


From Zero to z/Hero



Workshops 2012

CICSplex Workload Management



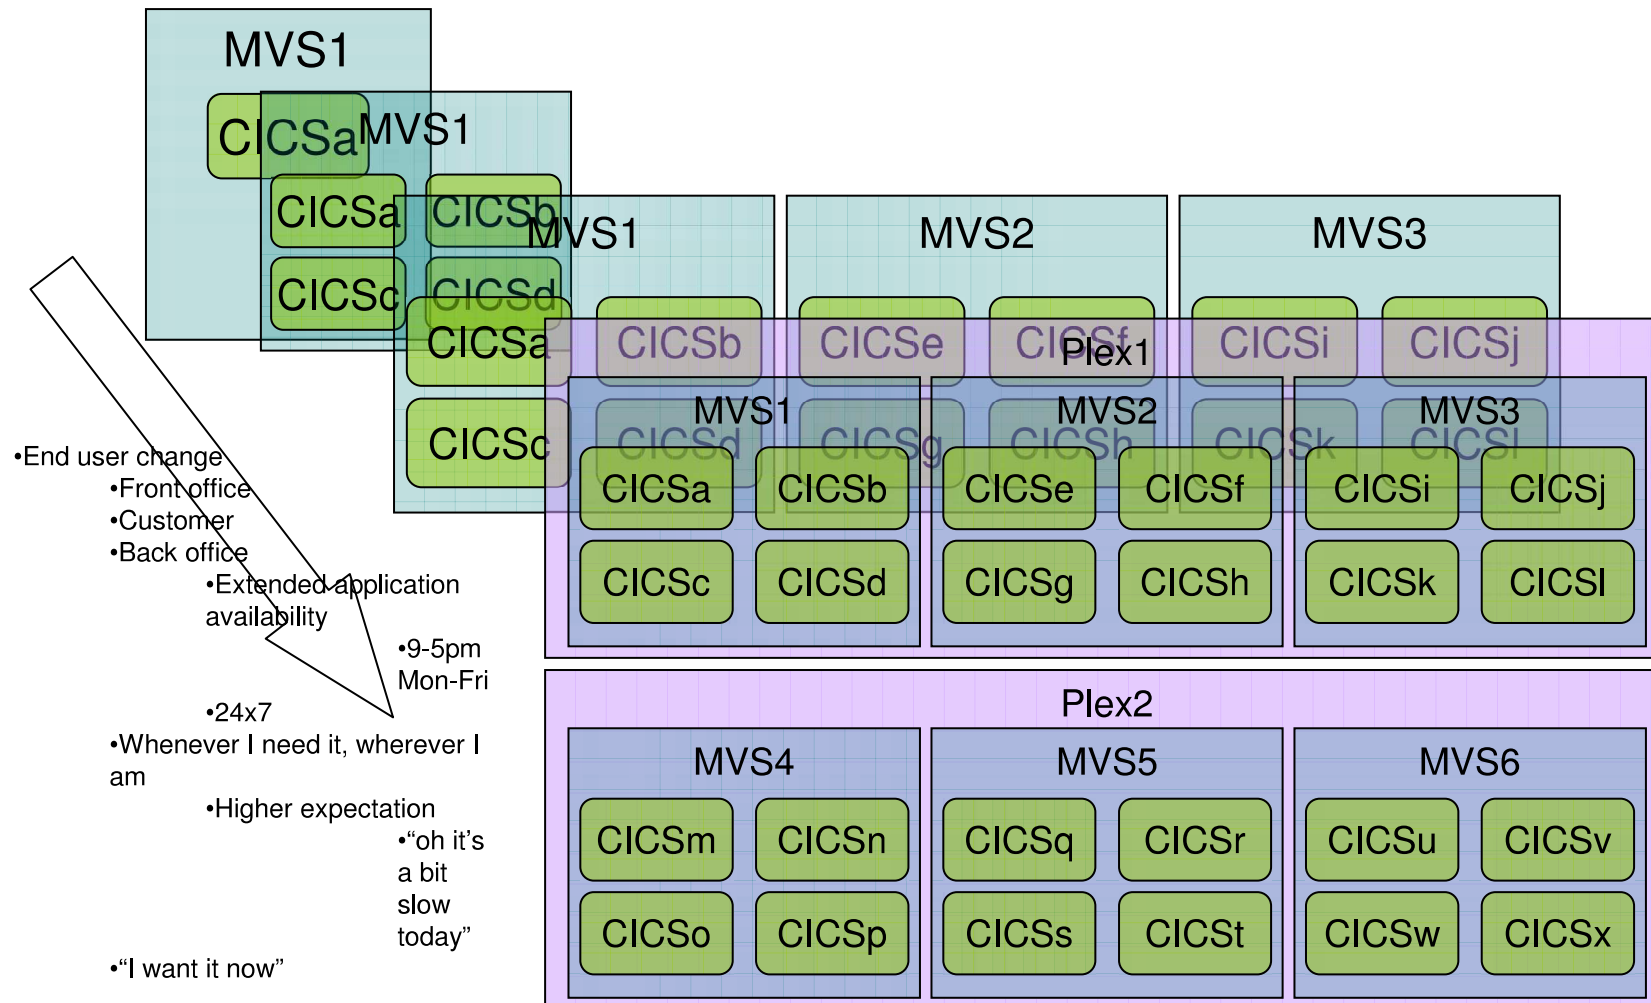
Marcel Amrein
IBM Deutschland Software Group
marcel.amrein@de.ibm.com

Agenda und “Disclaimer”

Diese Präsentation entspricht weitgehend der Session
“Dynamic Workload Management With CICSplex SM”
von Paul Johnson, Hursley Lab, bei der IMPACT 2009 in Las Vegas

