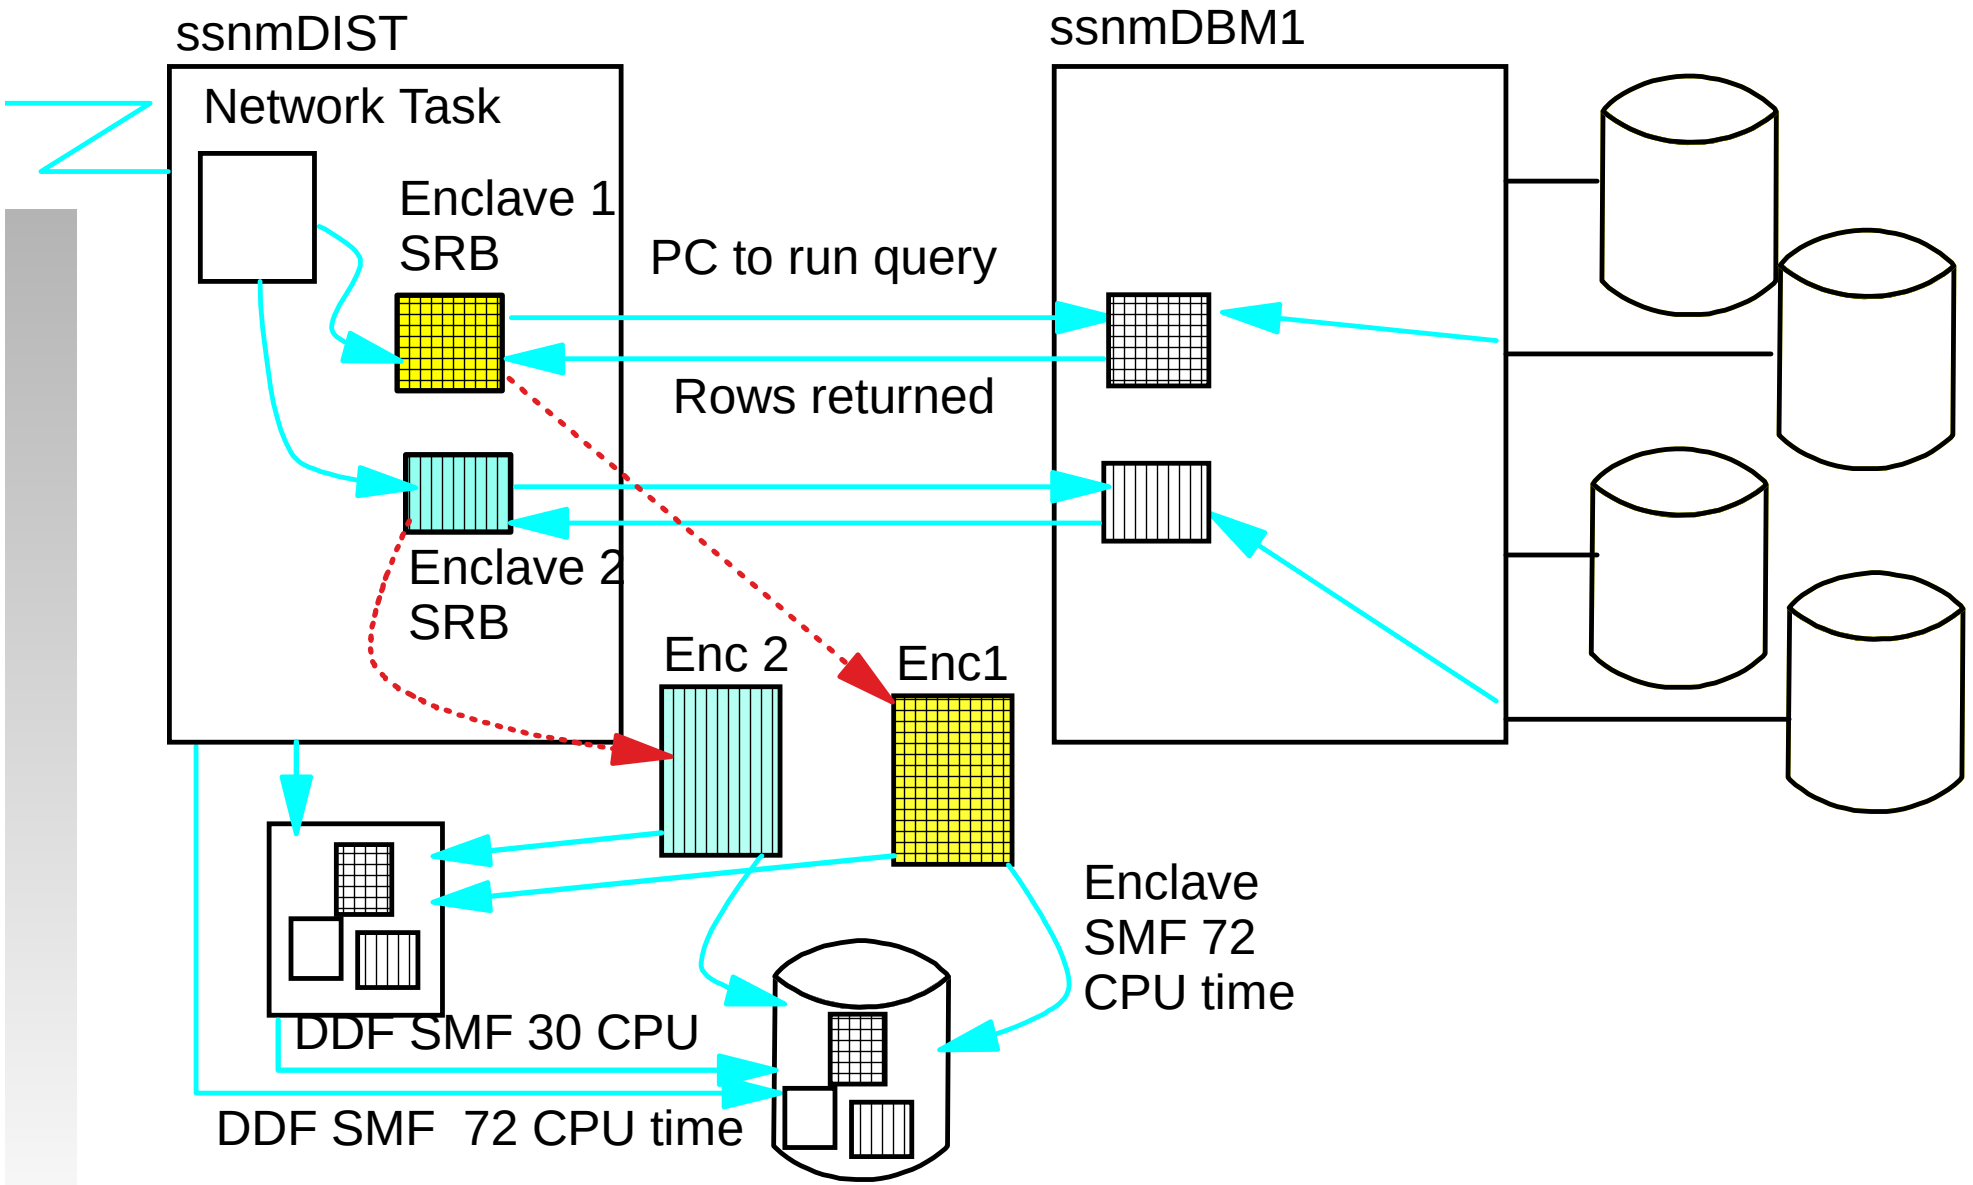

Classifying DDF Transactions

- All enclaves, including DDF transactions, are classified using the active WLM policy in both compatibility and goal modes
- In compat mode the ICS can be used to map a service class to a performance group, and likewise for reporting
- Classification attributes: Accounting Information, Collection Name, Connection Type, Correlation Information, LU Name, Netid, Package Name, Plan Name, Procedure Name, Subsystem Instance, and Userid
- **Defaults** if you do not classify in WLM policy:
  - **Goal mode:** enclaves default to the SYSOTHER service class which has a discretionary goal!
  - **Compat mode**: any enclaves run in the pgn/rpgn of the owning address space (DDF) as they did before

Managing DDF Work (Enclaves)

- Transactions are subject to period switch

- Goal mode:
  - All goal types allowed
  - WLM manages enclaves with its own dispatch priority, etc.

