SEGUS

White Paper

How much CPU did we really save with DB2 10?

Fabio Massimo Ottaviani - EPV Technologies

1 Introduction

IBM's DB2 10 announcement included the following statement:

> *Customers can achieve "out-of-the-box" DB2 CPU savings of up to 5%–10% for traditional workloads and up to 20% for specific workloads, when compared to running the same workloads on DB2 9.*

As you can imagine, the possibility of reducing CPU usage (and therefore software costs) created such high expectations, that it has become the single most important reason for customers to migrate to DB2 V10.

The reality is, as always, a little bit different. After completing their migrations, some sites are happy while others are not completely satisfied.

However, the real problem is that many sites are not even quite sure about the results they achieved and have difficulties when they have to answer the inevitable question: "How much CPU did we really save with DB2 V10 ?".

In this paper we will discuss a methodology, based on real life experiences, which could help you to accurately answer this question at your site.

2 Setting the right expectations

The first important step is to set realistic expectations. Every manager knows just how much capacity (in MIPS or MSU) they purchased, so they run the risk that they will instinctively apply the expected DB2 10 CPU reduction factor to the whole of that capacity.

The graph in Figure 1 shows the hourly profile of the CPU utilization on a peak day, on a CEC rated at 2,700 MIPS (CEC capacity is represented by the red dotted line)[1].
As you can see, the total CPU utilization in the peak hour (at 16[2]) is about 2,303 MIPS. So there are about 400 MIPS of spare capacity[3].

1 SMF 70 records are used to produce the graph.
2 In all graphs and tables, the start of the hourly interval is indicated; so 16 means between 16:00:00 and 16:59:59).
3 If you want to evaluate a possible reduction of the software costs, you must consider the MSU used in the monthly peak of the 4-hour rolling average.

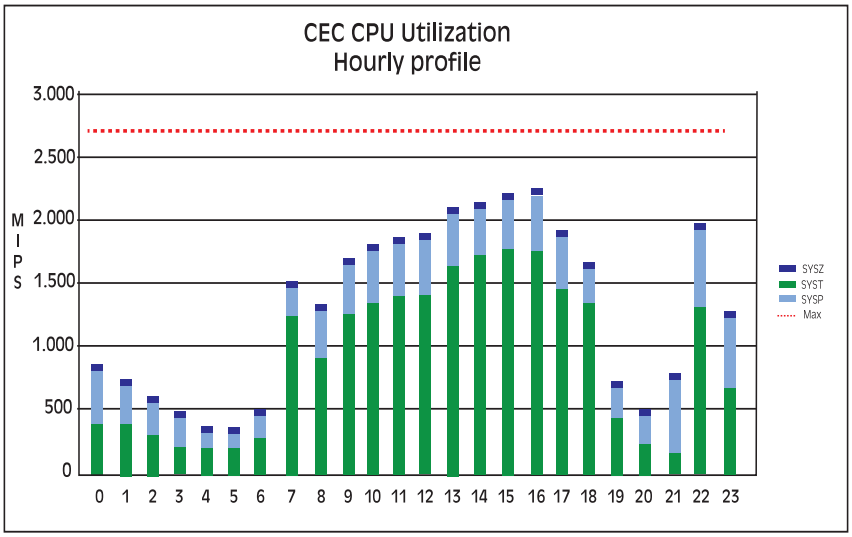

Figure 1

Only three systems are hosted on this particular CEC; their contribution to the CPU utilization peak hour, represented here by the stacked bars, is :

- SYSP (production = pale blue) — 1,849 MIPS,
- SYST (development = green) — 421 MIPS,
- SYSZ (test = dark blue) — 33 MIPS.

SYSZ is a sandbox system; it doesn't run any DB2 workload, so it should not be taken into consideration when evaluating the total amount of workload where a reduction may be possible by migrating to DB2 V10.
By summing SYSP and SYST only, we therefore get 2,270 MIPS - which is still a very optimistic view.