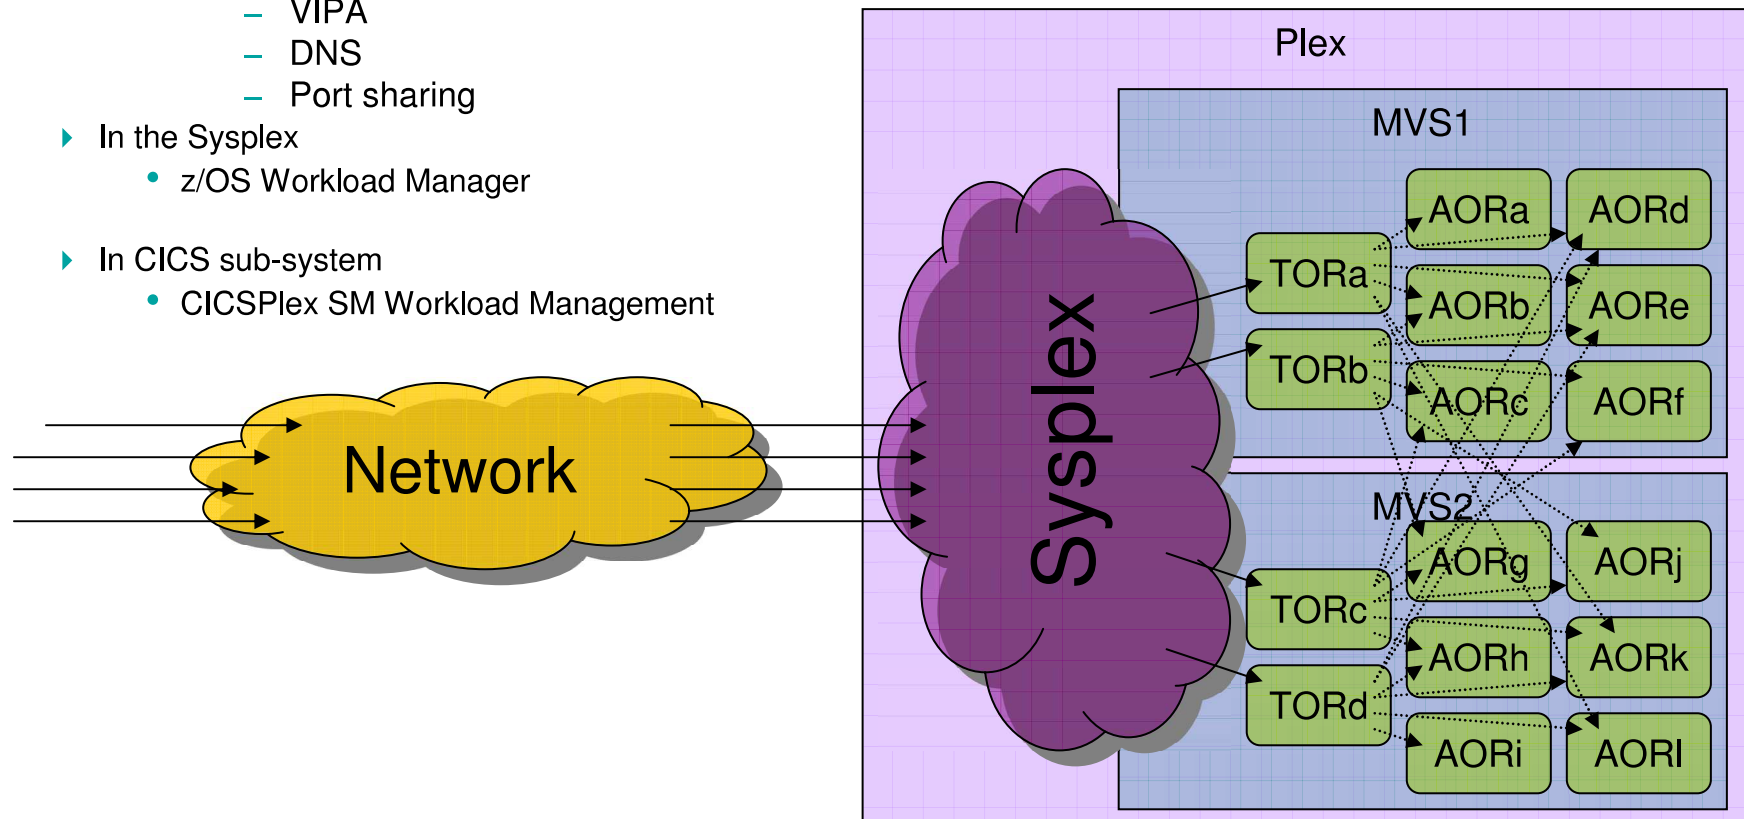
- Warum und wozu Workload Management?
- z/OS WLM und CICS
- die Mechanismen der CPSM WLM Funktion
 - Health Check
 - Workload
 - Workload Balancing und – Separation
- CPSM WLM definierbare Objekte und Direktiven
 - WLM Specification
 - WLM Definition

What's the problem?



What's the solution?

- Workload Management ...
 - ▶ In the Network
 - VTAM Generic Resource
 - TCP/IP load balancing
 - Sysplex distributor
 - VIPA
 - DNS
 - Port sharing
 - ▶ In the Sysplex
 - z/OS Workload Manager
 - ▶ In CICS sub-system
 - CICSplex SM Workload Management



Workload management in the sysplex z/OS WLM overview

- Works with z/OS SRM (System Resource Manager)
- Dynamically allocates resources
 - ▶ Processor Storage, IO Priority
- Service Definition
 - ▶ Contains 1 or more Service Policies
- Policy defined in terms of “Goals”, not resources
 - ▶ Service Policy
 - One per Sysplex
 - Can switch Policy when required
 - ▶ Service Class
 - describes performance objectives for part of workload
- Report Class
- Goals
 - ▶ Response time – How quick to run work
 - Average Response Time
 - E.g. Average transaction response should be 2 seconds
 - Percentile Response Time
 - E.g. 95% of transactions should complete in 3 seconds
 - ▶ Velocity
 - How fast work should run as a % of the time it's ready
 - High Value – When work is ready, run it quickly
 - Low Value – When work is ready, can wait to run
 - Used to get CICS active
 - ▶ Discretionary – Work with no goals
- Classification Rules
 - ▶ How work get classified to Service Class
 - ▶ Some Classification Rules available for CICS:
 - SI – SubSystem (Applid)
 - UI – Userid
 - TN – Transaction Name
 - LU – LUName
- Goal Importance
 - ▶ Associated with Service Class
 - ▶ 1-5 : 1 = Very important, 5 = Not important
- Performance Index (PI)
 - ▶ $PI = \text{averageResponseTime} / \text{Goal}$
 - ▶ Used to compare Goals
 - ▶ How well is work meeting it's Goal
 - $PI = 1$: Meeting the Goal
 - $PI < 1$: Over-achiever (beating the goal)
 - $PI > 1$: Under-achiever (Not meeting the goal)

z/OS WLM and CICS

- When CICS starts, as job or Started Task,
 - ▶ uses JES or STC classifications
- Once CICS is active,
 - ▶ CICS connects to z/WLM
 - ▶ CICS Sub-System related Service Classes are used
- At CICS startup, CICS will allocate a Pool of Performance blocks based on Maxtasks
 - ▶ Sampled at regular intervals, e.g. every 250ms
 - ▶ Communicates transaction state to z/OS WLM
- When a Task starts:
 - ▶ Performance block associated with Transaction
 - ▶ Transaction Classified (New, or passed via MRO)
 - ▶ Set-up Performance Block
- A Task executes
 - ▶ Dispatcher updates Performance Block
- Task ends:
 - ▶ Report transaction Complete or ...
 - ▶ Notify of partial transaction complete (e.g. in AOR via MRO)
- z/OS WLM controls which Requesting (TOR) CICS Regions receive the work
- z/OS WLM can also affect which AOR is chosen when using CICSplex SM
- Once in CICS routed under CICSplex SM criteria
 - ▶ Regions in bigger/more powerful LPARs may be favoured
 - ▶ CICSplex SM will favour local routing
 - ▶ Can give the impression of an “Unbalanced” Workload
- ▶ In reality more work goes to the System(s) best able to deal with the work
 - An uneven distribution may be an accurate reflection of the capabilities of the environment
- ▶ CICSplex SM uses z/OS WLM ‘Average Transaction Response Times’ to influence the ‘LOAD’

What can CICS do?

