SEGUS

White Paper

Saving money on the mainframe:
The most important things a manager should check

Fabio Massimo Ottaviani - EPV Technologies

1 Introduction

IT Budgets are tight and companies running z/OS applications are always looking for ways to save money and increase profitability.

There are many ways to achieve savings: some require deep, technical analyses, while others can involve complex organizational issues.

However, there are also simple and very important steps that a manager can take in order to keep costs for software products under control - always one of the biggest components of mainframe expenses - and hopefully reduce them in the process.

Most of the expensive software products on the market, (including z/OS itself), are usually priced using the WLC policy (Workload License Charges), where license fees are determined, and charged, using the monthly peak of the 4-hour rolling average value of the MSUs used.

In an *ideal* world, the monthly peak would always:

- occur on business-critical days;
- occur in business-critical hours;
- be due to business-critical systems;
- be due to business-critical workloads;
- be due to workloads which could only run on standard CPUs.

Unfortunately, due to the complexity of current systems, this is not always the case in the real world. But, by using effective reporting methods, a manager could easily perform checks on a daily basis, in order to ensure that the money spent for software products is minimized.

It's important to note, that while ex-post controls can normally only reduce future costs, in WLC it may also be possible to actually reduce the current month's bill too.

In this paper we will discuss these issues, giving examples based on the real-life experiences of some of our customers.

2 Not a "beautiful" day?

The first thing to examine is: on which day of the month did the MSU monthly peak occur?

It would be expected to take place on a working day, and probably on a business-critical day, (e.g.: month-end or the start of the month).

A simple report such as the one in Figure 1[1], can provide a quick and easy check.

It refers to February 2012; the machine was a 2097-717 valued at 1329 MSUs.

As you can see in the z/OS column[2], not only did the peak day occur on a Sunday, but it exceeded all the other values by 354 MSUs.

DATE	DAY	TYPE	MODEL	MSU	z/OS
02/26/2012	SUN	2097	717	1.329	1.309
02/22/2012	WED	2097	717	1.329	955
02/08/2012	WED	2097	717	1.329	896
02/24/2012	FRI	2097	717	1.329	882
02/02/2012	THU	2097	717	1.329	877
02/23/2012	THU	2097	717	1.329	862
02/17/2012	FRI	2097	717	1.329	859
02/14/2012	TUE	2097	717	1.329	849
02/20/2012	MON	2097	717	1.329	845
02/27/2012	MON	2097	717	1.329	843
02/09/2012	THU	2097	717	1.329	837
02/15/2012	WED	2097	717	1.329	830
02/21/2012	TUE	2097	717	1.329	824
02/28/2012	TUE	2097	717	1.329	818
03/01/2012	THU	2097	717	1.329	813
02/16/2012	THU	2097	717	1.329	811
02/07/2012	TUE	2097	717	1.329	805
02/13/2012	MON	2097	717	1.329	804
02/10/2012	FRI	2097	717	1.329	786
02/06/2012	MON	2097	717	1.329	756

Figure 1

As Saturday and Sunday are not business days at this particular site, the high value on a Sunday has to be considered completely abnormal.

In this actual case, it was found that the peak had been caused by a long, resource-heavy recovery activity, which was required in order to fix a data corruption issue following the migration to new storage processors on the previous day.

1 Most of the examples in this paper present standard views using the product EPV for z/OS
2 This column provides the total MSUs used by all the LPARs in the machine.

As can be seen in Figure 2, WLC rules foresee the possibility of managing "unusual situations" by excluding some hours or days from bill calculations. This is the reason why it is important to be able to identify and document situations which are out of the ordinary.

> *Customers have the primary responsibility for preventing uncontrolled loops, operator errors, or unwanted utilization spikes. However, IBM understands that, occasionally, situations that could not be prevented (especially situations related to disaster recovery) might cause exceptional utilization values.*
>
> *In these situations, IBM does not normally expect customers to pay for the increased utilization associated with the unusual situation.*
>
> *Use your best judgement to determine if an unusual situation has occurred. IBM does not publish a list of unusual situations because, by their nature, they will be unpredictable.*
>
> **From the "Using the Sub-Capacity Reporting Tool" manual.**

Figure 2

In Figure 3, we present a somewhat different situation based on the report presented in Figure 1.

The red bars represent Saturdays, and the blue bars all the other days in May 2011.

