Managing Workload Manager: Multiple Sysplexes and Asymmetric Sysplexes

Scott Chapman

Enterprise Performance Strategies, Inc.

Scott.Chapman@EPStrategies.com

EPS Webinar
July 2026

Contact, Copyright, and Trademarks

Questions?

Send email to performance.questions@EPStrategies.com, or visit our website at https://www.epstrategies.com or http://www.pivotor.com.

Copyright Notice:

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check®, Reductions®, Pivotor®**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM®, z/OS®, zSeries®, WebSphere®, CICS®, Db2®, S390®, WebSphere Application Server®, and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract

Using WLM to manage the performance of your z/OS workloads is of course a huge topic. But we sometimes get asked about the more pedestrian, but equally important, part of how you manage your WLM policies when you have multiple sysplexes. Should each plex have its own separate policy or should you have one policy that you copy between plexes? What are the implications of those different answers? Somewhat related are the issues that arise when you have a sysplex that has asymmetric systems: systems with different access to capacity or maybe even run on quite different hardware. If either of those situations apply to your environment, then this webinar presented by **Scott Chapman** may give you some new ideas for how to handle those sometimes “sharp” edges of that situation.

EPS: We do z/OS performance…

- Pivotor - Reporting and analysis software and services
  - Not just reporting, but analysis-based reporting based on our expertise
- Education and instruction
  - We have taught our z/OS performance workshops all over the world
- Consulting
  - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
  - We present around the world and participate in online forums

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
  - March 30 – April 3, 2026 (4 days, excl Wednesday the 1st)
- WLM Performance and Re-evaluating Goals
  - July 27 – 31, 2026 (4 days, excl Wednesday the 29th)
- Parallel Sysplex and z/OS Performance Tuning
  - May 12-13 2026
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor™

- If you don't see them in your performance reporting tool, or you just want a free cursory performance review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html

- We also have a free Pivotor offering available as well
  - 1 System, SMF 70-72 only, 7 Day retention
  - That still encompasses over 100 reports!

Agenda

- The original sysplex vision
- Why have multiple sysplexes?
- What do we mean by asymmetric sysplexes?
- What are the implications of the above for z/OS performance management?

Parallel sysplex: Historical Vision vs. Reality

Ancient History

- 1990: IBM introduces SYStems comPLEX, aka “sysplex”
- Initially just a “base” sysplex:
  - Cross system communication via XCF via CTC (channel to channel) links
  - Shared couple datasets for managing sysplex
  - Sysplex timer for coordinating the time across multiple CECs
- Parallel Sysplex introduced 1994 with the CMOS processors
  - Added Coupling Facility (CF) to facilitate data sharing across multiple machines
  - CFs were even a separate machine type (9674 vs 9672)
  - Was partially a result of CMOS processors being much slower than bipolar
    - So needed a lot more of them compared to the older, faster, more expensive bipolar machines

Parallel Sysplex design goals

- The ability to automatically balance transactions across parallel systems according to their available capacity.
- The ability to run work on those systems whose capability is appropriate for the work.
- The ability to keep workloads running even when there are hardware or software interruptions on a processing facility within the Parallel Sysplex. This avoids outages that are costly to businesses. The Parallel Sysplex can be configured so there are no single points of failure.
- Reduced cost in systems management and operations.
- Single system image for users, operators, and systems management

IBM Announcement ZA94-0160 dated April 6, 1994

Parallel Sysplex vision

- Multiple “small” machines interconnected via CF machines
- Application work distributed dynamically across those machines
- Any system run any work for the application
- System outage transparent to the application whose work just routes to available systems
- Managed as “a single system image”
  - One configuration for all machines, simplifying management