- Dynamic Routing
 - ▶ Transactions associated with a terminal
 - Physically 'entered' at a terminal
 - Via an EXEC CICS START
TERMID command
 - ▶ Dynamic Program Link (DPL)
 - EXEC CICS LINK PROGRAM
command
- Distributed Routing
 - ▶ Non-Terminal initiated STARTs (NTIS)
 - EXEC CICS START with no
TERMID
 - ▶ CICS BTS
 - EXEC CICS RUN
ACTIVITY(name)
<TRANID(abcd)> ASYNCH
 - ▶ Enterprise Java Beans (EJBs) & IIOP
 - ▶ Inbound WebServices
 - if modified to make them
routable
- CICS uses 'User Replaceable' Modules
 - ▶ One for Dynamic Routing
 - ▶ One for Distributed Routing
- CICS ...
 - ... calls the appropriate program to
decide where to route to
 - ... passes the program the DFHDYPDS
commarea
 - ... Routes based on the applid/sysid
provided by the URM
- CICS GLUEs (Route reject processing)
 - XICERES exit for dynamic STARTs
 - XPCERES exit for dynamic DPLs

When is the Dynamic Routing called?

- Dynamic Routing Program called for
 - ▶ Program
 - If defined:
DYNAMIC(YES)
 - If not defined:
Program not defined and the program is not autoinstalled
 - Don't forget to point the tranid to DFHMIRS!
 - ▶ Transaction
 - If defined:
 - DYNAMIC(YES)
 - If not defined:
 - SIT - DTRTRAN={CRTX|name}
- Distributed Routing program called for
 - ▶ Transaction Starts (Non-Terminal Initiated)
 - TRANSACTION defined ROUTABLE(YES)
 - ▶ CICS BTS
 - PROCESS or ACTIVITY started by RUN ASYNCHRONOUS
 - TRANSACTION associated with PROCESS or ACTIVITY has DYNAMIC(YES) and REMOTESYSTEM is blank
 - ▶ Enterprise Java Beans (EJBs)
 - REQUESTMODEL TRANSID defined ROUTABLE(YES)
 - ▶ Inbound WebServices
 - Target program defined DYNAMIC(YES)
 - A program in the PIPELINE changes contents of
 - Container DFHWS-USERID
 - Container DFHWS-TRANID & referenced TRANID defined DYNAMIC(YES)

So what part does CICSplex SM play?


- It supplies a program to act as both...
 - ... the Dynamic Routing Program
 - SIT - DTRPGM
 - SPI - SET SYSTEM DTRPROGRAM
 - ... the Distributed Routing Program
 - In SIT - DSRTPGM
 - SPI - SET SYSTEM DSRTPROGRAM
 - **EYU9XLOP**
- Once called...
 - ▶ CICSplex SM Workload Management facilities create a list of suitable, candidate, target CICS Regions, based on:
 - the transaction
 - the terminal id, luname, user id, or process type.
 - Note:
 - CICSplex SM does not do the routing,
 - CICS does
 - ▶ The list of candidate Target Regions is based upon the Workload to which the Requesting/Routing Region belongs.

What rules does it play by?

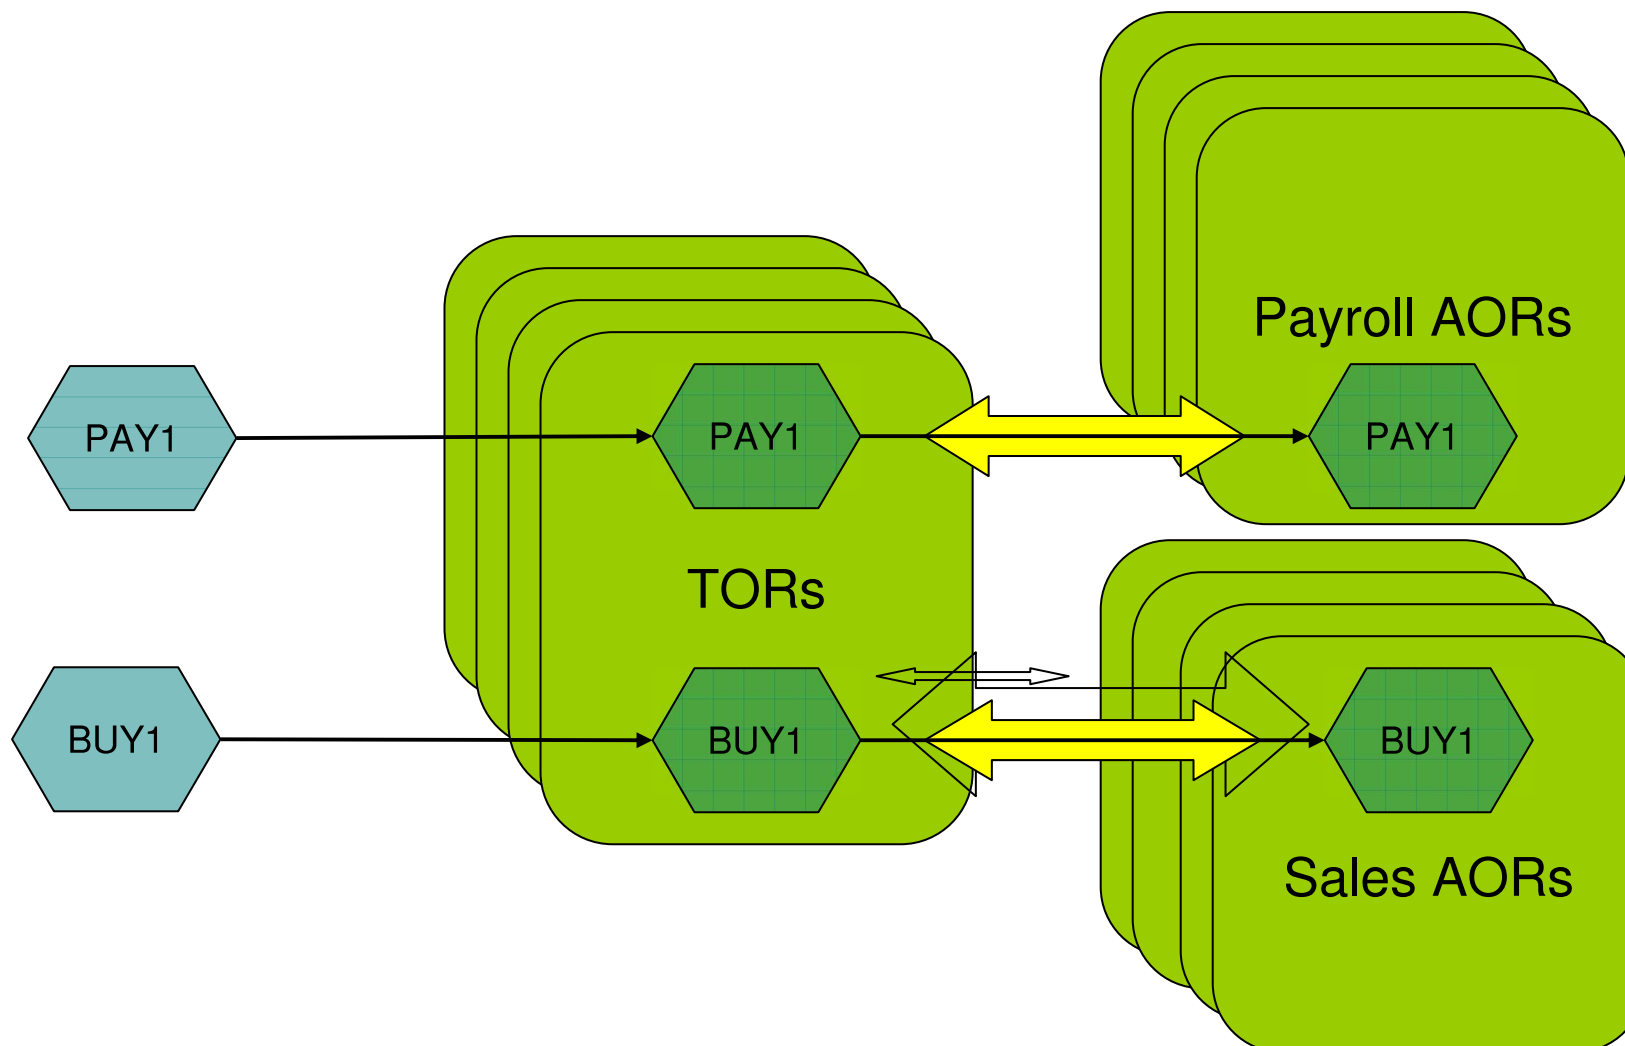
- CICSplex SM Workload Management can be divided the following
 - ▶ Workload Balancing
 - The process which decides which of the Target Regions is considered to be the most suitable, the 'best', candidate to route to, assuming:
 - Work 'could' be sent to any of the Target Regions
 - Work does not have an 'affinity' with a specific Region
 - ▶ Workload Separation
 - How CICSplex SM Workload Management allocates specific work
 - to specific set(s) of Target Regions
 - for specific reasons
 - ▶ Affinity management
 - Routing workloads which have existing affinities declared for them