If Saturday was a business-critical day for this customer, this behavior could be considered normal. Unfortunately this was not the case.

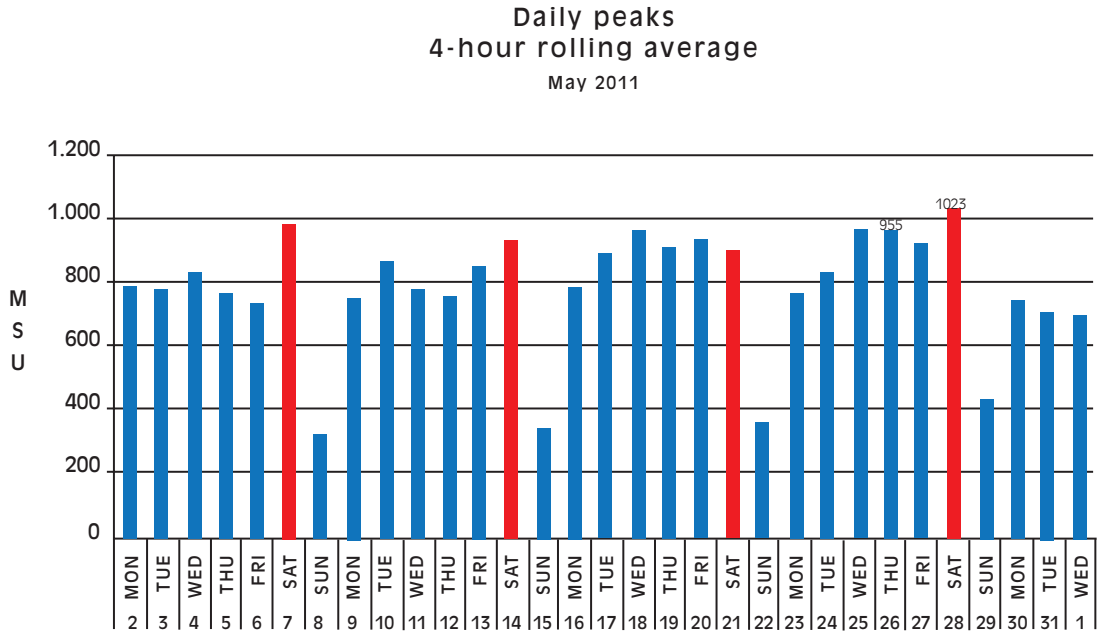

Figure 3

This odd behavior came as a surprise to the customer. However, after a couple of phone calls, the mystery was solved: Application stress tests had been run on every Saturday in May.

System parameters had been set in order to protect against excessive test consumptions, but one of them had been misinterpreted and therefore was set incorrectly. This then allowed test executions to encroach on the desired limits.

The consequence of this little misunderstanding, was an increased bill of about 68 MSUs (the difference between the 1023 MSUs highest value on Saturday and the 955 MSUs highest values on all the other days).

Depending on a company's relationship with IBM, IBM may or may not allow disallow these days from the WLC bill in a situation like that. As you can imagine, there is a good amount of discretion involved in the evaluation process.

3 (un)Happy Hour

Our next example is more typical, in that a monthly peak occurs on a business-critical day. In this case, the next step in the WLC check should be to verify in which hours of the day the peak occurred.

What caught the customer's attention in Figure 4 below, is the large difference between the first peak day and all of the other days (79 MSUs).

DATE	DAY	TYPE	MODEL	MSU	z/OS
11/14/2011	MON	2817	710	1.191	928
11/11/2011	FRI	2817	710	1.191	849
11/10/2011	THU	2817	710	1.191	844
11/18/2011	FRI	2817	710	1.191	837
11/30/2011	WED	2817	710	1.191	832
11/03/2011	THU	2817	710	1.191	830
11/23/2011	WED	2817	710	1.191	821
11/22/2011	TUE	2817	710	1.191	820
11/02/2011	WED	2817	710	1.191	809
11/17/2011	THU	2817	710	1.191	798
12/01/2011	THU	2817	710	1.191	777
11/28/2011	MON	2817	710	1.191	754
11/25/2011	FRI	2817	710	1.191	743
11/21/2011	MON	2817	710	1.191	741
11/15/2011	TUE	2817	710	1.191	737
11/24/2011	THU	2817	710	1.191	734
11/09/2011	WED	2817	710	1.191	732
11/29/2011	TUE	2817	710	1.191	731
11/07/2011	MON	2817	710	1.191	714
11/16/2011	WED	2817	710	1.191	705

Figure 4

By analyzing the peak day at the hourly level, the customer realized that something strange had happened. The workload peak usually occurs in the morning, so they expected to have the peak of the MSU used in the 4-hour rolling average just after noon.

