Workload Manager—Interpreting New WLM Measurements

Peter Enrico

IBM Corporation
522 South Road
Poughkeepsie, NY 12601-5400

Abstract

MVS/ESA SP5 provides performance measurement tools with a variety of new measurements useful in evaluating the performance of workloads running in an MVS/ESA environment. With the introduction of the MVS Workload Manager, MVS can now manage workloads according to user-specified goals. This includes goals not only for traditional workloads such as TSO, batch, and started tasks, but also goals for CICS and IMS transactions! With these new measurements users can determine if the goals for their CICS and IMS transactions are being met both on a single system and on the sysplex as a whole. In addition, users can also gain an understanding of where their CICS and IMS transactions spent their time processing. This presentation will help users understand how to use the new measurements made available by the Workload Manager (WLM) component of the MVS/ESA operating system. Discussed will be measurements that could be used to track the performance of CICS and IMS transactions, as well as batch jobs, TSO transactions, and long running address spaces such as started tasks.

Introduction

MVS/ESA SP5 simplifies the definition, control, and reporting of the performance requirements for MVS workloads. It provides a direct way to specify performance goals for work with the functions of the MVS Workload Manager.

To reduce the complexity of managing system resources, MVS workload management provides goal-oriented dynamic resource management. If using the MVS Workload Manager (WLM) an installation defines performance goals for CICS, IMS, JES, APPC/MVS, TSO/E, Distributed DB2, and OpenEdition MVS work on the basis of business importance. A key objective of workload management is to use these goals for dynamically adjusting access to processor and storage resources.

Each release of the MVS/ESA operating system provides additional workload measurements that further assist system analysts in diagnosing performance in their operating systems. MVS/ESA SP5, which introduces the new MVS Workload Manager, is no exception. This paper discusses new performance measurements that are made available by the MVS Workload Manager.

This paper assumes a general knowledge base of the MVS Workload Manager and its concepts and facilities.

MVS/ESA Workload Manager measurements, as discussed in this paper, refer to those measurements that are collected directly from the SRM and WLM components of the MVS/ESA operating system and are reported by products such as RMF to provide performance feedback on system performance. Please note that it is not the intent of this paper to sell the RMF product, but to provide a better understanding of the workload management measurements provided by SRM and WLM. There are other products available that collect and report SRM and WLM measurements. RMF, however, is used throughout this paper because it is the measurement tool used at the author's shop.

This paper is structured as follows:

- Brief overview of workload manager external controls
- Overview of new measurements in the RMF workload activity report to help answer the following questions:
    - What were the desired goals for the work?
    - Were the goals met?
        - What is a Performance Index and why is it so valuable?
        - What was the Actual achieved value versus the Goal value?
    - What response time did the transactions achieve?
        - What is a transaction?
        - What were the components of the average response time?
        - What is the distribution of response times achieved?
    - If a goal was not met, what got in the way of success?
        - Where did the address space transactions spend their time processing?
        - Where did the CICS and IMS transactions spend their time processing?
- A discussion of some possible uses for this data.

2.0 Overview of WLM External Controls

Since WLM is still such a relatively new component of the MVS/ESA operating system, a review of the basic concepts is now provided.

An installation can gather MVS work together into a new grouping called a service class. Users can think of a service class as being similar to the current construct of performance group. Like performance groups, service classes have periods. However, this is where the similarity stops.

Performance group periods are assigned various resource-based controls for such resources as processor, storage, and MPL. These static controls, as specified in the installation's IEAIPSxx and IEAOPTxx members of PARMLIB, can be thought of as all the 'knobs' and 'dials' that allowed an installation to control how and to whom certain resources are allocated.

A primary objective of the MVS Workload Manager is to simplify the management of MVS systems by providing externals to reflect an installation's expectation for work being processed in the system. To allow this, the MVS Workload Manager enables the installation to explicitly state to MVS the service objective, or goal, towards which the work should be managed. These goals are assigned to the periods of each service class.

2.1 Goal types

The MVS Workload Manager supports four different goal types. As stated previously, goals are assigned to service class periods. Each goal type has a unique meaning and implication if used. The four goal types are as follows:

- Response time goal
  - Average response time goal
  - Percentile response time goal
- Discretionary
- Velocity
- System Goal

Response Time Goal
A response time goal is the simplest goal to understand. This goal type can be based on either the average response time of ended transactions, or a response time target for a given percentage of completions (for example, 80% of CICS banking transactions complete within 1 second).

The primary difference between an average response time goal and a percentile response time goal is that an average response time goal is HEAVILY influenced by 'outliers'. That is, a single transaction which has gone amiss can really skew an average. For example, assume 99 transactions all end in 1 second, but one transaction runs for 2 minutes. The average is over 3.5 seconds, even though 99% completed in 1 second. When considering using an average response time goal, decide whether you really want WLM to manage your work based on the worst behaving transactions. Percentile goals are much less influenced by these 'outliers' since they will not fall into the group of transactions used to consider if the goal was met.

As will be discussed later in this paper, you can use the new Workload Manager measurements collected by RMF will help to identify this condition. These measurements can also be used to assist in converting an average response time goal to a percentile response time goal.

Discretionary Goals
A discretionary goal means that an installation wants MVS to provide resources to a particular workload only when those resources are not needed by any workloads with non-discretionary goals. Work assigned a discretionary goals will be the first to have resources taken away if those resources are needed by any workloads with non-discretionary goals.

Velocity Goals
There is MVS work which is not discretionary, yet it cannot be given a response time goal due to the infrequent number of completions. To address this there is a need for a goal that basically states "When this work is ready, be sure it runs without delays", or "When that work is ready, keep it plodding along to ensure it will eventually finish".

The velocity goal type supports both of these needs, as well as gradations in between. Velocity is a measure of the acceptable processor and storage delays holding up work that is ready to run.

It should be noted that the delays considered in the velocity calculation are only those delays that WLM has some control over. I/O delays at a control unit or device, mount delays, and operator delays are **not** part of the WLM velocity.

System Goals
The MVS Workload Manager can handle some work by default, without requiring a customer to bother setting externally specified goals. These 'system' goals simply provide static ways for MVS to manage certain recognized types of work. There are 3 predefined service classes that are managed according to these system goals. These service class names are protected from explicit use when assigning service classes to work, but are instead service classes to be assigned to certain types of work by default.

SYSTEM
When selected system address spaces are created, they are assigned the highest dispatching priority (255) and are excluded from storage isolation controls. These include MASTER, SMF, CONSOLE, CATALOG, GRS, RASP, XCFAS, SMXC, IOSAS, DUMPSRV, ANTMAIN, JESXCF, ALLOCAS, IXGLOGR and WLM. It is best not to assign a service class to these critical address spaces, but to allow them to be managed within the SYSTEM service class.

SYSSTC
This service class is for all started tasks not otherwise associated with a service class. Address spaces managed in the SYSSTC service class are given a high dispatching priority of 253.

SYSOTHER
This service class is intended as a 'catcher' for all address spaces other than started tasks that an installation has not bothered to classify. Address spaces managed in the SYSOTHER service class **are assigned a discretionary goal**.

2.1.1 Importance
When multiple goals are defined, it is necessary to have a way to prioritize which of those goals are really critical, and which are only wishful thinking. MVS supports this through an importance value associated with a goal. Each goal can be rated as very important to the business (1), down to a goal that is desirable but can be sacrificed readily (5). The absolute value specified is less meaningful than the relative value of one importance compared to that of other goals.

Importances have several purposes:

- They identify the critical goals to WLM. WLM attempts to satisfy all importance '1' goals before going after the goals at importance '2', then '3' etc.
- They help prevent the workloads from getting into trouble when goals that are too aggressive are set. The importance allows WLM to make trade-offs that protect the really critical work over the less critical work.
- They allow WLM to react to changing capacity. At many installations today, reacting to an outage involves the cancellation of some work, re-prioritization of other work, and reallocation of the remaining resources. WLM will use the importance of the goals to decide immediately which of the remaining work can donate resources needed by the work with higher importance goals.

  If scarce resources are preventing work from achieving goals, WLM will not just select one type of work to pick on. It will try to achieve the goals of higher importance by degrading equally the work whose goals have lower importance.

These previously mentioned characteristics of importance are critical to remember when examining the WLM workload measurements collected by RMF.

2.2 Goals for CICS and IMS Transactions

A primary feature of the MVS Workload Manager is the ability to specify goals for on-line transactions, like CICS and IMS/TM user transactions.