A Theoretically Ideal Parallel Sysplex

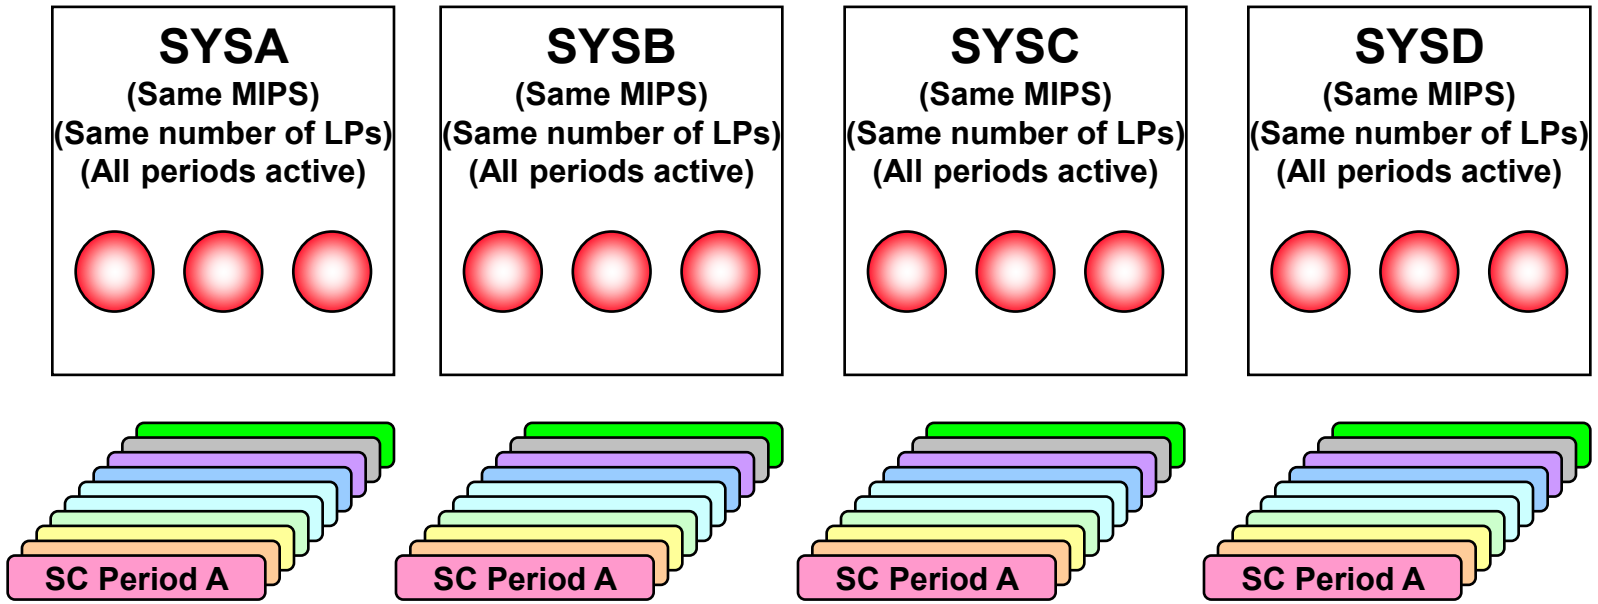

- 4 production systems all running the same work, all about the same size

A more likely reality today

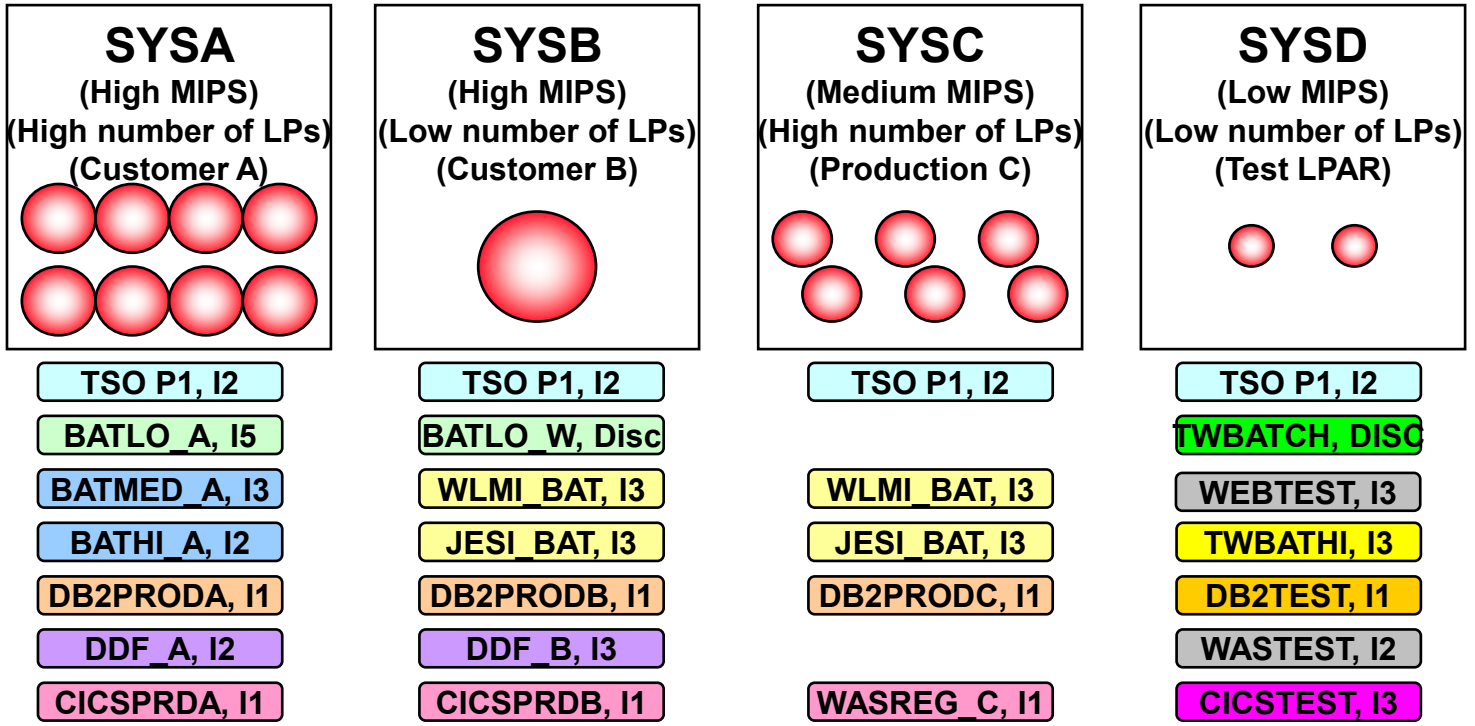

- Variety of system sizes
- Different work on different systems
- Mix of production and test

Of course there's an infinite variety of variations between these extremes!

How many sysplexes should you have?

“Single System Image”

- If the goal is to only have one “system” to manage, that means one sysplex
  - I.E. development, test, production all in one sysplex
- Idea of mixing production and non-production was always a debate point
  - If the sysplex is resilient to a single system outage, the development system failing should have no impact on production
- Separate sysplexes for production and non-production are common
  - System outages may be momentarily disruptive to performance
  - Other shared resources in a sysplex may impact performance:
    - Locks/enqueues
    - Security
    - Disk
- Multiple production parallel sysplexes (for separate applications) also common

Scott's thoughts

- How many parallel sysplexes you need/want is very situationally dependent!
- There is a cost component: ICF engines aren't free
  - Can share them though, especially for smaller & non-production plexes
- Because things like the security and locks are shared your risk management may dictate separating production and non-prod and/or by application
- But there's nothing inherently wrong with having prod/non-prod together
  - One sysplex is easier to manage then multiple
  - One sysplex may end up being cheaper
- For dev/test plexes, depending on your situation, running CF LPARs on spare GP capacity may be fine too
  - Performance will be impacted, but if for dev/test that's probably fine
- You can change your mind, but splitting or merging sysplexes is not fun
  - Been there, done that.

What is an asymmetric sysplex?

Sysplex Capacity Concerns

- Machines in a sysplex may be of different generations and capacities
  - Generally, n-2 generations can coexist in the same sysplex
- Most sysplexes with multiple machines do tend to be of similar processor "speed" and perhaps total capacity
  - I.E. You may have a 620 with a 617, but something like a 620 with a 515 is unusual
  - More common with "BC" machines where you may have a J05 and a N06
- Where there is a "speed" difference, there will be a greater difference in the SU/sec "constant" for each LPAR
  - This impacts the service accumulation for WLM period switching
  - Usually this is a minor point, but it can be a consideration in some situations