In fact, a high value (889) was reported at "13", (between 13:00:00 and 13:59:59), but in the late afternoon, the MSU utilization climbed up to a much higher value (928).

CEC : XXXX - WORKLOAD : z/OS - 4 HOUR MOVING AVG BY HOUR - MON, 14 NOV 2011																										
SYSTEM	TYPE	MODEL	MSU	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
PRD1	2817	710	1.631	31	31	31	31	27	28	34	46	56	63	70	72	73	76	76	74	76	81	79	76	64	49	47	55
PRD2	2817	710	1.631	84	69	58	68	92	111	140	178	206	295	432	551	739	752	704	709	727	752	804	808	790	760	695	593
ZTST	2817	710	1.631	4	4	4	4	4	4	4	5	5	5	5	6	6	6	6	6	6	6	6	6	5	5	5	5
UDEV	2817	710	1.631	16	15	15	13	13	13	14	20	26	32	45	51	53	54	48	42	41	41	38	38	35	33	37	35
TOTAL				135	119	108	117	136	156	192	250	294	395	553	680	872	889	834	831	850	880	927	928	894	847	784	688

Figure 5

Looking at the different systems' contributions, it became clear that the peak was due to something running on the PRD2 system (see Figure 5 above). The customer asked the technical team for a deeper analysis.

MSU BY HOUR - PRD2 - MON, 14 NOV 2011																	
WKL	ADDRESS SPACE	SRVCLASS	0	1	2	...	14	15	16	17	18	19	20	21	22	23
JOB	USER0001	OMVS-P2						4.2	65.9	65.4	64.8	66.1	66.0	42.8		
JOB	JOB1	BATCHOPC												57.2		
JOB	JOB2	BATCHOPC									47.2	51.2				
JOB	JOB3	BATCHOPC										34.1				
JOB	JOB4	BATCHOPC										31.4	26.2	8.8		
JOB	JOB5	BATCHOPC												17.8	48.5	

Figure 6

It turns out that the late afternoon peak was caused by a user running into a loop.

As you can see in Figure 6,[3] where all the highest utilization hours are shown in red, USER 0001 used about 65 MSUs continuously for about 5 hours[4]. This caused an increase of 39 MSUs in the WLC bill for November 2011.

As clearly stated in the IBM manual, (see Figure 2), the customer has the primary responsibility for preventing uncontrolled loops. However, in situations like this, when a loop has occurred on one day only, in one system, and only for a few hours in the month, it is very likely that IBM would disallow the extra hours.

By conducting checks in this way, anomalies can easily be spotted and documented, thus helping to reduce and minimize costs.

3 Only the highest MSU address spaces running in late afternoon are presented here.
4 This report shows the MSUs used by each address space in each hour. Reported values are not 4-hour rolling averages.

4 The system you don't expect

Every site runs systems which are dedicated to different workloads and activities, such as production, application development, application test, system test, etc.
Normally you would expect the WLC monthly bill to be mostly related to the amount of MSUs used by production systems to support company business.

A report such as the one in Figure 7 can be very useful to show the contribution of each system to the monthly peak of the number of MSUs used.

DATE	TYPE	MODEL	MSU	PROD	TST1	TST2	ZNET	TOTAL
2011-10	2097	714	1.139	988	10	8	86	1092
2011-09	2097	714	1.139	975	10	9	45	1039
2011-08	2097	714	1.139	912	9	9	40	971
2011-07	2097	714	1.139	873	7	7	39	926
2011-06	2097	714	1.139	965	10	10	44	1029
2011-05	2097	714	1.139	943	10	10	45	1008
2011-04	2097	714	1.139	967	9	9	42	1027
2011-03	2097	714	1.139	961	10	10	44	1025
2011-02	2097	714	1.139	956	10	10	45	1021
2011-01	2097	714	1.139	940	9	9	41	999
2010-12	2097	714	1.139	944	8	8	43	1003
2010-11	2097	714	1.139	923	10	10	44	987

Figure 7

Only four systems were hosted on the reported machine. The production system (PROD) used most of the MSUs while two test system (TST1 and TST2) contributed up to a maximum of 10 MSUs each. The last system (ZNET) was dedicated to host network component such as VTAM and TCP/IP.

