WSC Experiences with z/OS 3.2 Performance Enhancements

February 23rd, 2026

9:45am-10:05am

Stephanie DeLuca
IBM Z Performance Specialist
Worldwide Systems Center (WSC)
stephanie.deluca@ibm.com

About the IBM Worldwide Systems Center

IBM Z Worldwide Systems Center (WSC)

Drive IBM Z platform adoption, growth and retention
with deep technical expertise in new and existing IBM Z capabilities

The WSC experts provide SME support and act as consultative, trusted advisors to IBM Z clients. Our experts establish best practices and develop scalable, reusable demos, tools, methodologies and enablement content for IBM Z. We showcase and use these assets to upskill the IBM Z community by mentoring, instructing, and advising clients, business partners and IBMers. Our SMEs are key systems partners for product development including Early Support Programs and critical situation support. We also perform economics analysis, execute competitive performance benchmarks and comprehensive health checks.

To engage the IBM Z WSC
contact your local IBM account team
or IBM Business Partner

Agenda

z/OSMF WLM Policy Advisor

SYSSTC Health Checker Rule

New SRB Use Case for Dynamic I/O Activate

SRM Lock Contention Relief

Appendix

- CPENABLE Default
- New SMF Records

Z/OS 3.2 WLM POLICY ADVISOR ENHANCEMENTS

The Idea

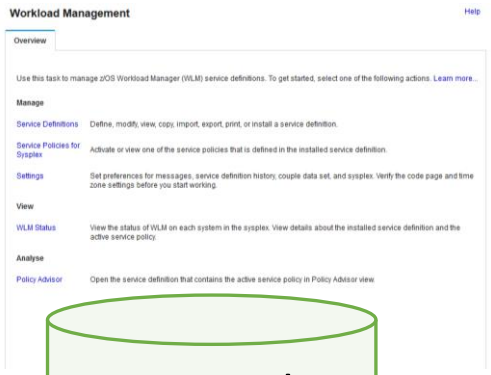

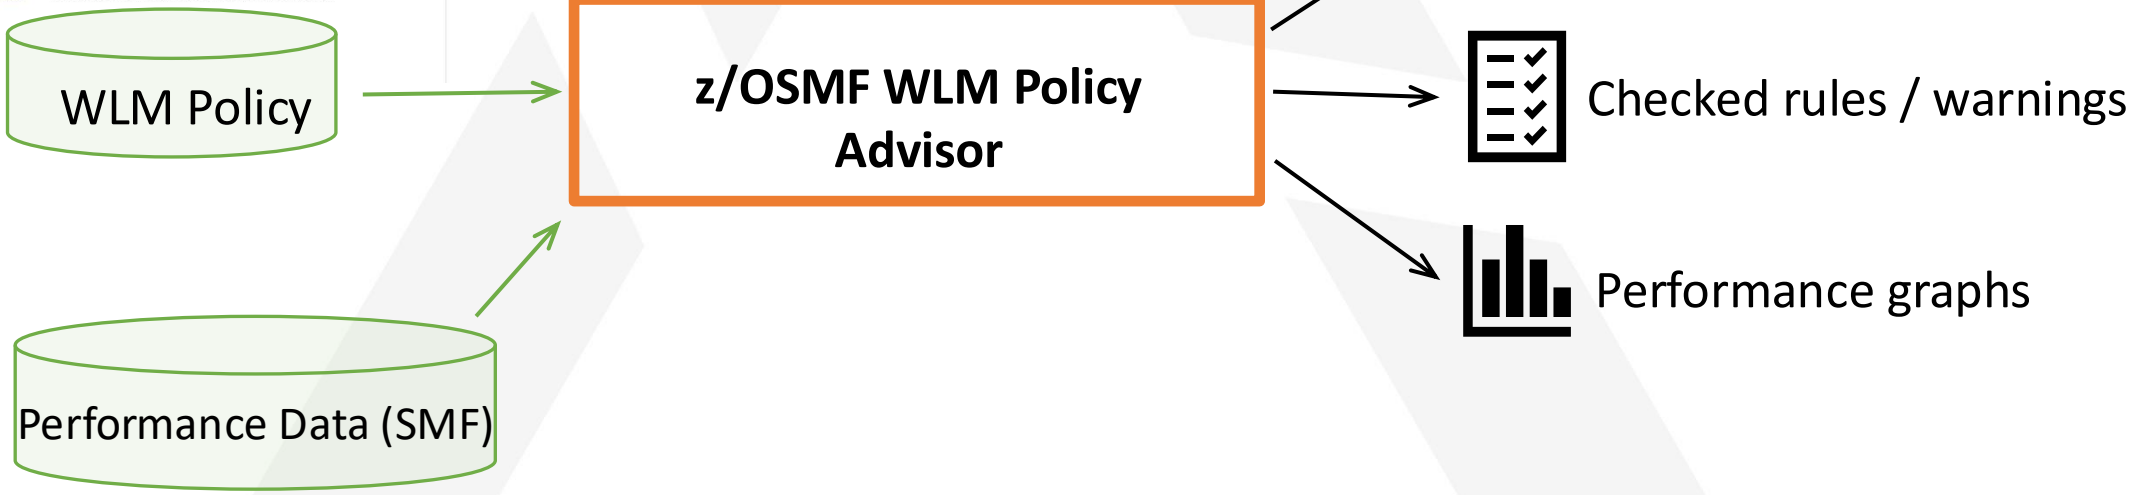

z/OSMF Workload Management Task

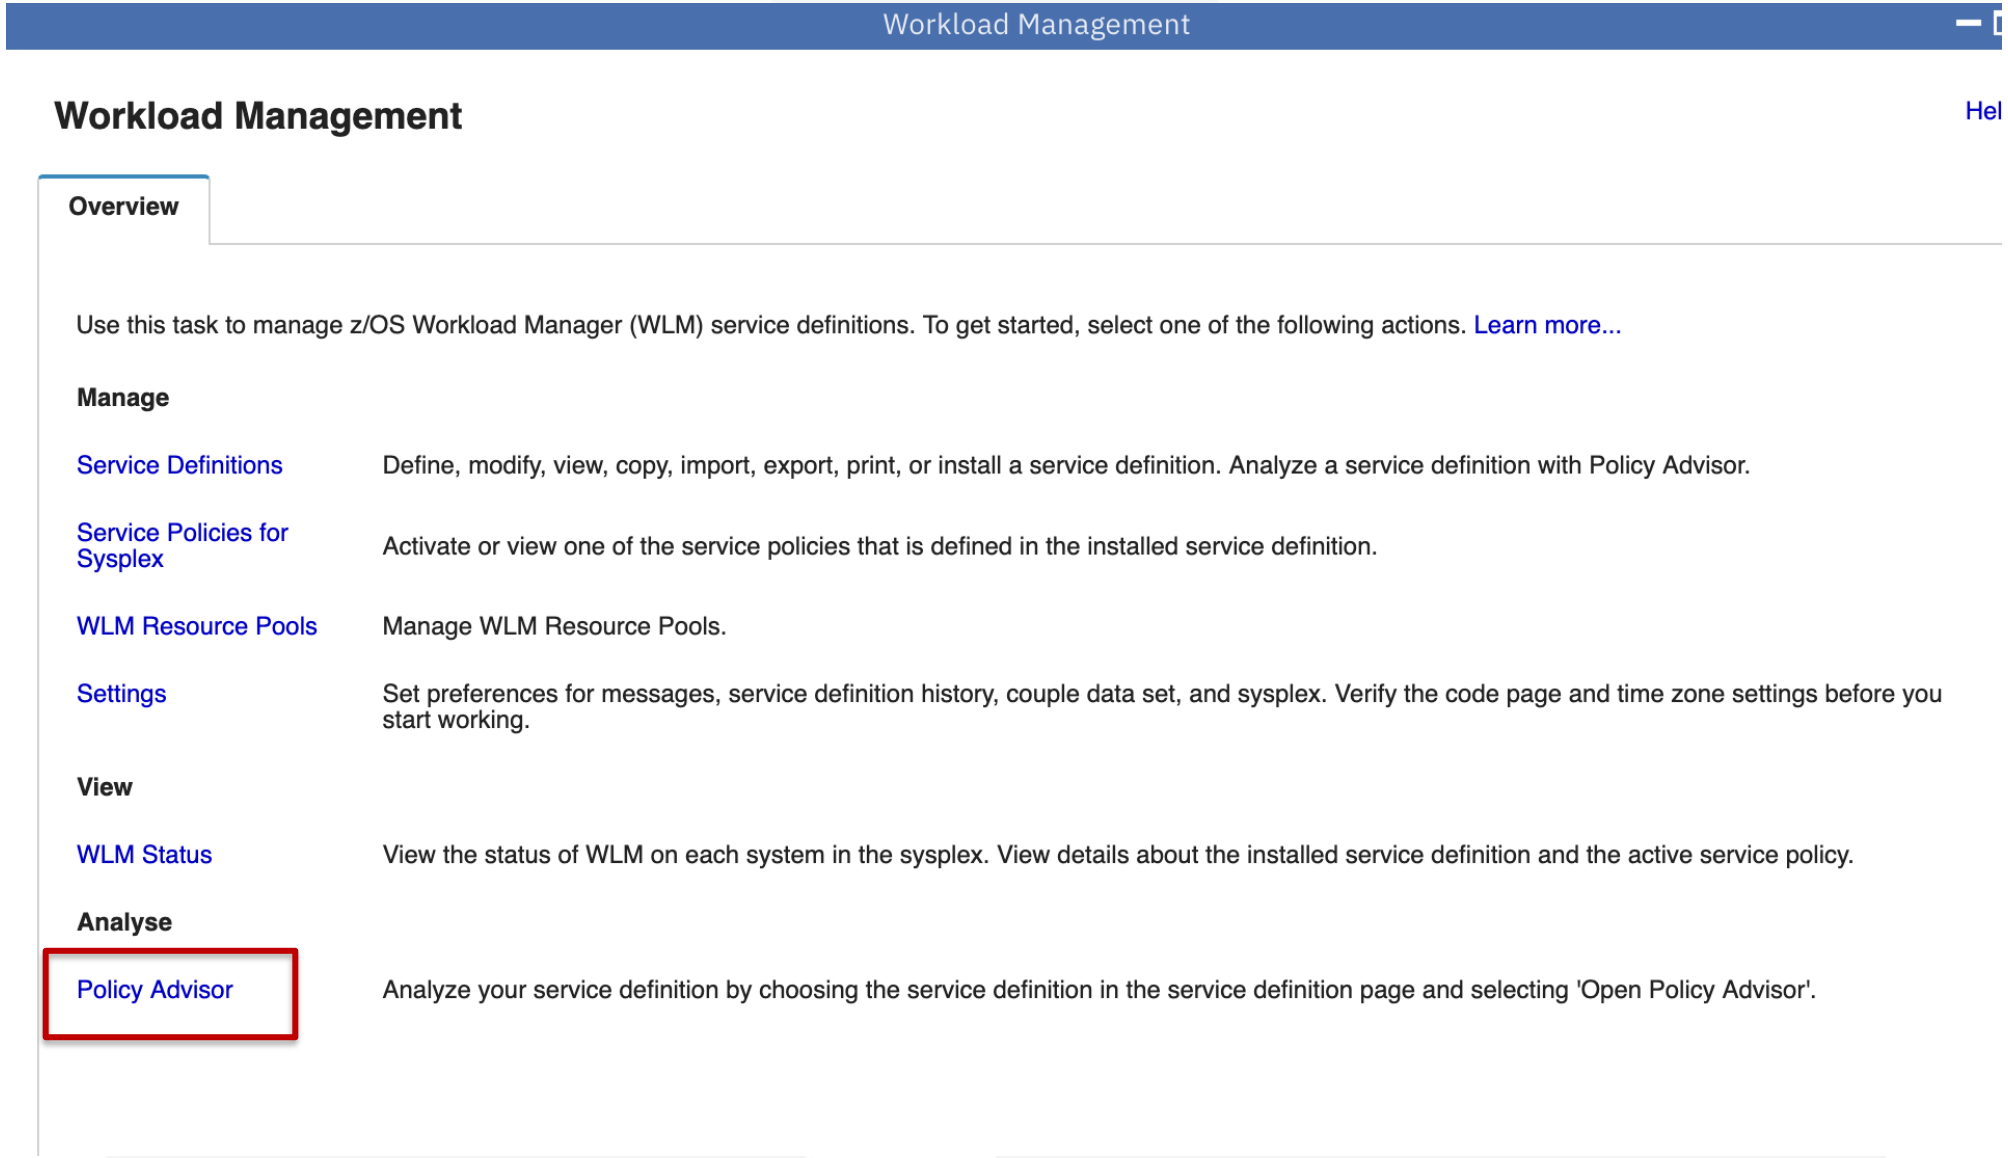

Initial z/OS 3.1 WLM Policy Advisor

New support to analyze a WLM service definition through the z/OSMF Workload Management task

- Provides analysis and data visualizations to review service definition set up and service class performance
- Offers several recommendations for the specified service definition based on best practices
- Simplifies WLM policy definition tasks
- New direction for working with WLM

Requires a dataset with SMF 72.3 records for the active service policy

- Importance tab
- Performance Index tab
- Goal tab

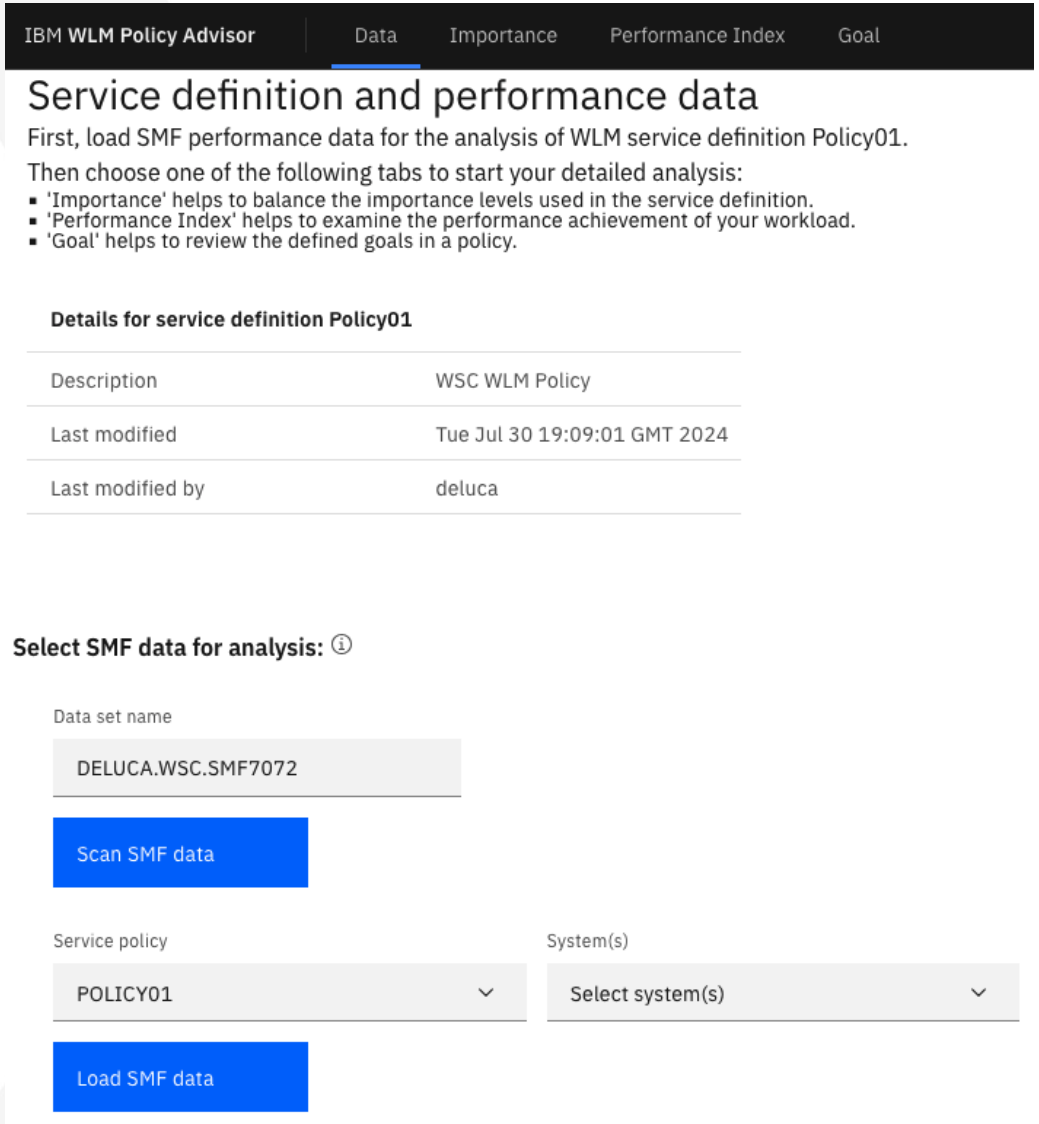

Policy Advisor Enhancements in z/OS 3.2

- Redesigned user interface
- Streamlined workflow and enhancements to load SMF data
  - Load data from multiple datasets
  - Search for dataset names from the UI
  - Choose specific LPARs and dates
- New overview page with actionable recommendations
- Improved filtering on Performance Index and Importance tabs
- All pages augmented with additional findings and recommendations