Prior to the MVS Workload Manager, customers influenced the response time for these transactions by controlling the resources provided to the transaction and resource managers. So if a performance manager wanted to achieve a better response time for a certain subset of the on-line transactions, the manager had to figure out which regions processed the transactions of interest. Then the next task was to analyze the resource usage by those regions and work at improving this resource usage in hopes of achieving a better response time for the transactions of interest. This level of indirectness

often made it difficult to tune on-line transaction processing workloads and often resulted in an inefficient use of system resources.

Examples of such address spaces are CICS regions, IMS control regions and message processing regions, and DB2 address spaces.

With the advent of the MVS Workload Manager, you can now specify response time goals for the on-line transactions. The MVS Workload Manager will manage the transaction and resource managers to work towards meeting the goals of the transactions they serve.

This feature has some limitations. The MVS Workload Manager is only able to manage these transaction and resource managers if they exploit specific MVS workload management services provided in MVS/ESA SP 5.1. The following products are exploiting these services (as of the date of this paper):

- CICS/ESA V4.1 and higher
- IMS/ESA TM V5.1 and higher
- IMS/ESA DB V5.1 and higher
- DB2 V4.1 and higher (for distributed DB2 transactions)

It should be noted that CICS and IMS transactions can only be assigned a response time goal (either average or percentile).

Resource consumption controls
In addition to the 4 goal types above, one further type of control is available with WLM. That is the ability to tell MVS to explicitly control the CPU access for a given collection of work. This can be stated as either a maximum or a minimum amount of CPU resource per second, which should be made available to all the work combined into a Resource Group.

The units of capacity are the same units that have been familiar for years as the SRM constant for various processors. Therefore it is very easy to tell MVS: "Don't let the service units consumed by this workload exceed the rate that could be captured using half of a model 9672-R11". A customer may actually be running that work across 3 LPAR images on 2 separate CECs, neither of which are 9672s!

3.0 Workload Manager Measurements

The WLM measurements collected by RMF can be broken down into 6 primary categories.

1. Workload manager definition information
2. Goals versus actuals data
3. Resource consumption data [1]
4. Execution state data
5. Response time distribution data
6. Transaction state data for CICS and IMS transactions

Most of this data is available only when the system is running in WLM goal mode, and much of this data is the same sort of data WLM uses during its policy and resource adjustment routines to manage the work on the system. Since the data being reported is very similar to the data WLM bases its management decisions on, performance analysts can effectively use the data to evaluate a workload's performance according to that workload's goals. In effect, the reports can be used to help understand the answers to the following questions:

- What were the desired goals for the work?
- Were the goals met?
    - What is a Performance Index and why is it so valuable?

[1] The resource consumption data reported in the upper portion of the RMF Workload Activity report will not be discussed in this paper. Most of this measurement data has not changed in support of the MVS Workload Manager. The author of this paper has written another paper entitled "Understanding MVS/ESA RMF Workload Measurements" and this paper was publish in the USA CMG 1992 Proceedings. Please feel free to contact the author for a copy of this paper.

    - What was the Actual achieved value versus the Goal value?
- What response time did the transactions achieve?
    - What is a transaction?
    - What were the components of the average response time?
    - What is the distribution of response times achieved?
- If a goal was not met, what got in the way of success?
    - Where did the address space transactions spend their time processing?
    - Where did the CICS and IMS transactions spend their time processing?

It is the intent of this paper to help individuals analyzing the RMF Workload Activity report to allow them to become familiar with the Workload Manager measurements. It is **not** within the scope of this paper to educate the reader to be able to complete a detailed analysis of these reports, or to provide tuning recommendations. These are the subjects for a whole other paper, if not a whole book.

3.0.0.1 Reporting Options

The reports discussed in this paper contain those measurements that are collected by RMF into the SMF 72, subtype 3 records, and reported by the RMF Post Processor Workload Activity Report. Most of the same measurements can be collected by RMF Monitor III and reported real-time by the RMF Workload Delay Monitor. The RMF Workload Activity report can be generated:

- For each individual system in the sysplex to gain an understanding of how well work ran on each system.
- On a sysplex wide basis to gain a single view of how well work within the sysplex ran.

In addition, users of the RMF post processor can control the granularity of the data being reported. The following lists some of the report summaries available:

- Summary by Policy. Conceptually, this is similar to the compat mode 'Summary by System' report.
- Summary by Workload
- Summary by Service Class. Conceptually, this is similar to the compat mode 'Summary by Performance Group' report.
- Summary by Service Class Period. Conceptually, this is similar to the compat mode 'Summary by Performance Group Period' report.
- Summary by Report Class. Conceptually, this is similar to the compat mode 'Summary by Report Performance Group' report.

All reports discussed in this paper are reports that summarize the data by Service Class Period. Report examples that have been provided for the reader to refer to are as follows:

- For a TSO workload with an average response time goal report examples refer to Figure 1
- For a Batch workload with a velocity goal report example refer to Figure 2
- For a CICS 4.1, IMS/DB 5.1 workload with a percentile response time goal report example refer to Figure 3
- For a CICS 4.1 VSAM workload with an average response time goal report example refer to Figure 5
- For a CICS 4.1, IMS/DB 4.1 workload with an average response time goal report example refer to Figure 6

```
                                         W O R K L O A D   A C T I V I T Y
                                                                                                             PAGE  21
      MVS/ESA                 SYSPLEX PLEX1             DATE 06/17/1995           INTERVAL 05.00.319   MODE = GOAL
      SP5.1.0                 RPT VERSION 5.1.0         TIME 14.54.58

                                            POLICY ACTIVATION DATE/TIME 06/17/1995 14.46.04

------------------------------------------------------------------------------------------------------------- SERV. CLASS PERIOD(S)

REPORT BY: POLICY=ALLCENT    WORKLOAD=TSO        SERVICE CLASS=TSO        RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=HIGH

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT   INTERVAL SERVICE   --SERVICE RATES--   -PAGE-IN RATES--   ----STORAGE----
AVG        0.96  ACTUAL             000.00.00.076   IOC       52735    ABSRPTN     9001    SINGLE      0.60   AVG      302.92
MPL        0.77  QUEUED             000.00.00.000   CPU      660727    TRX SERV    7209    BLOCK       0.00   TOTAL    232.46
ENDED      2498  EXECUTION          000.00.00.076   MSO     1321473    TCB         23.0    HSP         0.00   CENTRAL  232.46
END/SEC    8.32  STANDARD DEVIATION 000.00.00.136   SRB       39414    SRB          1.4    HSP MISS    0.00   EXPAND     0.00
#SWAPS     2494                                     TOT     2074349    TCB+SRB%     8.1    EXP SNGL    0.00
EXECUTD       0                                     PER SEC    6907                        EXP BLK     0.00

        ---RESPONSE TIME---  EX   PERF   AVG   USING  ------------------ EXECUTION DELAYS % ------------------ ---DELAY  %---
        HH.MM.SS.TTT         VEL  INDX  ADRSP  CPU %  TOTAL  CPU CAPP SWIN  MPL  -AUXILIARY PAGE DELAY % FROM-  OTHR IDLE QUIE
GOALS   00.00.00.100   AVG                                                       PRIV COMM XMEM VIO  SHSP EHSP
ACTUALS 00.00.00.076        11.2%  0.8  145.9   0.1     0.4  0.2  0.0  0.1  0.0   0.1  0.0  0.0  0.0  0.0  0.0   0.2 99.3  0.0

-------------------------------------------------RESPONSE TIME DISTRIBUTION------------------------------------------------
   ----TIME----    ---NUMBER TRANSACTIONS---          ----PERCENT----  0    10   20   30   40   50   60   70   80   90   100
   HH.MM.SS.TTT    BUCKETS            TOTAL           BUCKETS   TOTAL  |....|....|....|....|....|....|....|....|....|....|
<  00.00.00.050      1766              1766              70.7    70.7  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.100       232              1998               9.3    80.0  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.200       253              2251              10.1    90.1  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
>  00.00.00.200       247              2498               9.9     100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>

REPORT BY: POLICY=ALLCENT    WORKLOAD=TSO        SERVICE CLASS=TSO        RESOURCE GROUP=*NONE      PERIOD=2 IMPORTANCE=HIGH

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT   INTERVAL SERVICE   --SERVICE RATES--   -PAGE-IN RATES--   ----STORAGE----
AVG        0.24  ACTUAL             000.00.00.419   IOC       18440    ABSRPTN    10501    SINGLE     10.92   AVG      326.46
MPL        0.24  QUEUED             000.00.00.000   CPU      229115    TRX SERV   10501    BLOCK       0.40   TOTAL     79.34
ENDED       187  EXECUTION          000.00.00.419   MSO      511323    TCB          8.0    HSP         0.00   CENTRAL   79.34
END/SEC    0.62  STANDARD DEVIATION 000.00.00.361   SRB        7508    SRB          0.3    HSP MISS    0.00   EXPAND     0.00
#SWAPS      187                                     TOT      766386    TCB+SRB%     2.7    EXP SNGL    0.00
EXECUTD       0                                     PER SEC    2551                        EXP BLK     0.00

        ---RESPONSE TIME---  EX   PERF   AVG   USING  ------------------ EXECUTION DELAYS % ------------------ ---DELAY  %---
        HH.MM.SS.TTT         VEL  INDX  ADRSP  CPU %  TOTAL  CPU CAPP SWIN  MPL  -AUXILIARY PAGE DELAY % FROM-  OTHR IDLE QUIE
GOALS   00.00.00.500   AVG                                                       PRIV COMM XMEM VIO  SHSP EHSP
ACTUALS 00.00.00.419        20.7%  0.8    0.2  11.2    42.9 29.4  0.0  0.0  0.0  11.9  0.0  1.3  0.3  0.0  0.0  45.9  0.0  0.0

-------------------------------------------------RESPONSE TIME DISTRIBUTION------------------------------------------------
   ----TIME----    ---NUMBER TRANSACTIONS---          ----PERCENT----  0    10   20   30   40   50   60   70   80   90   100
   HH.MM.SS.TTT    BUCKETS            TOTAL           BUCKETS   TOTAL  |....|....|....|....|....|....|....|....|....|....|
<  00.00.00.250        90                90              48.1    48.1  >>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.500        61               151              32.6    80.7  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.01.000        26               177              13.9    94.7  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
>  00.00.01.000        10               187               5.3     100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
```