Based on the customer's experience, the ZNET load had always been very stable. This was certainly true up to October 2011 when its load suddenly doubled.

A deeper analysis at the daily level showed something happened on October 24th (see Figure 8 below). It was Monday, so the first idea was to check for maintenance activities performed over the week end.

DATE	DAY	MSU	SYSA	TST1	TST2	ZNET	TOT
10/02/2011	Sun	1.139	432	5	5	12	454
10/03/2011	Mon	1.139	931	7	7	42	987
10/04/2011	Tue	1.139	904	10	8	43	965
10/052011	Wed	1.139	805	8	7	42	862
10/06/2011	Thu	1.139	767	7	6	45	825
10/07/2011	Fri	1.139	702	6	8	41	757
10/08/2011	Sat	1.139	656	5	5	23	689
10/092011	Sun	1.139	357	5	5	14	381
10/10/2011	Mon	1.139	840	7	7	44	897
10/11/2011	Tue	1.139	887	10	6	43	946
10/12/2011	Wed	1.139	932	10	8	43	993
10/13/2011	Thu	1.139	934	8	8	44	994
10/14/2011	Fri	1.139	889	7	6	40	942
10/15/2011	Sat	1.139	712	5	5	25	747
10/16/2011	Sun	1.139	395	5	5	18	423
10/17/2011	Mon	1.139	886	7	7	43	942
10/18/2011	Tue	1.139	896	8	7	43	954
10/19/2011	Wed	1.139	869	9	8	43	928
10/20/2011	Thu	1.139	851	8	7	45	910
10/21/2011	Fri	1.139	796	7	7	41	850
10/22/2011	Sat	1.139	684	5	5	24	718
10/23/2011	Sun	1.139	376	5	5	16	402
10/24/2011	Mon	1.139	863	7	7	79	955
10/25/2011	Tue	1.139	891	9	7	78	985
10/26/2011	Wed	1.139	900	10	8	78	996
10/27/2011	Thu	1.139	892	8	8	79	987
10/28/2011	Fri	1.139	842	7	7	75	931
10/29/2011	Sat	1.139	698	5	5	40	748
10/30/2011	Sun	1.139	385	5	5	38	433
10/31/2011	Mon	1.139	979	7	7	84	1077
11/01/2011	Tue	1.139	988	9	7	86	1089

Figure 8

At the same time, they analyzed the ZNET workload in more detail and they found a corresponding increase of CPU consumption by the session manager address space.

The good news about software, is that there is always a new version to correct current problems; the bad news is that the new version may introduce new (and sometimes even bigger) problems.

In this case, it was a new version of the session manager causing such a big increase (about 40 MSUs) of CPU consumption.
In the end, most of these MSUs were recovered - thanks to some PTFs and parameter tuning actions.

Being able to measure and report these issues in detail gave the customer the opportunity to discuss the October and November monthly bills with IBM in order to reduce them.

5 The importance of "importance"

Once it has been verified that WLC monthly peaks occur on business-critical days, in business-critical hours and that they are due to business-critical systems, the next check to perform is for the "business importance" of the workloads which used most of the MSUs and determined the monthly software bill.

It is not by chance that one of the most important z/OS components, the WLM, decides which workload can use the system resources, (CPU, memory, etc.), based on a parameter called "importance".

There are five levels of importance that can be assigned to user workloads; they range from 1 (highest) to 5 (lowest). In addition, there are system-related workloads which are considered more important and are referred to as importance 0 workloads.

Finally, there are discretionary workloads which are less important than any other workload and are referred to as importance 6 workloads.

It's easy to understand why the WLM "importance" has to be carefully matched to the "business importance" of the workload, in order to provide the appropriate support to the company business.

