



Product Update Bulletin 2005-1

This product update forwards updates to both software and documentation. The software update is Release 15.1 of CPEXPERT. The release number indicates that it is the first release of 2005.

The development work in this update primarily concerns the WLM Component. While updates have been made to the other CPEXPERT component, the WLM Component has been updated to: (1) provide an analysis of PR/SM LPAR definition and PR/SM performance problems; (2) provide additional coupling facility performance analysis; (3) provide a new user interface to allow controlling the performance areas being analyzed by the WLM Component, and (4) provide new or revised rules in other areas.

Additionally, Release 15.1 provides (1) updates to the CICS Component, the DB2 Component, the WebSphere MQ Component, and the DASD Component; (2) provides user links to the SAS Output Delivery System (ODS) html output optionally produced by CPEXPERT, (3) provides updated documentation on CD; and (4) corrects errors that have been reported.

- **Updates to the WLM Component.** The WLM Component has been updated in the following areas.
 - **Provide analysis of LPAR definition and performance in a PR/SM environment.** IBM requires PR/SM with z890 and z990 processors, and supports up to 30 LPARs in a single Central Processor Complex (with an IBM Statement of Direction to support up to 60 LPARs). CPEXPERT Release 15.1 analyzes potential performance problems with LPAR definitions and performance in a PR/SM environment.

CPEXPERT Release 15.1 provides *initial* analysis of PR/SM performance problems. With CPEXPERT Release 15.2 (to be issued in October 2005), I plan to expand this initial PR/SM analysis, and also to analyze Intelligent Resource Director (IRD) performance problems.

The following new rules have been added to the WLM Component¹ as the initial CPEXPERT analysis of PR/SM performance constraints:

- Rule WLM670 (*Coupling facility LPAR used shared processors*) reports that a coupling facility LPAR used a processor that was shared between multiple logical partitions.

¹At present, the PR/SM analysis applies only to users with a MXG performance data base. Within 2 weeks after receiving test data from a NeuMICS performance data base, I will be able to support PR/SM analysis with NeuMICS.

- Rule WLM671 (*Coupling facility LPAR with shared processors had low weight*) reports that a coupling facility LPAR used a processor that was shared between multiple logical partitions, and that the LPAR was assigned a low weight.
- Rule WLM672 (*Coupling facility LPAR weight was too low for CF exploiter*) reports that a coupling facility LPAR used a processor that was shared between multiple logical partitions and that the weight was too low for the CF exploiter (e.g., IMS, DB2, etc.).
- Rule WLM674 (*Coupling facility LPAR was capped*) reports that a coupling facility LPAR was capped.
- Rule WLM800 (*WAIT COMPLETION=YES was specified for LPAR*) reports that an LPAR was assigned time-driven PR/SM dispatching rather than event-driven PR/SM dispatching.
- Rule WLM802 (*LPAR share is higher than assigned number of central processors*) reports that the LPAR weight resulted in a capacity share that was higher than the assigned number of logical central processors, with the result that the LPAR can never use its assigned capacity share.
- Rule WLM803 (*LPAR share is higher than assigned number of special processors*) reports that the LPAR weight resulted in a capacity share that was higher than the assigned number of logical special² processors, with the result that the LPAR can never use its assigned capacity share. This finding is particularly important to installations that have added a zAAP (IFA) special processor or an Integrated Facility for Linux (IFL) special processor.
- Rule WLM805 (*LPAR weights are too similar for LPARs assigned central processors*) reports that two or more LPARs were assigned similar (but not identical) weights, with the result that PR/SM cannot properly manage dispatching of logical central processors to comply with the weights.
- Rule WLM806 (*LPAR assigned weights are too similar for special processors*) reports that two or more LPARs were assigned similar (but not identical) weights, with the result that PR/SM cannot properly manage dispatching of logical special processors to comply with the weights.
- Rule WLM807 (*Capped LPAR has very low percent of a central processor*) reports that the weight assigned to a capped LPAR resulted in a low share of the central processor pool ("Pool 1"). After considering the number of logical processors assigned, the individual logical processors had a very low percent of a central physical processor.
- Rule WLM808 (*Capped LPAR has very low percent of a special processor*) reports that the weight assigned to a capped LPAR resulted in a low share of the special processor pool ("Pool 2"). After considering the number of logical special

