zTSU

z/OS V1R7
WLM
Server Specific Load Balancing Services for Communication Server and DB2

Table of Contents

Trademarks

- The following names are trademarks of the IBM: see url http://www.ibm.com/legal/copytrade.shtml

- Other company, product and service names may be trademarks or service marks of others

Overview

Sysplex Routing

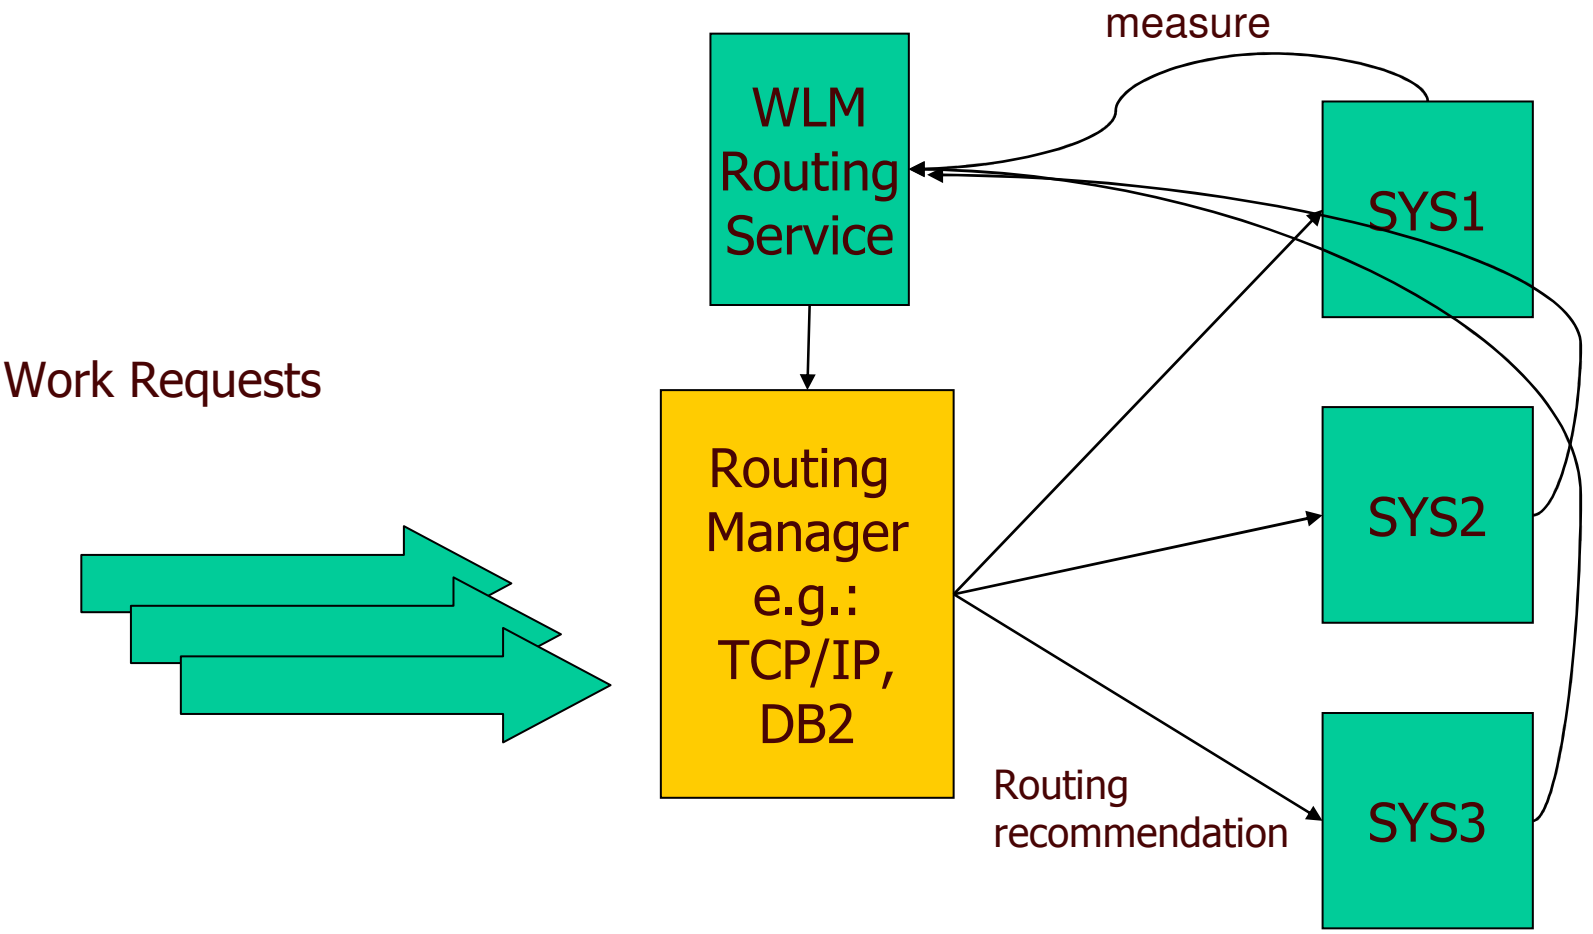

Overview

WLM Routing Service Today

Step 1: Registration of Servers by IWMSRSRG service

Register each Server on its local system as routing candidate

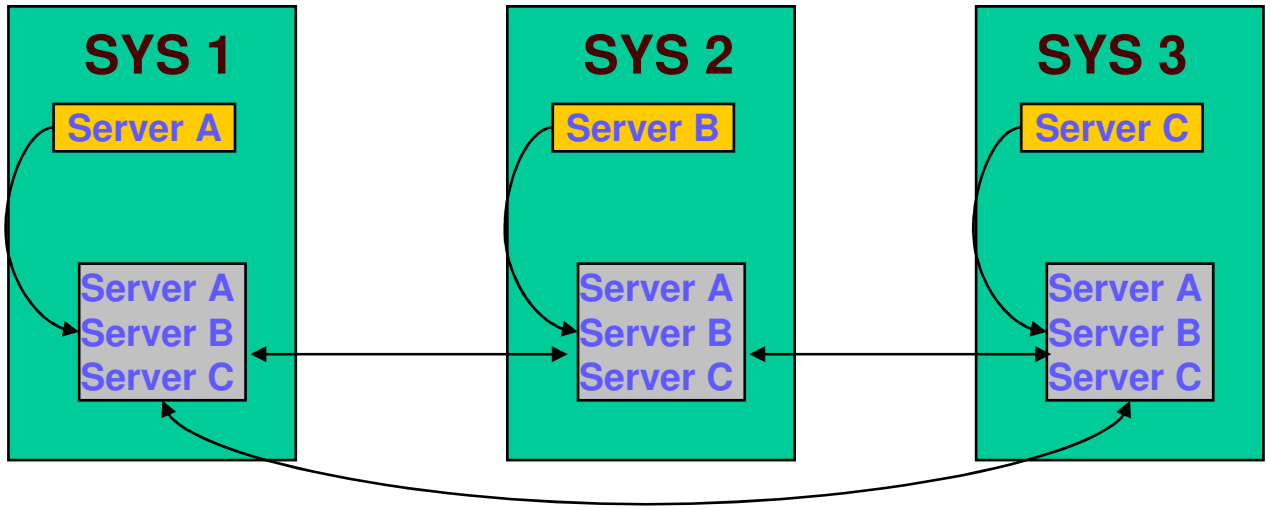

WLM communicates registered servers and capacity data between systems

Overview

WLM Routing Service Today

Step 2: Ask WLM for routing recommendations by IWMSRSRS service