Figure 1 Example of RMF 5.1.0 Workload Activity Report: TSO periods 1 and 2 are being reported

4.0 Question: What were the desired goals for the work?

Some of the most important pieces of information that the RMF Workload Activity report provides are the policy information and the goals for the work that ran during the measurement interval. A service class period's policy information can be found at the top of each report and what goal was set for that service class period can be found in the middle of each report.

4.1 Policy Information

The **REPORT BY** line at the top of each report provides a summary of the work whose performance data is being summarized. This line reports:

POLICY=
: The policy in effect during the measurement interval. The date and the time the policy was activated is reported in the heading of each report.

WORKLOAD=
: The workload that the service class was assigned to.

SERVICE CLASS=
: The service class of the period being reported on.
RESOURCE GROUP=
The resource group, if any, that this service class was assigned to. If a service class is not assigned to a resource group then '*NONE' will be reported.

PERIOD=
: The service class period being reported on. This is a number from 1 to 8.

IMPORTANCE=
: The importance level assigned to the period being reported on.
HIGHEST
Importance 1
HIGH
Importance 2
MEDIUM
Importance 3
LOW
Importance 4
LOWEST
Importance 5
NONE
If work being reported on is assigned either a system or discretionary goal.

4.2 Assigned Goal Information

Use the **GOALS** versus **ACTUALS** lines in the middle section of the report to determine what the goal for a service class period was. This section appears as follows:

```
          ---RESPONSE TIME---  EX
          HH.MM.SS.TTT         VEL
 GOALS    00.00.00.100   AVG
```

The first line is labeled **GOALS** and on this line the value of the goal that was assigned to the service class period is reported.

4.2.0.0.1 Average Response Time Goal

If the work were assigned an average response time goal, the **GOALS** line would contain the goal response time in terms of **HH.MM.SS.TTT** with the term 'AVG'. next to it.

4.2.0.0.2 Percentile Response Time Goal

If the work were assigned a percentile response time goal, the **GOALS** line would contain the goal response time in terms of **HH.MM.SS.TTT** with a goal percentage from 1% to 99% next to the goal response time.

4.2.0.0.3 Velocity Goal

If the work were assigned a velocity goal, then the field in the column labeled **HH.MM.SS.TTT** would be blank. Instead the first line of the column labeled **EX VEL** would contain the goal execution velocity that the service class period was assigned.

4.2.0.0.4 System Goal

If the work were assigned a System goal then the appropriate term, 'SYSTEM', 'SYSSTC', or 'SYSOTHER', would be reported in the **HH.MM.SS.TTT** column.

4.2.0.0.5 Discretionary Goal

If the work were assigned a Discretionary goal then the term 'DISCRETIONARY' would be reported in the **HH.MM.SS.TTT** column.

```
                                W O R K L O A D   A C T I V I T Y
                                                                                                                PAGE   1
      MVS/ESA                 SYSPLEX PLEX1             DATE 06/17/1995          INTERVAL 05.00.319   MODE = GOAL
      SP5.1.0                 RPT VERSION 5.1.0         TIME 14.54.58

                                      POLICY ACTIVATION DATE/TIME 06/17/1995 14.46.04

------------------------------------------------------------------------------------------------------------------- SERV. CLASS PERIOD(S

REPORT BY: POLICY=ALLCENT    WORKLOAD=BATCH      SERVICE CLASS=BATCHHI    RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=LOWEST

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT   INTERVAL SERVICE   --SERVICE RATES--   -PAGE-IN RATES--   ----STORAGE----
AVG        6.57  ACTUAL             000.09.04.720   IOC      141953    ABSRPTN     6678    SINGLE      7.00   AVG      565.68
MPL        6.57  QUEUED             000.01.09.733   CPU     3936925    TRX SERV    6678    BLOCK      57.00   TOTAL   3717.76
ENDED         2  EXECUTION          000.07.54.987   MSO     9083196    TCB        137.3    HSP         0.00   CENTRAL 3717.76
END/SEC    0.01  STANDARD DEVIATION 000.00.00.000   SRB       18942    SRB          0.7    HSP MISS    0.00   EXPAND     0.00
#SWAPS        0                                     TOT       13181K   TCB+SRB%    45.9    EXP SNGL    0.00
EXECUTD       0                                     PER SEC   43890                        EXP BLK     0.00

        ---RESPONSE TIME---  EX   PERF   AVG   USING  ----------------- EXECUTION DELAYS % ----------------- ---DELAY  %---
        HH.MM.SS.TTT         VEL  INDX  ADRSP  CPU %  TOTAL  CPU CAPP SWIN  MPL  -AUXILIARY PAGE DELAY % FROM-  OTHR IDLE QUIE
                                                                                 PRIV COMM XMEM VIO  SHSP EHSP
GOALS                       14.0%
ACTUALS                      8.8%  1.6    6.6   6.4    65.7 49.1  0.0  0.0  0.0   0.7  0.1  0.2 15.7  0.0  0.0  27.9  0.0  0.0
```

Figure 2 Example of RMF 5.1.0 Workload Activity Report:

Batch running in a service class period with an execution velocity

5.0 Question: Were the goals met?

The next important bit of information that the WLM workload measurements provide is to determine how well the goals for the work that ran during the measurement interval were achieved. Data reported in the RMF workload report enables you to compare how well work in a service class period did compared to its goals and thus allowing for a quick insight into where problem areas may exist. It also allows you to compare how well one workload is performing compared to another workload.

The most important field to examine to determine if a goal was met is the field under the column labeled **PERF INDX**. This is the field that reports the *Performance Index* (also known as 'the PI').

5.1 Performance Index (PI)

Since there are several types of goals, both WLM and performance analysts need some simple way to compare how well or how poorly one service class period is doing compared to its assigned goal and how well it is doing compared to other work with differing goals or resource requirements. This comparison is made possible through the use of a *Performance Index*.

The PI reported in the workload activity report is a calculated value reflecting how well the work in each service class is meeting its defined goal over the reporting interval. A PI of 1.0 indicates the service class period is exactly meeting its goal. A PI greater than 1 indicates the service class period is missing its goal, and a PI less than 1.0 indicates the service class period is exceeding its goal.

The PI also makes it possible to compare how well or how poorly one service class period is doing compared to other service class periods of differing goals. It is possible, for example, to compare how well a batch workload assigned a velocity goal is doing compared to a TSO period assigned an average response time goal and to a period having a set of CICS transactions and assigned a percentile response time goal, and figure out at a glance which periods are nearest to meeting their goals (PI <= 1) and which is farthest away from its goal (has the largest PI greater than 1).