- Compat mode:
  - Performance groups and report performance groups can be assigned using SUBSYS=DDF
  - Limit of one rpgn
  - Time slicing is ignored
  - Domain is ignored

What is a Transaction?

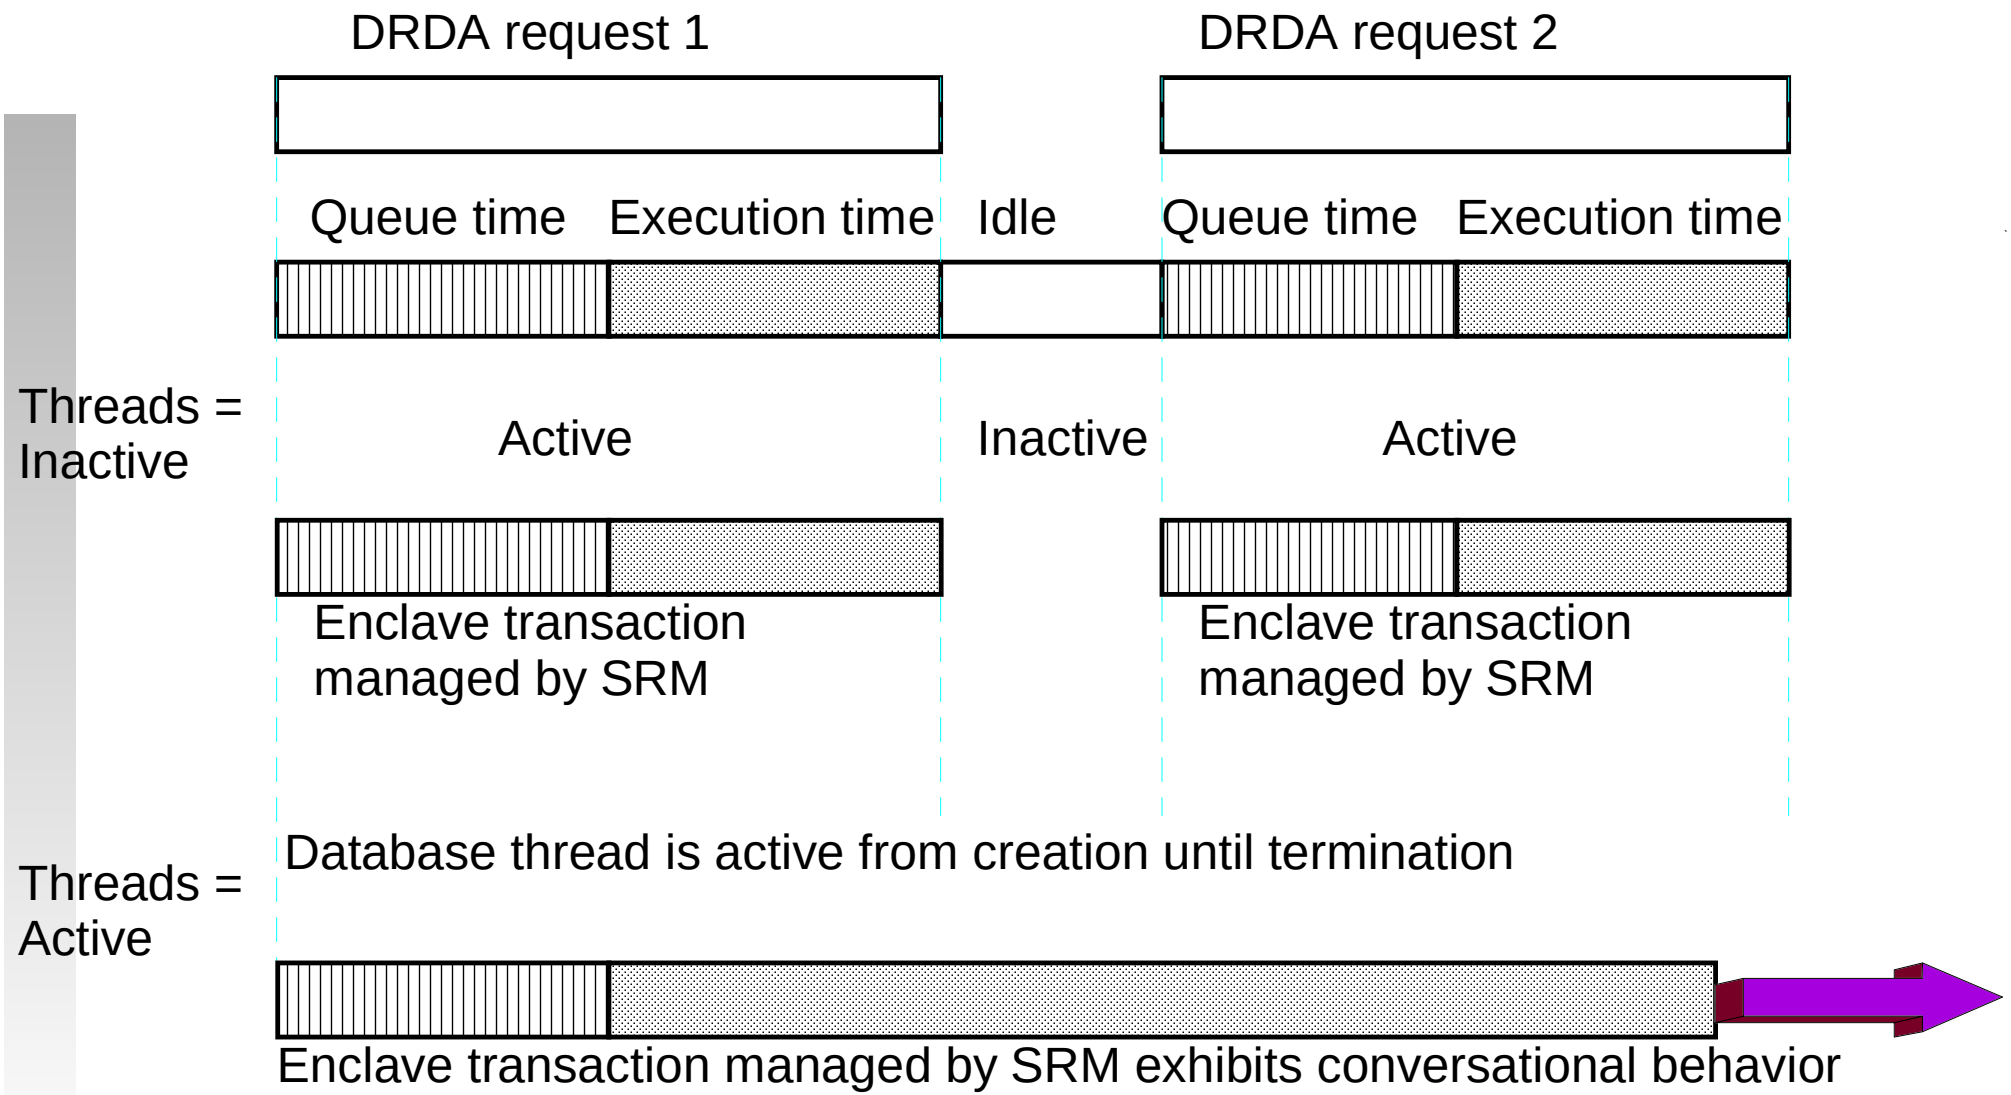

What Goals Should I Use?