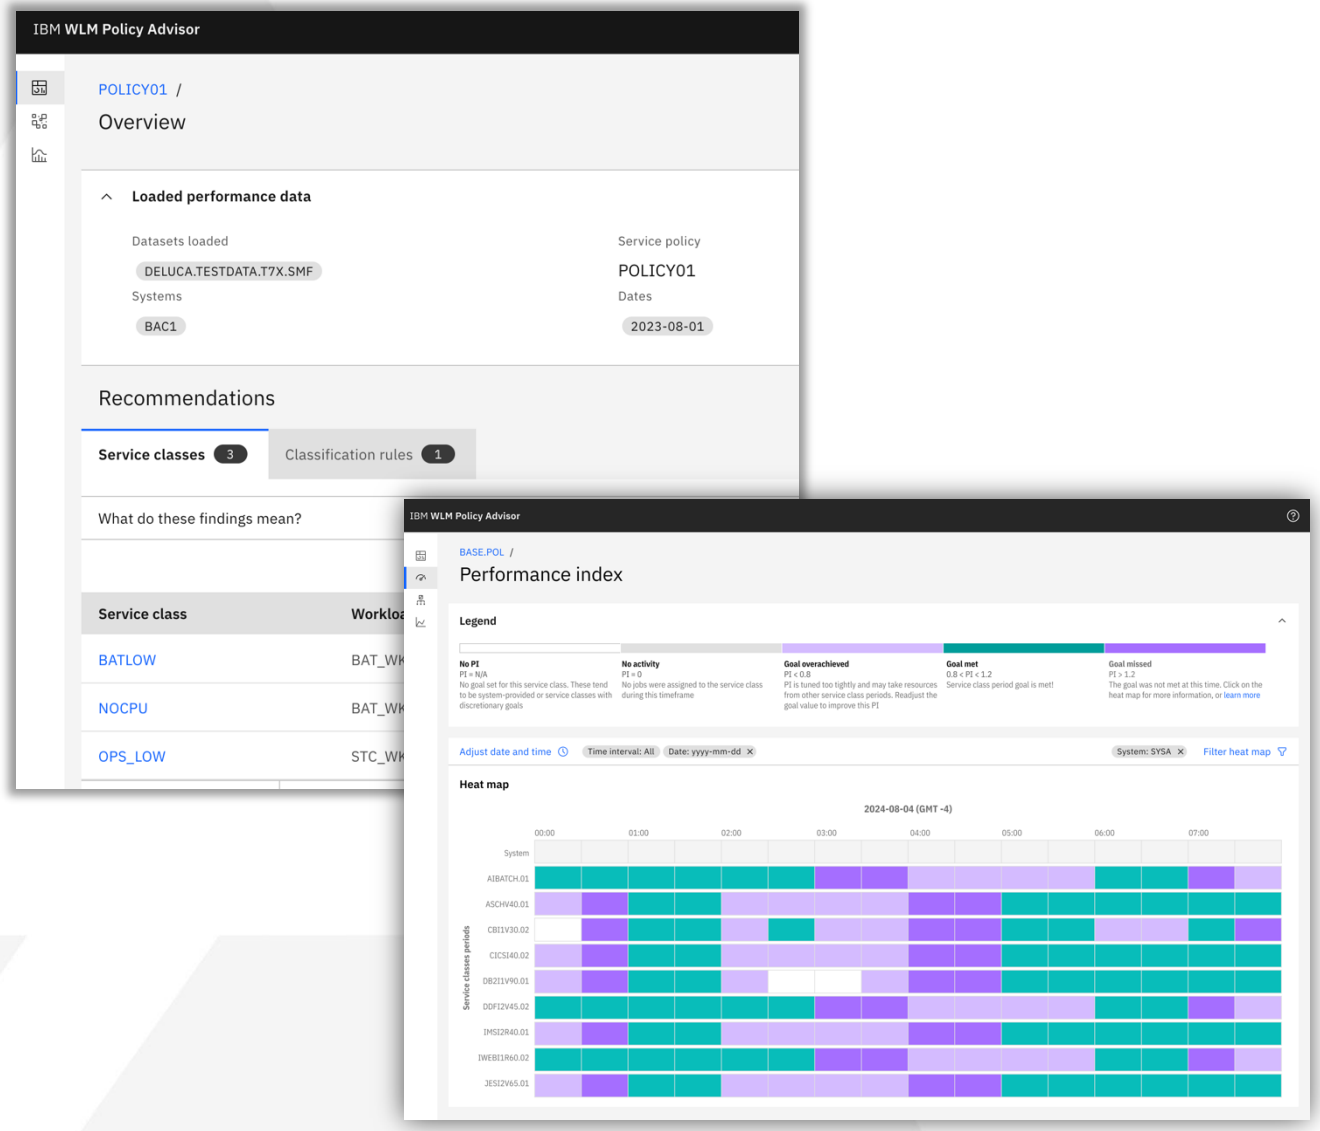

Streamlined Workflow to Load SMF Data

Allows for searching of datasets, loading multiple datasets, and selecting specific systems and dates