²Special processors are Internal Coupling Facility (ICF) processors, Integrated Facility for Linux (IFL) processors, or Integrated Facility for Applications (IFA) processors. PR/SM manages special processors separately from central processors.

processors assigned, the individual logical special processors had a very low percent of a special physical processor.

- Rule WLM809 (*LPAR weight resulted in less than 5% of a central processor*) reports that the weight assigned to an LPAR resulted in a processor share that was less than 5% of a central processor. Such a low share of a central processor could cause the LPAR to experience serious performance degradation.
- Rule WLM810 (*LPAR weight resulted in less than 5% of a special processor*) reports that the weight assigned to an LPAR resulted in a processor share that was less than 5% of a special processor. Such a low share of a special processor could cause the LPAR to experience serious performance degradation. This problem is common when zAAP processors are added to a configuration, because the ZAAP logical processors inherit LPAR characteristics associated with the LPAR's logical central processors.
- Rule WLM811 (*PR/SM time slice was not dynamically computed*) reports that the PR/SM time slice (or dispatch interval) was not dynamically computed by PR/SM.
- Rule WLM820 (*PR/SM Ready Queue caused significant delay to work in LPAR*) reports that queue delays on the PR/SM Ready Queue caused significant delays to dispatching logical central processors to a physical central processor. These queue delays could be from logical central processors recovering from a normal Wait State, or could be from logical central processors having been involuntarily intercepted by PR/SM. This finding applies only with z/Architecture systems.
- Rule WLM821 (*PR/SM Ready Queue caused significant delay to work in LPAR*) reports that queue delays on the PR/SM Ready Queue caused significant delays to dispatching logical central processors to a physical central processor. These queue delays could be from logical central processors recovering from a normal Wait State, or could be logical central processors having been involuntarily intercepted by PR/SM. This finding applies only with systems that are not z/Architecture, and there is insufficient information to distinguish between recovering from a normal Wait State or PR/SM involuntarily intercepting the logical central processor.
- Rule WLM825 (*PR/SM Ready Queue caused significant delay to work in LPAR*) reports that queue delays on the PR/SM Ready Queue caused significant delays to dispatching logical special processors to a physical special processor. These queue delays could be from logical special processors recovering from a normal Wait State, or could be logical special processors having been involuntarily intercepted by PR/SM. With the special processors, there is insufficient information at present to distinguish between recovering from a normal Wait State or PR/SM involuntarily intercepting the logical special processor.
- Rule WLM830 (*Possible "Short Engine" effect exists with LPAR*) reports that the "short engine" effect might exist with the LPAR. In the case of a single TCB (such as a CICS region quasi-reentrant TCB), an application can be seriously delayed when (1) the CPC becomes busy such that PR/SM enforces weights, and (2) an

LPAR has a large number of logical processors assigned such that a relatively low physical processor share is “guaranteed” to any specific logical processor. The IBM Washington System Center staff has described this situation as the “short engine” effect.

- **Added additional new rule.** The following new rule has been added to the WLM Component to analyze additional performance issues:
 - Rule WLM091³ (*FICON Director statistics were collected by multiple systems*) notifies users that the FICON Director switch statistics collection was run on multiple systems in a sysplex, causing unnecessary overhead and potentially causing serious I/O conflicts.
- **Revised the user interface to the WLM Component to require explicit selection of the analysis to be done⁴.** When the WLM Component was developed in 1994, its primary function was to analyze the WLM Service Definition for common problems, to analyze the actual performance of service class periods (i.e., detect when Performance Goals were missed), and to identify the reasons the performance goals were missed.

