

# Dynamic Workload Management using CICSplex System Manager


**Dan Donnelly**

Donnelly@uk.ibm.com

Session 4139A

## CICSplex SM Review

- Enterprise systems management for CICS Transaction Server via:
  - Web User Interface
  - Batch repository update
  - Application Programming Interface
- Single System Image & Single Point of Control for:
  - Resource definition
  - Operations
  - Monitoring
  - Threshold Analysis
  - Automation
  - Workload Management
  - System Management Integration




## Terminology

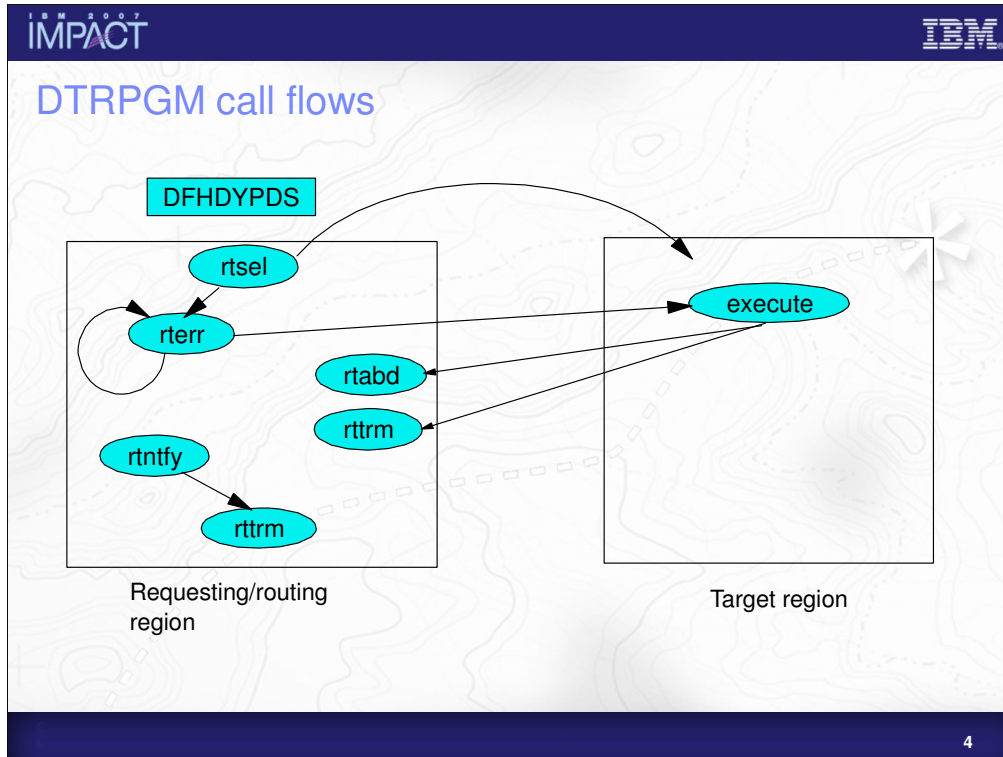
- Workload management

- Separation
- Balancing
- Affinity management

- Roles

- Requestor
  - Initiates routing request
- Router
  - Makes routing decision
- Target
  - Processes work





The DTRPGM SIT parameter specifies the dynamic routing program to be used for all requests that are NOT:

- non-terminal-related START requests
- BTS activities
- Enterprise beans

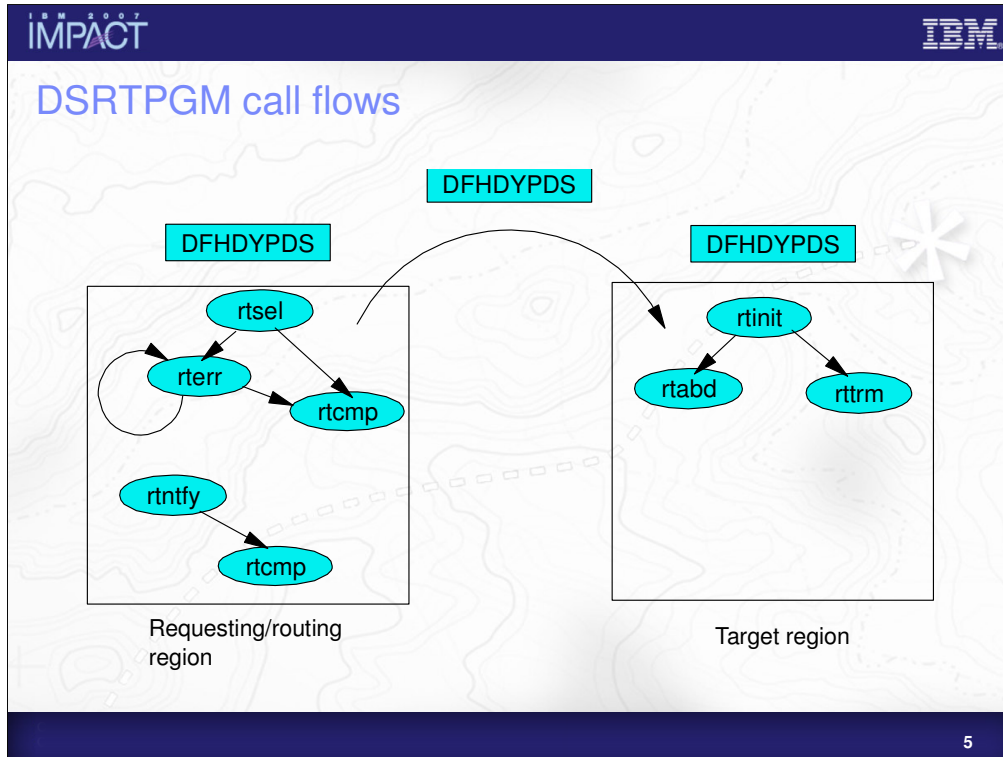
CICS passes routing requests to the program via a COMMAREA named DFHDYPDS.

The routing program first selects a target region (rtsel). This may result in either the program executing in the target region or in a Route Selection Error (rterr). This could occur when links between regions are down.

If we have a Route Selection Error then we try to select another region (depending on if we need to handle affinities). This will either result in another error or an execution.

If execution abends in the target region we return to the routing region with a Route Abend (rtabd). If execution completes without abending we return with a Route Termination (rtrm).

We may be notified of a transaction being statically routed (rtntfy). In this case we do nothing and simply call route termination (rtrm).



The DSRTPGM SIT parameter specifies the dynamic routing program to be used for all requests that ARE:

- non-terminal-related START requests
- BTS activities
- Enterprise beans

CICS passes routing requests to the program via a COMMAREA named DFHDYPDS.

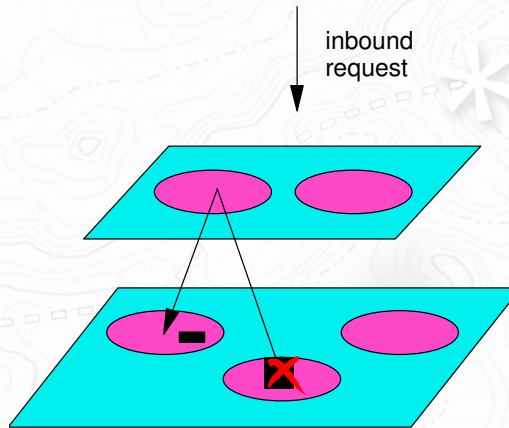
The routing region performs a Route Selection as before. If this fails then, like DTRPGM, the Route Error routine may be called. As before this routine may try and select another region.

If a target region is successfully selected the Route Complete (rncmp) routine is called.

Meanwhile the work is routed to the target region. The COMMAREA is passed on the routing request and the target region performs an rinit call which can result in the abend or terminate. As the COMMAREA is passed through to the target these calls can be performed in the target region.

## Affinities

- Inbound request routed
- Creates data in region
- Next route in different region
- Fails (hopefully!)