SU/Sec Constant - For Each System Over Time

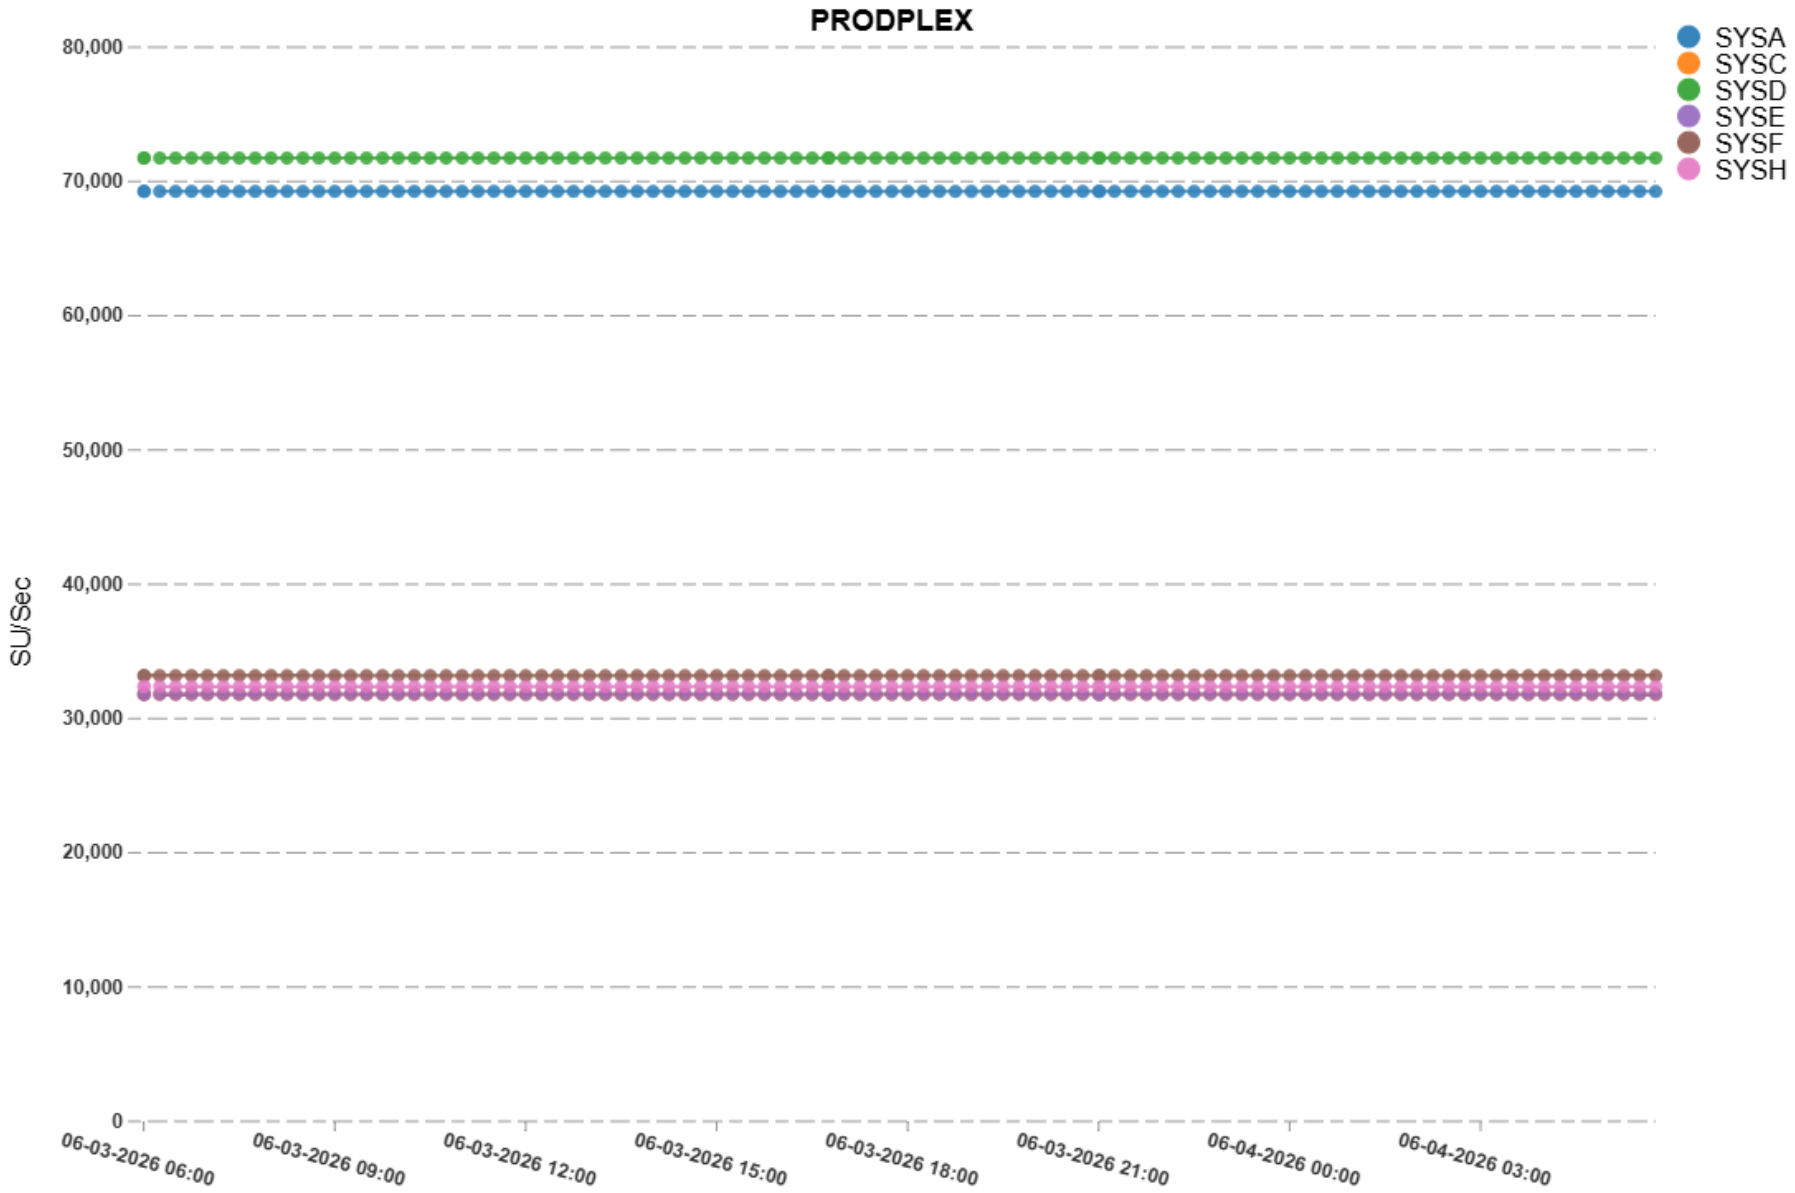

Here SYSA and SYSD have very different SU/Sec ratings to the other LPARs.

Software asymmetry

- A relatively common issue is having certain workload pinned to a subset of the systems in the sysplex instead
  - These affinities can reduce flexibility and degrade availability
- Application design affinities
  - Applications sometimes reference specific system/subsystems instead of using the generic names and allowing the sysplex to route between them
  - Generally, should discourage applications from this!
- License affinities
  - Sometimes software might not be licensed on all systems, requiring work using that software to be routed to wherever it is licensed
  - For IBM software, moving from R4HA to TFP can help here

Work Type Asymmetry

- Does the plex do just production work or is there dev/test systems too?
- Some dev/test work isn't applicable to production and vice-versa
- Production and non-prod probably have different:
  - Activity volumes
  - Activity timings
  - Performance requirements
  - Security requirements

z/OS Performance Management Implications

SU/sec differences: multiperiod goals

- SU/sec will often be different on different systems
  - Most directly impacts period aging
  - Capacity differences can impact routing though