As you know, in each system there are many different workload types, and some of them have little, or nothing, to do with DB2.

In general, there are no benefits derived from DB2 10 for:

- System components;
- Non-DB2 STC;
- Any workload not connecting to DB2.

So a further step towards our goal of setting realistic expectations can be achieved by splitting the system CPU utilization by workload.

The graph in Figure 2[4] shows the contribution of each workload in SYSP[5] to the total system CPU utilization.

4 SMF 30 records are used to produce the graph; SMF 72 can also be used.
5 From here on we will discuss only the SYSP production system, but a similar analysis should also be conducted for SYST.

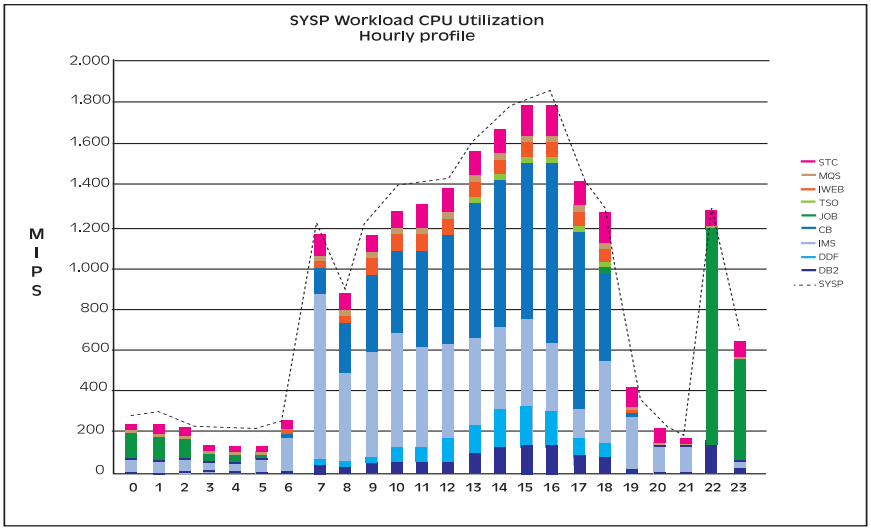

Figure 2

You can see here, that the sum of the CPU used by all of the workloads is always less than the SYSP total CPU (represented by the black dotted line in figure 2). This is due the fact that some system activities, (such as paging, some SRM activities part of the I/O processing, etc.), are not attributed to any workload[6].

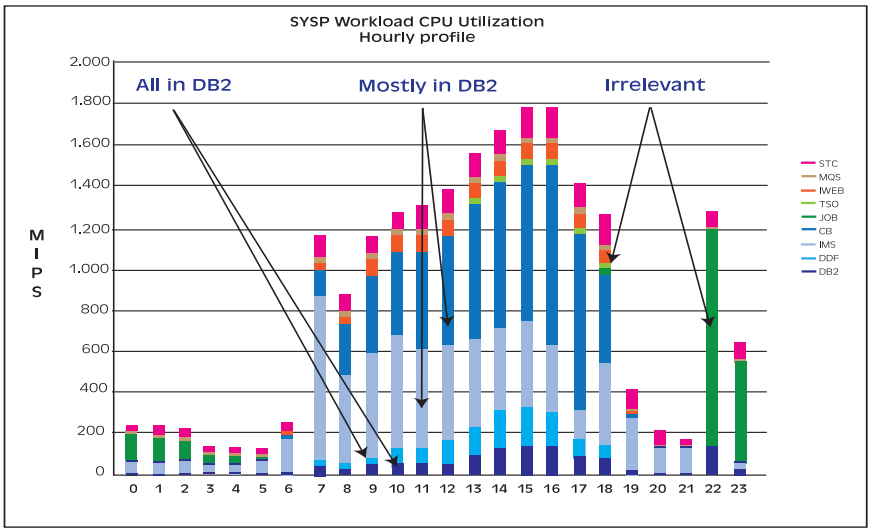

Figure 3

Starting from the bottom of the stacked bars in Figure 3 we find the following:

- 100% DB2 workloads; they are DB2, (which accounts for the consumptions of all the DB2 system address spaces), and DDF, (which accounts for DIST address space consumptions).
- Mostly DB2 workloads; they are IMS, for IMS MPP and BMP workload consumptions, and CB

6 The difference between system CPU utilization, as measured in SMF 70, and the sum of all workloads CPU utilization, as measured in SMF 30 or 72, is called uncaptured CPU time.