5.1.0.1 Average Response Time Goal PI

The calculation for the performance index for a service class period is simply the ratio of the defined goal for the period to what the work in the period actual achieved. For instance, the PI formula for a period with a response time goal is as follows:

$$\text{Average Response Time Goal PI} = \frac{\text{Actual Average Response Time}}{\text{Goal Average Response Time}}$$

For example, say TSO third period were assigned an average response time goal of 6 seconds. If TSO third period was actually exceeding its goal by achieving an average response time of 3 seconds then the PI would equal 0.5. If TSO third period was actually missing its goal by achieving an average response time of 9 seconds then the PI would equal 1.5.

5.1.0.2 Execution Velocity Goal PI

The PI formula for a period with a velocity goal is opposite from the PI formula for a period with an average response time goal. This is because an achieve velocity that is higher than the assigned goal means that the period would be exceeding its goal.

$$\text{Execution Velocity PI} = \frac{\text{Goal Execution Velocity \%}}{\text{Actual Execution Velocity \%}}$$

For example, say the test batch workload were assigned a velocity goal of 60%. If the test batch period was actually exceeding its goal by achieving a velocity of 80% then the PI would equal .75. If the test batch period was missing its goal by achieving a velocity of 30% then the PI would equal 2.0.

5.1.0.3 Percentile Response Time Goal PI

The PI calculation for a percentile response time goal is very similar to the PI calculation for an average response time goal. The primary difference is that rather than using the 'Actual Average Response Time' the actual response time that was achieved for the percentage of the goal is used.

$$\text{Percentile Response Time Goal PI} = \frac{\text{Actual Response Time At Percentile}}{\text{Goal Response Time}}$$

For example, say TSO second period has a percentile goal of 80% to complete in 2 seconds, and during the measurement interval 140 transactions completed in this TSO second period. Since 80% of 140 is 112 then this means exactly 112 transactions have to complete within 2 seconds and the remainder of the transactions have to complete in greater than 2 seconds for the goal to be met exactly and the PI to equal 1. If more than 112 transactions complete within 2 seconds then the goal would be exceeded and the PI would be less than 1. If less than 112 transactions completed within 2 seconds then the goal would be missed and the PI would be greater than 1.

To properly calculate the PI for work with a percentile response time goal a response time distribution for completed transactions is needed. Later in this paper response time distribution data is discussed.

6.0 Question: Why is the Performance Index so valuable?

As discussed previously (2.1.1 Importance) WLM will try to achieve the goals of higher importance work before attempting to meet goals of lower importance work. WLM will do this by selecting as receivers of resources the highest importance period with the largest PI greater than 1. WLM looks for donors of resource by first utilizing unused resources. If there is no unused resource to draw from then WLM will take resources from the lowest importance periods with the lowest PIs. Refer to Table 1 for an example of this search order. WLM will help periods from the highest importance down, and take from the lowest importance up.

Table 1 WLM Policy Adjustment Receiver Selection Order

Period	Importance	PI
A	1	1.4
B	1	1.0
C	1	0.8
D	2	1.1
E	2	0.7
F	3	2.1
G	3	1.5
H	3	1.0
I	4	3.4
J	4	2.0
K	4	0.5
Discretionary	none	none

This is an important concept to understand. Although the PI reported is not the same one WLM uses it can still provide insight into how well a workload is doing compared to its own goal and the goal of other periods. [2]

- If a user calls up complaining of slow response times, the PI could be used as a quick indicator to determine if the user's goals are being achieved. If the goals for that user are not being met, the state samples (as explained in 13.2.1 Using the State Data) may provide some insight into the delays the user may be experiencing.
- One would expect to see higher PIs on the lower importance work before seeing large PI values in higher importance work. Tracking the PIs over time will help performance analysts to monitor changes in a workload and in capacity. If the PIs start indicating that lower importance work is consistently missing goals it could be an early warning sign that the system is running out of capacity. It could also be the case where all/most of the goals are just too aggressive.
- It is very possible and reasonable to see some lower importance periods with PIs less than 1 and higher importance work with PIs greater than 1. If the higher and lower importance workloads have different resource requirements. WLM may find that taking resources from the lower importance work will not benefit the higher importance work. The lower importance work would be allowed to continue to run well. Performance analysts could use the data in the rest of the RMF report to investigate why the higher importance workloads are missing their goals.

7.0 Question: What was the Actual achieved value versus the Goal value?

```
         ---RESPONSE TIME---  EX   PERF
         HH.MM.SS.TTT         VEL  INDX
GOALS    00.00.00.100   AVG
ACTUALS  00.00.00.076        11.2%  0.8
```

Below the **GOALS** line there is another line labeled **ACTUALS** and on this line the value what the work running in the service class period actually achieved. The possible contents of these fields are as follows:

Average Response Time Goal
If the work in the period being reported on was assigned an average response time goal, then the average response time actually achieved by the work that completed in the period would be reported below the goal value on the **ACTUALS** line.

Percentile Response Time Goal
If the work in the period being reported on was assigned a percentile response time goal, then the **ACTUALS** line would report the actual average response time achieved by the work that completed in the period. Next to this value a percentage from 0% to 100% would be be reported. This field reports the percent of transactions that achieved less than or equal to the actual response time achieved. It is simply the number of transactions that completed in in less than or equal to the goal response time divided by the number of completed transactions.

$$\text{Actual Percentage} = \frac{\text{Number Transactions Completed in \&lt.= 100\% of Goal}}{\text{Total Ended Transactions}}$$

.

Velocity Goal
If the work in the period being reported on was assigned a velocity goal, then the column labeled **EX VEL** would contain the actual execution velocity that the work in the period achieved.

System Goal
If the work in the period being reported on was assigned a System goal then no actual achieved value is reported since there is no actual goal value to be used for comparison.

[2] WLM calculates and analyzes PIs and attempts to help at most 1 period every 10 seconds. RMF only reports the PI on an RMF interval basis which can be from 1 minute to 60 minutes. Naturally a difference will exist between the PI values used by WLM and those reported by RMF. The PIs on the RMF reports are still representative of how well a goal is met.

Discretionary Goal
If the work in the period being reported on was assigned a Discretionary goal then no actual achieved value is reported since there is no actual goal value to be used for comparison.

Actual Execution Velocity
It should be noted that the **EX VEL** will sometimes report an actual execution velocity even if the work running in the period was not assigned a velocity goal. Any period that contains address spaces and thus accumulating state samples will have an execution velocity reported. Periods with response time goals for CICS or IMS transactions will not contain an actual execution velocity.

The actual execution velocity (**ACTUAL EX VEL**) field is calculated by using the state samples accumulated in the period and reported in the state samples percentages section of the report. These state samples are discussed in the next section of this paper.

$$\text{Actual Execution Velocity} = \frac{\text{Using CPU\%}}{\text{Using CPU\% + Total Delay\%}}$$

It should be noted that velocity goals are more sensitive to the delay and using states since those states are used by WLM to evaluate a period's performance. An achieved response time includes many more states than the delay and using states and is thus a more accurate measurement of the quality of service the transactions are actually achieving. A period could be achieving a very high velocity, but if most of the time the transactions in that period are in a state that WLM does not recognize then the velocity achieved can be misleading.

8.0 Question: What response time did the transactions achieve?

The **GOALS** versus **ACTUALS** section of the report showed the actual average response times achieved by the work that ran in the period during the measurement interval. However, it is also important to know the components of a transaction's average response time and the response time distribution of the completed transactions. This information will help the performance analyst to understand

- how appropriate an assigned response time is.
- if some work should be assigned a velocity goal rather than a response time goal.
- If the transactions completing in a period are of like durations.

These issues will be addressed by examining the answers to three additional questions:

- What is a transaction?
- What were the components of the average response time?
- What is the distribution of response times achieved?

9.0 Question: What is a transaction?

Before interpreting the actual average response times reported by RMF on the Workload Activity report, it is important to understand the definition of a transaction for the type of work being examined.

Definitions of a transactions when interpreting WLM response time measurements
TSO Transactions
A TSO transaction for WLM's policy is identical to a TSO transaction as reported via RMF for the past 20 years. The installation still has a choice via IEAOPTxx whether a CLIST should be counted as a single transaction or a collection of individual transactions.

JES Transactions
A batch job is one transaction. It starts when the job is submitted (when the JES reader processes the job), and completes when the initiator finishes executing the job. That means it includes the time queued by JES waiting for an initiator, but it does not include output processing.

CICS Transactions
A CICS transaction begins when the initial CICS region receives a message, generally from VTAM, and ends when that region returns the result to VTAM. If the transaction is routed to another CICS region (AOR, FOR, etc) the time spent processing by those other regions accumulates for the original transaction.