- WLM compatibility mode
  - Goal does not matter. Service class used only to associate an enclave with a performance group.
- WLM goal mode
  - THREADS=INACTIVE **and** RELEASE(COMMIT)
    - DDF creates one enclave per active interval
    - Response times do not include think time
    - Response time goals and multiple periods can be used
  - THREADS=ACTIVE **or** RELEASE(DEALLOCATE)
    - DDF creates one enclave for the life of the thread
    - Enclave response time includes think time
    - Response time goals should not be used
    - Multiple periods should not be used

Enclave System Effects (DDF)

- (Lots) more transactions. Each DDF transaction used to be part of the DDF address space SRM transaction, now SRM sees each DDF transaction.
- Increased active time since more transactions exist
- Large decrease in DDF SRB time/service in SMF 30 record
- Corresponding increase in CPU time/service in the service classes/performance groups where enclaves are running
- MSO and I/O service is unchanged
- DDF logons are not held up by existing DDF work

Sysplex Query Parallelism

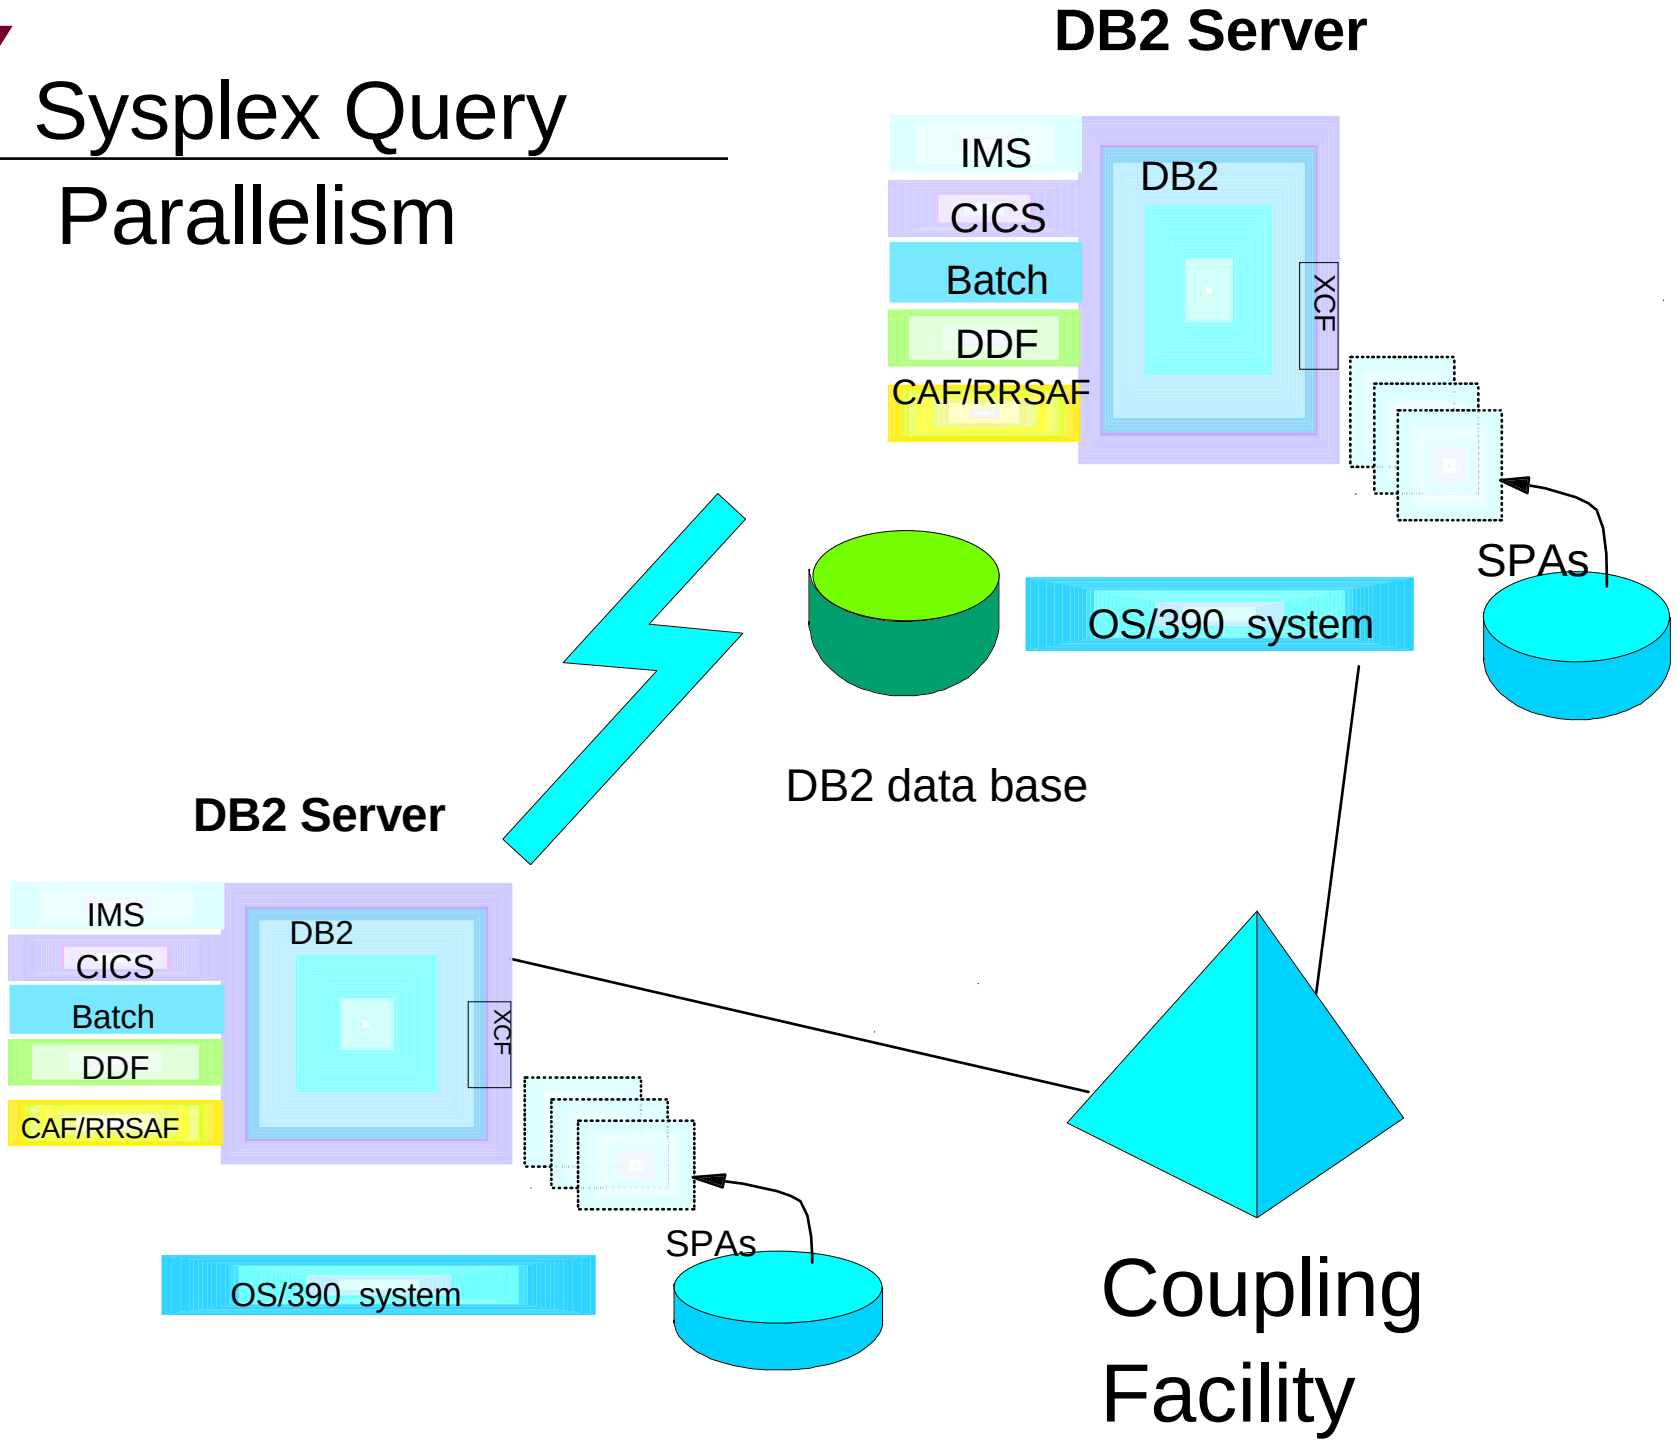

Sysplex Query Parallelism

- Supports data mining applications
- DB2 query can be split across parallel sysplex
- Requires DB2 V5.1 and DB2 data sharing
- Classification of originator doesn't change
- Classification of participants is under 'DB2' rules
  - ► Classification attributes inherited from originator
- Accounting is done on the system where enclave runs
- Each remote piece starts in 1st period