DATE	IMP 0-6	IMP 0-5	IMP 0-4	IMP 0-3	IMP 0-2	IMP 0-1	IMP 0
01-2012	975	843	640	617	476	370	173
12-2011	913	741	509	481	365	275	127
11-2011	873	811	536	509	355	258	166
10-2011	865	777	534	517	413	296	127
09-2011	913	831	507	483	367	291	174
08-2011	868	813	507	484	387	292	116
07-2011	861	714	508	493	387	310	134
06-2011	856	798	552	527	436	329	134
05-2011	879	763	567	549	448	355	170
04-2011	728	656	472	452	364	267	115
03-2011	824	764	554	538	406	304	188
02-2011	883	826	513	489	390	283	132
01-2011	963	880	533	508	435	322	118

Figure 9

In the report above, each IMP 0-n column shows the MSUs used in the 4-hour rolling average by all the workloads having an importance less than or equal to n.

So IMP 0-6 accumulates the MSUs used by all the workloads while IMP 0-5 accumulates the MSUs used by all the workloads excluding importance 6 workloads, (running with a discretionary goal).

The difference between these two columns is the contribution of discretionary workloads to the software bill. Note that, at this customer site, this difference was very high in every month.

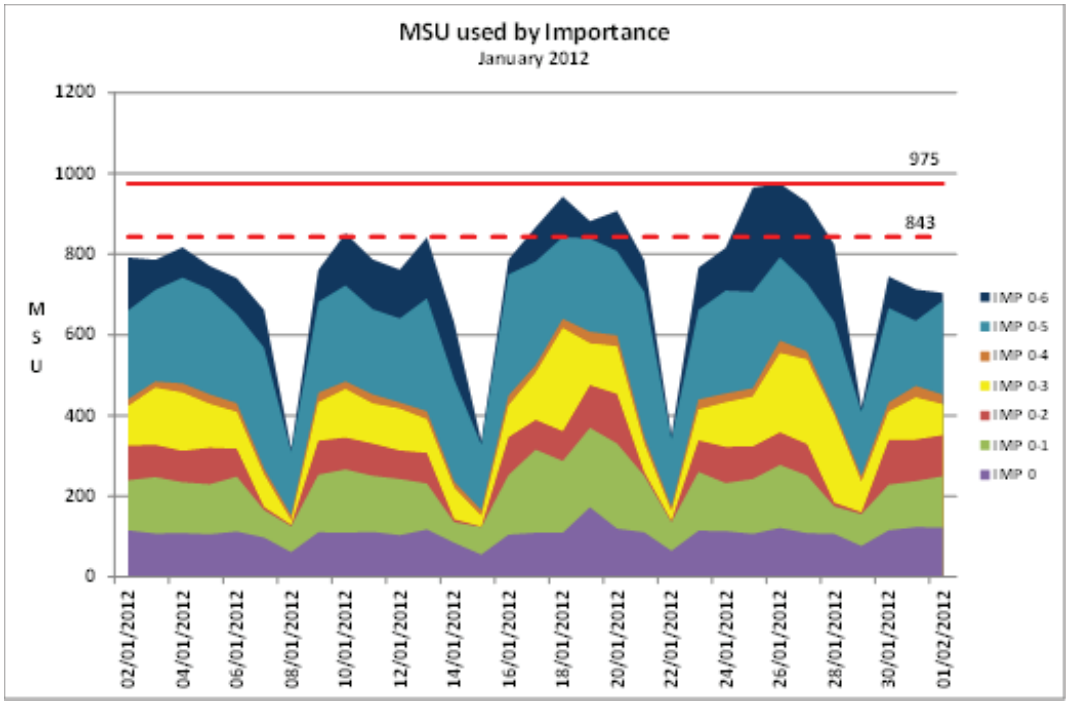

Figure 10

As you can see in Figure 10 above, in January 2012 the difference was 132 MSUs (975 – 843)[5].

This means that a lot of money could have been saved if these discretionary workloads (batch jobs in this case) had been scheduled in non-peak hours.

By re-scheduling these discretionary workloads, companies can reduce, or avoid, sharp "peaks" in MSUs used.

5 t's important to consider the highest values of IMP 0-6 and IMP 0-5 in the month, even if they occur on different days.

6 Could we save more money with zAAP and zIIP?

The final important check required, regards the level of exploitation of zAAP and zIIP engines in the environment. Moving as much work as possible to these specialized processors is, in fact, a key point in the strategy of any customer who really wants to reduce hardware and software costs.