## Relations and Lifetimes

- **Relations**

- Userid – First use of user ID
- Luname – First use of Luname
- Global – First use of work unit
- Bappl – Business application

## Relations and Lifetimes

### ▪ Lifetimes

- Delimit – Maintain until end/start transaction
- Logon – Duration of VTAM logon
- Signon – Duration of CICS signon
- System – Whilst AOR is active
- Perm – Duration of workload
- Pconv – Duration of psuedo conversation
- Activity – Duration of CICS BTS activity
- Process – Duration of CICS BTS process

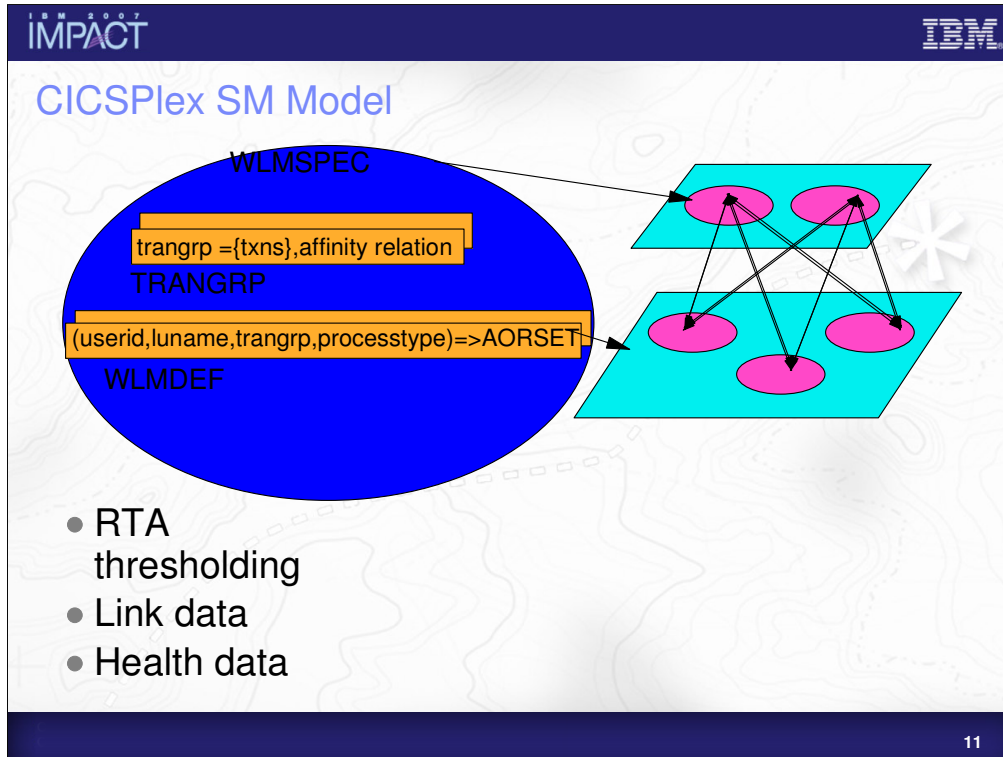


## Affinity determination

- CICS Interdependency Analyzer for z/OS (CICS IA)
- Run-time and batch reporting tool
- Captures resource relationships such as
  - which resources (Programs, Files, TSQs TDQs etc) are required by a transaction
  - what resources are no longer used
  - the sequencing of transactions within an application
  - Includes CICS, DB2, WebSphere MQ and IMS DB resources

## Affinity determination

- Captures affinity relationships
  - what programs have potential affinities due to the CICS API commands within them
  - which transactions have affinities to other transactions and the type and lifetimes of these affinities
  - which transactions have affinities to particular CICS regions
- Relationship data loaded onto a DB2 data base
- Ad hoc analysis
- Loadlib Scanner and VSAM File Reporter



Workload Management Specifications (WLMSPEC) – These allow you to specify the characteristics of a workload, such as: the workload algorithm, default affinity criteria and default target scopes for the workload. A WLMSPEC is associated with the routing regions for the workload.