  - Capacity, logical CPs, and workload can impact velocity attainment
- Small differences generally can be ignored
- Larger differences may need to be considered
  - May need different SCs with different durations / goals

```
          ---Period---  ----------------------Goal---------------------
Action    #  Duration    Imp.  Description
  __
  __      1  7000         2    90% complete within 00:00:00.150
  __      2  14000        3    90% complete within 00:00:00.500
  __      3               4    80% complete within 00:00:01.000
```

- If SU/sec = 70,000 then 7000 SUs is 0.1 CPU seconds
- If SU/sec = 35,000 then 7000 SUs is 0.2 CPU seconds
- A transaction taking 0.2 CPU second is not going to complete in 0.15 seconds
- That's not necessarily a huge problem, but it is something we would question

Asymmetric Sysplex Velocity Goals

- Goals are Sysplex wide
  - How do you assign a velocity goal in an asymmetrical environment?
  - Example: CICSPROD class assigned VEL 50 goal
    - Too aggressive for SYSA
    - Too easy for SYSB
    - Just right for SYSD
  - System classification
    - If SYSA then classify to CICSA with VEL 20 goal
    - If SYSB then classify to CICSB with VEL 70 goal
    - If SYSD then classify to CICSD with VEL 40 goal
  - System works for STCs and TSO, for other work may need to use other qualifiers