Workload Balancing

- Workload Balancing is ...
 - ▶ Providing CICS with the 'best' Target Region, at the moment the request was made, from all of the possible candidates.
- ITS NOT EVENLY DISTRIBUTING WORK



Workload Separation



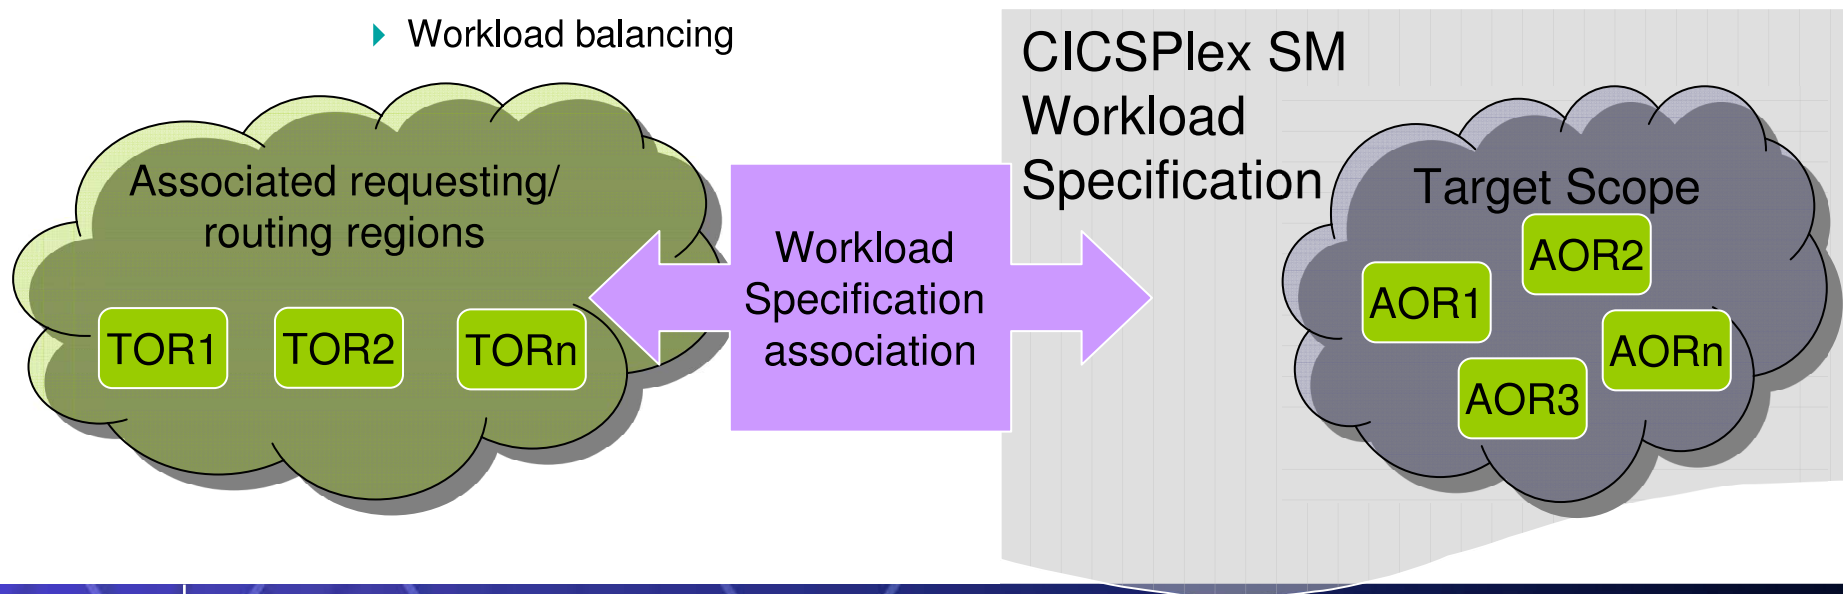
What are 'Affinities'?