Transaction Groups (TRANGRP) – These allow you to specify a set of transactions and any affinity information that should be associated with them.

Workload Management Definitions (WLMDEF) – In association with TRANGRPs these define workload separation and affinity requirements for a group of transactions. This is the finest level of control, i.e. if a transaction is covered by the WLMDEF the information in the WLMDEF overrides what is in the WLMSPEC. For example, if a WLMSPEC says we can route work to all regions in the CICSplex we can override that and route a transaction to only a subset of regions in the CICSplex.

## Health Indicators

- **Basic Health indicators**

- Maxtask
- SOS
- Dumping
- Stalled

- **Region indicators**

- Down
- Inactive
- Quiescing

### RTA event

- VHS/VLS
- HS/LS
- HW/LW

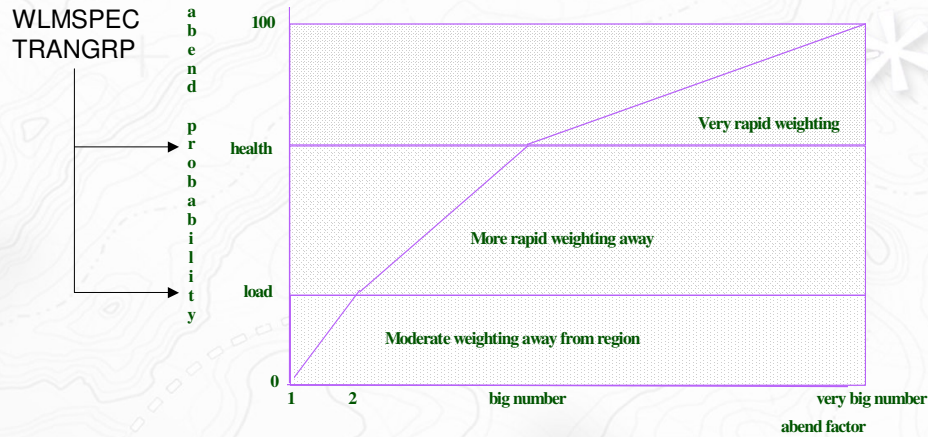
### Link indicators

- Link down
- Lost contact
- No link

### Abend health



## Abend factor calculation



If the abend probability is below the load value in the WLMSPEC increases, the abend factor increases at a base rate.

If it is above the load value then it increases at a more rapid rate. If it is above the health value it increases at an even more rapid rate.

This allows you to control how sensitive WLM is to changes in abend probability.

## Queue

- Calculate weight
  - $(\text{load} * \text{link} * \text{abend} * 100) + \text{health}$
- Load
  - $\text{curtasks} / \text{maxtask}$
- Link
  - 2 sets based on % load
    - local
    - MRO
    - MRO/XCF
    - ISC
- Tiebreaker
  - Toss a coin

## Queue example

**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

weight = (link\*load\*Abnd\*100) + 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*



## Goal Mode

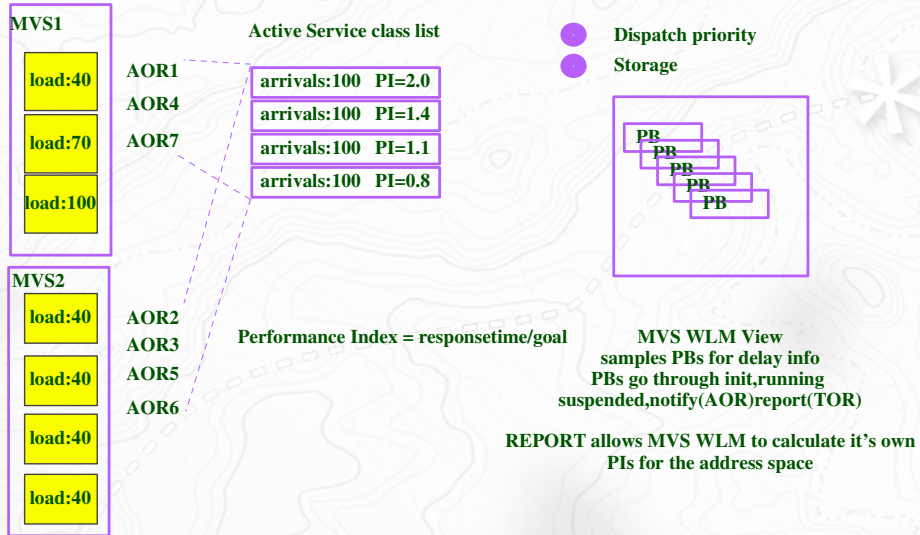
- Utilizes MVS WLM
  - You must define average response time goals to MVS WLM
  