- (Component Broker) which accounts for the WebSphere Application Server workload consumptions.
- Workload which shows irrelevant consumptions during peak hours, such as JOB for batch and TSO for TSO users;
- Workload not using DB2, such as STC (started tasks), MQS (WebSphere MQ address spaces) and IWEB (IBM HTTP Server address spaces).

Considering only the DB2, DDF, IMS and CB workloads, the CPU utilization peak can be evaluated now as 1,492 MIPS at 15 (at 16 the peak is 1,478).

Unfortunately, this is still an optimistic view. The DB2 10 benefits only refer to activities in DB2. So even for workloads which are always connected to DB2, such as IMS and CB in our example, only a part of the CPU they use may be considered, when evaluating the DB2 consumptions to be reduced by migrating to Version 10.

To estimate this part, you must analyze internal DB2 measurements which can be produced by activating accounting trace classes 1 and 2 and then collecting the SMF 101 records.

3 How much CPU is really used in DB2

The total amount of CPU used by DB2 can be split into two major components:

- The CPU used by DB2 system address spaces; which can be considered as DB2 subsystem overhead (it is the DB2 workload shown in Figures 2 and 3);
- The CPU used by DB2 applications - which is the real productive work[7].

3.1 DB2 system address spaces

DB2 system address spaces CPU consumption is presented in the report shown in Figure 4[8].
Peak hours for the system component are highlighted in red. Highest values are around 115 MIPS.

A.S. CPU USAGE DETAIL MIPS BY HOUR - SYSP - Mon, 18 Jun 2012											
DB2ID	SYSTEM	ADDRESS SPACE	8	9	10	11	12	13	14	15	16	17
DB2P	SYSP	DB2PMSTR	13.2	15.8	16.1	14.0	13.8	14.7	16.0	14.8	13.6	5.7
DB2P	SYSP	DB2PDBM1	37.3	56.0	67.5	75.7	75.8	95.0	101.5	101.8	100.0	81.0
DB2P	SYSP	DB2PIRLM	0.5	0.2	0.4	0.3	0.4	0.3	0.5	0.4	0.4	0.3
		TOTAL	51.0	72.0	84.0	90.0	90.0	110.0	118.0	117.0	114.0	87.0

Figure 4

As already discussed, this information is provided by SMF 30 and 72 records. (It can also be obtained from SMF 100 if the DB2 statistics trace is activated.)

7 DDF consumptions can also be estimated by using the DIST address space as shown in Figures 2 and 3. However that would not easily allow us to evaluate the benefits of DB2 10, as we will see in the last chapter.
8 Only the busiest hours are reported here.

3.2 DB2 applications

While the CPU used by DB2 system address spaces can be considered as wholly DB2 CPU, the same can't be said for the DB2 applications.

In fact, most of DB2 applications alternate activities within DB2, (accessing DB2 objects), and activities outside of DB2, (to perform calculations, access sequential files, etc).

Information provided by SMF 101 records allows you to distinguish between these different activities: class 1 CPU includes all the consumption, starting from the opening of the DB2 thread up to its end; class 2 CPU includes only the consumption when accessing DB2 objects.

The report in Figure 5[9] shows only CPU class 2 consumptions.

It provides a very clear picture of the amount of CPU used inside DB2 by the different connection types.
The biggest workload is RRSAF, which is the connection type used by WebSphere Application Server workloads.
The other most relevant workloads are IMS and DDF.