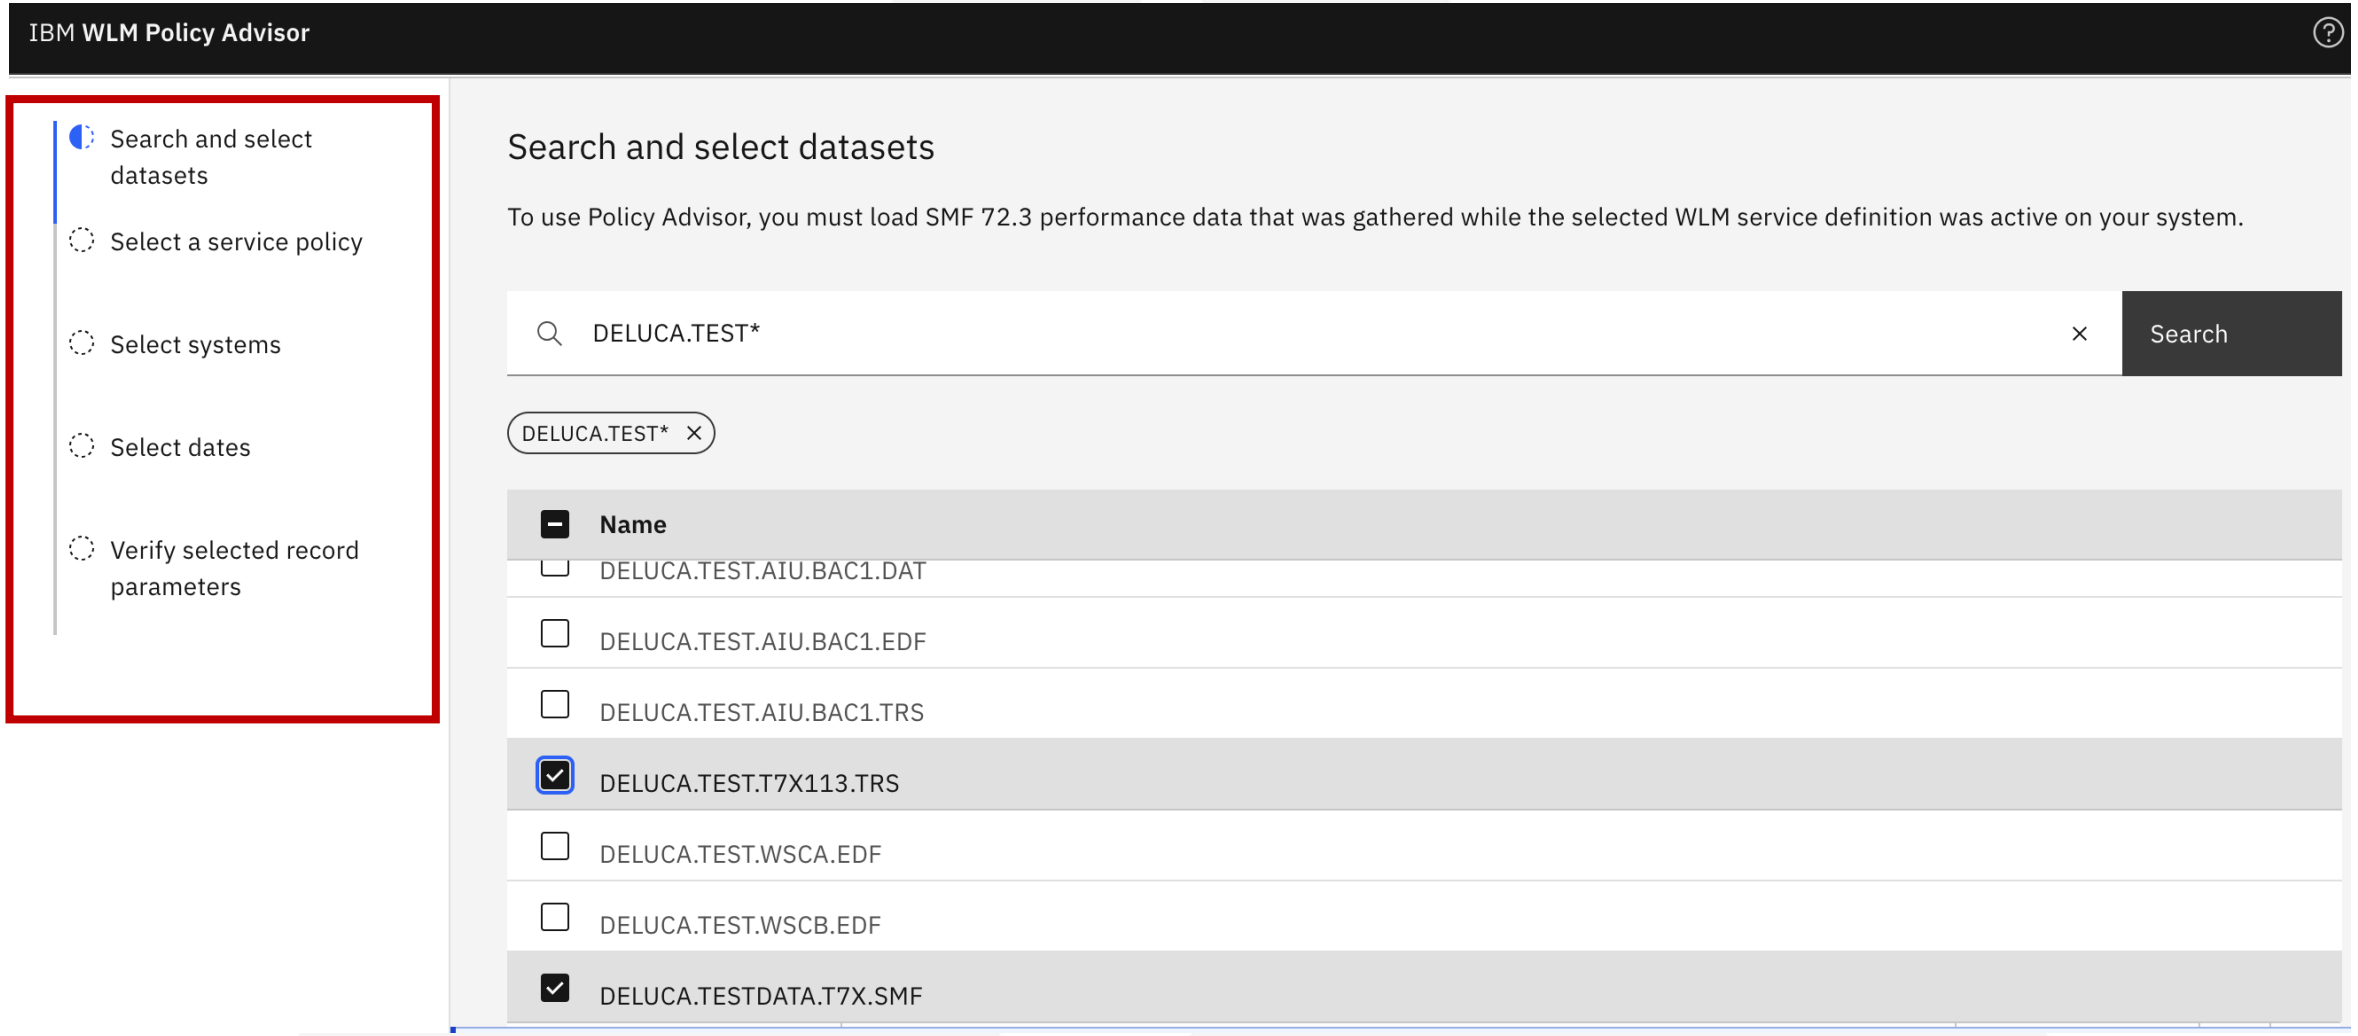

New Overview Page

Shows which data is loaded and provides immediate and actionable recommendations

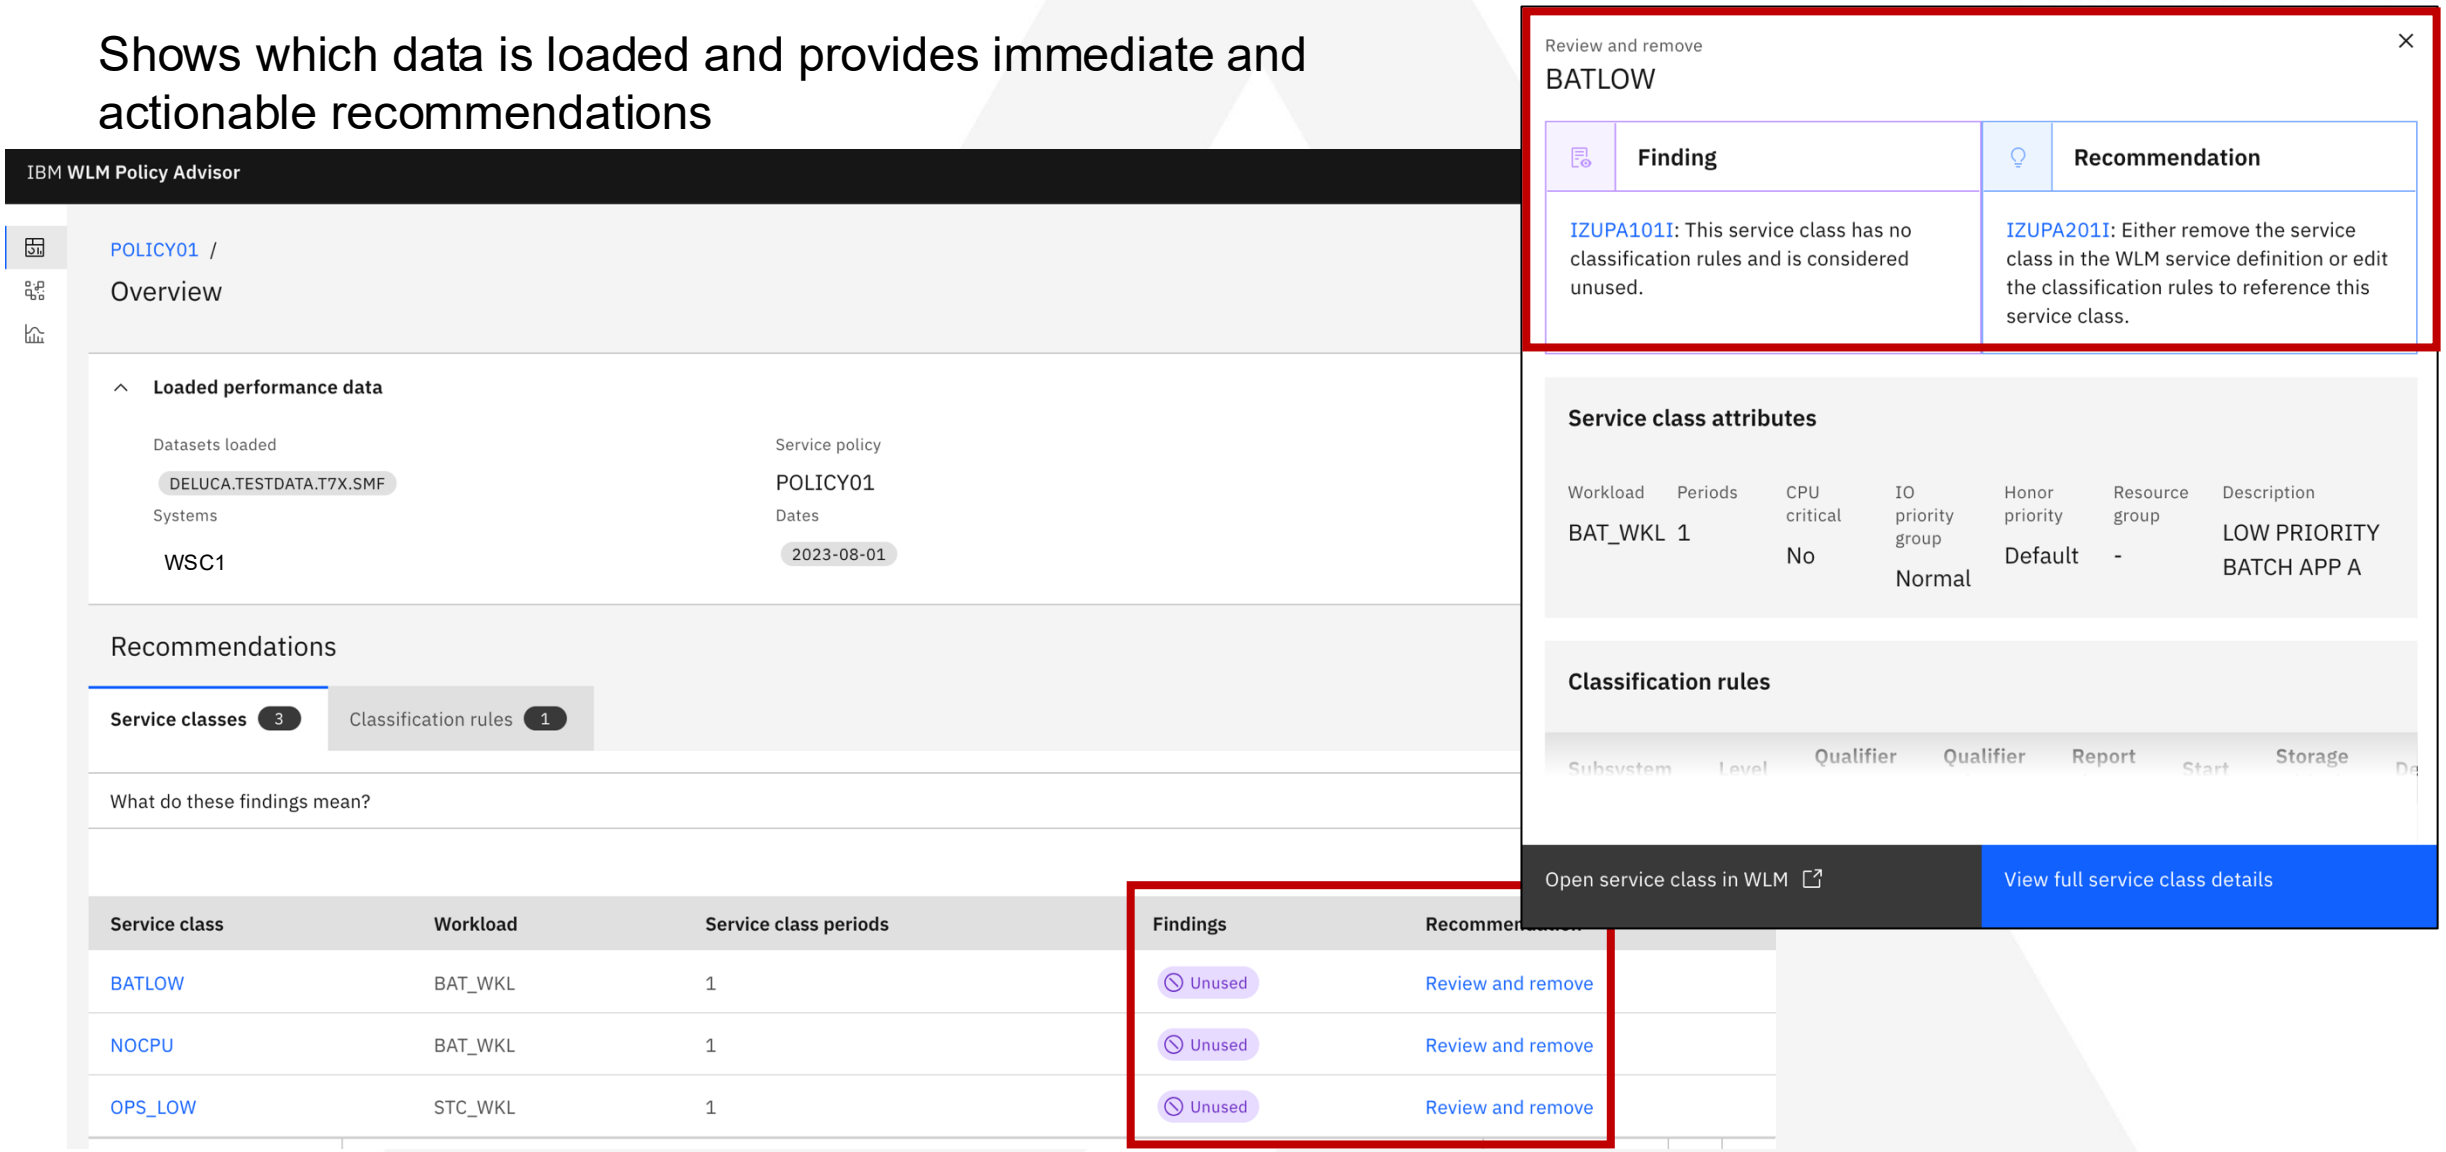

Additional Findings and Insights

<u>Performance Index (PI) Page</u>

- Updates to PI heat map
- More filtering options
- Findings and detailed recommendations