In the 10 years that have lapsed since the WLM Component was introduced, I’ve continually added new analysis to the WLM Component. I added analysis of coupling facility and XCF problems, analysis of system logger problems, analysis of paging subsystem problems, analysis of IOP/SAP problems, analysis of zAAP problems, and (with this release) analysis of LPAR and PR/SM problems.

The new analysis was developed because that is how I spend my time: researching and developing analysis in areas where there seems to be performance problems that CPEXpert can identify. This new analysis was added to the WLM Component mostly because I felt that the analysis needed to be done, and the WLM Component seemed a suitable place to put the code and documentation.

One result of adding this analysis is that the WLM Component has grown significantly in terms of what it analyzes. Unfortunately, another result is that the amount of output that can be generated has become burdensome to review by CPEXpert users. This amount of output is irksome if all a CPEXpert user wants to do is to use the WLM Component to analyze problems in a specific system area (such as the coupling facility area) to see whether CPEXpert will identify performance problems in that specific area.

To solve this problem, I have restructured the WLM Component internally to analyze different areas, based on user-provided guidance. The following new guidance variables in USOURCE(WLMGUIDE) control which system areas the WLM Component will analyze in any specific execution of the WLM Component:

³Thanks to **San Knutson** (GEICO, MD) for suggesting this analysis.

⁴Thanks to **Kris Ferrier** (State of Washington, WA) and **William “Nick” Nicholson** (Bank of America, VA) for suggestion this feature.

VARIABLE	DESCRIPTION
CHKGOAL	Analyze service class periods to determine reasons goals were missed?
CHKINIT	Analyze problems with WLM-managed initiators?
CHKIOP	Analyze problems with IOP/SAP imbalance?
CHKLOGGER	Analyze problems with System Logger?
CHKPAGE	Analyze problems with paging subsystem?
CHKPLCY	Analyze problems with WLM service policy definition?
CHKRPT	Analyze report class periods to determine reasons goals were missed?
CHKPRSM	Analyze problems with LPAR definitions and other PR/SM issues?
CHKSYS	Analyze problems with SYSTEM/SYSSTC excessive CPU use?
CHKXCF	Analyze problems with coupling facility and XCF?
CHKZAAP	Analyze problems with zSeries Application Assist Processors (zAAP)?

NOTE: The default specification with this release of CPEXpert is for the WLM Component to analyze **only** the WLM Service Definition and **only** reasons that service class periods missed their performance goals. I decided on this as the initial analysis because new users sometimes become overwhelmed by the number of problems that are found in their environment. Also, this was the initial analysis provided with the WLM Component when it was developed.

You must explicitly specify **%LET variable = Y** (where “variable” is one of the above guidance variables) if you wish the WLM Component to analyze other areas. Any combination of analysis can be specified. The WLM Component will provide a listing identifying any analysis area that is not selected, so you will be alerted that the analysis was not done.

If you suppress analysis of any system area, CPEXpert will not process RMF data related to that area and will suppress all analysis of data related to the area. Since the analysis is suppressed, no findings related to that area will be saved in the CPEXpert HISTORY file. Please send me a note if you don't like this approach.

- **Changed rules.** Rule WLM066 and Rule WLM068 (Excessive CPU time in SYSTEM and SYSSTC, respectively) have been changed to compute the percent CPU as a function of the total CPU Busy time of all engines in the LPAR, rather than as a function of the RMF interval time. This approach is a much better way of analyzing CPU Used by the SYSTEM and SYSSTC service class than the approach I had originally implemented, since the value is relative to CPU time actually used.

My original method of analyzing the percent CPU used by SYSTEM and SYSSTC did not consider the number of engines assigned to the LPAR, and did not consider the

PR/SM dispatch time to the engines. Analyzing the CPU Used as a function of the amount of CPU used by the LPAR eliminates⁵ these problems completely.