CICS transaction information is passed to WLM beginning with CICS/ESA V4.1.

IMS Transactions
An IMS transaction begins when the Control Region receives a message from VTAM, and ends when the control region passes the response back to the network.

A customer has the option to allow inserted programs (program to program switch) to be considered as new transactions, or to just be considered a continuation of an existing transaction.

IMS transaction information is passed to WLM beginning with IMS/ESA V5.1.

Distributed DB2 Transactions
A goal can be specified for a distributed DB2 transaction. This can be a request arriving remotely across the network via Distributed Relational Database Architecture. Once again, the start time is the arrival time from the network and the transaction ends at the commit point when DB2 completes the distributed request.

Distributed DB2 transaction information is passed to WLM beginning with DB2 V4.1

10.0 Question: What were the components of the average response time?

This second column on the RMF Workload Activity report, titled **TRANSACTION TIME**, reports average transaction response times information. This portion of the report shows the total average elapsed time of the ended transactions and then report the two components of this elapsed time. Those components are queued time and execution time.

```
TRANSACTION TIME   HHH.MM.SS.TTT
ACTUAL             000.00.00.076
QUEUED             000.00.00.000
EXECUTION          000.00.00.076
STANDARD DEVIATION 000.00.00.136
```

The fields in this section of the report are defined as follows:

ACTUAL
Reported in this field is the average system transaction time. This actual average response time includes all time for a transaction as specified in the 'Definitions for Transactions' section of this paper.

$$\text{ACTUAL} = \frac{\text{Total Elapsed Time of Ended Transactions}}{\text{Number of Ended Transactions}}$$

QUEUED
The Queue time field has different meanings dependant upon what kind work is running in the service class. This field is simply calculated by dividing the difference of the total system transaction time and the total elapsed time by the number of ended transactions. For Started Task transactions this field is meaningless since these types of transactions cannot incur queue time. For TSO transactions queue time can be considered logon delay time. For APPC and JES transactions this field is the time spend waiting on an initiator queue.

Queue time for CICS transactions depends upon the level of CICS and the region topology. Transactions which either (a) are not processed by an AOR within the sysplex or (b) are received by a TOR that is not at CICS V4 or above have zero queue time. When a service class contains a complete picture of a transaction, its queue time is the difference between its response time and the sum of all its execution phases.

Queue time for IMS transactions consists of the portion of its response time that the transaction was not being processed in a message processing region. This includes time queued within the control

region, control region processing, and time waiting for a message processing region to begin executing the transaction.

$$QUEUED = \frac{\text{Total Elapsed Time of Ended Transactions - Total Transaction Execution Time}}{\text{Number of Ended Transactions}}$$

If a response time goal for a batch or APPC period is consistently being missed the **QUEUED** time should be examined. It could be that the period is missing its response time goal because it is spending significant time queued waiting for an initiator.

EXECUTION
This field reports the average transaction processing time of transactions that ended during the measurement interval, and can be thought of as MVS response time. For TSO transactions, this does not include processing time of VTAM or BTAM, or line delay time when running remotely. For APPC and Batch transactions, this only includes the time the jobs were executing and not waiting for an initiator.

The meaning of this field as it pertains to CICS and IMS transactions will be discussed later in this paper under heading 13.3 Question: Where did the CICS and IMS transactions spend their time processing?

This execution time is calculated by dividing the total system transaction time of ended transactions by the number of ended transactions.

$$EXECUTION = \frac{\text{Total Transaction Execution Time}}{\text{Number of Ended Transactions}}$$

STANDARD DEVIATION
This field reports the standard deviation of the actual transaction processing time (**ACTUAL**). Its value is somewhat limited because the distribution on which it is based is not always known/reported. It helps indicate whether most of the response times for the transactions were consistently close to the average response time value reported in the **ACTUAL** field. A low response time standard deviation value would indicate that the response times were fairly consistent. A high response time standard deviation value indicates inconsistent response times.

For periods with response time goals, it is recommended that the reported response time distribution data be used to determine the response time consistency of the transactions that completed during the measurement interval. Response time distribution data is discussed next in this paper.

11.0 Question: What is the distribution of response times achieved?

The bottom of the Workload Activity report shows the **RESPONSE TIME DISTRIBUTION** data for service class periods that were assigned a response time goal. For examples, see Figure 1, Figure 3, or Figure 5 This section provides response time distribution data for transactions that completed during the measurement interval. This data is not collected for service class periods with velocity, discretionary, or system goals. It is also not collected for report classes.

RMF collects into the SMF 72, subtype 3 records the response time distribution data as a set of 14 buckets. The width of each bucket is automatically scaled based on the response time goal for the period for which the data is collected. Each bucket contains the number of transactions that completed within some percentage of the response time goal.

- Bucket 1 - less than 50% of goal
- Bucket 2 - from 51% to 60% of goal
- Bucket 3 - from 61% to 70% of goal
- Bucket 4 - from 71% to 80% of goal
- Bucket 5 - from 81% to 90% of goal
- Bucket 6 - from 91% to 100% of goal
- Bucket 7 - from 101% to 110% of goal
- Bucket 8 - from 111% to 120% of goal
- Bucket 9 - from 121% to 130% of goal

- Bucket 10 - from 131% to 140% of goal
- Bucket 11 - from 141% to 150% of goal
- Bucket 12 - from 151% to 200% of goal
- Bucket 13 - from 201% to 400% of goal
- Bucket 14 - greater than 400% of goal

For example, if the average response time goal for a service class period was 1 second. As shown previously, the 9th bucket contains the number of transactions that completed between 121% and 130% of goal, then the 9th bucket will contain the number of transactions that completed between 1.21 seconds and 1.30 seconds.

Although RMF collects the data for all 14 buckets into the SMF 72 subtype 3 record, it summarizes the data on the workload activity report into only 4 buckets.

1. Bucket 1 - less than 50% of goal
2. Sum of Buckets 2-6 - less than or equal to 100% of goal
3. Sum of Buckets 7-12 - less than or equal to 200% of goal
4. Sum of Buckets 13&14 - over 200% of goal

This data is presented on the RMF report in 3 different forms in 3 separate columns.

The first column of this section is labeled **NUMBER TRANSACTIONS** and reports the response time distribution data in terms of number of transactions that completed during the measurement interval in each of the 4 previously discussed buckets. Both an individual count of the number of transactions that completed within each of the 4 buckets, as well as a running total of the number of transactions that completed are reported. The total transactions that completed within the 'over 200% of goal' bucket should equal the number of ended transactions reported in the TRANSACTIONS column at the top of the report.

The second column, which is labeled **PERCENT**, reports the same data as in the previous column, but in terms of percent of completed transactions rather than as the number of completed transactions.

The final column provides a simple visual display of the accumulated percentage of completions within each of the 4 response time distribution buckets that RMF reports. This visual representation allows for a quick overview to which percentages of the transactions have met the goal.

11.0.0.1 Using the Response Time Distribution Data

Response time distribution data is extremely helpful for understanding how well response time goals are being met. Some of the many ways this data can be used include the following.

- If a response time goal is being missed and the response time distribution reports a large deviation of transaction response times this could indicate different types of transactions (ie. long and short, routed and non-routed, etc.) are classified to the same service class. **This could be perfectly acceptable since goals should be based on business requirements and not on reporting requirements. Using report classes is encouraged.** If this is deemed to be a problem, then some areas could be investigated further:
    - Duration values may need to be examined and adjusted for a multi-period service class.
    - Separate different types of transactions into different report classes. Consider continuing to allow the transactions to be managed in the same service class.
- If the PI indicates that a response time goal is consistently not being met, then the response time distribution data will help to understand the response time requirements of the work running in the service class period. If a large portion of the transactions are completing with a response time greater than the goal, then the goal may be too aggressive or additional capacity may be needed.
- If a period is missing its average response time goal and a great majority of the transactions are completing in less then the goal value, then the period probably contains a few 'outlier' transactions that are not typical for the period. This indicates that the period is a good candidate to convert the goal from an average response time goal to a percentile response time goal. Use the response time distribution data to help choose a proper percentile value.

12.0 Question: If a goal was not met, what got in the way of success?

A great benefit of the Workload Manager measurements is performance analysts can use the new data to help determine where transactions spent their time processing. For service classes containing address spaces, state sampling by WLM allows RMF to report using and delay states for the address spaces. See Figure 1 or Figure 2 for report examples of this type of state data. For service classes containing OLTP workloads, a different type of state sampling by WLM allows RMF to report where the OLTP transactions (currently only CICS and IMS) spent their time processing. See Figure 3, Figure 5, or Figure 6 for report examples of this type of state data.