- MVS WLM provides
  - Service definition (1 per sysplex)
    - Service policies (1 active)
      - Workload definitions
      - Service classes
      - Report classes
      - Goals (Velocity,discretionary,responsetime(average or percentile)
      - Classification for CICS is
      - (substype,subsysinstance,userid,tranid,trangrp,LU)=>service class



## Performance Indexes

- Performance index =  $\text{curAvg}/\text{goal}$
- Response time exponentially faded
  - $\text{newAvg} = \text{totalRespTime}/\text{totalCount}$
  - $\text{curAvg} = \text{curAvg} * \alpha + (1 - \alpha) * \text{newAvg}$
- Arrival rate exponentially faded
  - $\text{curRate} = \text{curRate} * \alpha + (1 - \alpha) * \text{newRate}$
- Service classes eliminated if no arrivals in an interval

# Goal Example

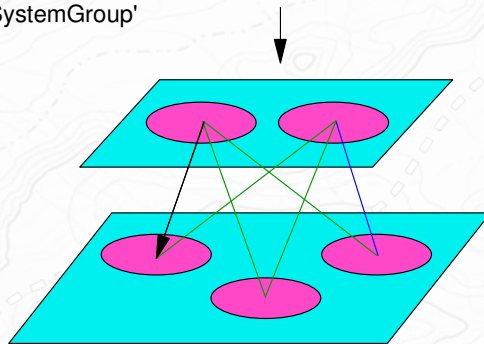


## A final word....

- **Balancing is probably a bad word**
  - Statistical in nature
  - Latency is allowed to offset performance cost
  - NOT round robin
  
- **Low transaction rates**
  - Serialised arrival rates
    - Basically random distribution into local AORS

## Dynamic transaction routing

trangrp(x) = {trandid}  
(userid,luname,trangrp,processtype)->SystemGroup  
Associated with SystemGroup'




trandid(abcd)  
dynamic(yes)

trandid(abcd)  
dynamic(no)


## Start TERMID

trangrp(x) = {trandid}  
(userid,luname,trangrp,processtype)->SystemGroup  
Associated with SystemGroup'



# Start

trangrp(x) = {tranid}  
(userid,luname,trangrp,processtype)->SystemGroup  
Associated with SystemGroup




## CICS BTS

```

trangrp(x) = {trandid}
(userid,luname,trangrp,processtype)->SystemGroup
Associated with SystemGroup'
    
```

```

EXEC CICS RUN
ACTIVITY(name)
<TRANID(abcd)>
ASYNCH
    
```




```

define tranid(abcd)
dynamic(yes)
routable(yes)
    
```

## Dynamic DPL - inbound

trangrp(x) = {trandid}  
 (userid,luname,trangrp,processtype)->SystemGroup  
 Associated with SystemGroup'

Web, EXCI, ECI  
 specifying prog(xyz)  
 <trandid(abcd)>



```
define program(xyz)
  <trandid(abcd)>
  dynamic(yes)
```


```
define program(xyz)
  dynamic(no)
```

```
define trandid(abcd)
  program(dfhmirs)
```