- **Added new option to analyze performance aspects of report classes.** The WLM service definition can identify work via the workload classification scheme and assign the work to report classes (in addition to service class). In past, CPEXpert analyzed performance problems only of service class periods. With this new option, CPEXpert can analyze performance problems of report classes (either all report classes, or selected report classes). CPEXpert can identify causes of poor performance experienced by specific workloads, as identified by report classes.

NOTE: For this technique to work properly, report classes must be homogeneous (that is, a report class must describe a subset of a service class). If a report class is heterogeneous (that is, a report class describes more than one service class), the Workload Manager cannot properly identify the Performance Goal and Goal Importance associated with the report class. This is because the underlying service class periods can have different goal specifications and goal achievement. Consequently, a Performance Index might not be properly computed and CPEXpert's analysis could be flawed.

- **Added an option to analyze service class periods based on Goal Importance.** This option allows users to specify the Goal Importance for which they wish CPEXpert to analyze the performance of service class periods. With this option, CPEXpert can be directed to analyze only important service class periods.
- **Reduced run time for WLM Component.** I've redesigned some of the data extraction and analysis modules to reduce the run time of the WLM Component. This new design takes advantage of SMF variables available in later versions of z/OS. If the z/OS release is V1R3 or later, CPEXpert can skip processing some large SMF files. These new algorithms should reduce the run time considerably (a 71% reduction in run time was achieved when processing data from a very large site).
- **Updates to the DB2 Component.** The main changes to the DB2 Component for CPEXpert Release 15.1 are:
 - Added Rule DB2-227 (*Buffer pool "hit" ratio was less than expected for buffer pool*)⁶ to allow users to specify a buffer pool hit ratio that signals potential problems for specific applications.
 - Added Rule DB2-237 (*GETPAGE rate was higher than expected for buffer pool*)⁶ to allow users to specify a GETPAGE rate that signals potential problems for specific applications.

⁵Thanks to **Craig Collins**, State of Wisconsin, for suggesting this revision to CPEXpert's analysis approach.

⁶Thanks to **Gautam Chakravarty** (Lockheed Martin, MD) for suggesting this analysis.

- Added Rule DB2-270 (*Available virtual storage was dangerously low for DBM1*)⁷ to report that the DB2 DBM1 address space was in danger of exhausting virtual storage. The result of running out of virtual storage could mean an abend of DB2 transactions, could mean that DB2 enters Recovery, or could mean that the DB2 subsystem terminates. This analysis uses the new DB2 IFCID 225 trace records available with DB2 Release 7
- Added an option to analyze DB2 performance problems for only those data sharing members belonging a specific DB2 data sharing group⁸.
- **Updates to the CICS Component.** With Release 15.1, a “wild card” feature⁹ has been added to allow users to select CICS regions based on leading characters. This allows users who have CICS region naming standards to select production or test regions (for example), by specifying the leading characters associated with production regions.
- **Updates to the DASD Component.** With Release 15.1, the documentation for CPEXpert MXG modification has been revised to conform with MXG current standards for user or vendor modifications.
- **Updates to the WebSphere MQ Component.** With Release 15.1, the WebSphere MQ Component has been updated only by correcting various errors identified by users. I’m eagerly awaiting suggestions for additional analysis (and I still need test data from WebSphere MQ users running shared queues).
- **Updates that apply to all CPEXpert components.** The following updates apply to all CPEXpert components:
 - **Added an option to select information to analyze based on CEC serial number.** This option allows users to specify a CEC serial number to restrict analysis to only the specified CEC. This option is important with large installations with more than one central processor complex (CPC), and is particularly useful with the WLM Component’s analysis of PR/SM.
 - **Added %xxxHTML macros to USOURCE.** I added a macro for each CPEXpert component to USOURCE. This allows users to make any modifications they wish to the SAS ODS HTML statement for any CPEXpert component. Users can concatenate USOURCE ahead of SOURCE to override the default %xxxHTML macros that are distributed in SOURCE.