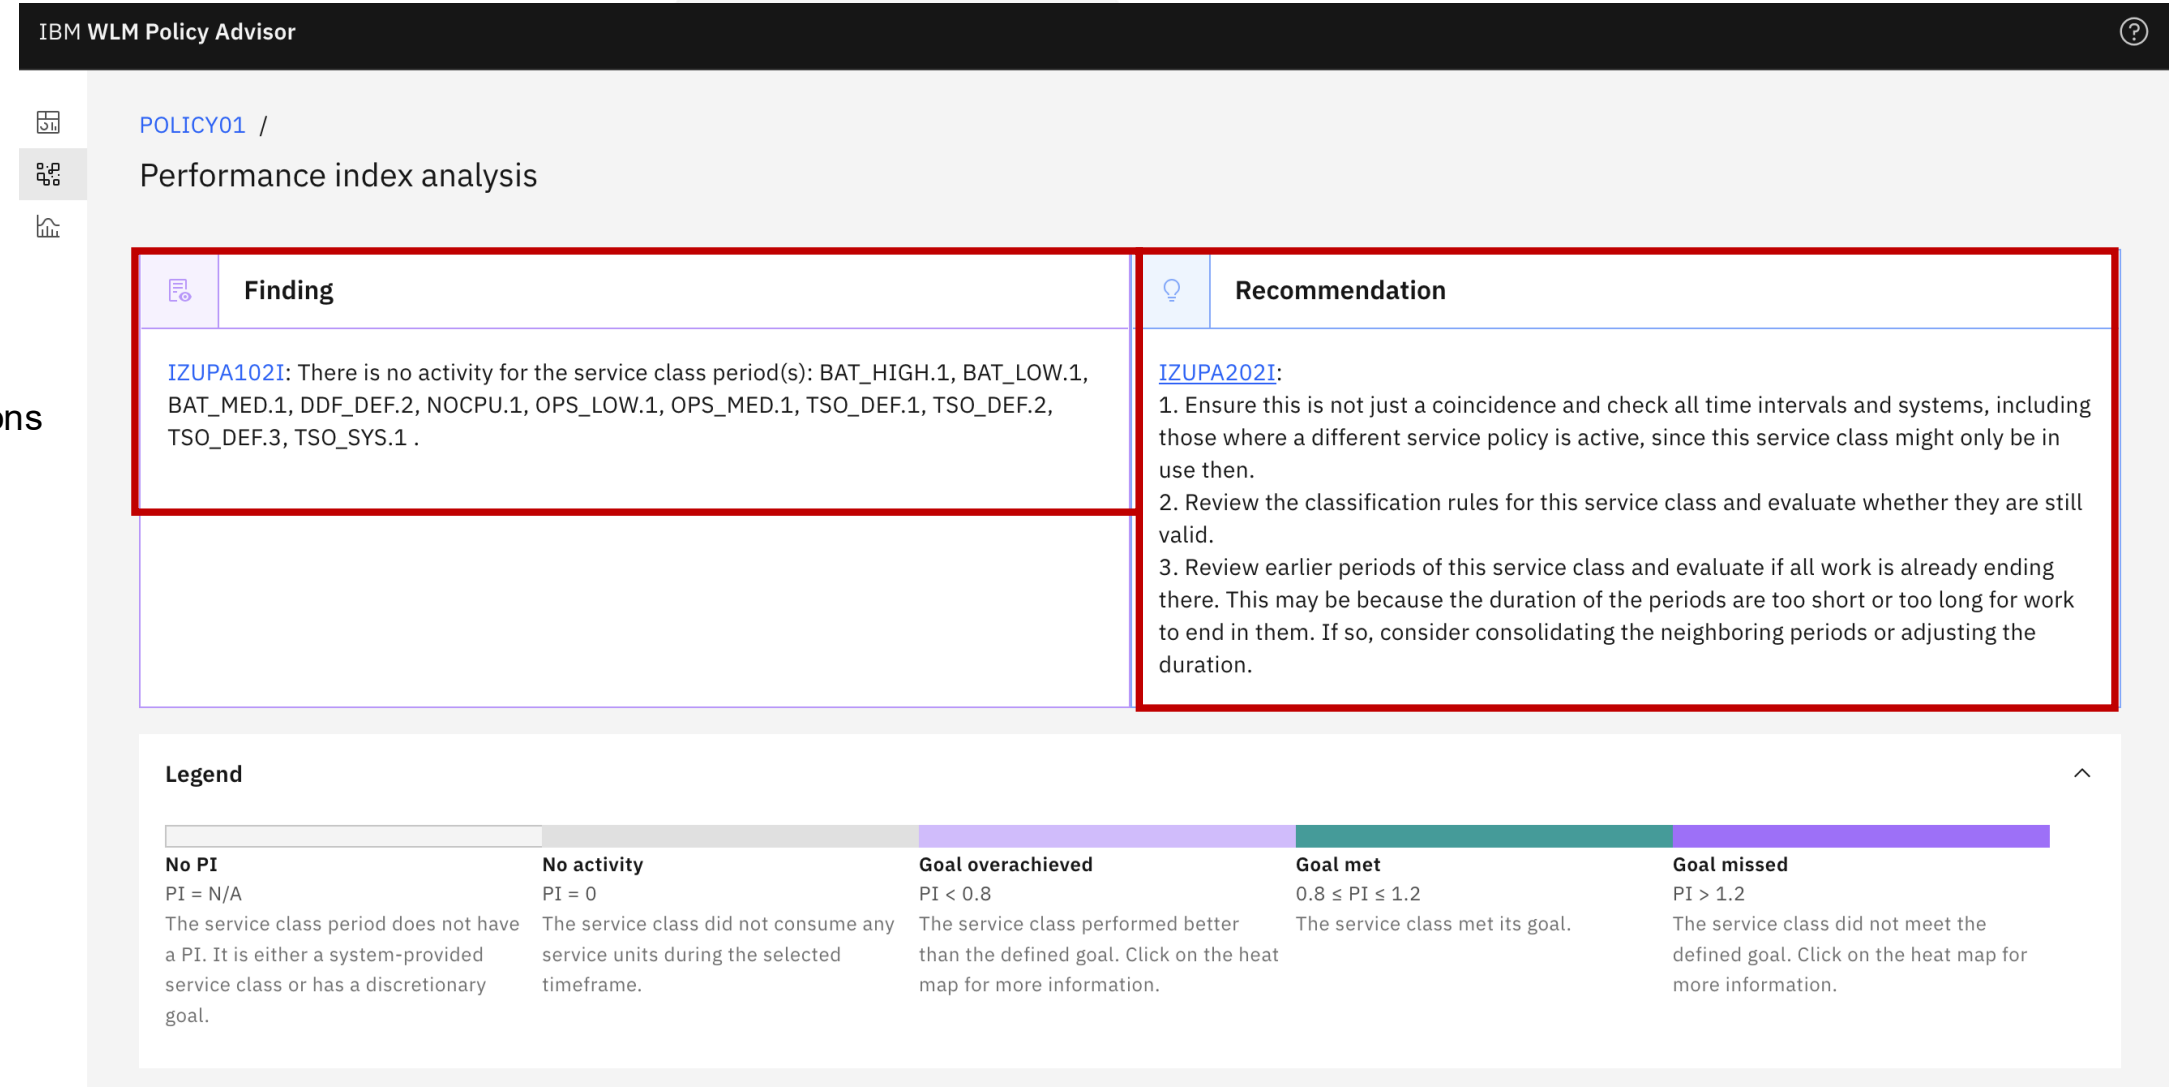

Additional Findings and Insights

<u>Importance Page</u>

- Service class distribution across Importance Levels (defined)
- Actual distribution of work using Service Units
- Detailed findings and recommendations

IBM **WLM Policy Advisor**

POLICY01 /

Importance level distribution

Finding	Recommendation
IZUPA104I: There are no service class periods defined for importance level 5 .	IZUPA204I: Review the service class periods of neighboring importance levels and determine which ones can be moved into importance level 5 .

Skip periods without work classification

Workload distribution (Service policy: POLICY01)

The workload distribution across importance levels shows how balanced the actual amout of work was running across the supported range of importance levels.

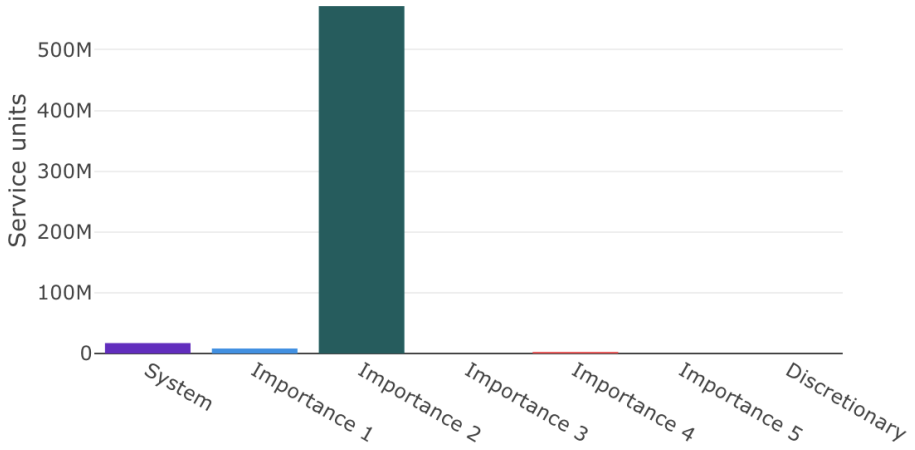

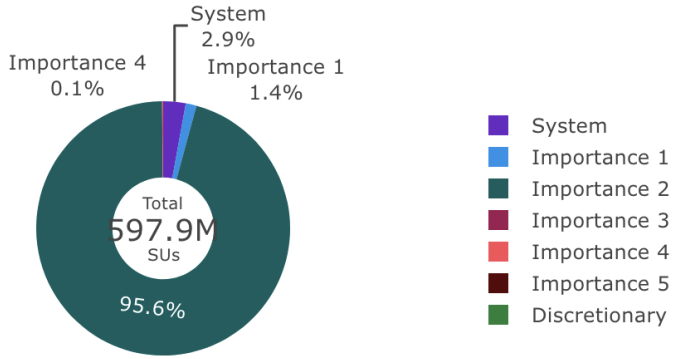

WSC Experience and Recommendations

- Extremely valuable to review WLM policies
  - Since z/OS 3.1, Policy Advisor has been a key tool in WLM engagements at the WSC
- Modern UI along with additional guidance and findings allow new WLM administrators to quickly visualize service class performance, make meaningful updates, and add value to the enterprise
- In many cases z/OSMF Workload Management Task along with WLM Policy Advisor can be an all-in-one solution
- Try it out! Feedback encouraged

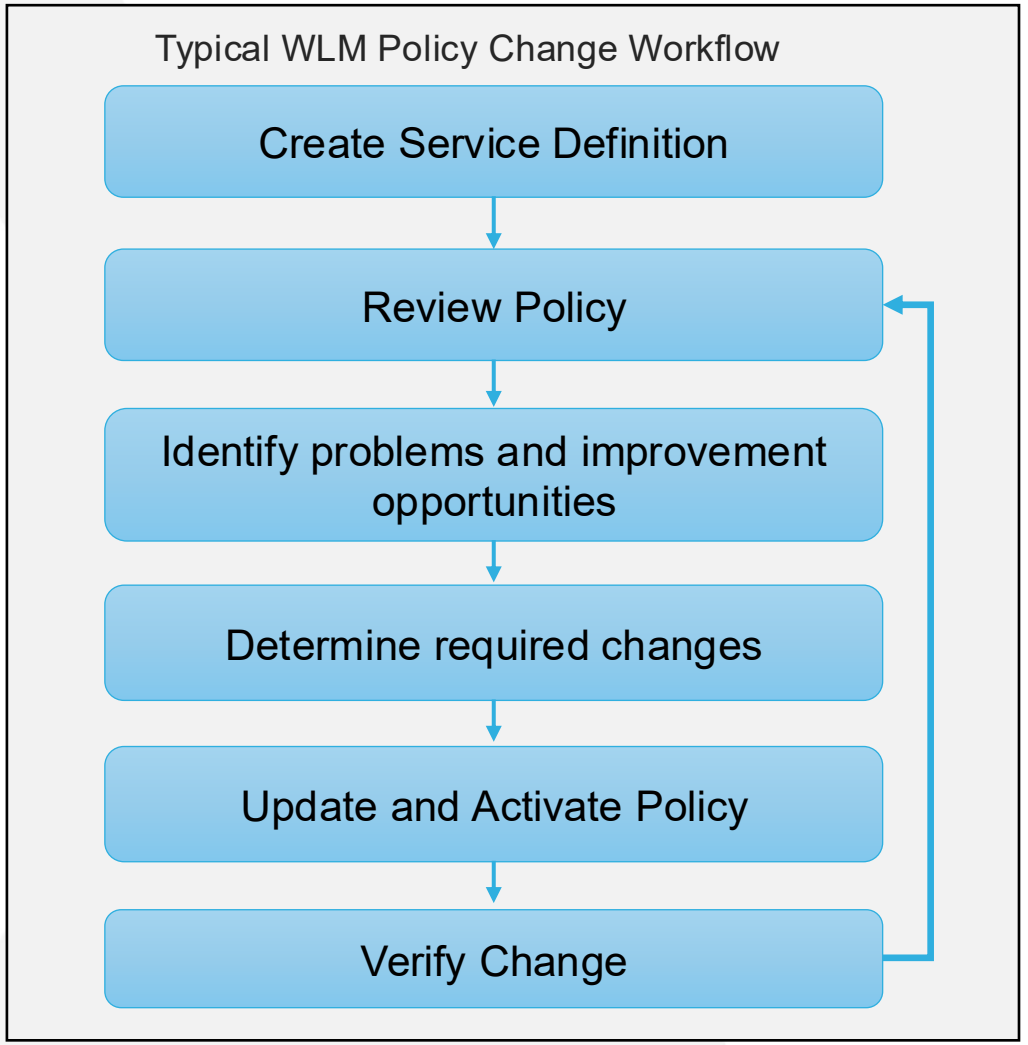

SYSSTC CPU Z/OS HEALTH CHECKER RULE
WLM_SCLASS_SYSSTC

z/OS Health Checker Rule for SYSSTC CPU Usage

IBM Health Checker for z/OS is a z/OS component that provides a set of system checks to help uncover potential problems before they impact availability