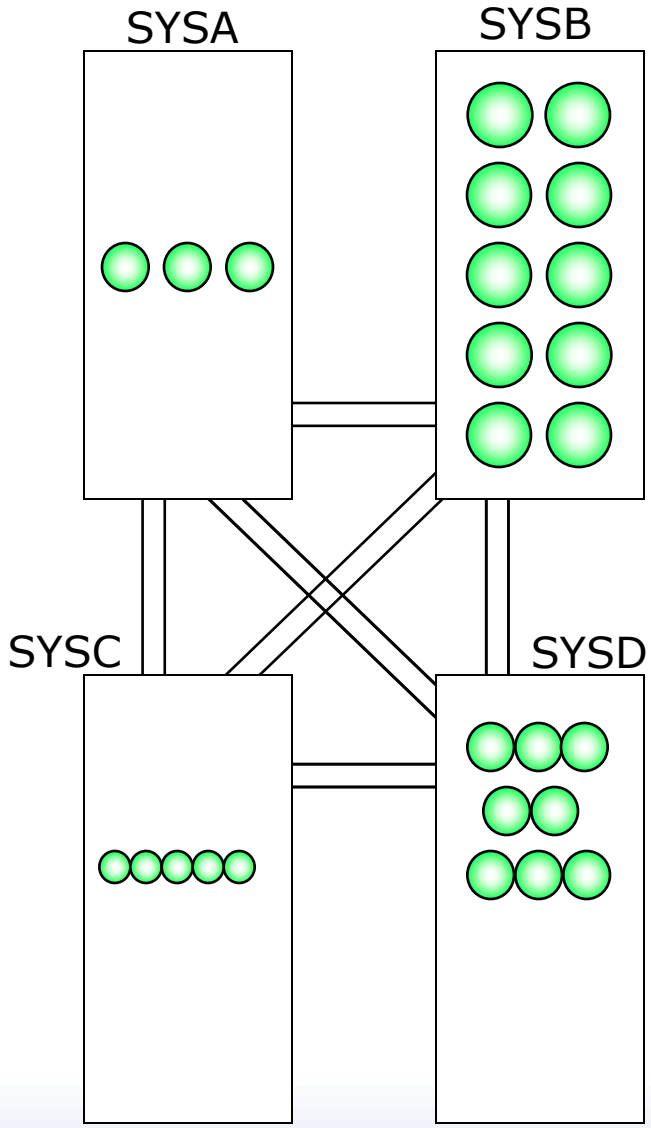

WLM: Sysplex PI

- WLM evaluates sysplex PI before system PI
  - Intent is to prioritize overall performance across the sysplex
  - The decision to help or not help work on one system can be influenced by how the work is performing on other systems
- This was a good intention in the mid 90s when WLM came out and the parallel sysplex vision was that ideal standard
- Today reality for many sysplexes is:
  - Composed of heterogenous systems with different work and/or capacity
    - E.G. dev, test, and production may all be in the same sysplex
    - Different systems may have access to different capacity
    - Not all workloads distribute all their work dynamically across the sysplex

If you're mixing Prod and Non/Prod

- Be careful of reusing production SCs for application work on the dev/test systems
  - E.G. CICSPROD is used as the goal for both production CICS regions and dev/test regions on different LPARs in the same plex
  - In this case, performance of CICSPROD on the dev/test LPARs might influence the management of CICSPROD on the production LPAR
- In many cases, this won't be a problem… but it can be
- Better to use different SCs for production and non-production work by qualifying the work on the classification rules by:
  - System name: STC, TSO
  - Subsystem instance: CICS, DDF, IMS, etc.
  - Transaction/job name: most

WLM PA - Top Sysplex PI Receiver Candidates

(Policy Adjustment Decisions)

