



Measuring IMS Transactions (Part 3)

Giuseppe Giacomodonato – EPV Technologies
Fabio Massimo Ottaviani – EPV Technologies
Enzo Rossi – EPV Technologies

February 2014

7 Input queue time

The input queue time is the time a message spends in the input queue or in the message queue buffers waiting for a message region to become available to process it. In a busy system, this time can become the major portion of the IMS transaction response time.

Based on what we discussed in the previous chapters, measuring the input queue time might look fairly simple. We know the time when the message arrived in the queue (arrival time) and when the transaction started (start time) so the formula to be used should be:

$$\text{input queue time} = \text{start time} - \text{arrival time}$$

Unfortunately this formula works for some transactions but not for all of them.

When a transaction is switched by another one both transactions have the same arrival time and this is true for all the IMS transactions in a chain.

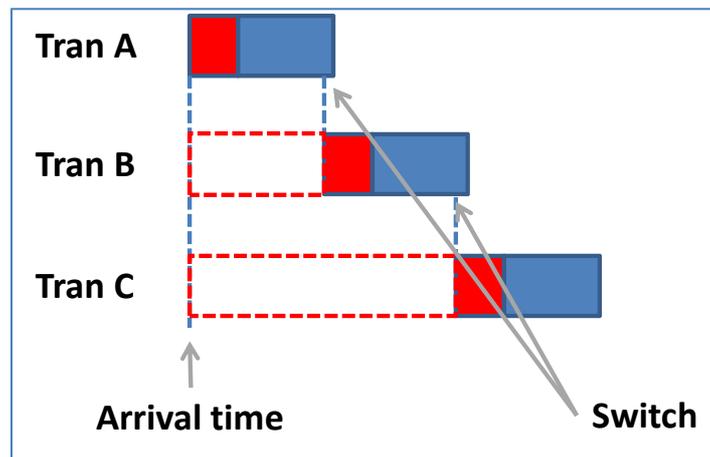


Figure 15

In the example in Figure 15 we have three transactions: Tran A arrives in the queue and after waiting some time (input queue time represented by the red bar) starts its processing (elapsed time represented by the blue bar). Just before ending Tran A switches Tran B.

After some time Tran B starts, however in this case we have the real input queue time, which is the time from the switch to the Tran B start, and an inflated input queue time (represented by both the



red bar and the red dotted line) which is the time from the Tran B arrival (set at the same time as Tran A) and the Tran B start. The same is true for Tran C.

It's important to note that every switched transaction includes in its input queue time both the input queue time and the elapsed time of previous transactions in the chain; so the longer the chain the bigger the error in the input queue time calculation.

Solving this problem would be very easy if the switch time was available in the current metrics provided by 56FA and FA records. Unfortunately this is not the case.

A good solution to this issue is suggested and used in MXG.

```

/* WHEN TRANSACTIONS ARE CHAINED, THE IMF RECORD FOR THE SECOND */
/* AND SUBSEQUENT TRANSACTIONS CONTAIN THE ARRIVAL TIME OF THE */
/* ORIGINAL TRANSACTION, WHICH PRODUCES INCORRECT MEASURES OF */
/* INPUT QUEUE TIME. THIS ALGORITHM SORTS THRU THE CHAIN AND */
/* SET THE ARRIVAL TIME OF THE SECOND TRANSACTION TO THE END */
/* TIME OF THE PRECEDING, LEADING TO VALID INPUT QUEUE TIMES */
/* FOR CHAINED TRANSACTIONS. */
    
```

from VMACCIMS MXG macro

While this method provides much better results than simply making the difference between start time and arrival time, it tends to provide an optimistic measure of the input queue time.

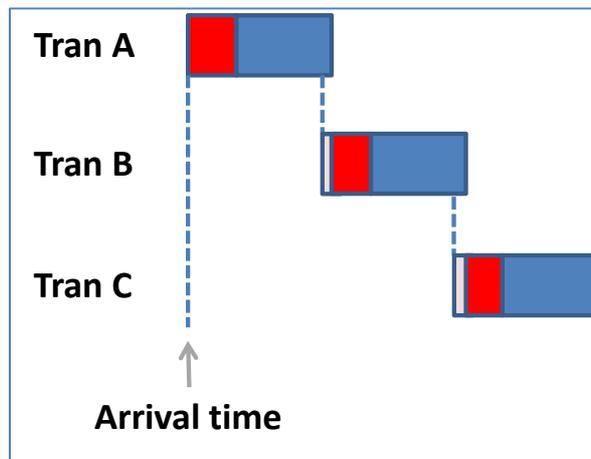


Figure 16

As highlighted in Figure 16, using the end of the previous transaction in the chain will underestimate the input queue; the error (represented by the pink bar) is normally small and corresponds to the remaining elapsed time of the previous transactions after the new transaction switch.

Another important issue to take care about is the identification of switched transaction chains. It can be done essentially by sorting the transactions by arrival time and checking if the arrival time of the current one is exactly the same as the previous occurrence.

Of course it is always possible to have the same arrival time for transactions which are not chained, especially on very busy IMS subsystems, so the precision of the arrival time stamp has to be very high.



As discussed in Chapter 6.2, both 56FA and FA log records provide metrics reaching microsecond precision. It's very important to select them when the method described above is used to calculate the input queue time.

In our tests we got practically the same input queue time values by using 56FA and FA log records.

8 New performance metrics

New performance metrics have been added to IMS 56FA and FA log records.

In the following we will only discuss zAAP/zIIP related metrics; for a complete discussion of all the others please refer to the IBM and BMC standard documentation.

8.1 New zAAP/zIIP metrics in 56FA

Starting from IMS V12¹, application standard CPU and zIIP/zAAP times are collected separately in the 56FA log records. This also means that before V12, and prior to the required maintenance, IMS log records reported zAAP/zIIP² time as part of the overall transaction CPU time.

The zAAP/zIIP time is now collected in the TPEZAAP field while the standard CPU time is collected in TPEXTIME.

TPEZAAP time is normalized to the standard processor speed. So, for knee capped machines, it can show a larger value than the transaction actual elapsed time.

To correctly decompose IMS transactions elapsed time you have to de-normalized TPEZAAP by using the following formula to get the "real" zAAP/zIIP time:

$$TPEZAAP_REAL = TPEZAAP * 256 / R723NFFS$$

R723NFFS is the normalization factor for zIIP available in SMF 72.

Information about zAAP/zIIP eligible time on standard CPU is not provided in 56FA log records.

8.2 New zAAP/zIIP metrics in FA

The situation is similar for the FA log records produced by Mainview for IMS.

Total CPU time, including zAAP/zIIP time, is provided in the TRXZTCPU field while the standard CPU time is provided in the TRXZONCP. So you need to use the following formula to get the zAAP/zIIP time:

$$zAAP/zIIP\ time = TRXZTCPU - TRXZONCP$$

Also in this case the zAAP/zIIP time is normalized to the CPU speed.

Information about zAAP eligible time and zIIP eligible time on standard CPU is provided in the TRXZAACP and TRXZIOCP fields.

¹ Appropriate maintenance (see PM36273) is required; this support is also provided in IMS log record 7.

² IBM announced that zAAP will not supported anymore in future machines; zAAP eligible work will be run on zIIP.



9 Conclusions

In this presentation we discussed the most important metrics available that you can use to measure IMS applications' performance.

We compared the results obtained by using 56FA and FA records showing all the issues you should face when using each one of them.

Finally we briefly described the new metrics made available to measure zAAP/zIIP utilization of IMS applications.