The question of *'if a goal was not met, what got in the way of success'* can be answered by 1 of 2 other questions.

- Where did the address space transactions spend their time processing?
- Where did the CICS and IMS transactions spend their time processing?

13.0 Question: Where did the address space transactions spend their time processing?

To understand why a particular goal for a group of address spaces was not met, it is important to understand where the transactions that ran in the service class period spent their time processing.

13.1 WLM's Use of State Sampling

Performance monitors frequently use state sampling as a way of determining where work spends its time. With MVS/ESA V5, WLM will periodically sample the state of every dispatchable unit. These samples are accumulated for each address space, and then further accumulated into the service class period associated with the address space. These state samples are used by WLM to help make policy adjustment decisions. They are also provided to reporting products, such as RMF, to help allow performance analysts to understand the delays the workload is experiencing and to provide possible explanation as to why a goal may have been missed.

The states that WLM cares about, and that RMF reports, are the states reflecting usage of, or delay for, a resource that WLM can allocate. Those resources are the processor (CPU) and storage. A list of the states WLM samples include:

- Using the CPU
- Waiting for access to the CPU
- Waiting for a page fault (separate states are sampled depending on the type of page fault).
- Waiting for a swap-in to be started
- Waiting for a swap-in to complete

The collection of state samples are used in several ways. First the velocity achieved by an address space or service class period comes directly from these samples. As mentioned earlier, the calculation for an execution velocity is as follows:

$$\text{Execution Velocity} = \frac{\text{Using CPU\%}}{\text{Using CPU\% + Total Delay\%}}$$

Besides allowing a direct comparison to a velocity goal, WLM uses the sampling to determine where work is spending its time. This information allows WLM to determine what resource is the primary bottleneck for the work and allows WLM to assess the impact of some possible action it might take. This same information can be used by performance analyst to help understand those delays that WLM attempted to alleviate and provide insight into resource contention among the workloads.

13.2 State Sample Percentages Section of Report for Address Spaces

Average Number of Address Spaces Sampled
The first field of interest in this section reports the average number of address spaces that were sampled during the measurement interval. This value is reported below the column labeled **AVG ADRSP**. This is very different from the AVG or the MPL values reported in the upper section of the report. Those fields report the average number of address spaces that wanted to run or did run at any given time during the measurement interval. The **AVG ADRSP** column reports the average number of address spaces that were sampled, regardless of whether or not those address space actually executed or wanted to execute during the measurement interval.

```
                Number of Address Spaces Sampled
AVG ADRSP =  ---------------------------------
                Number of Times WLM Sampling Ran
```

Using Percentages
This section of the report shows the using states that the address spaces in the period were found using some given resource. This section of the report only reports on CPU using sample percentages since that is the only using state sampled by WLM. For more exact information on both CPU and storage usage, refer the the resource consumption sections at the top of the report. These sections will not be discussed in this paper.

USING CPU
Percent of the measurement interval work was using the CPU. This value is used in the calculation for the execution velocity being achieved (see Actual Execution Velocity). This is a useful field to help determine what workloads are using the CPU as compared with those workloads that are delayed for CPU.

Execution Delay Percentages
This section of the RMF report shows the different delay states that the workload is experiencing. The delays included in this section are only those delays for resources that WLM has has control over. An example of a resource that WLM has control over is CPU delay and is thus included in this section. An example of a delay that WLM does not currently have control over is I/O delay. Thus any delays incurred due to I/O are not included in this section.

The data reported in this section provides some insight into what delays a workload is experiencing and that WLM is trying to alleviate for the workload. They provide a good insight as to why a particular workload may be missing its goals.

TOTAL:
Percent of the measurement interval work was delayed by a state that WLM controls. This value is used in the calculation for the execution velocity being achieved (see Actual Execution Velocity).

CPU
Percent of the measurement interval the work was delayed because it was waiting to use the CPU.

CAPP
Percent of the measurement interval the work was delayed because it wanted to use the CPU but could not because a resource group maximum was being enforced.

SWIN
Percent of the measurement interval the work was delayed because it was waiting to be swapped in.

MPL
Percent of the measurement interval the work was delayed because it was out and ready to run but swap in could not be started due to lack of MPL.

AUXILIARY PAGE DELAY % FROM:
PRIV
Percent of the measurement interval the work was delayed due to a auxilary storage paging for private area pages.

COMM
Percent of the measurement interval the work was delayed due to common area paging.

XMEM
Percent of the measurement interval the work was delayed due to cross memory paging. This means that an address space in the period being reported on did a cross memory access to another address space and that other address space took a page fault.

VIO
Percent of the measurement interval the work was delayed due to VIO paging.

SHSP
Percent of the measurement interval the work was delayed due to scroll hiperspace paging.

EHSP
Percent of the measurement interval the work was delayed due to an expanded storage only hiperspace read miss.

States Not Controlled by WLM:
The section of the report, with the heading **DELAY %** shows those execution states that WLM sampled address spaces in, but these are states that the WLM does not manage resources for. These states are not included in **TOTAL** for **EXECUTION DELAYS %**.

OTHR
Percent of the measurement interval the work being sampled was in a state unknown and unmanaged by WLM. This includes time waiting on an I/O, DFHSM, DB2 latches, enqueues, operator replies, tape mounts, and long waits.

IDLE
Percent of the measurement interval the work was idle. This includes any STIMER wait, TSO terminal input and output wait, APPC wait,

QUIE
This is the percent of the measurement interval that work in the period was quiesce by an operator. It should be noted that an address space that is quiesced could be incurring delay and using samples in one of the other sampled states.

13.2.1 Using the State Data
Analysts will find many ways to use the state data reported by RMF. If a PI indicates that a service class period is consistently missing its goals, the state data could be examined to help understand the areas of delay.

- If a goal is being missed and most of the time was spent in **OTHR** then maybe the goal is being missed due to delay that WLM does not manage. The goal should probably be re-evaluated to determine if it is too aggressive.
- If a period is consistently missing goals due to storage delays, then it would be useful to know who is using the needed storage. One would expect lower importance period to suffer from storage delay before higher importance period.
- If a period is consistently missing goals due to CPU delays then it would be useful to know the service class periods that are using the processor. One would expect lower importance work to suffer from processor delay before higher importance work.
- Use the delay states as an early warning to help determine if the system is running out of either CPU capacity or storage capacity. Again, the lower importance workloads should show up as suffering before the higher importance workloads. If the higher importance workloads are suffering then the system may be out of capacity or the goals may be too aggressive.
- If a goal is being missed and work is incurring capping delays (as shown by the **CAPP** field) then a resource group max is being enforced. Resource group maximums take priority over goals.

When reviewing the RMF report the following formula may be helpful.

```
100% of Samples = Total Using + Total Delay + Other + Idle
```

13.3 Question: Where did the CICS and IMS transactions spend their time processing?

As discussed previously, it is possible to set response time goals for both CICS and IMS transactions. When this is done, RMF is able to collect and report execution state data for the subsystems work managers that exploit MVS Workload Manager services. The collection and reporting of this data provides the performance analyst with the following information.

- Response time data for ended transactions
- Insight into what 'phases' the transactions spent time processing.
- Insight into where the transactions spent their time during each 'phase' of processing.

Phases of Processing.

The easiest way to explain the *phases of processing* is to discuss an actual example. The RMF report shown in Figure 3 is for a set of CICS transactions classified to a service class named CICSPS. The transactions in this service class are assigned a percentile response time goal of 80% to complete within 0.6 seconds. These transactions are part of a CICS/ESA V4.1 IMS/DB V5.1 workload. Both CICS and IMS in this example are exploiting the new Workload Manager services.

The goals versus actuals section of the report indicates that the transactions in this example actually achieved an average response time of 0.332 seconds and 86% of the ended transactions completed within the 0.6 second goal value. Note the PI is less than 1.0 since the goal was exceeded. During their execution the transactions spent time in the **begin-to-end** phase of processing and/or the **execution phase** of processing.

Begin-To-End Phase

The begin-to-end phase of processing begins when the initial CICS transaction manager receives a message, generally from VTAM, and ends when that manager returns the result to VTAM. In the example in Figure 3 the begin-to-end phase takes place in the terminal owning region (TOR). The TOR is responsible for starting and ending the transactions. The **ACTUAL** time shown at the top of the report is the average response time for ended transactions and is from when the transactions enter the TOR until the transactions end. Thus, the **ACTUAL** time of 0.332 seconds is the time spent in the begin-to-end phase. The **ENDED** field in the top section of the report shows the number of transactions that completed begin-to-end phase processing during the measurement interval. The **ACTUAL** time only applies to **ENDED** transactions.