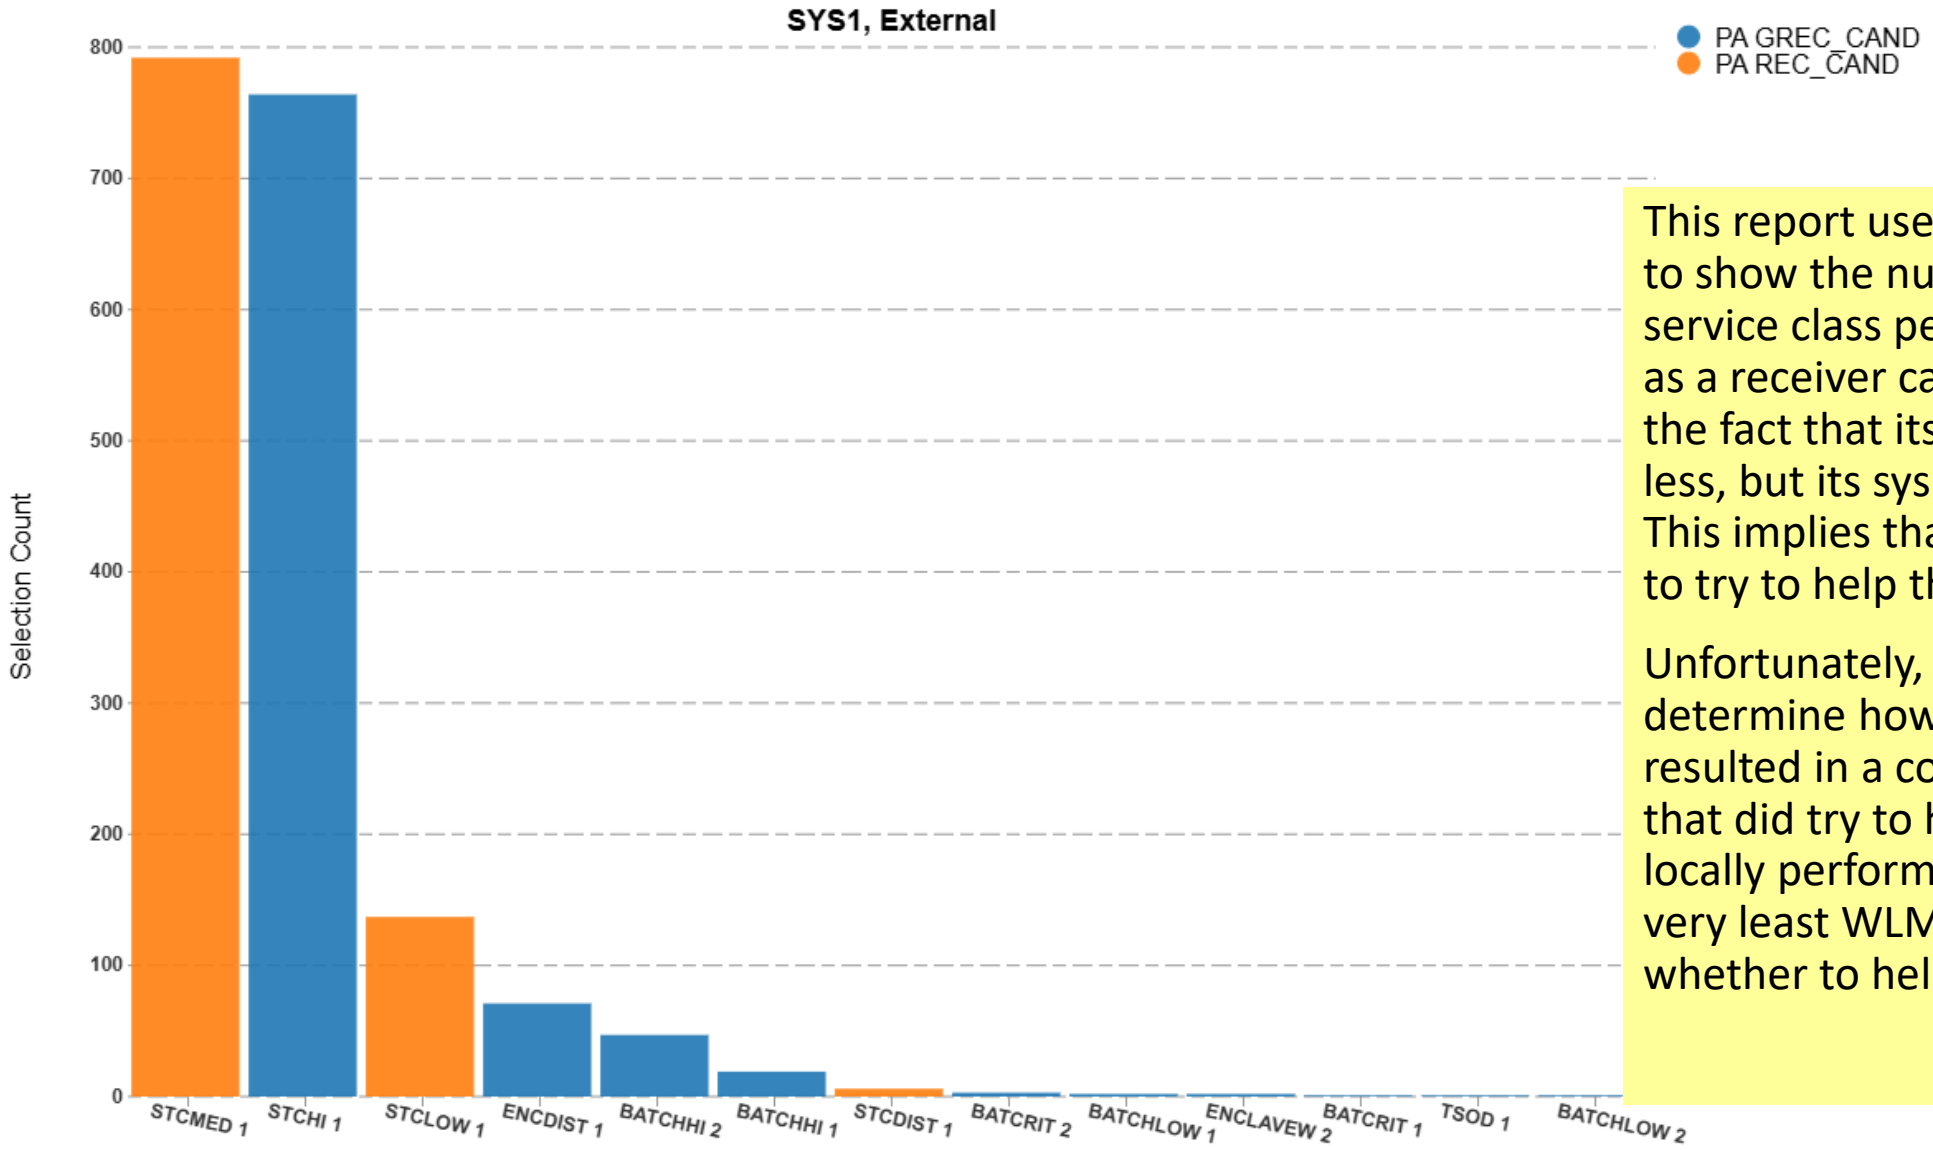

This report uses the 99.1 records to show the number of times a service class period was selected as a receiver candidate despite the fact that its local PI was 1 or less, but its sysplex PI was over 1. This implies that WLM selected it to try to help the sysplex PI.

Unfortunately, it's not easy to determine how many of these resulted in a committed change that did try to help that already locally performant SCP. But at the very least WLM was investigating whether to help them.

One policy or multiple?

- The WLM policy boundary is the sysplex
  - The policy covers all systems in the sysplex
  - The policy cannot be directly shared between sysplexes
  - If you have multiple sysplexes, what do you do?
- Manage separate policies individually
  - PRODPLEX and TESTPLEX are manually updated separately
- Copy a policy between sysplexes
  - Export as XML from wherever you changed it then import on the other plex(es)
  - This is what I did: 1 policy copied between 4 plexes

Policy management

Individual Policies

+ Mentally aligns with how WLM actually works: the policies are independent
+ More easily allows for targeted goals for different plexes

- Requires multiple manual updates when things should be managed the same
- May use different SC/RCs for the same work on different plexes, increasing complexity of cross-plex reporting

One policy, replicated

+ Can better enforce SC/RC consistency across the entire environment
+ Reduces chance of error by making manual changes across multiple policies

- Makes tuning goals for different plexes more complex
- If you want to use different SC/RCs for different work, it will be more difficult
- May lead to thinking of the policies as a single entity

Scott's opinion…

- Consider replicating a single policy if the plexes reside on the hardware and do somewhat similar work

Fun fact: "Sysplex name" is the 2nd most widely available workload classification qualifier (behind container qualifier)

- Consider individual policies if the plexes are on different hardware or do very different work

- But there's no right or wrong answer here

Wrap-up

Summary

- Asymmetry is not unusual: the originally envisioned ideal of a sysplex is rare and systems in a sysplex often differ in:
  - Their available capacity
  - Whether they run production or dev/test
  - The work that's allowed to run on them
- It may be necessary to assign specific and different SCs to the "same" work running on different systems
- If you have multiple sysplexes, you can manage multiple WLM policies or replicate one policy to multiple plexes
  - Either way is fine, but there are things to consider either way

Thanks!
Questions?

Email: scott.chapman@epstrategies.com