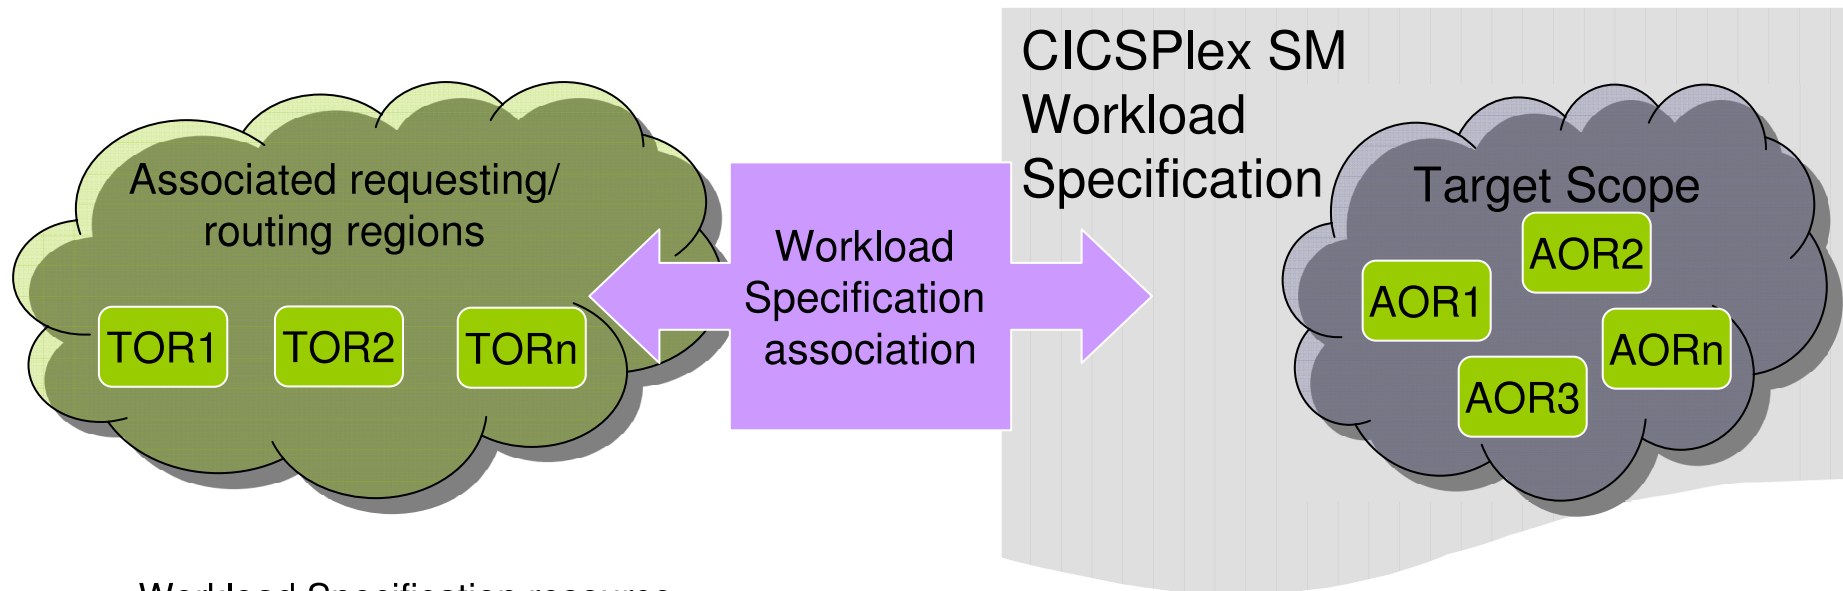
- A dependency to return to the same Target Region caused by a program leaving a requirement to return there, e.g. for local data
 - Affinity Relationships
 - ▶ USERID
 - ▶ LUsername
 - ▶ Global
 - ▶ BAPPL
 - Must be set if Affinities processing required
 - Each affinity relationship supports different lifetimes
- Affinity lasts for specified Lifetime
 - ▶ Lifetime specified on WLMSPEC (default) or Transaction Group
 - ▶ DELIMIT - Mode END (e.g. for a Menu transaction)
 - ▶ PCONV - Until end of pseudo-conv using EXEC CICS RETURN TRANSID
 - ▶ ACTIVITY/PROCES - Until CICS BTS Activity or Process is no longer active
 - ▶ SIGNON - Duration of a user Signon -> Signoff
 - ▶ LOGON - Duration of a VTAM logon -> Logoff
 - ▶ SYSTEM -Duration of the region (Avoid these!)
 - ▶ PERMANENT - Duration of the Workload (Avoid these!)

What is a CICSplex SM Workload?

- At it's most basic
 - ▶ One or more Requesting/Routing Regions
 - ▶ One or more Target Regions
 - ▶ One Workload Specification
 - ▶ Workload balancing



Workload Balancing – Workload Specification



- Workload Specification resource (WLMSPEC)...

- ▶ Defines a Workload

- The Workload must be “associated” with ...
 - ... Requesting/Routing Region(s) (Dynamic & Distributed)
 - ... Target Region(s) (Distributed only)
 - ... via systems (CSYSDEF) and system groups (CSYSGRP)

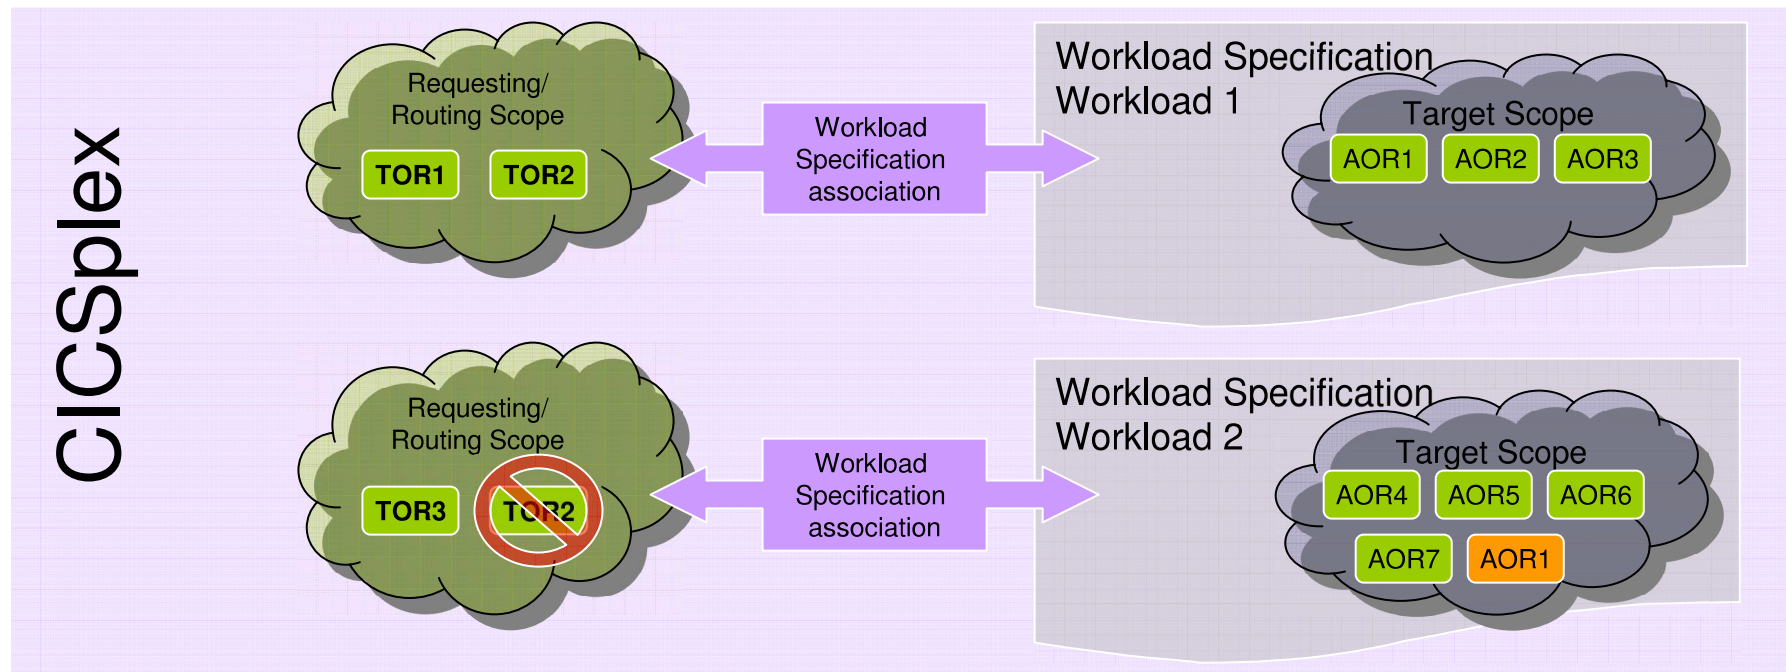
- ▶ Specifies default target scope

- a CICS Region (CSYSDEF)
 - a Group of CICS Regions (CSYSGRP)

- ▶ Specifies the algorithm type


- QUEUE or GOAL

Workload Balancing – Workload Specification ...