Execution Phase

The execution phase of processing takes place when the TOR switches control to another subsystem work manager. In the example in Figure 3 this is the application owning region (AOR) and eventually the IMS/DB region. The **EXECUTION** time shown at the top of the report is the average time for transactions to complete the execution phase of processing and is included in **ACTUAL**. Thus, the **EXECUTION** time of 0.251 seconds is the average time the transactions spent in the AORs and the IMS/DB regions. The **EXECUTD** field at the top of the report shows the number of transactions that completed the execution phase of processing during the measurement interval. The **EXECUTION** time only applies to **EXECUTD** transactions.

```
                                  W O R K L O A D   A C T I V I T Y
                                                                                                          PAGE   5
      MVS/ESA                 SYSPLEX UTCPLXJ8           DATE 08/16/1995           INTERVAL 30.10.727   MODE = GOAL
      SP5.1.0      RPT VERSION 5.1.0 TO 5.2.0 CONVERTED  TIME 07.30.00

                                             POLICY ACTIVATION DATE/TIME 08/11/1995 17.58.14
------------------------------------------------------------------------------------------------------------- SERV. CLASS PERIOD(S)

REPORT BY: POLICY=WLMPOL01   WORKLOAD=CICS       SERVICE CLASS=CICSPS     RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=HIGH

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT
AVG        0.00  ACTUAL             000.00.00.332
MPL        0.00  QUEUED             000.00.00.081
ENDED    733475  EXECUTION          000.00.00.251
END/SEC  407.75  STANDARD DEVIATION 000.00.00.603
#SWAPS        0
EXECUTD  734006

           -------------------------------RESPONSE TIME BREAKDOWN IN PERCENTAGE-----------------------------   ------STATE------
SUB    P   TOTAL  ACTIVE  READY   IDLE   ----------------------------WAITING FOR----------------------------   SWITCHED TIME (%)
TYPE                                     LOCK    I/O   CONV   DIST  LOCAL  SYSPL  REMOT  TIMER   PROD   MISC   LOCAL SYSPL REMOT
CICS  BTE   93.1     5.0    6.7    0.1    0.1    0.0   81.2    0.0    0.0    0.0    0.0    0.0    0.0    0.0     0.0  87.6   0.0
CICS  EXE   44.7     6.8   12.5    0.0    0.0    2.9    0.0    0.0    0.0    0.0    0.0    0.0   22.6    0.0     0.0   0.7   0.0
IMS   EXE   30.0    30.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0     0.0   0.0   0.0

        ---RESPONSE TIME---  EX   PERF
        HH.MM.SS.TTT         VEL  INDX
GOALS   00.00.00.600  80.0%
ACTUALS 00.00.00.332  86.0%  N/A   0.8

                                        ----------RESPONSE TIME DISTRIBUTION----------
   ----TIME----    ---NUMBER TRANSACTIONS---          ----PERCENT----  0    10   20   30   40   50   60   70   80   90   100
   HH.MM.SS.TTT    BUCKETS            TOTAL           BUCKETS   TOTAL  |....|....|....|....|....|....|....|....|....|....|
<  00.00.00.300        512K            512K              69.8    69.8  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.600        119K            631K              16.2    86.0  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.01.200      74568             705K              10.2    96.2  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
>  00.00.01.200      28019             733K               3.8     100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
```

Figure 3 Example of RMF 5.1.0 Workload Activity Report

- CICS IMS/DB workload with CICS/ESA 4.1 and IMS/ESA 5.1 running on a 9 system parallel sysplex.

13.4 State Sample Percentages Section of Report for CICS and IMS Transactions

Response Time Breakdown [3]
In addition to reporting the time spent in each phase of processing RMF reports the states the transactions are experiencing in the **RESPONSE TIME BREAKDOWN IN PERCENTAGE** section of the report. This section is valid only for transactions whose work managers exploit the necessary WLM services. Unlike the address space state samples discussed earlier, WLM has no control over the states for CICS and IMS transactions. They are collected and reported for the interests of performance analysts. There is one line per phase per type of work manager. In the example in Figure 3 there are 3 lines reported for 3 phases of processing.

1. **CICS BTE** for time spend in TORs
2. **CICS EXE** for time spend in AORs
3. **IMS EXE** for time spend in IMS/DB

```
|-Transaction Response = 0.332 ---------------------------|

|-ACTUAL Time = 0.332 = 100%  ---------------------------|

    |-CICS BTE Phase = 93.1% ----------------------------|

      |-Switched Time = SYSPL = 87.6% --------------------|

        |-Execution Time = 0.251 = 75.6%  ----------------|

            |-CICS EXE = 44.7% ----|--IMS EXE = 30.0% ----|
```

Figure 4 Response time break down of CICSPS transactions

The above response time breakdown was obtained from the RMF report in Figure 3.

The **CICS BTE** field shows that the TORs have information covering 93.1% of the response time. RMF does not have information covering 100% of the response time because it takes some time for the system to recognize and assign incoming work to a service class before it can collect information about it.

For most of the 93.1% of the time, the transactions were not being run in the TOR, but had been routed locally, to an AOR and eventually to an IMS/DB region in either th same or different MVS image with the sysplex. This can be seen in the **SWITCHED TIME % SYSPL** field, which is 87.6%. This value is close, but not equal, to the **WAITING FOR CONV** field. This indicates that there is some delay in the TOR once the AOR has returned the transactions.

The total execution time is some percentage of the total response time. It is the **EXECUTION** transaction time (0.251) divided by the **ACTUAL** transaction time (0.332) which is 75.6%. The CICS execution phase (**CICS EXE** field) covers 44.7% of the response time. Some of that time the work is active in the AOR, sometimes it is waiting behind another task in the region, but 22.6% (**PROD** field) of the response time is found being spent outside of the CICS subsystem, waiting for IMS/DB to provide some service to these transactions. The IMS execution phase (**IMS EXE** field) covers 30.0% of the response time and all of that time IMS is active processing the transactions.

[3] It should be noted that the values reported in the **RESPONSE TIME BREAKDOWN IN PERCENTAGE** section of the report are not always as easy to interpret as the examples provided in this paper. Individuals analyzing these reports should refer to "RMF V5 - Getting Started on Performance Management" and the "CICS/ESA V4 Performance Guide" manuals to understand these reports better. These manuals are referenced in the bibliography section of this paper.

The **WAITING FOR LOCAL, SYSPL**, and **REMOT** states percentages show the percentage of the response time the service class was delayed when CICS was waiting to establish a session. The STATE percentages **LOCAL, SYSPL**, and **REMOT** show the percentage of the response time during which transactions were routed via MRO, MRO/XCF, or VTAM connections.

The meaning of the fields in the **RESPONSE TIME BREAKDOWN IN PERCENTAGE** section of the report will vary since each state has a unique meaning for each product that exploits the WLM services. The meaning of these fields for CICS are as follows:

TOTAL
Total percent of time that the transactions spent in states that are shown in this section of the report. These states are not a complete breakdown of the response time shown in the **ACTUAL** field. There will always be a gap due to states that are not sampled.

ACTIVE
Percent of time the CICS transactions spent running.

READY
Percent of time the CICS transactions were dispatchable but did not run.

IDLE
Percent of time the CICS transactions spent in a long wait.

WAITING FOR
The following fields document the time the CICS transactions spent waiting for some event while processing during the phase being reported on (ie. BTE or EXE).

LOCK
Waiting for a lock. For example, waiting for:

- A lock on a CICS resource
- A record lock on a recoverable VSAM file or exclusive control of a record in a BDAM file.
- An application resource that has been locked by a EXEC CICS ENQ command.

I/O
Waiting for an I/O request or I/O related request to complete. For example:

- File control, transient data, temporary storage, or journal I/O.
- Waiting for I/O buffers or VSAM strings.

CONV
Waiting for a conversation between work manager subsystems. For example, TOR waiting for a response back from an AOR.

DIST
Waiting for a distributed request. This state is not used by CICS.

LOCAL
Waiting for the establishment of a session with another CICS region in the same MVS image in the sysplex. (eg. MRO XM or IRC)

SYSPL
Waiting for the establishment of a session with another CICS region in a different MVS image in the sysplex. (eg. MRO XCF)

REMOT
Waiting for the establishment of a session with another CICS region (which may or may not be in the same MVS image). (eg. ISC LU6.2 or LU6.1)

TIMER
Waiting for a timer event of an interval control event to complete.

