
Rule WLM602: XCF message buffer length may be too small

Finding: CPExpert has determined that a large percent of the cross system coupling facility (XCF) messages were larger than the value specified in the CLASSLEN associated with the transport class.

Impact: This finding can have a LOW IMPACT or MEDIUM IMPACT on the signalling performance of the sysplex.

Logic flow: This a basic finding. There are no predecessor rules.

Discussion: The XCF component of MVS/ESA allows authorized programs on one MVS system in a sysplex to communicate with programs on the same system or on other systems. A typical example of this communication is between CICS regions; CICS regions often communicate with other CICS regions in the same system or with CICS regions on other systems in the sysplex.

Within the XCF terminology, authorized programs are termed *XCF members*, and the XCF members are logically a part of specific *XCF Groups*. XCF group members communicate with each other using the XCF *signalling* mechanism.

Optimal signalling performance requires that XCF groups have access to adequate signalling resources. These resources consist of signalling paths and buffers. A *transport class* is the mechanism used by MVS to allow resources to be assigned to XCF groups. Resources (signalling paths, buffers, etc.) are assigned to one or more transport classes, and XCF groups are assigned to the transport classes. Thus, resources can be made available to the XCF groups as they are needed.

The two major transport class resources to be tuned are (1) the message buffers assigned to transport classes and (2) the number of signalling paths assigned to transport classes.

The following discussion relates to the *message buffers*. Other rules in the WLM600(series) relate to the signalling paths.

Please refer to the discussion associated with Rule WLM601 for additional information about XCF concepts.

Message buffers are assigned to transport classes¹ in two ways: (1) the basic assignment to the transport class via the CLASSLEN and MAXMSG parameters on the CLASSDEF statement and (2) the MAXMSG parameter on the PATHOUT statement.

- The CLASSLEN parameter defines the message length for the transport class. MVS allocates *fixed-length buffers* at the size specified in the CLASSLEN parameter for the transport class.

If no CLASSLEN parameter is specified, MVS uses the value of the CLASSLEN parameter specified on the COUPLE statement (with a default value of 956 bytes).

The message length specified by the CLASSLEN parameter should be large enough to accommodate most messages, but not so large as to waste storage. Selecting the correct buffer length is a tradeoff between (1) overhead incurred by having buffers too small, (2) wasted storage incurred by having buffers too large, and (3) the performance implications of mixing large and small messages in the same transport class.

- If the fixed-length buffers are too small to hold a message, MVS acquires additional buffers to accommodate the message. Increased system overhead is caused when MVS must acquire additional buffers.

In order to minimize this overhead, MVS may dynamically increase the length of the buffers if (1) the number of over-sized messages message traffic warrants the increase and (2) the increase in buffer length would not exceed the maximum buffer space specified on the receiving system.

- If the buffers are too large for a message, the unused storage remaining in the buffer is wasted. This is an inefficient use of storage. Additionally, MVS could exhaust the supply of buffer space associated with a transport class if the space is wasted by specifying a buffer length that is too large for most messages. In the later case, XCF messages would be rejected if the supply of buffer space is exhausted.
- If large and small messages are mixed in the same transport class, the small messages tend to be delayed simply because the large messages take longer to process.

¹Message buffers are assigned to transport classes only for **outbound** traffic since only outbound traffic can be separated into transport classes. Inbound traffic cannot be separated by transport classes; buffers are assigned to inbound traffic based on the total buffer space defined on the PATHIN statement.

-
- The MAXMSG parameter defines the amount of message buffer space allocated for messages sent in the transport class. The MAXMSG parameter can be specified on the PATHOUT or PATHIN statements, or on the CLASSDEF statement.

If no MAXMSG value is specified for the paths associated with a transport class or for the CLASSDEF statement, MVS uses the value of the MAXMSG parameter specified on the COUPLE statement (with a default of 750K bytes of buffer space).

SMF Type 74 (Subtype 2) provides statistics about the number of messages sent by XCF groups in a transport class, where the messages are sent, how many messages were too small for the defined buffer size, how many messages fit the defined buffer size, how many messages were too big for the defined buffer size, and how many messages were over the message length for which XCF was optimized.

CPEXpert analyzes this information to determine whether the correct buffer allocation has been defined. CPEXpert computes the total outbound message traffic for a transport class. CPEXpert concludes that the message length specified for the transport class is too small when a significant percent of the messages were **larger** than the buffer length specified for the transport class and a significant percent of these messages caused overhead.

The value considered a "significant percent" of the large messages is controlled by the **PCTBIG** guidance variable. Please refer to Section 2 of this document for a discussion of the PCTBIG guidance variable.

Additionally, CPEXpert applies a "reality check" by ensuring that a reasonable number of messages were sent in the transport class.

When the above conditions are met, CPEXpert produces Rule WLM602 to alert you that there is a mismatch between the buffer length specified for the transport class and the lengths of messages sent in the transport class. Rule WLM602 is not produced for measurement intervals in which Rule WLM601 is produced.