I decided on this approach because there was becoming an unwieldy proliferation of user-requested modifications to the default SAS ODS HTML statement. I became concerned that describing all these options would simply confuse new CPEXpert users. Consequently, the SAS ODS HTML statement has been placed into a separate

⁷Thanks to **FE (Fred) Nijdam** (Rabobank, Netherlands) for suggesting this analysis!

⁸Thanks to **FE (Fred) Nijdam** (Rabobank, Netherlands) for suggesting this feature!

⁹Thanks to **Marty Stahl** (Social Security Administration, Maryland) for suggesting this feature!

module for each CPEXpert component, and users can make whatever modifications they wish to those modules in USOURCE. See the *CPEXpert Installation Guide* for details.

- Added the **%LINK2** user macro¹⁰, so users can optionally create html code as an addition to the html link placed in CPEXpert output that links to the CPEXpert rule documentation. The code in the %LINK2 macro can be used to link to user-generated graphics, for example, that provide graphical information related to the CPEXpert finding. See the *CPEXpert Installation Guide* for details (this feature applies to all CPEXpert Components).
- Updated the *CPEXpert Installation Guide* to provide installation and operation instructions for users running CPEXpert on a Windows platform. Some CPEXpert users are moving their SMF processing and analysis from a mainframe to a Windows platform. The *CPEXpert Installation Guide* has been updated to provide documentation on how to run CPEXpert in this Windows environment.
- **Updated documentation on CD.** This release provides all updated documentation to reflect new or revised rules. The CPEXpert documentation is accessible via Adobe Acrobat Reader. A free version of Adobe Acrobat Reader is available at <http://www.adobe.com/products/acrobat/readstep2.html>.
- **Add log of changes to code.** I've added a log of changes to the CPEXpert code for Release 15.1. This log is titled CHANGES and shows each module changed (except when the change is minor). Interested users can review the CHANGES log to see what code changes have been made, and can easily identify new guidance variables. I've followed the MXG model for releases (that is, the CHANGES log show changes in the new release, while the CHANGESS log is an accumulation of changes).
- **Add log of errors reported.** I've added a log of all errors reported by users. This error log is in the GENER142 module in SOURCE.

FTP download

CPEXpert Release 15.1 (both code and documentation) is available via ftp download from my ftp server. If you prefer this method of obtaining the update, please send an email to Don_Deese@cpexpert.com to obtain instructions for ftp download, obtain your user id, and obtain your password.

Installation

I suggest that you use the following steps to install Release 15.1:

¹⁰Thanks to **FE (Fred) Nijdam** (Rabobank, Netherlands) for suggesting this feature. Fred has implemented code to link CPEXpert output to amazing graphics, displaying "in context" graphical information for management!

- Create a new PDS titled "prefix.CPEXPERT.V151.SOURCE".
- Create a new PDS titled "prefix.CPEXPERT.V151.USOURCE".
- Install CPEXpert into the "prefix.CPEXPERT.V151.SOURCE" using the normal installation procedures described in the *CPEXpert Installation Guide*.
- Copy your **old** USOURCE members into "prefix.CPEXPERT.V151.USOURCE". This step should be done so you do not have to recreate all of your unique parameters¹¹.

Please remember that you **MUST** change any of the CHKxxxx guidance variables if you wish the WLM Component to analyze areas other than the basic analysis of the Service Policy and service class period performance constraints.

- If you elected to receive updates to CPEXpert on CD, the software is distributed on the CD that contains the CPEXpert user documentation. With this release, I have adopted the MXG style of providing CPEXpert updates on CD:
 - **Installing CPEXpert from CD to MVS.** Copy the *ebc151.ebc* file from `cd:\source\ebc151.ebc` as binary, to CPEXPERT.TAPE.REL151, allocated as `RECFM=FB,LRECL=80,BLKSIZE=0,AVGREC=M,SPACE=(80,2,1)`. The *ebc151.ebc* file is in standard IEBUPDTE format.