Unfortunately, because of the complexity of today's systems, it is not uncommon to find customers with a good portion of the workloads which could run on zAAP or zIIP, (called zAAP or zIIP eligible), running on standard CPUs and so, unnecessarily, increasing the monthly software bill.

The most common reasons for this are:

- wrong PRSM and system parameters;
- missing PTFs;
- loops;
- not enough zAAP or zIIP available due to wrong estimates in the Capacity Planning process.

DATE	CURR	NO AAPCP	NO IIPCP
11-2011	193	193	160
10-2011	185	185	150
09-2011	190	190	160
08-2011	188	188	156
07-2011	189	189	155
06-2011	192	192	158
05-2011	198	198	152
04-2011	199	199	165
03-2011	200	200	164
02-2011	201	201	164
01-2011	197	197	166
12-2010	203	203	170

Figure 11

In the report above (Figure 11), the CURR column shows the MSUs used in the 4-hour rolling average. NO AAPCP is an estimate of what the CURR value would be if enough zAAP was available. NO IIPCP is an estimate of what the CURR value would be if enough zIIP was available.

It is worth noting that at this particular customer site, NO AAPCP showed the same values as CURR in every month; this means that no MSU could have been saved by adding more zAAPs.

On the other hand, NO IIPCP was always substantially less than CURR.

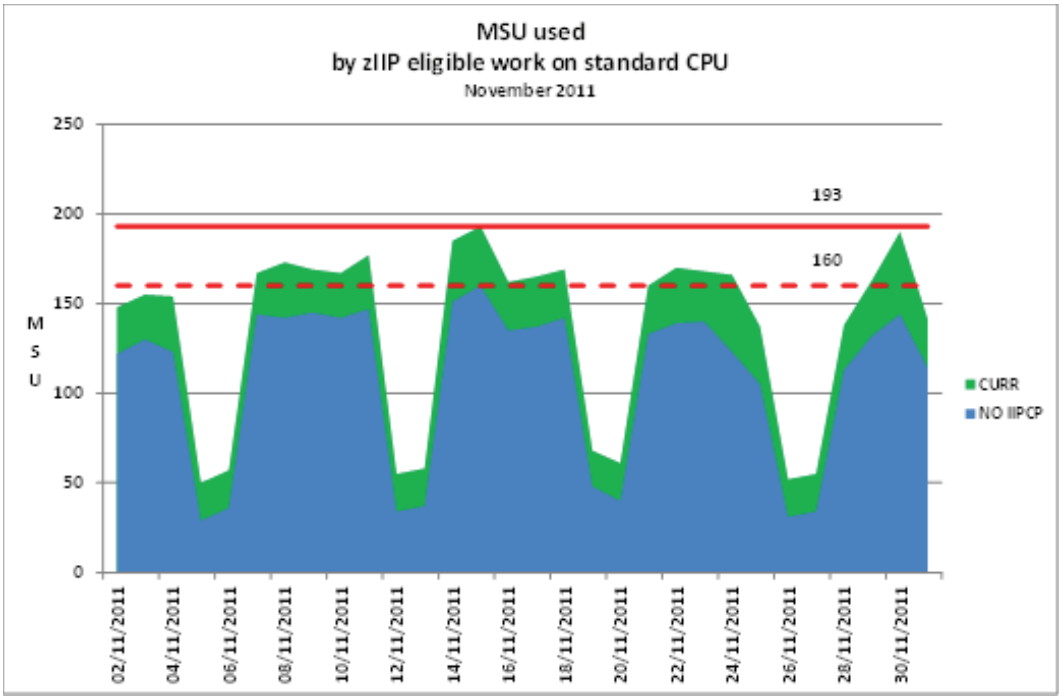

Figure 12

As you can see in Figure 12 above, in November 2011, the difference between CURR and NO IIPCP was 33 MSUs (193 – 160)[6].

This means that a substantial amount of money could have been saved if more zIIPs, (probably only one in this case), had been made available.

6 It's important to consider the highest values of CURR and NO IIPCP in the month even if occurring on different days.

7 Conclusions

In this paper we presented 5 important checks that should be performed in order to keep WLC-based costs under control.

We also showed how effectively-designed WLC reports enable these checks to be performed very quickly and easily.

Being able to clearly identify and document issues makes it possible to not only reduce next month's costs, but, depending on the issue, to also reduce the current month's bill.

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