The following example illustrates the output from Rule WLM602:

RULE WLM602: XCF MESSAGE BUFFERS MAY BE TOO SMALL

The XCF message buffer length may be too small for the DEFAULT transport class. Unnecessary overhead is incurred when XCF must expand the buffers to fit a message. The CLASSLEN parameter was specified as 16,316 for this transport class. You should consider increasing the message length for this transport class or you may wish to split the transport class, depending upon actual message lengths. This finding applies to the following RMF measurement intervals:

MEASUREMENT INTERVAL	SENT TO	SMALL MESSAGES	MESSAGES THAT FIT	MESSAGES TOO BIG	TOTAL MESSAGES
10:00-10:30,26MAR1996	J80	2,654	0	462	3,116
11:30-12:00,26MAR1996	J80	3,006	1	381	3,388
12:00-12:30,26MAR1996	J80	3,481	0	493	3,884
12:00-12:30,26MAR1996	Z0	2,943	6	472	3,421

Suggestion: If Rule WLM602 is regularly produced, CPExpert suggests that you consider the following alternatives²:

- You should evaluate the message length specified for the transport class and the message lengths of the XCF groups assigned to the transport class. You should consider using the CLASSLEN parameter of the CLASSDEF statement to increase the message length for the transport class.

With z/OS Version 1 Release 2 (V1R2), Message IXC344I has been changed to provide more insight into the requirements of transport classes. In response to a DISPLAY XCF,CLASSDEF command, Message IXC344I displays detailed data for specific transport classes. With z/OS V1R2, the message has been enhanced to provide counts of messages sent at **each different signal size that was used**. By examining the count of messages sent at the appropriate signal size, you can determine whether the transport class should be split, and what the new sizes should be.

If most of the outbound messages do not fit the buffer lengths, it normally is better for the buffer lengths to be slightly larger than the outbound messages. A small amount of wasted storage usually has less performance impact than the unnecessary overhead caused by messages being larger than the buffer length.

- You should evaluate whether the XCF groups are properly assigned to transport classes. XCF groups are assigned to transport classes via the GROUP parameter on the CLASSDEF statement.

²**WARNING:** There exists little practical experience with analyzing coupling facility data and with selecting proper values for the controlling parameters. The CPExpert analysis and suggestions are based on (1) the information contained in the referenced documents and (2) our analysis of data provided by IBM or CPExpert users. Please keep this paucity of knowledge in mind when considering the alternatives. Additionally, **please** provide Computer Management Sciences with feedback!

-
- XCF groups can be assigned to more than one transport class. When evaluating which transport class to use (when XCF groups are assigned to more than one transport class) XCF will select the transport class with the smallest buffer that will hold the message being sent. You potentially can "optimize" the buffer space used by assigning XCF groups to more than one transport class.

All groups assigned to a transport class have equal access to the signalling resources of that class. Consequently, you should make sure that you do not assign "low priority" groups to transport classes that have high performance requirements if the "low priority" groups could cause performance degradation to the "high priority" groups.

Fortunately, SMF Type 74 (Subtype 2) records contain information about the XCF groups and XCF members, including the number of signals sent and received by each member. This information is in the **Member Data Section** of the Type 74 records, and can be analyzed to assess the impact of message traffic of the XCF members and XCF groups.

- Alternatively, it may be preferable to reassign XCF groups to transport classes. In practice, this situation is unlikely to occur as most installations will have a relatively small number of transport classes.
- You can adjust CPEXpert's analysis by altering the value specified for the **PCTBIG** guidance variable in USOURCE(WLMGUIDE). The default value for PCTBIG is intended to cause Rule WLM602 to be produced when more than a modest number of the messages cause MVS to incur unnecessary overhead for over-sized messages.
- If Rule WLM602 occurs frequently and there is no action you wish take, you can exclude the transport class from CPEXpert's analysis, using the **EXCLASSn** guidance variables. The EXCLASSn guidance variables allow you to exclude one or more transport classes from analysis.

Reference: MVS/ESA: Setting Up a Sysplex (GC28-1449)
Section 5: Planning Signalling Services in a Sysplex

MVS/ESA: Initialization and Tuning Reference (GC28-1452)
COUPLExx (Cross-System Coupling Facility Parameters)

OS/390: Setting Up a Sysplex (GC28-1779)
Section 5: Planning Signalling Services in a Sysplex

OS/390: Initialization and Tuning Reference (GC28-1752)
COUPLExx (Cross-System Coupling Facility Parameters)

z/OS: Setting Up a Sysplex (SA22-7625)
Section 5: Planning Signalling Services in a Sysplex

z/OS: Initialization and Tuning Reference (SA22-7592)
COUPLExx (Cross-System Coupling Facility Parameters)

"Parallel Sysplex Performance: tuning tips and techniques,"
Kelley, Joan (IBM, Poughkeepsie, NY), SHARE 86, February 1996.

z/OS V1R2: MVS System Messages, Volume 10 (IXP-IZP), SA22-7640

z/OS V1R3: MVS System Messages, Volume 10 (IXP-IZP), SA22-7640

z/OS V1R4: MVS System Messages, Volume 10 (IXP-IZP), SA22-7640 |