- ▶ There can be multiple Workloads within a CICSplex
- ▶ A Region can be a Target for multiple Workloads
- ▶ A Region can be BOTH a Requesting and a Target Region
- ▶ But... Requesting/Routing Region can ONLY be associated with **1** Workload
- ▶ While you might not be using distributed today, don't have your target regions as target scopes in more than one workload

Factors affecting routing




- 'Queue Algorithm'
 - ▶ $LOAD = \text{Current Tasks} / \text{Maxtasks}$
 - ▶ Also defined to CICSplex SM (via EYUPARMS)
 - WLMLOADCOUNT – for “Current Tasks” count
 - ALLQUEUED
 - All queue tasks waiting for MAXTASK or TRANCLASS
 - MXTQUEUED
 - Exclude tasks waiting for TRANCLASS
 - WLMLOADTHRSH
 - 1-100%
 - Threshold as a % value of LOAD before CICSplex SM will route off the local CEC (all other things ... Health, Link, Affinities ... being equal)
- 'GOAL' Algorithm
 - ▶ 'QUEUE' + MVS WLM average transaction response time goals



- Combination of ...
 - ▶ Target Region Health
 - SOS, Dumping, Maxtasks
 - CICSplex SM notified by CICS
 - Stalled
 - CICSplex SM Stall detection process
 - ▶ Target Region “Quiesced” from Workload
 - ▶ RTA Event associated with WLMSPEC/TRANGRP
 - ▶ Abend Compensation

Abend Compensation

- ▶ Used to help avoid an AOR if Transaction Abends
- ▶ Specified on WLMSPEC and TRANGRP definitions
- ▶ If handled, routing program doesn't know about it
 - Acceptable level of abend probability (Abend Health)
 - 0 – Abend Compensation OFF
 - 2-99% chance of abend, consider region unhealthy
 - Region considered unhealthy if this value is exceeded
 - Abend load threshold (Abend Load)
 - 0 – Required if Abend Health is 0
 - 1-98% chance of abend
 - Double Target Load when reached



Link Factors



- Type of Communications Link to Target...
 - No Link
 - Run Locally,
 - Target Region = Routing Region
 - MRO (same LPAR)
 - MRO (same SYSPLEX)
 - IPIC for Dynamic Program Link between CICS TS 3.2 Regions
 - APPC
- Modifying routing behaviour
 - ▶ Route reject processing
 - XICERES and XPCERES CICS Exits
 - ▶ Customisation of routing code (EYU9WRAM)

Queue

AOR1

Link : MRO/XM
 Load :55
 PAbnd(ABCD):2.0
 SOS :No
 Dump : No
 Stall :No
 RTA Severity :None

AOR2

Link : MRO/XM
 Load :60
 PAbnd(ABCD):6.0
 SOS :No
 Dump : No
 Stall :No
 RTA Severity :None

AOR3

Link : MRO/XM
 Load :70
 PAbnd(ABCD):0.0
 SOS :No
 Dump : No
 Stall :yes
 RTA Severity :None

AOR4

Link : ISC
 Load :80
 PAbnd(ABCD):0.0
 SOS :No
 Dump : No
 Stall :No
 RTA Severity :None

All regions maxtask 100; abend load=2.0; Abend health=6.0

$\text{weight} = (\text{link} * \text{load} * \text{Abnd} * 100) + \text{Health}$