Use the JCL described in Exhibit 2-1 of the *CPEXpert Installation Guide* to install the CPEXpert software. Copy your old USOURCE members as described above.

- **Installing CPEXpert from CD to PC:** Copy the SOURCE directory on the CD to `prefix\cpxpert\rel151\source`. Copy your old USOURCE members as described above. Make sure that you update your USOURCE(GENGUIDE) member to specify **%LET SOURCE =user_directory;** statement to point to the new SOURCE directory.
- If you have installed the CPEXpert option to produce output to SAS Output Delivery System (ODS) and if you have exercised the LINKPDF option so you can "click" on the rule to see the documentation, **please remember to load the new documentation onto your system.**
- If you chose to implement the xxxHTML option (e.g., WLMHTML) to provide your own SAS ODS HTML statement parameters, please place your own macros into USOURCE and concatenate the USOURCE pds ahead of the SOURCE pds.

Special note of SAS/ITRM users of CPEXpert: With CPEXpert Release 15.1, you must have **MXG TYPE747P** file structure (whether populated or not) in your performance data base. SAS gives an invalid return to the `%sysfunc(exist(file))` macro with SAS/ITRM, so

¹¹The exception to this statement is if you chose to use the "yesterday" option in selecting data. The "yesterday" calculation is done by a PREVIOUS macro contained in USOURCE(GENGUIDE). The PREVIOUS macro is placed in USOURCE so you can select "yesterday" as the immediate previous day, or can define "yesterday" as any other previous day you chose. Placing the PREVIOUS macro in USOURCE(GENGUIDE) allows you to make such a modification, without your altering a SOURCE member. Consequently, if you copy your old USOURCE(GENGUIDE) member to CPEXPERT.V14.USOURCE, you may wish to append the PREVIOUS macro from the Release 15.1 USOURCE into your USOURCE(GENGUIDE) member.

TYPE747P must be in your pdb, even if it is a null file. If you chose to “turn off” Rule WLM091 (*FICON statistics collection was run on multiple systems*), MXG TYPE747P will not be required.

Thanks

I would like to say "**Thank you**" to the following individuals. These folks have discovered errors, proposed new features, graciously sent me test data, or suggested documentation changes since the last Update Bulletin:

Markus Bansemir (HUK-Coburg, Germany)
James Barton (Department of Veterans Affairs, TX)
Gerard Bosker (Rabobank, Netherlands)
Glenn Bowman (Wakefern Food Corporation, NJ)
Craig Collins (State of Wisconsin, WI)
Michael Duffy (Lloyds TSB, United Kingdom)
Kris Ferrier (State of Washington, WA)
Wayne Hammond (Bank of America, GA)
Sam Knutson (GEICO, MD)
Hugh Lapham (Royal Canadian Mounted Police, Canada)
Barry Merrill (Merrill Consultants, TX)
Dean Montevago (Visiting Nurse Service of New York, NY)
William (“Nick”) Nicholson (Bank of America, VA)
Alex Torben Nielsen (Computer Sciences Corporation, Denmark)
Fred Nijdam (Rabobank, Netherlands)
Bryant Osborn (Bank of America, VA)
John Os den (Rabobank, Netherlands)
Keli Prado (Farmers Insurance Group, CA)
Harald Seifert (HUK-Coburg, Germany)
Wayne Schumack (United Health Technologies, MN)
Al Sherkow (I/S Management Strategies, Ltd., MN)
Marty Stahl (Social Security Administration, MD)
Andreas von Imhof (Rabobank, Netherlands)
Kathy Walsh (IBM, Washington Systems Center, MD)
Bruce Widlund (Merrill Consultants, TX)

Please send me an email or phone if you have suggestions, you want new features, or you would like to see more or different reporting done by CPExpert.

Best regards,

Don Deese

Computer Management Sciences, Inc.
6076-D Franconia Road
Alexandria, Virginia 22310
(703) 922-7027 FAX: (703) 922-7305
www.cpexpert.com
Don_Deese@cpexpert.com