Stored Procedures in DB2 V5 and OS/390 R3

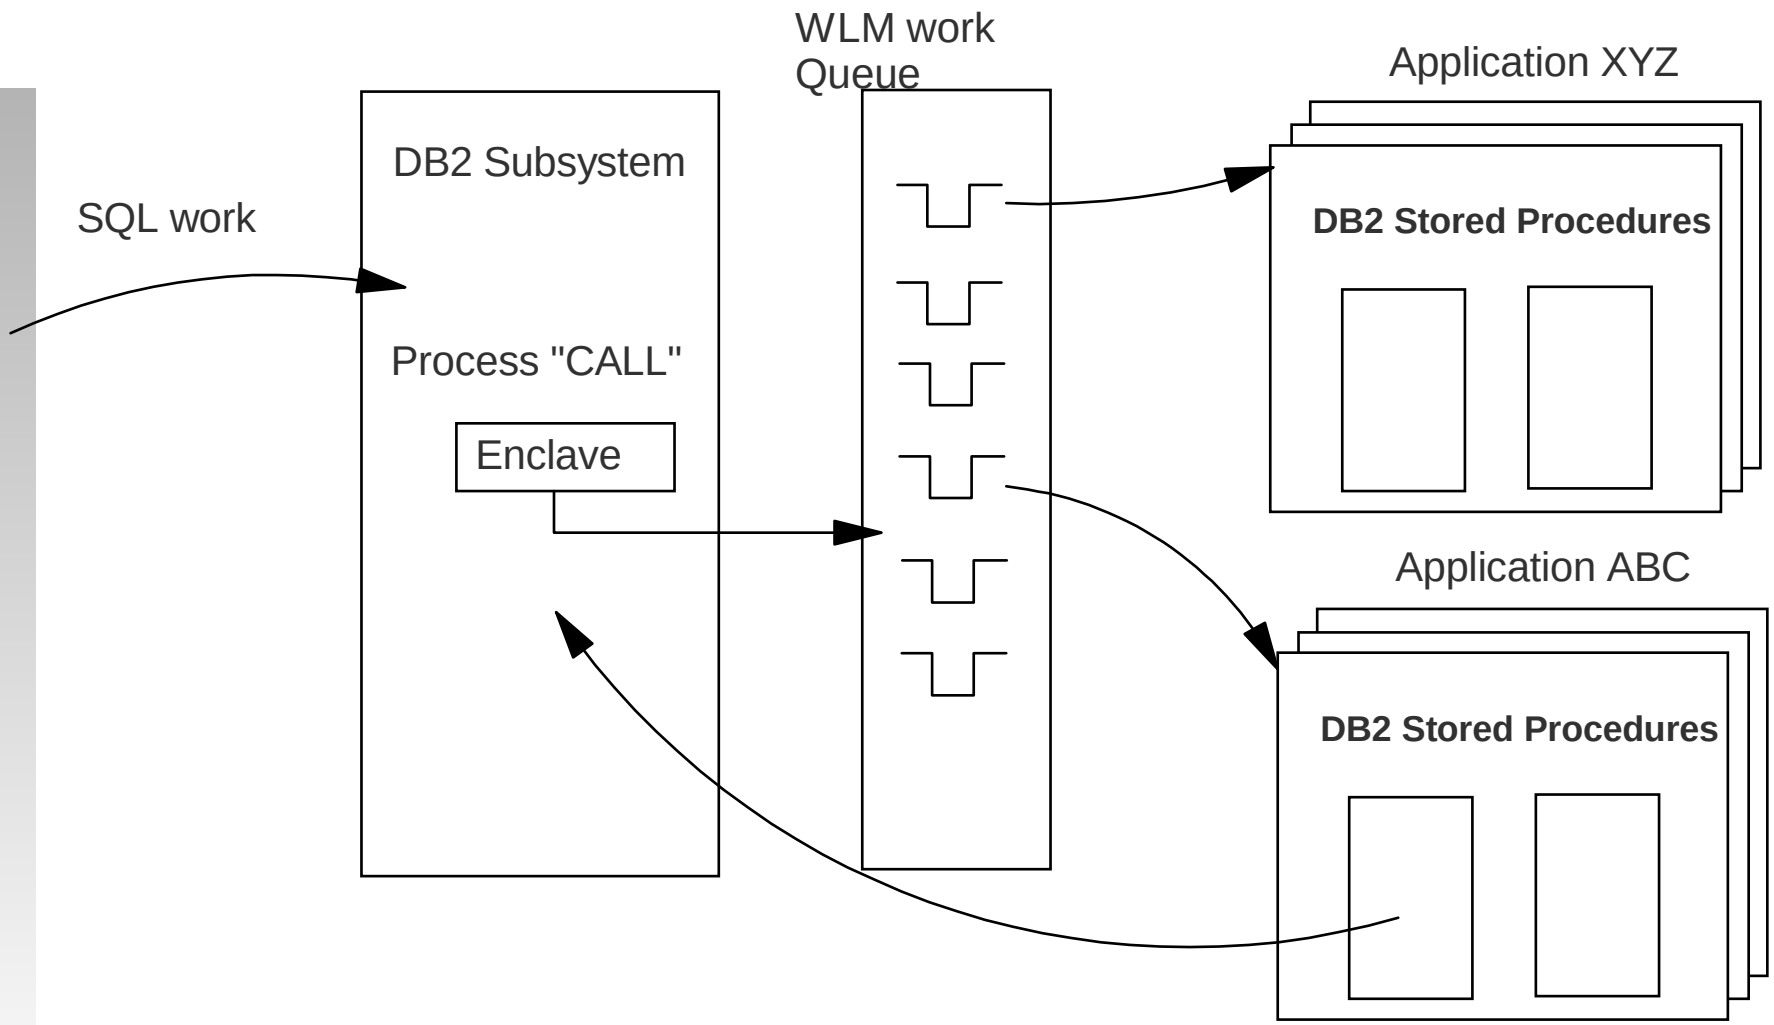

DB2 Spaces