When there's an exception, each rule indicates a severity level -> Low, Medium or High

New rule: WLM_SCLASS_SYSSTC

- Dynamic severity check when the CPU consumption of address spaces in the system supplied WLM service class SYSSTC exceeds one of the specified threshold percentages
- Inactive by default
- z/OS 2.5 and above with APAR OA66312

```
SDSF HEALTH CHECKS  MVSA                                  LINE 186-210 (243)
COMMAND INPUT ===>                                          SCROLL ===> CSR
NP    NAME                                 State              Status
      VSM_CSA_LARGEST_FREE                 ACTIVE(ENABLED)    SUCCESSFUL
      VSM_CSA_LIMIT                        ACTIVE(ENABLED)    SUCCESSFUL
      VSM_CSA_THRESHOLD                    ACTIVE(ENABLED)    SUCCESSFUL
      VSM_PVT_LIMIT                        ACTIVE(ENABLED)    SUCCESSFUL
      VSM_SQA_LIMIT                        ACTIVE(ENABLED)    SUCCESSFUL
      VSM_SQA_THRESHOLD                    ACTIVE(ENABLED)    SUCCESSFUL
      WLM_OPT_PARM_CPENABLE                ACTIVE(ENABLED)    SUCCESSFUL
      WLM_SCLASS_SYSSTC                    INACTIVE(ENABLED)  INACTIVE
      XCF_CDS_MAXSYSTEM                    ACTIVE(DISABLED)   ENV N/A
      XCF_CDS_SEPARATION                   ACTIVE(ENABLED)    SUCCESSFUL
      XCF_CDS_SPOF                         ACTIVE(ENABLED)    EXCEPTION-LOW
      XCF_CF_ALLOCATION_PERMITTED          ACTIVE(ENABLED)    SUCCESSFUL
      XCF_CF_CONNECTIVITY                  ACTIVE(ENABLED)    SUCCESSFUL
```

Why SYSSTC Usage Should be Included in Regular System Checks

Work running in SYSSTC has the second highest dispatch priority on the system

- Dispatch priority is fixed (xFE)
- It should be reserved for work that requires extremely fast access to CPU, but is not and does not have the potential to be CPU intensive
- At times, we see environments where inappropriate work runs in SYSSTC (Db2 DBM1, MSTR, DIST, MQ MSTR and CHIN, etc.) which risks delaying more critical OS and system tasks

Using the SYSSTC health check allows you to detect high CPU usage in SYSSTC, identify the address spaces causing it, and evaluate whether they are appropriate to be in SYSSTC

Example Structure of Workload Importance Levels	
SYSTEM	Predefined system address spaces
SYSSTC	Important monitors, automation, IRLM, everything which consumes very little CPU and requires fast and instantaneous access to CPU
Importance 1	Critical server address spaces, IMSDBCT, DBM1, MSTR, TORs when managed towards BOTH, monitors
Importance 2	Most important production work, your loved ones
Importance 3 and 4	Less important production work, TSO, very critical batch
Importance 4 and 5	Batch, TSO last period, all default service classes which do not need any attention
Discretionary	Batch

Implementation – Evaluate Threshold Values

The rule allows for three threshold values to indicate severity (Low, Med, High)

- At least one of these must be specified
- Threshold is the percent of total CPU time that all work can consume on the system

WSC Recommendation:

- Thresholds can and should be customized
- Review what runs in SYSSTC today and remove anything which does not absolutely require this high dispatching priority
- Identify typical CP consumption at peak times in SYSSTC to determine appropriate thresholds

Default: `THRESHOLD_LOW(20%),THRESHOLD_MED(40%),THRESHOLD_HIGH(60%)`

LPAR MVSA Captured CP Usage by Service Class

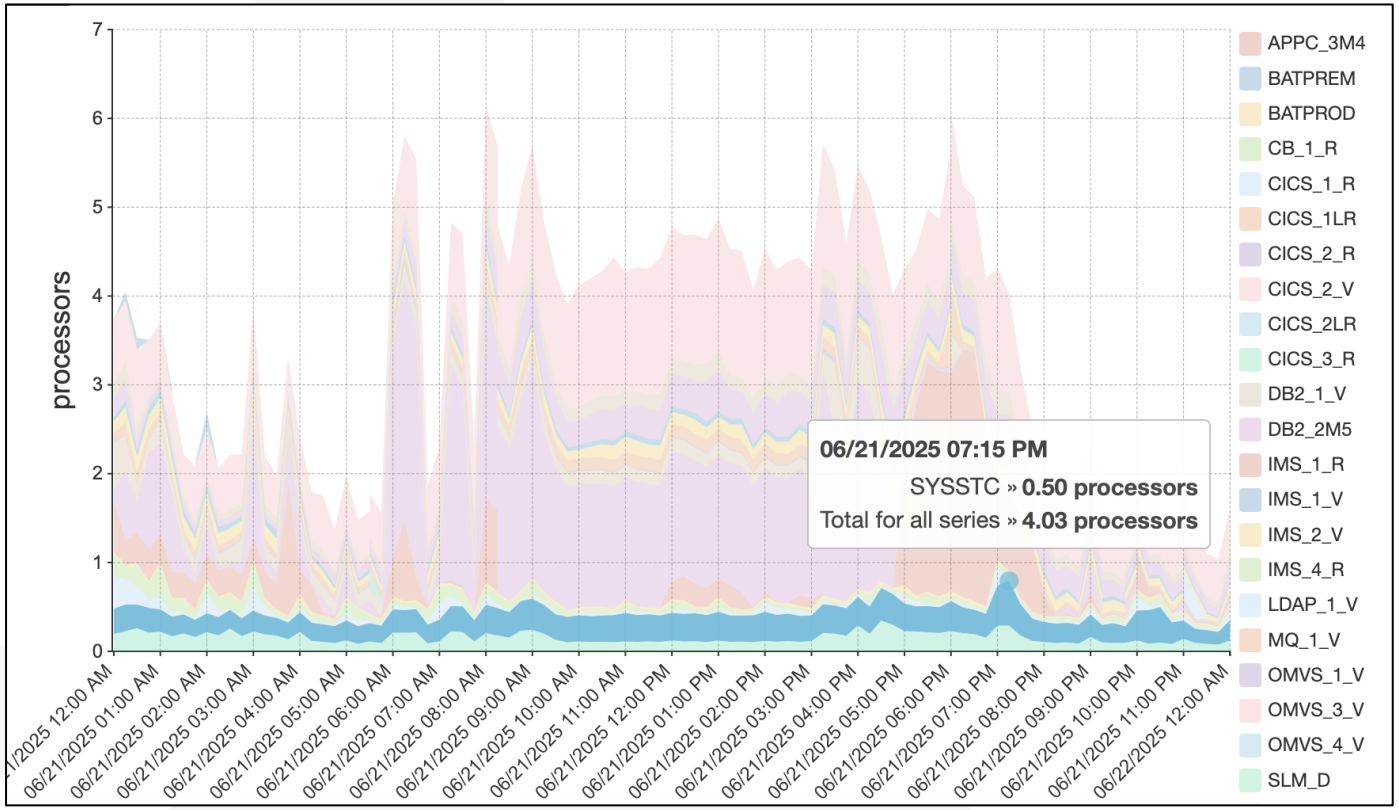

Implementation – Enable the Rule

Temporary check changes and enablement can be made using MODIFY HZSPROC commands

- Ex. F HZSPROC, UPDATE CHECK(IBMWLM,WLM_SCLASS_SYSSTC),ACTIVE

Make permanent customizations and keep the rule past the next refresh using HZSPRMxx parmlib member

- Evaluate and update defaults based on current SYSSTC usage
- Used at IBM Health Checker for z/OS startup