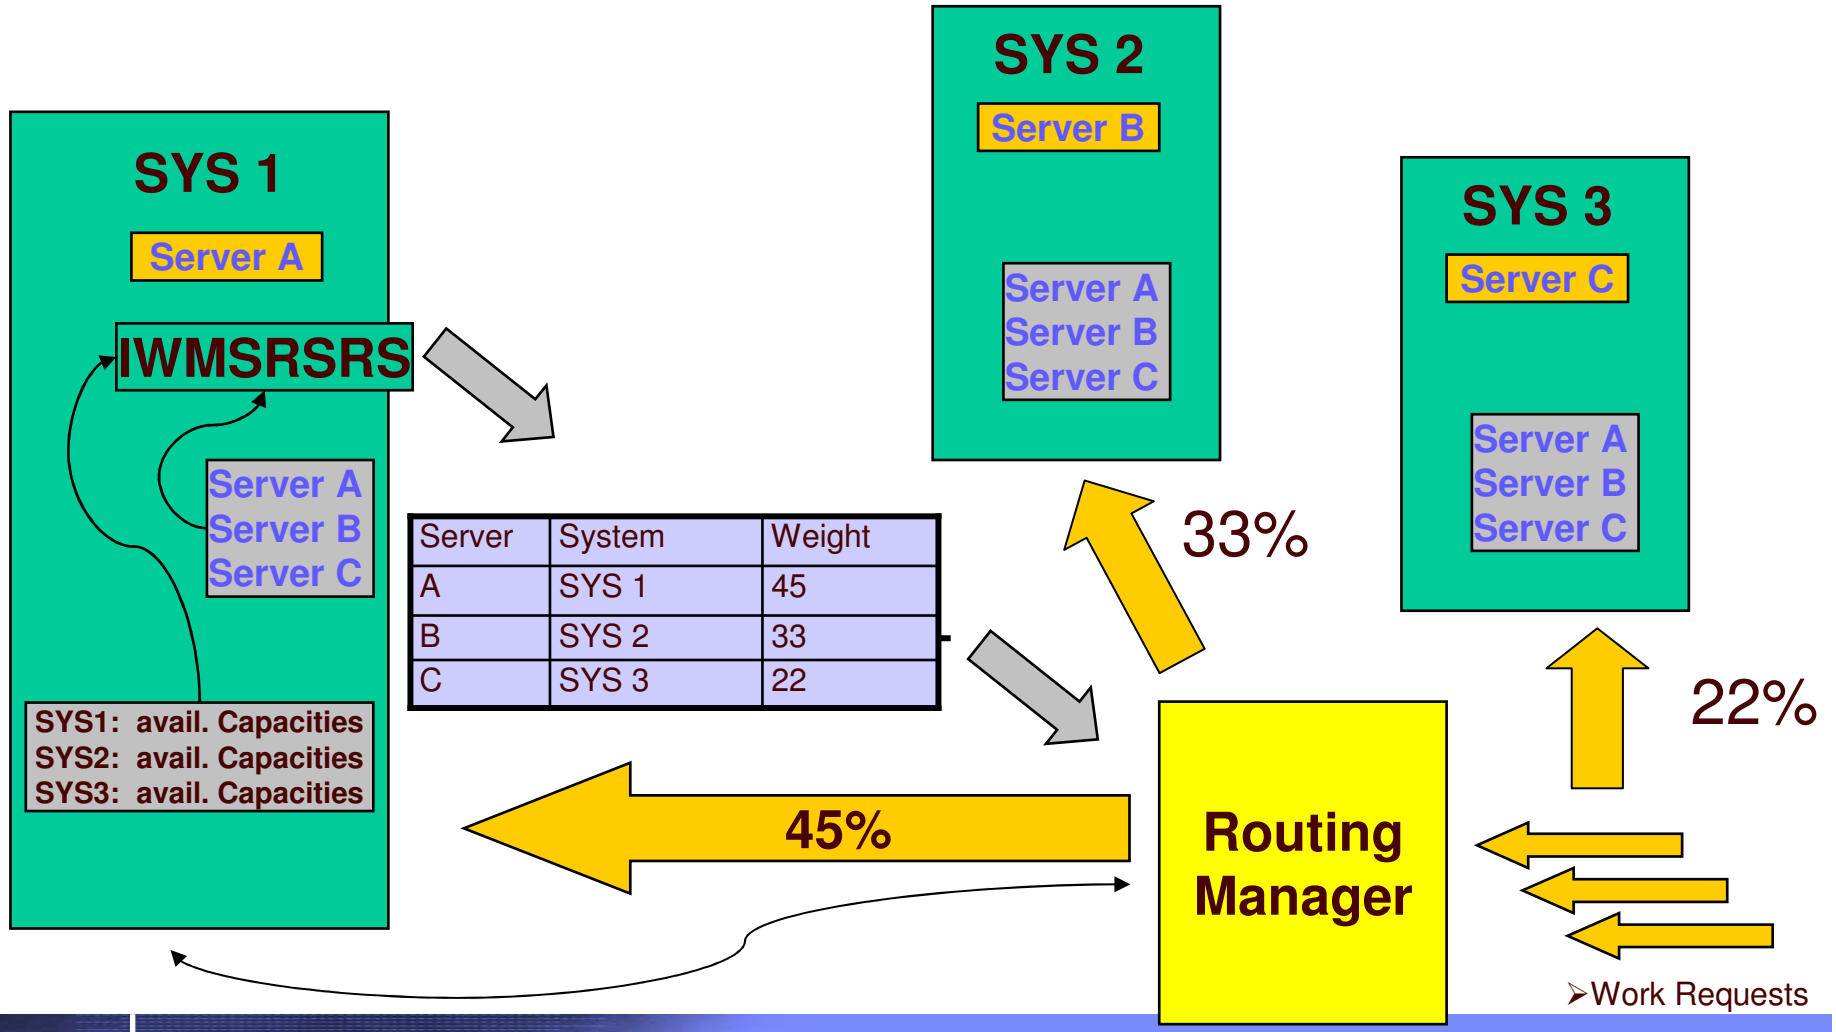

Overview WLM Routing Service Today

Select Importance Level that consumes > 5% of System CPU Capacity at least at one System (SUs = Service Units)

Calculate System Weight for each System:

$$\text{System weight} = \frac{\text{SUs for this system at selected level} * 64}{\text{total SUs for all systems at selected level}}$$

Calculate Server Weights:

$$\text{Server weight} = \frac{\text{system weight}}{\text{\# of servers on system}}$$

Overview

WLM Routing Service Today

Example:

Two Servers per System
Selected Level: 5

Sum:

200+400+300 = 900

	System 1		System 2		System 3	
Level	SUs	%	SUs	%	SUs	%
0	2000	100	2000	100	2000	100
1	1800	90	1900	95	1840	92
2	1600	80	1500	75	1600	80
3	1100	55	1500	75	800	40
4	400	20	1200	60	800	40