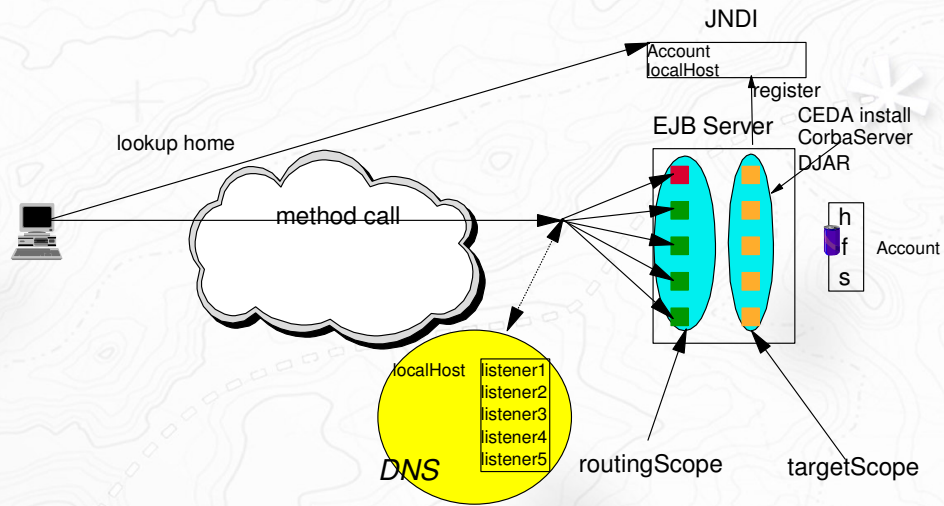


## Peer to Peer Dynamic DPL

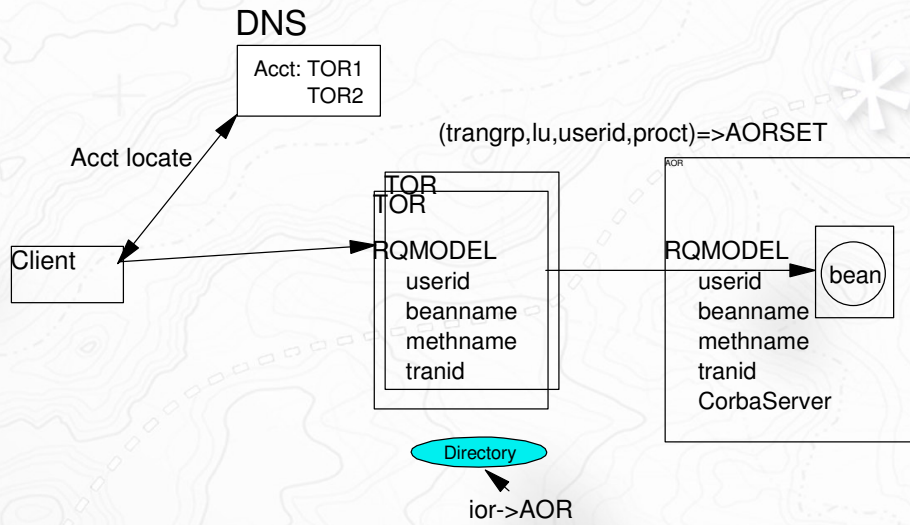
trangrp(x) = {trandid}  
 (userid,luname,trangrp,processtype)->SystemGroup  
 Associated with SystemGroup'