PROD
Waiting for another product to complete its function. For example, when the work requests have been passed to a DBCTL subsystem.

MISC
Waiting for a resource that does not fall into any other categories.

SWITCHED TIME (%)
The percentage of time that transactions during this phase of processing spent routed to another region for processing. For CICS the values in these fields are more of a function of how the transactions were routed and not necessarily a function of where the transactions were routed.

LOCAL
Percent of time transactions spent time processing in another CICS region in the same MVS image in the sysplex (example: via MRO XM or IRC).

SYSPL
Percent of time transactions spent time processing in another CICS region in either the same or different MVS image in the sysplex (example: via MRO XCF).

REMOT
Percent of time transactions spent time processing in another CICS region which may or may not be in the same MVS image (example: via ISC LU6.2 or LU6.1).

14.0 Recommendations for obtaining the best reporting data

To obtain the best reporting data the following recommendations should be followed.

- **Most important 1** - Keep the number of service classes to a minimum to allow WLM to be responsive to the goals
- **Most important 2** - Set goals based on business needs and **not** to make the reports look good.
- **Most important 3** - Make use of report classes.
- To obtain the best CICS/ESA V4 and IMS/ESA V5 execution state data assign transactions with different characteristics to different classes.
    - Do not mix CICS-supplied transactions with user transactions.
    - Do not mix routed and non-routed transactions.
    - Do not mix conversational with pseudo-conversational transactions.
    - Do not mix long-running and short running transactions.
- Refer to paper by Peter Enrico and Ed Berkel of IBM S/390 Division entitled **"Effective Use of MVS Workload Management Controls"** for recommendations on setting goals. This paper can be obtained from
    - US CMG '95 Leading Edge Papers
    - Cheryl Watson's Tuning Letter, March/April 1995 issue, Watson & Walker Inc., Sarasota, FL 34236
    - US CMG '95 Proceedings
    - Contact the authors at 522 South Road, Poughkeepsie, NY, 12602

15.0 Summary

This paper has discussed the MVS/ESA SP5 WLM workload measurements that are made available to performance products. To help achieve this goal, utilization of the RMF Workload Activity report as the mechanism to explain the WLM measurements was used. The primary purpose of this paper is to assist the reader in becoming familiar with workload measurements available as of MVS/ESA SP5. This was done by providing sample reports and detailed descriptions of the workload measurements, and then by explaining some of their uses.

This paper was written for the general interest of experienced MVS/ESA performance analyst, and for the greater interest of those that are still learning the basics.

16.0 References

1. Understanding MVS/ESA RMF Workload Measurements" - Peter Enrico, IBM, CMG '92 Proceedings

2. "Effective Use of MVS Workload Management Controls" - Ed Berkel, Peter Enrico, IBM, CMG '95 Leading Edge Papers
3. "MVS/ESA Planning: Workload Management" - IBM GC28-1493
4. "CICS/ESA Performance Guide, Version 4, Release 1" - IBM SC33-1183
5. "Resource Measurement Facility Version 5 Getting Started on Performance Management" - IBM LY33-9176
6. "Cheryl Watson's Tuning Letter" - May/June 1995 Vol. 5, No. 3" - Watson & Walker, Inc.
7. "CICS/ESA 4.1 and MVS/ESA 5.1 Workload Management and Performance Tuning in a Parallel Environment" - Chris Baker, IBM, 1995 Share summer conference

17.0 Acknowledgements

The author wishes to thank the following people for their constructive review and suggestions for this paper:

- Ed Berkel
- Steve Grabarits
- Gary Hall
- Kristine Logan
- Terry Stache

18.0 RMF Workload Activity Report Examples

```
                                        W O R K L O A D   A C T I V I T Y
                                                                                                                PAGE   8
      MVS/ESA                 SYSPLEX PLEX1             DATE 06/21/1995          INTERVAL 04.59.595   MODE = GOAL
      SP5.1.0                 RPT VERSION 5.1.0         TIME 13.02.01

                                     POLICY ACTIVATION DATE/TIME 06/21/1995 12.51.18
------------------------------------------------------------------------------------------------------------------- SERV. CLASS PERIOD(S

REPORT BY: POLICY=EQUALHRD   WORKLOAD=CICS       SERVICE CLASS=CICUSRTX   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=HIGH

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT
AVG        0.00  ACTUAL             000.00.00.086
MPL        0.00  QUEUED             000.00.00.015
ENDED     14326  EXECUTION          000.00.00.077
END/SEC   47.82  STANDARD DEVIATION 000.00.00.063
#SWAPS        0
EXECUTD   13069

           ------------------------------RESPONSE TIME BREAKDOWN IN PERCENTAGE----------------------------   ------STATE------
SUB    P   TOTAL  ACTIVE  READY   IDLE   ---------------------------WAITING FOR---------------------------   SWITCHED TIME (%)
TYPE                                    LOCK    I/O   CONV   DIST  LOCAL  SYSPL  REMOT  TIMER   PROD   MISC   LOCAL SYSPL REMOT
CICS  BTE   96.0     1.4    0.9    0.0    0.1    4.0   89.5    0.0    0.0    0.0    0.0    0.0    0.0    0.0    89.9   0.0   0.0
CICS  EXE   75.8     5.7    7.7    0.0    0.1   62.3    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    50.6   0.0   0.0

         ---RESPONSE TIME---  EX   PERF
         HH.MM.SS.TTT         VEL  INDX
GOALS    00.00.00.090   AVG
ACTUALS  00.00.00.086         N/A   1.0

------------------------------------------------RESPONSE TIME DISTRIBUTION-------------------------------------------------
   ----TIME----    ---NUMBER TRANSACTIONS---          ----PERCENT----  0    10   20   30   40   50   60   70   80   90   100
   HH.MM.SS.TTT    BUCKETS            TOTAL          BUCKETS   TOTAL  |....|....|....|....|....|....|....|....|....|....|
<  00.00.00.045       5597             5597             39.1    39.1  >>>>>>>>>>>>>>>>>>>>
<= 00.00.00.090       3866             9463             27.0    66.1  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.180       3633            13096             25.4    91.4  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
>  00.00.00.180       1230            14326              8.6     100  >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
```

Figure 5 Example of RMF 5.1.0 Workload Activity Report

- CICS VSAM workload. Note that waiting on I/O is the largest component of the CICS execution phase of processing. In this example the I/O is to a VSAM file.

```
                                    W O R K L O A D   A C T I V I T Y

                                                                                                                    PAGE   5
      MVS/ESA                 SYSPLEX PLEX1             DATE 06/25/1995          INTERVAL 04.59.000   MODE = GOAL
      SP5.1.0                 RPT VERSION 5.1.0         TIME 09.00.25

                                       POLICY ACTIVATION DATE/TIME 06/21/1995 12.51.18
------------------------------------------------------------------------------------------------------------------- SERV. CLASS PERIOD(S

REPORT BY: POLICY=HPTSPOL1   WORKLOAD=PRODWKLD   SERVICE CLASS=CICSHR
                                                          Hotel Reservations

-TRANSACTIONS--  TRANSACTION TIME   HHH.MM.SS.TTT
AVG        0.00  ACTUAL             000.00.00.114
MPL        0.00  QUEUED             000.00.00.036
ENDED       216  EXECUTION          000.00.00.078
END/SEC    0.24  STANDARD DEVIATION 000.00.00.270
#SWAPS        0
EXECUTD     216

           ------------------------------RESPONSE TIME BREAKDOWN IN PERCENTAGE-----------------------------   ------STATE------
SUB    P   TOTAL  ACTIVE  READY   IDLE   ---------------------------WAITING FOR---------------------------   SWITCHED TIME (%)
TYPE                                     LOCK    I/O   CONV   DIST  LOCAL  SYSPL  REMOT  TIMER   PROD   MISC   LOCAL SYSPL REMOT
CICS  BTE   93.4    10.2    0.0    0.0    0.0    0.0   83.3    0.0    0.0    0.0    0.0    0.0    0.0    0.0    83.3   0.0   0.0
CICS  EXE   67.0    13.2    7.1    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0   46.7    0.0     0.0   0.0   0.0
```

Figure 6 Example of RMF 5.1.0 Workload Activity Report

- CICS DBCTL workload with CICS/ESA 4.1 and IMS/ESA 4.1. Note that waiting on another product is the largest component of the CICS execution phase of processing. In this example the other product is IMS/ESA 4.1.