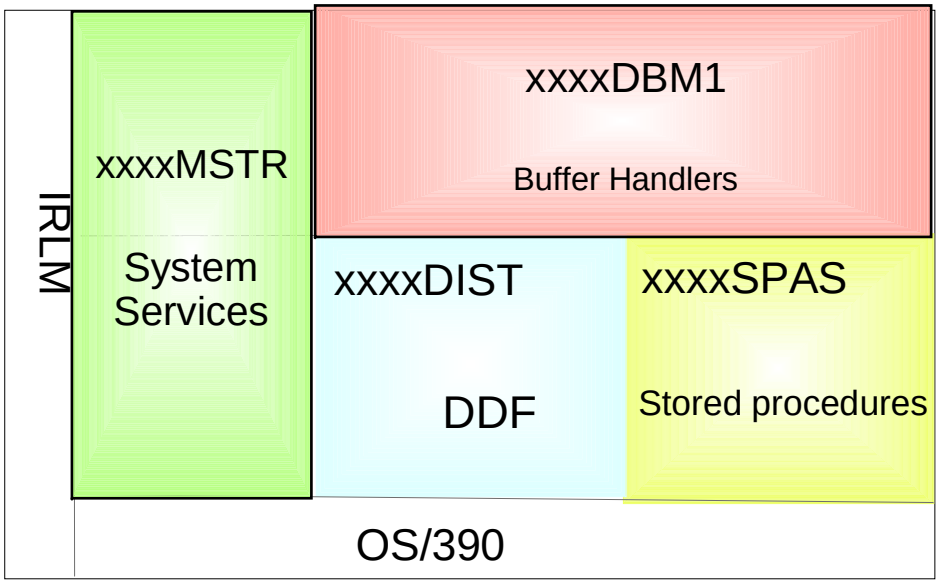

IRLM - SYSSTC

MSTR - hi velocity/imp.
DBM1
*DIST
*WLM SPAS

*service tasks

Last page