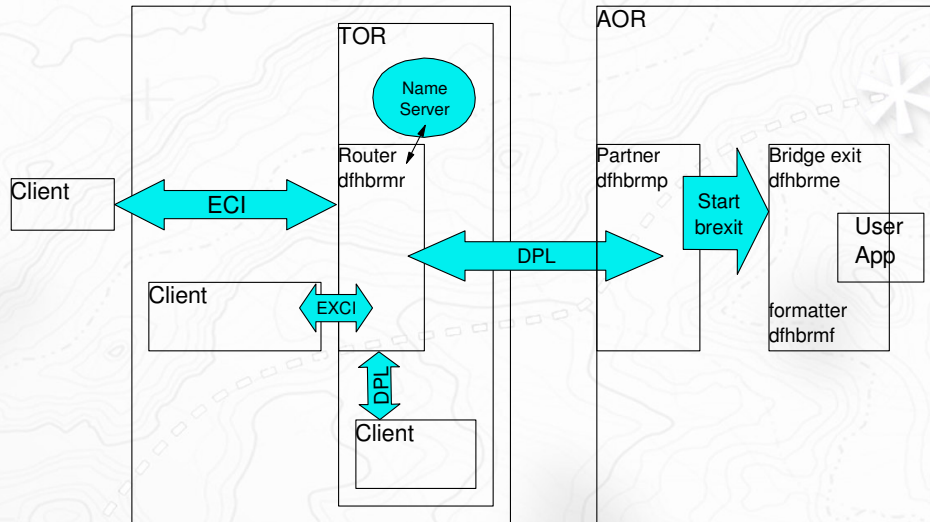
# Dynamic WLM of EJBs



# EJB target determination



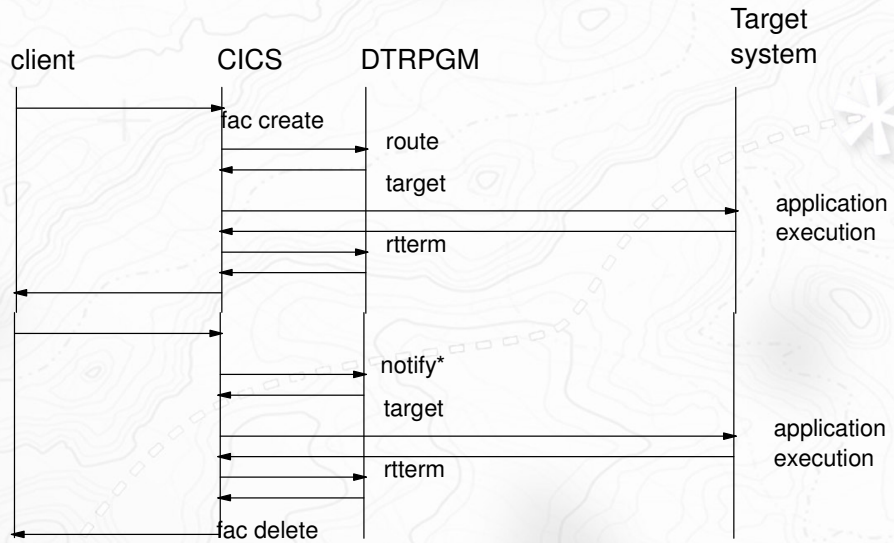
### Link3270 workloads



## Workload separation and affinities

- Separation allowed
  - (userid,transgrp,luname,processtype)=>targets
    - Processtype not relevant
    - LUname only when LU predictable
      - User exit - OK
      - Normal execution – unpredictable
  
- BRIDGE affinities managed by CICS code
  - CICSplex SM only called for
    - route select
    - route notify

### Typical flow



## Route reject processing

- Target regions can reject routing request
- Provided via user exits
  - XICERES exit for dynamic STARTs
  - XPCERES exit for dynamic DPLs
  - If rejected, router tries another target.
- Typical usage
  - resource availability in region e.g. files
  - Resource status e.g. disabled transaction

## What it can't cope with

- **Soft failures**
  - Some support via DB2
  - Failure returned to TOR who processes it as route error
  - Detectable via TIOA, COMMAREA, TCTUA then customise WRAM
  
- **Killer transactions**
  - If you think you have them
    - use workload separation to limit affect
    - use storage isolation if overwrite problems



## Exit Customization – Why?

- You know your transactions!
  - Examination of TIOA
  - Examination of COMMAREA
- Modification of TIOA data via EXEC CICS RETURN INPUTMSG
- Direct modification of COMMAREA
- Using above on invocation/return to control affinity dynamically via
  - CREAFF
  - DELAFF

## Routing data outside routing exit

- Dynamic routing information from CICS environment
- Invoked via EXEC CICS LINK PROGRAM(EYU9XOP)  
COMMAREA(EYURWTRA)
  - EYURWTRA contains
    - DFHDYPDS data
    - Simulated transaction info
    - Workload classification data
    - Returned applid, sysid and API fields
    - No support for Signoff/Logoff affinity
    - Will not work in goal mode