5	200	10	400	20	300	15
6	80	4	20	1	0	0
7	0	0	0	0	0	0

System 1 weight = 200 *64 / 900 = 14 **Server weight** = 14 / 2 = **7**

System 2 weight = 400 *64 / 900 = 28 **Server weight** = 28 / 2 = **14**

System 3 weight = 300 *64 / 900 = 21 **Server weight** = 21 / 2 = **10**

Weakness of WLM Routing Service Today:

The routing recommendation – the weight – for a server is only based on the available capacity of the LPAR

This can result in unexpected results:

- If the LPAR is loaded with low important work, but the server has high importance
  –> weight too low
- If the LPAR is low utilized, but the server has a bad performance index or is waiting for resources
  –> weight too high

Solution

Two new server specific routing services are made available with z/OS 1.7:

- **IMW4SRSC**
- **Function SPECIFIC for IWMSRSRS**

1.) New routing service IWM4SRSC

- Returns routing recommendation for one given server
- LPAR scope
- Used by Communication Server for z/OS 1.7
- No registration of servers necessary

2.) New Function code "SPECIFIC“ in old routing service IWMSRSRS

- Returns routing recommendations for all registered servers
- Each recommendation is server specific
- Syplex scope
- Used by DB2 V.8 (thru APAR to be announced)

New routing service IWM4SRSC

Weight calculation: Product of two factors, scaled by 64

The PI Factor (Performance Indicator Factor)

This gives an indication of how good this server, with respect to the work that is related to this server, is achieving its goals as defined in the active WLM policy.