TOTAL CPU MIPS BY HOUR - Mon, 18 Jun 2012											
CONNTYPE	**GROUP**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**
BATCH	SYSPDB2P		0.00								0.13
DB2CALL	SYSPDB2P	1.20	2.40	1.47	2.17	2.34	22.36	23.59	24.46	18.42	26.17
IMS BMP	SYSPDB2P	0.08	0.11	0.11	0.11	0.08	0.08	0.08	0.07	0.07	0.07
IMS MPP	SYSPDB2P	140.35	172.11	172.46	149.71	150.22	150.78	160.08	161.65	141.97	80.40
IMS TRAN BMP	SYSPDB2P	2.99	3.05	3.12	3.12	3.02	3.06	3.18	3.10	3.06	2.62
REMOTE UOW	SYSPDB2P	37.24	50.08	70.85	75.51	90.04	118.14	152.54	173.47	157.89	74.49
RRSAF	SYSPDB2P	161.61	267.48	350.70	425.16	446.74	533.97	582.36	620.05	668.00	674.32
TSO	SYSPDB2P				0.01				0.00		
UTILITY	SYSPDB2P		0.01								
	TOTAL	343.47	495.24	598.71	655.79	692.44	828.39	921.83	982.80	989.41	858.20

Figure 5

Peak hours, for the application components shown in figure 5, are highlighted in red. The highest values are around 980 MIPS at hour 15.

9 Only the busiest hours are reported here.

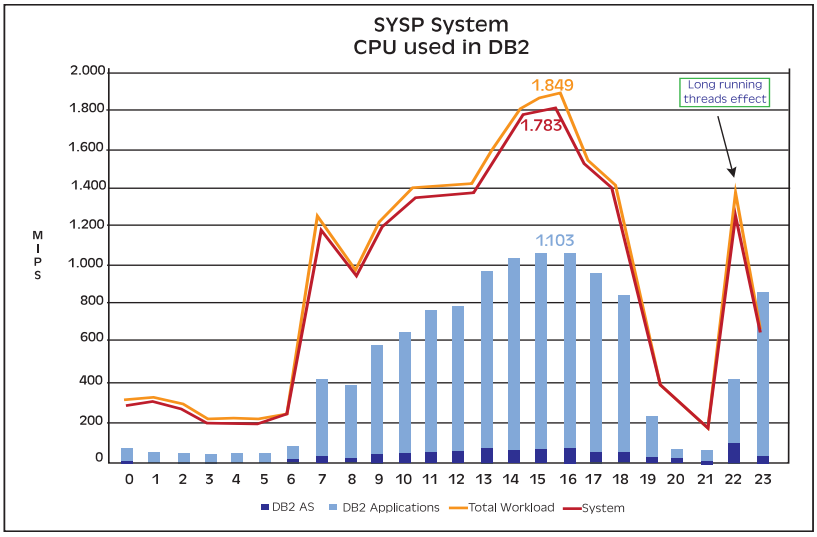

Figure 6

Figure 6 above shows the in-DB2 CPU utilization due to System AS component (dark blue bars) and DB2 applications component (light blue bars).

You can see that the sum of these components has a peak of 1,103 MIPS, compared with a total system CPU utilization peak of 1,849 MIPS, represented here by the orange line. Captured consumptions of all the workloads is represented by the red line; the peak value is 1,783 MIPS.

So assuming that migrating to DB2 10 should give a 10% benefit in terms of CPU reduction, a realistic expectation should be to save about 110 MIPS on the SYSP system.

It's important to remember that SMF 101 records, used to evaluate the CPU utilization due to DB2 Applications, are written at thread termination and not at interval expiration (as SMF 70, 30, 72 and 100 are). So a long running thread could impair the accuracy of the hourly profile. You can see this occurring at hours 22 and 23 in Figure 6. However, this is not an issue in this case, because these are not peak hours.

SEGUS Inc is the North American distributor for EPV products

For more information regarding EPV for z/OS, please visit www.segus.com or call (800) 327-9650