
Rule WLM706: DASD staging data set high threshold was reached

Finding: The SMF Type 88 data showed that the DASD staging data set high threshold was reached.

Impact: This finding has a LOW IMPACT, MEDIUM IMPACT, or HIGH IMPACT on the performance of your computer system. The level of impact depends on the applications using the log stream, and the extent to which log stream delays effects the performance of these applications.

Logic flow: This is a basic finding, based on an analysis of the SMF Type 88 system logger data. The finding applies only to log streams that are defined to use a coupling facility.

Discussion: The system logger is an MVS component that allows an application to log data from a sysplex. You can log data from one system or from multiple systems across the sysplex.

Please refer to Rule WLM701 for more general information about the MVS system logger.

Data in a log stream is contained in two kinds of storage: (1) *interim storage*, where data can be accessed quickly without incurring DASD I/O, and (2) *DASD log data set storage*, where data is “hardened” for longer term access. When the interim storage medium for a log stream reaches a user-defined threshold, the log data is offloaded to DASD log data sets.

There are two types of log streams: coupling facility log streams and DASD-only log streams. The main difference between the two types of log streams is the storage medium system logger uses to hold interim log data:

- In a coupling facility log stream, interim storage is coupling facility list structures.
- In a DASD-only log stream, interim storage is contained in local storage buffers on the system, as an MVS data space areas associated with the system logger address space.

Additionally, for data integrity there exists duplexed storage, so that if one system or component fails, the log stream can be recovered from the duplexed storage. These concepts differ, depending on whether the log stream is defined for a coupling facility or for DASD-only.

-
- If the primary storage is defined as a list structure in a coupling facility, the duplexed data can be retained in another coupling facility, or can be retained in *staging data sets*. Staging data sets are used when the coupling facility is in the same CPC, or uses volatile storage.
 - If the primary storage is defined as DASD-only, the duplexed data is retained in *staging data sets*.

When a log stream in a coupling facility is duplexed to staging data sets, the system logger automatically makes a duplicate copy of the data every time data is written to a log stream. This is done to protect against data loss due to coupling facility problems or due to system failure. The duplicate copy is kept in the staging data sets until the data is off-loaded from the coupling facility structure to DASD log data sets. After the data is off-loaded to DASD log data sets, the system logger discards the duplicate copy of the log data.

Interim storage in a coupling facility structure normally is “offloaded” to DASD log data sets based on two parameters associated with each log stream: the HIGHOFFLOAD and LOWOFFLOAD parameters. The values for these parameters are expressed as a percent of the interim storage being filled. For log streams defined in coupling facility list structures, these parameters apply as follows:

- When the coupling facility structure is filled to the **high offload threshold** point or beyond, the system logger begins offloading data from the coupling facility to the DASD log stream data sets. For example, if the HIGHOFFLOAD parameter is specified as 80% (this is the default value), the system logger normally would begin offloading interim storage to DASD log data sets when 80% or more of the structure is used.
- The **low offload threshold** is the point in the coupling facility structure, as a percent of space consumed, where the system logger stops offloading coupling facility log data to log stream DASD data sets. The default LOWOFFLOAD parameter value is 0%, indicating that the system logger will offload all the log stream to DASD log data sets once offloading has commenced.

From the above description, the amount of data that normally is offloaded is the difference between HIGHOFFLOAD and LOWOFFLOAD, as percentages of the coupling facility list structure size. For example, if the HIGHOFFLOAD value was specified as 80% and LOWOFFLOAD value was specified as 60%, 20% ($80\% - 60\% = 20\%$) of the structure would be offloaded once offloading commenced.

For log streams in a coupling facility that are duplexed to staging data sets, the values of the HIGHOFFLOAD and LOWOFFLOAD parameters **apply**

to the staging data sets as well as to the coupling facility structure. This is simply because if the staging data sets become full, MVS would not be able to continue duplexing data and there would be a data integrity exposure in case of failure. Consequently, if a staging data set fills up **before** an offload of a log stream in a coupling facility structure is triggered by the high threshold specification, an offload will be triggered because of the full staging data set.

When a staging data set reaches the high threshold, the system logger immediately offloads data from the coupling facility to DASD log data sets, even if the coupling facility usage for the log stream is below the high threshold. Thus, if the staging data sets are small in comparison to the coupling facility structure size for a log stream, the staging data sets will keep filling up and the system logger will frequently offload coupling facility data to DASD log data sets. This means that your installation would experience frequent (and unexpected) offloading overhead that could affect performance¹.

IBM's "Setting up a Sysplex" document (Section 9.4.5.6: Monitoring Staging Data Set Usage Log Streams) contains the following comments:

"Whether your staging data sets are defined by system logger or on the STG_SIZE parameter, you should carefully monitor your staging data sets. This applies to both coupling facility and DASD-only log streams, and is important because the consequences of having your staging data set fill up can be quite disruptive. When a system's staging data set fills up, system logger applications on that system will not be able to write to the log stream until log data can be offloaded to DASD, which frees up space in the staging data set. Thus, when your staging data sets are too small, system logger will perform coupling facility offloading more frequently than the HIGHOFFLOAD and LOWOFFLOAD thresholds defined for the log stream would otherwise require. This can negatively affect the performance of all the log streams in that structure."

The MVS system logger writes SMF Type 88 records containing statistics for each connected log stream. This information is available as MXG TYPE88 file.

CPEXpert examines the SMF88STN variable (the structure name) in the MXG TYPE88 data set to select records that apply only to coupling facility structures². For these records, CPEXpert examines the SMF88ETT variable (the number of times the system logger detected a Staging Data Set Threshold hit condition). CPEXpert produces Rule WLM706 when the

¹If your staging data sets are too small, you also run the risk of filling them up completely. If this occurs, system logger immediately begins offloading the coupling facility log data in DASD log data sets to harden it. System logger applications will be unable to log data until system logger can free up staging data set space. This serious situation is evaluated by Rule WLM702.

²The SMF88STN variable will be *DASDONLY* for log streams that are DASD-only log streams.

SMF88ETT value exceeds the **STDSHIGH** guidance variable in USOURCE(WLMGUIDE).

The default value for the **STDSHIGH** is 0, indicating that CPExpert should produce Rule WLM706 whenever a Staging Data Set Threshold was encountered during an RMF interval.

Suggestion: IBM suggests that you size the staging data sets larger than the coupling facility structure size for the log streams.

While you can modify CPExpert's analysis by altering the STDSHIGH guidance variable, you should not do so unless you have unusual circumstances.

Reference: OS/390 Setting up a Sysplex

- OS/390 (V2R2): Section 9.2.4
- OS/390 (V2R3): Section 9.2.5
- OS/390 (V2R4): Section 9.2.6
- OS/390 (V2R5): Section 9.2.6
- OS/390 (V2R6): Section 9.2.6
- OS/390 (V2R7): Section 9.2.6
- OS/390 (V2R8): Section 9.2.6
- OS/390 (V2R9): Section 9.2.6
- OS/390 (V2R10): Section 9.2.6
- z/OS (V1R1): Section 9.2.6
- z/OS (V1R2): Section 9.4.6
- z/OS (V1R3): Section 9.4.6
- z/OS (V1R4): Section 9.4.6