## Operational control

- Can display
  - Active workloads
  - Active routing regions
  - Active target regions
    - can quiesce targets
  - Active WLM definitions
    - can discard - look for affinities first
  - Active affinities
    - can discard - be aware of consequences!
  - Transaction and TRANGRPs associated with workload

## How to do it

- Applications already exist
  - with affinities which you don't know about
  - Lucky if you have the source
    - Luckier still if you have someone to change them to remove affinities
  - They must continue to execute during the transition to dynamic routing

## Understand your system configuration

- Understand your configuration
  - Operational procedures
  - failure procedures, backup procedures
- Design your new configuration
  - If possible add new regions for dynamic routing
  - Easier to diagnose problems
  - If you have “spaghetti” in combined regions, move towards TOR/AOR STATIC routing first

## Understand your applications

- Absolutely essential
  - Use affinities and interdependency utilities to identify relationships of parts
  - Try and move an application at a time
    - 'Baby steps' to get some success
  - Everything else being equal
    - Move shortest lifetime affinities first
      - balanced more frequently
    - Eliminate longest lifetime affinities first
      - balanced least frequently
    - Eliminate affinities that cause broadcasts
      - Reduced risk to CMAS failure
      - better performance

## Getting cute on performance

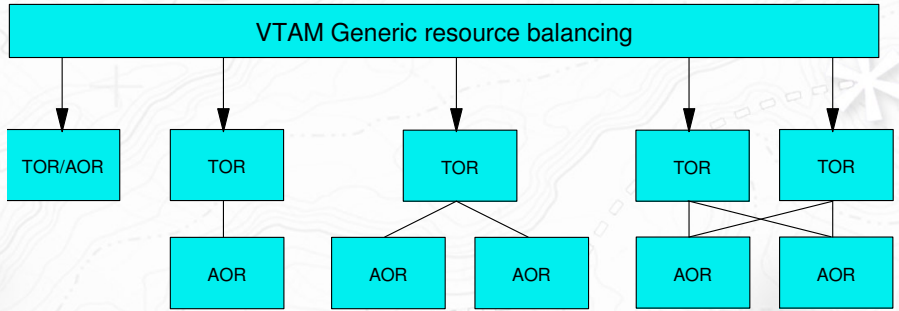
- Some affinities are worse than others
- Affinity Creation:
  - Broadcast affinities
    - lock obtained across all CMASes involved in workload.
    - If any such CMAS is down,
      - lock cannot be obtained so affinity cannot be created
      - transaction instance cannot execute.
    - If affinity exists, then can execute in absence of CMAS.
      - pconv, logon, signon, delimit must be same incarnation of AOR

## Getting cute on performance

| Broadcast Affinity | Permanent | System | Logon | Signon | Delimit | Pconv | activity | process |
|--------------------|-----------|--------|-------|--------|---------|-------|----------|---------|
| Global             | Yes       | Yes    | N/A   | N/A    | N/A     | N/A   | N/A      | N/A     |
| Luname             | Yes       | Yes    | No    | N/A    | No      | No    | N/A      | N/A     |
| UserID             | Yes       | Yes    | N/A   | No     | No      | No    | N/A      | N/A     |
| Bappl              | Yes       | Yes    | N/A   | N/A    | N/A     | N/A   | Yes      | Yes     |



### Steps to exploitation



existing system

static routing

Dynamic routing single TOR

Dynamic routing multi TOR

IBM 2007  
**IMPACT**

**IBM**

**Questions  
and  
Answers**

impact·venture\*

© IBM Corporation 2007. All Rights Reserved.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation 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 IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).  
AIX, CICS, CICSplex, DB2, DB2 Universal Database, i5/OS, IBM, the IBM logo, IMS, iSeries, Lotus, OMEGAMON, OS/390, Parallel Sysplex, pureXML, Rational, RCAF, Redbooks, Sametime, System i, System i5, System z, Tivoli, WebSphere, and z/OS.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.  
Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.  
Intel and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.  
UNIX is a registered trademark of The Open Group in the United States and other countries.  
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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