$$(1.0 * 0.55 * 2.0 * 100) + 0 = 110$$

$$(1.0 * 0.7 * 1.0 * 100) + 1000 = 1070$$

$$(1.0 * 6 * 2000.0 * 100) + 0 = 140000$$

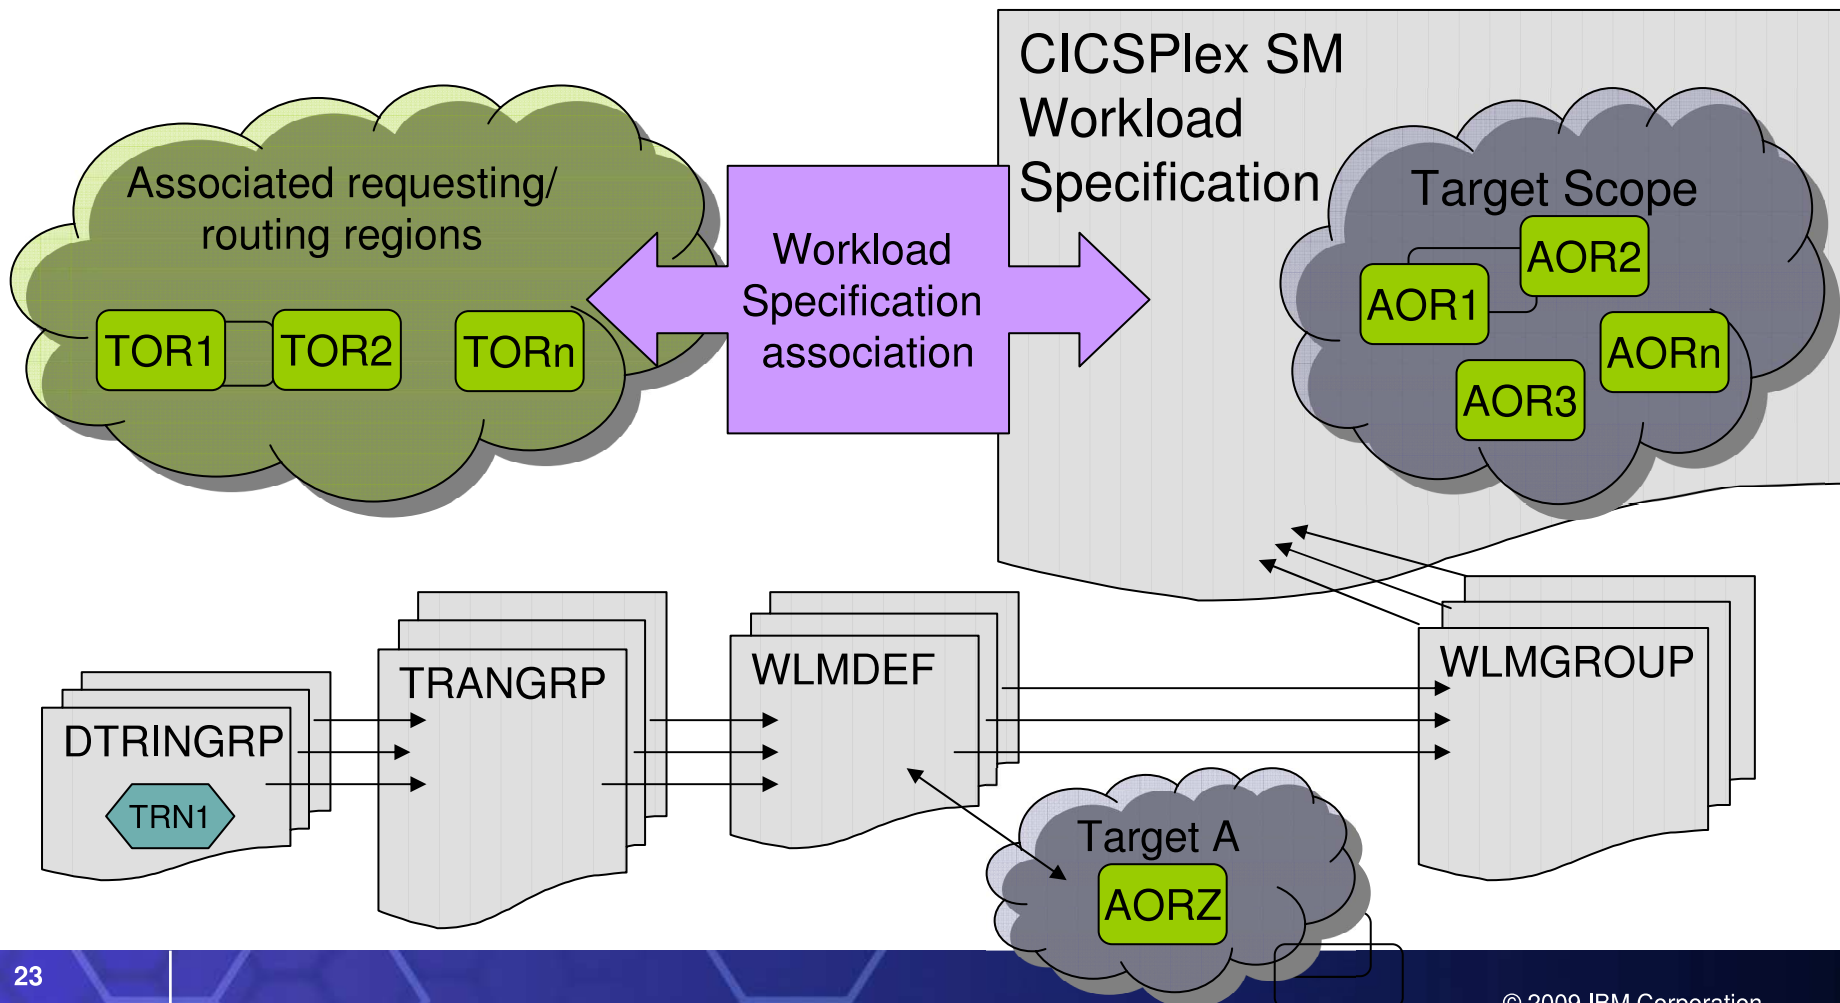
$$(1.3 * 0.8 * 1.0 * 100) + 0 = 104$$

VALUES ARE FOR ILLUSTRATIVE PURPOSES ONLY



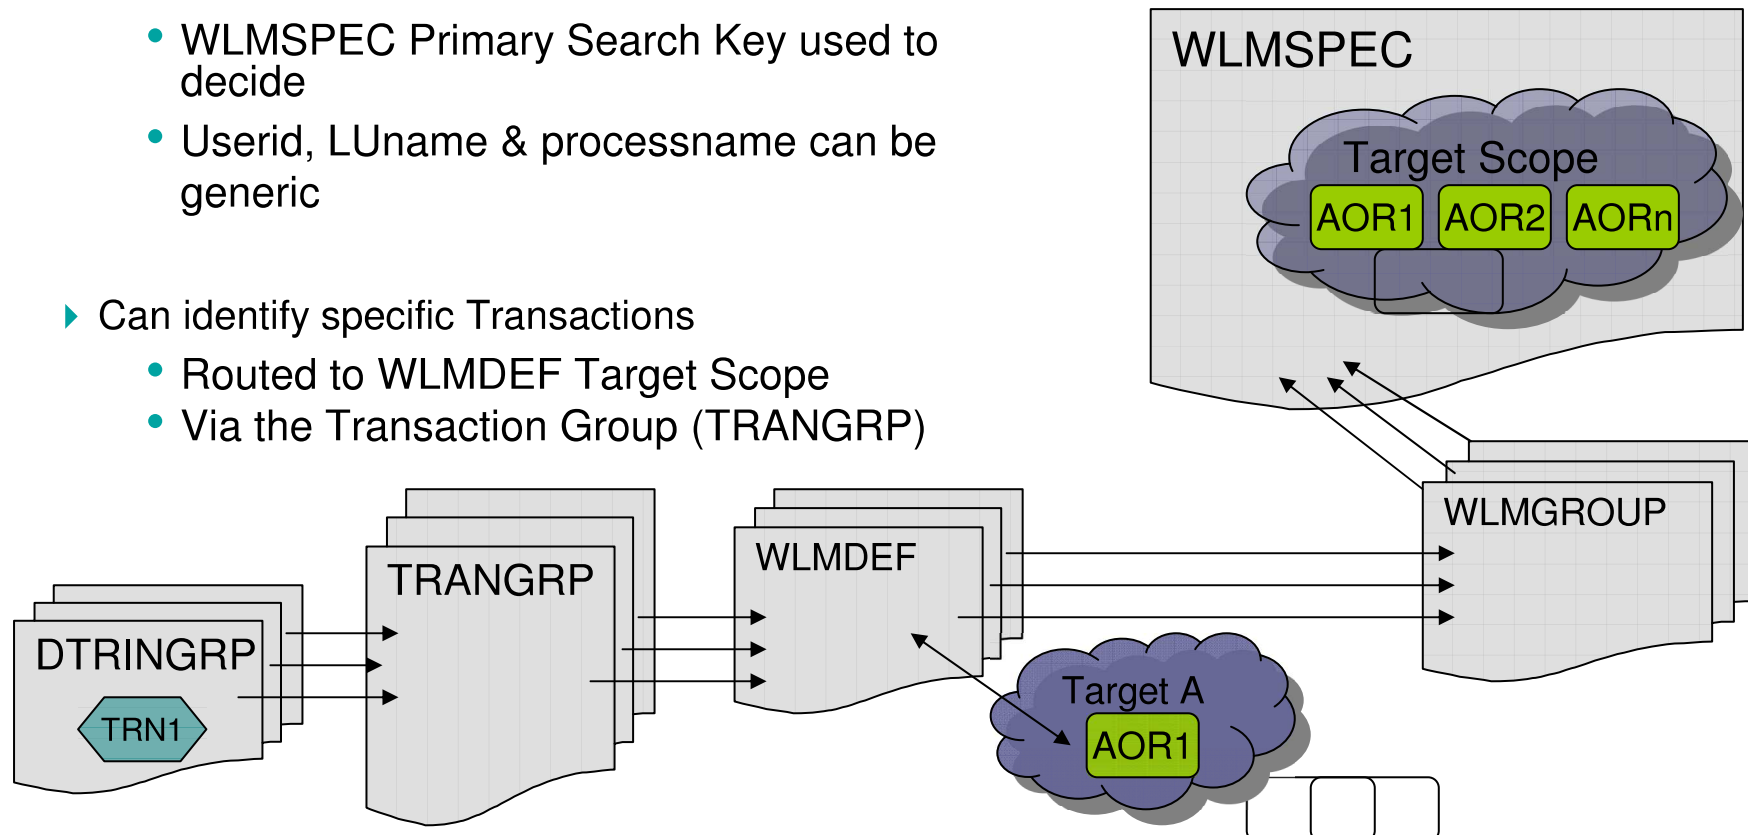
Workload Separation

- Routing specific parts of the work to different systems



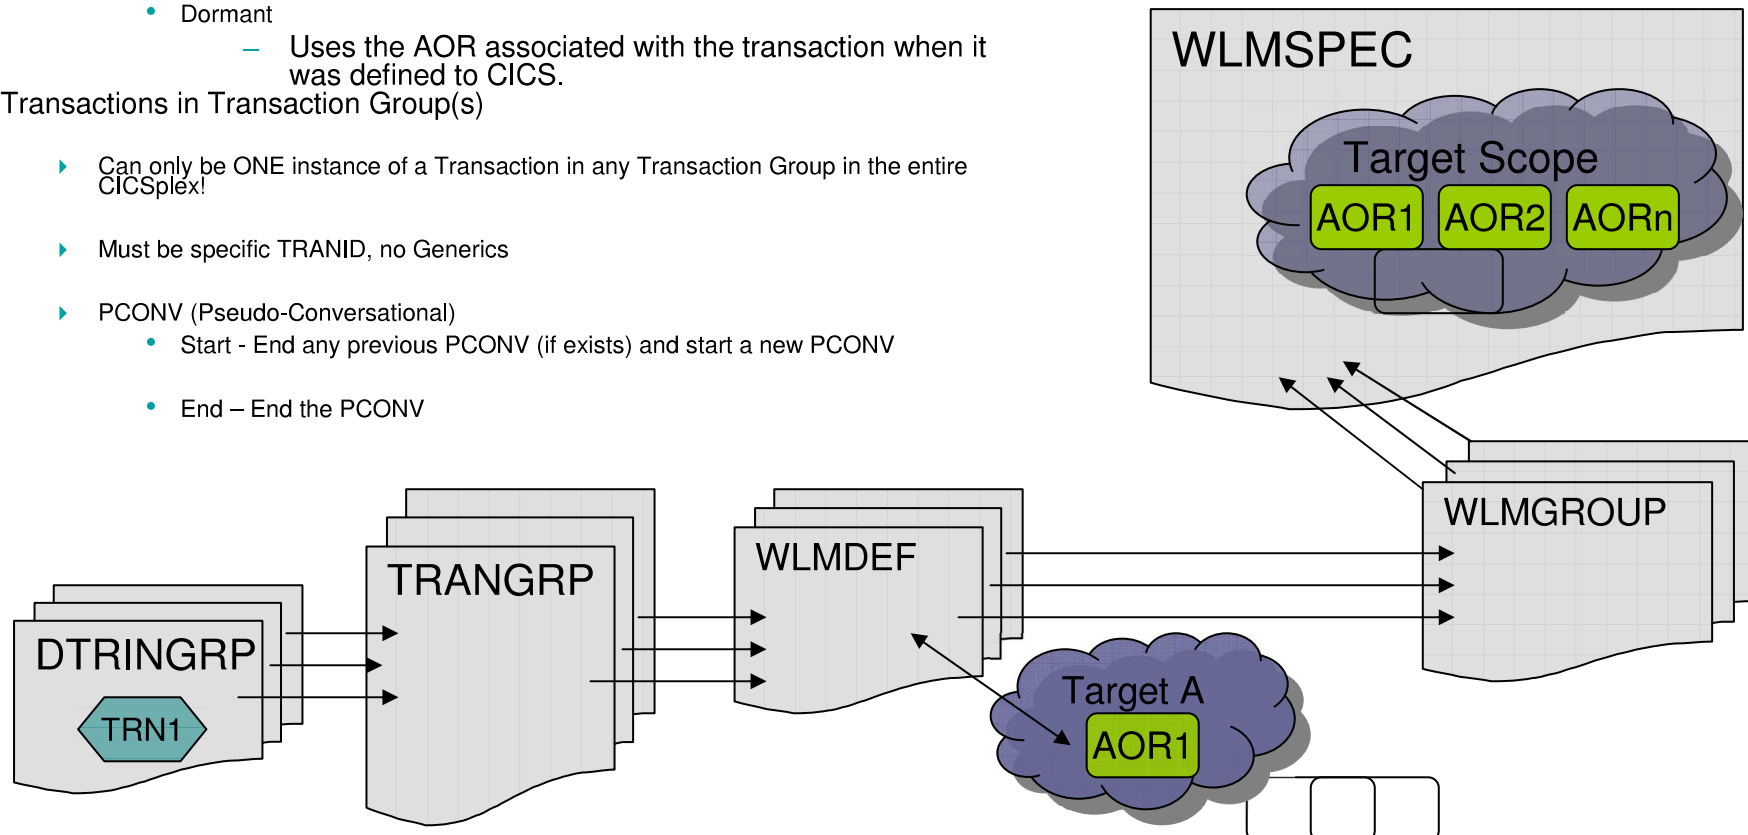
Workload definitions and groups

- WLMGROUPs are added to WLMSPECs
- WLMDEFs are added to WLMGROUPs
 - ▶ Specifies Target Scope
 - ▶ Based on Userid, LName or BTS Processname
 - WLMSPEC Primary Search Key used to decide
 - Userid, LName & processname can be generic
 - ▶ Can identify specific Transactions
 - Routed to WLMDEF Target Scope
 - Via the Transaction Group (TRANGRP)



Transaction Groups

- Transaction Group
 - ▶ Broadly same attributes as WLMSPEC
 - Match Key, Affinity Relation/Lifetime, Create Affinity, Abend Health/Load
 - ▶ Plus STATUS
 - Active
 - Use Affinity Settings
 - Dormant
 - Uses the AOR associated with the transaction when it was defined to CICS.
- Transactions in Transaction Group(s)
 - ▶ Can only be ONE instance of a Transaction in any Transaction Group in the entire CICSplex!
 - ▶ Must be specific TRANID, no Generics
 - ▶ PCONV (Pseudo-Conversational)
 - Start - End any previous PCONV (if exists) and start a new PCONV
 - End – End the PCONV



CPSM WLM Demo

- beruhend auf der FILEA (Uralt-) CICS Sample (IVP) Anwendung
- 1 TOR, 4 AORs
- reines Workload Balancing
- Workload Separation aufgrund Userid, Terminal ID
- Auswirkung eines “unhealthy” Region Status
- Affinitäten

Herzlichen Dank für Ihre
Aufmerksamkeit !

Feedback jeder Art gerne willkommen

bitte an:

marcel.amrein@de.ibm.com