Check reports for exceptions regularly

- Example output of successful rule check

HZSPRMxx

```
ADDREPLACE POLICY[(policyname)] [STATEMENT(name)]
UPDATE
CHECK(IBMWLM,WLM_SCLASS_SYSSTC),
ACTIVE,
VERBOSE(NO),
INTERVAL(00:01),
SEVERITY(LOW),
PARM('THRESHOLD_LOW(10%),THRESHOLD_MED(20%),THRESHOLD_HIGH(30%)'),
DATE('date_of_the_change'),
REASON('Your reason for making the update')
```

SDSF CK Panel Browse (S)

```
SDSF OUTPUT DISPLAY WLM_SCLASS_SYSSTC       LINE 0        COLUMNS 02
COMMAND INPUT ===>                                      SCROLL ===>
************* TOP OF DATA *************
CHECK(IBMWLM,WLM_SCLASS_SYSSTC)
SYSPLEX:    SC02PLEX  SYSTEM: SC02
START TIME: 02/04/2026 16:36:11.812840
CHECK DATE: 20240425  CHECK SEVERITY: MEDIUM-DYNAMIC
CHECK PARM:THRESHOLD_LOW(10%),THRESHOLD_MED(20%),THRESHOLD_HIGH(30%)


IWMH103I CPU consumption in SYSSTC service class is 1%. The low
threshold of 10% is not reached. No action is required.

END TIME: 02/04/2026 16:36:11.812851  STATUS: SUCCESSFUL
```

NEW SRB USE CASE WITH Z/OS 3.2
DYNAMIC I/O ACTIVATE

New SRB Use Case for Dynamic I/O Activate

SRB Enhancements

Boosts are: 1. Speed: subcap can run fullcap 2. zIIP: allowing workload onto zIIPs On by default in IEASYSxx BOOST=SYSTEM	SMP/E FIXCAT IBM.Function.SystemRecoveryBoost — Updated 10 April 2025										
	z16 and z17										
	z15							z16 and z17			
	System Recovery Boost			Recovery Process Boost* at MCL P46602.005 for IBM z15 Driver 41C (Bundle S29)				Recovery Process Boost*			
	IPL Startup	Standalone Dump (no zIIP boost)	Shutdown	Sysplex Partitioning – planned or unplanned removal	CF Structure Recovery – rebuild or duplex	CF Datasharing Member Recovery-disconnect or failure from locking resources	Planned/Unplanned HyperSwap	SVC Dump	STC start/restart	HyperSwap configuration load	Dynamic I/O Activate
Intended Duration	60 min	60 min	30 min	2 min	1 min	1 min	2 min	2 min	5 min	2 min	2 min
Basis of use	Auto	Auto	S IEASDBS	Auto	Auto	Auto	Auto	CHNGDUMP SET,SDUMP, RPBMINSZ threshold	WLM service definition BOOST attribute	Auto	Auto
z/OS 2.3	PTF	PTF	PTF	PTF	PTF	PTF	PTF				
z/OS 2.4	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	
z/OS 2.5	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	
z/OS 3.1	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	PTF
z/OS 3.2	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

*Recovery Process Boost limited to 30 min per day per LPAR in aggregate. In V2.4 and higher, can be enabled or disabled with S IEASRB,CLASS=RP,REQ=DISABLE|ENABLE

GDPS provides configuration and orchestration parallelization in GDPS V4R2 and higher.

New SRB Use Case for Dynamic I/O Activate

- New system recovery boost for dynamic I/O activation
  - Classified as Recovery Process Boost
  - Duration of 2 minutes
- Enabled by default via the BOOST=SYSTEM parameter in IEASYSxx
  - Automatically triggered based on certain conditions and criteria to accelerator dynamic I/O activate processing
- Use the display command D IPLINFO,BOOST,STATE to verify active boosts
- Requirements:
  - z/OS 3.2
  - z/OS 3.1 with PTF OA66837

SRM LOCK CONTENTION RELIEF

SRM Lock Contention Relief

- Currently, SRM services and timed algorithms run concurrently
- Larger systems with large transaction volumes have significantly larger amounts of SRM lock contention
  - Especially true for systems with large enclave transaction volume
  - Prior to z/OS 3.2, multiple recommendations for increasing RMPTTOM value
- Algorithm is restructured and enclave code path is reduced, providing performance improvements for system overall
  - The larger the system, and the larger the enclave volume, the bigger the improvement will be
  - Significantly reduces CPU time for systems with high concurrent Db2 trx. processing

Workload Management Update for IBM Z and z/OS – SRM Lock Contention Relief and More
Simon Flaig
1:15pm

Summary

- z/OSMF WLM Policy Advisor
- SYSSTC Health Checker Rule
- New SRB Use Case for Dynamic I/O Activate
- SRM Lock Contention Relief
- New SMF Records

Experience more with IBM

Visit us at the IBM Booth #113

After a full day of technical sessions, take a break with us!

Connect with our experts, snap a photo with the z17 Plexi or the latest Telum II, and get an up-close look at our Spyre Accelerator.

Come back each day for fresh topics and demos at our expert stations.

Think 2026

Join 5000+ senior business and technology leaders who are seizing the AI revolution to unlock unprecedented growth and productivity at Think 2026.

Find out more information using the QR code below.

IBM Digital Asset Haven

IBM Digital Asset Haven is the operational backbone for financial institutions and regulated enterprises entering the digital asset economy.

Find out more information using the QR code below.

Thank You

- Questions?

Your feedback is important!

SHARE
EDUCATE • NETWORK • INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation

APPENDIX

APPENDIX A
CPENABLE Default Change with z/OS 3.2

- Starting with z/OS 3.2, the default value for OPT parameter CPENABLE is changed from (0,0) to SYSTEM
- CPENABLE=SYSTEM automatically sets IBM recommended low and high thresholds according to the model of the CPC on which the LPAR is running
  - Recommendation for the past few generations is to set CPENABLE=(5,15) to limit the number of CPs handling interrupts
  - CPENABLE=SYSTEM as the default removes the need to have to update this parameter with hardware migrations as it will take the IBM recommended values

- Recommendation
  - If you were accepting the default of (0,0) in the past to enable all CPs for I/O interrupts, prepare for the z/OS 3.2 change by setting CPENABLE=SYSTEM
  - If you had installation specific requirements for having all CPs handle interrupts, be sure to explicitly set it as it is no longer the default

APPENDIX B
New and changed SMF fields for z/OS 3.2

Of performance note:

- New fields provided in SMF 30 to allow instrumentation on the high-water mark of virtual private memory usage
  - SMF30_HWMMemlimitMb is added to Storage and Paging section of SMF30
  - Issue before: Lack of instrumentation for the maximum amount of high virtual memory charged against the MEMLIMIT makes it difficult to assess an appropriate value for the MEMLIMIT
  - APAR OA67892: https://www.ibm.com/support/pages/apar/OA67892

Table 1. New and changed System Management Facilities (SMF) records for z/OS 3.2

SMF record	z/OS element or feature	Description	Documented in
Type 30	MVS	– Bit 6 in the SMF30_RAXFLAGS field is now defined in the Storage and Paging section. – The SMF30_HWMMemlimitMb field is added in the Storage and paging section. The record level in SMF30RLN is changed to X'1E0'.	MVS
Type 72 subtype 3	MVS	– The R723RST2G field is added at offset 816 in the Service/report class period data section. The subtype 3 record level indicator in the SMF72SRL field changes to X'90'.	MVS
Type 74 subtype 5	MVS	– The R7451CT7 field is added at offset 84 (formerly reserved space) in the RAID Rank/Extend Pool Data section.	MVS
Type 74 subtype 8	MVS	– Three new fields are added starting at offset 104 in the Synchronous I/O Link Statistics section. The subtype 8 record level changes to X'90'.	MVS
Type 88 subtype 1	MVS	– The descriptions of the SMF88LIB and SMF88LAB fields are updated in the Log stream section. – The SMF88LTDE field is added in the Log stream section.	MVS
Type 88 subtype 11	MVS	– The SMF88ATKE field is added in the Structure alter section.	MVS
Type 119, TCP/IP profile record (subtype 4)	Communications Server	This flag indicates that AIASSIST BATCHOUTBOUND has been configured on the GLOBALCONFIG statement of the TCP/IP profile, and the TCP/IP instance will begin recording training data for the AI model. **Reason for change:** AI-powered Network outbound packet batching	Communications Server
Type 1157–1161	z/OS Data Gatherer	– New record types to support the z/OS OpenTelemetry Emitter. (APAR OA66345, which also applies to z/OS 3.1)	MVS

APPENDIX C
SRB Enhancements for Transient zIIP Processing

- In installations with reserved zIIPs defined for an LPAR, those zIIPs are brought online at the start of a zIIP boost period to provide additional capacity and parallelism, and are taken offline after the boost ends

- With OA66837, if any zIIPs the boost brought online explicitly (transient zIIPs) are later configured online, they will become a normal zIIP and remain online after the end of the boost period
  - For recovery process boosts, at most two logical zIIPs can be brought online this way
  - https://www.ibm.com/docs/en/SSLTBW_3.2.0/pdf/izsb100_v3r2.pdf