The Importance factor

This is a measurement of how much CPU Capacity is displaceable by work of the server's importance, with respect to the work that is related to this server

New Function code "SPECIFIC“ in routing service IWMSRSRS

Weight calculation: Product of three factors

1.) System Utilization Factor:

Same as the resulting system weight for old "SELECT“ function (as described in the earlier foils)

2.) PI Factor

This gives an indication of how good this server, with respect to the work that is related to this server, is achieving its goals as defined in the active WLM policy.

New Function code "SPECIFIC“ in routing service IWMSRSRS

Weight calculation: Product of three factors (cont.)

3.) Queue Time Ratio:

If the server owns independent enclaves, the ratio of queue time to elapsed time of those enclaves

If more than one server is registered on the same system, the weight is divided by the number of those servers.

Interactions & Dependencies

- Hardware
  - no specific Hardware dependencies
- Software
  - IMW4SRSC can be used on any system with z/OS 1.7
  - IWMSRSRS with function SPECIFIC can be called when all servers registered under the given LOCATION run on systems with z/OS 1.7.
  - The services can be called as assembler macros by any application program
  - Current exploiters are Communication Server for z/OS 1.7 and DB/2 V.8

Support of Dynamic Processor Speed Changes

- Technically: Speed of CPs is changed to increase/decrease system performance
- Other types of processors like zAAP are not affected by speed change
- Speed is changed by applying a new LICCC record
- The HW currently supporting this function is z890
- A speed change dynamically adjusts all timing and performance related numbers (e.g. RMCTADJC)
- No IPL required when speed is changed
- z/VM 5.1 supports dynamic processor speed changes

Support of Dynamic Processor Speed Changes

- z/OS has no user interface for dynamic processor speed changes
- When a speed change is signaled to z/OS, SRM/WLM will do the following:
  1. Adjust all timing and performance related numbers
  2. Initiate a WLM policy reactivation
  3. Issue a new message

     IWM063I WLM POLICY WAS REFRESHED DUE TO A PROCESSOR SPEED CHANGE

Interactions & Dependencies

- Hardware
  - The hardware needs to support this new functionality. Currently z890 provides this hardware capability
- Software
  - Support of dynamic processor speed changes is part of z/OS V1R7
  - For z/OS V1R4 and above (HBB7707, JBB7717, HBB7708, HBB7709) this function is provided via APAR OA07510
    - OA07510 became available 11/04

Migration/Coexistence Considerations

- No particular migration / coexistence rules need to be observed
- When dynamic processor speed changes are intended, this new support should be installed on z/OS on all LPARs
- If this new support is not installed on z/OS, IPLing the system will set the proper timing and performance values as well

Installation

- This new support is part of z/OS V1R7 and no installation is required

- If this support is desired for releases V1R4 up to V1R6, APAR OA07510 is to be applied

Session Summary

- Workload Balancing by using old routing service took only the behavior of the whole system into account
- New Routing services take server specific behavior into account
- Additional indicators for routing recommendations: Performance index, Importance of server, Queue Time of owned Enclaves
- One service to get recommenation for one server, IWM4SRSC
- One service for sysplex wide registered servers, IWMSRSRS with function=SPECIFIC
- More precise workload balancing, less unexpected results
- This new z/OS V1R7 line item ‘Support of dynamic processor speed changes’ allows the processor speed to be changed while the system is running
- No IPL is required, SRM WLM will update performance and timing metrics
- Dynamic processor speed change needs appropriate hardware support. Currently z890 supports this new functionality.