



REALDOLMEN
to get there, together

IBM z Systems Newsletter

N°. 25 – 08/11/2016



www.realdolmen.com

Industriezone Zenneveld A. Vaucampsiaan 42 1654 Huizingen
Tel +32 2 801 55 55 Fax +32 2 801 55 99 BTW/TVA be 0429.037.235
rpr BRUSSEL/RPM BRUXELLES

Table of contents

A brief history of software pricing.....	3
General Introduction.....	3
General tendency.....	4
Starting off with machine based pricing.....	4
Enters Sub-capacity pricing	5
General price decreases – Technology dividend	6
General price decreases – New Workloads	7
New Workloads – from separated to integrated	7
New workloads today	8
Announcements.....	10
IBM z/OS Version 2 Release 2 enhancements and statements of direction	10
Price Changes effective January 1, 2017.....	10
Hardware withdrawal: IBM TS3500 Tape Library models L23 and L53 and select features	10
DB2 12 for z/OS and DB2 Tools	11
z/VM V6R4 announcement	11
DS8888 – the new all flash DS8xxx system.....	12
Hints and Tips.....	14
Hot Topics #30.....	14
Last call for MES upgrades on your z12 system(s).....	14
EOS dates Operating Systems.....	15
Agenda	16
Small print.....	17

A brief history of software pricing

General Introduction

It was about 13 years ago that they asked me if I was interested in doing the presales for mainframe business. I came from the 'other' side since I'd always been engaged in application development as a programmer, an analyst and a DB2 admin. So, at that moment I'd never seen any mainframe hardware whatsoever. But we had an experienced mainframe sales person at the time who learned me everything there was to learn about ESCON, CHPIDs, OSA-Express cards, CECs etc. etc. And then there was this other part as well: if you wanted (and this hasn't really changed) to work out a business case, you had to know something about software pricing too.

I still vividly remember my first steps in 'Software Pricing' on mainframe or should I rather say my first stumbles. I just looked up the word in a dictionary.

Stumble [stuhm buh l]

1. the act of stumbling (to strike the foot against something, as in walking or running, so as to stagger or fall)
2. a slip or blunder

I can assure you, both meanings were very appropriate at the time.

Just one year later, our experienced sales person took on another opportunity and, just as that, I was no longer the rookie but I was the most experienced mainframe (pre)sales guy in the company. A new sales was appointed for mainframe and of course I was the one introducing him to all this magnificent stuff. Don't worry, he pulled through with flying colors and we're still working together.

Still, it's a bit of a crazy story but we first met just a couple of hours before the official announcement of the z890 in April 2004 (yes, the 40th anniversary of the mainframe). There was a large IBM and BP event in Paris and he picked me up at my home for a 3-hour drive to Paris. Of course we started talking about the mainframe, comparing with the AS400 (now System i) he knew from his background. And, actually, we never stopped talking for the next three days. After each session new questions popped up and were, as good as possible, answered. Until finally we came to the subject of 'software pricing'. What could I say? The only thing there was to say: "Well, you know, here's where it gets a bit complicated". Talking about a euphemism.

And this remained a constant during the years to come. Talking to customers, the questions I most often heard about software pricing were: "Could you freshen up my memory about that, I'm a bit confused?" usually followed by "Could you explain that once again, I don't think I'm still following?"

I really would like to say, luckily things have gotten a lot less complicated over the years, but they haven't. So, why am I telling you all this. On November 15, we host a meeting for the GSE Young professionals Working Group at our HQ. These are all young mainframers we desperately want to get or to keep aboard of the mainframe ship. And just for the occasion, I'd like to tell them a bit about the evolutions we've seen over the years in software pricing on the mainframe.

General tendency

To put things into perspective of what we are talking about. We, as a business partner, are often engaged into negotiations on the hardware price of a new system. Mainframe is expensive, remember? But when the customer starts working out his business case, it's always played on the software part. Usually software represents 70% or more of the cost when e.g. mapped out on a four year basis. So any decrease in software cost when going to a new generation of mainframe might already make up for the hardware cost.

And this is the general tendency we have seen over the past years. Or should I say: tendencies. On the one hand IBM has always put effort into making the mainframe a competitive platform as compared to other, distributed environments. For a moment now, we're ignoring all the benefits you get from the platform as is, but we're purely focusing on the price, let's say, on the CFO part of the business. You can make it competitive by making sure that the customer is only paying for what he is really using. On the other hand, IBM has made sure that with every new generation of the mainframe, prices were more attractive on the new system. This meant investing in new hardware paid off in the long run with economically more interesting software prices.

And there's one more tendency, that is predominantly directing software pricing: it is IBM's intention that you will pay less for new workloads that you introduce to the mainframe rather than installing them on the 'cheaper' distributed systems. This is a key factor in remaining competitive with the distributed systems.

Starting off with machine based pricing

When browsing through my (old) documentation, I keep lingering at a document from 2004, the year of the introduction of the z890: 'z890 and z800 Software Pricing P-Guide'. One of the first pages headlines 'What is Sub-Capacity Pricing?'. In 2004 sub-capacity pricing is a new term for a pricing mechanism that was introduced for the z900. I know I'm more than generalizing throughout this summary in order to give you a clear view of the bigger picture.

Up to then, software pricing was machine based. This meant there was a flat monthly pricing amount according to the model of e.g. your z800 box. But that was always a problem. You don't buy a mainframe every year and so customers made a three or four year projection of what they needed in the next years and bought their systems accordingly. So, the first years, they only used e.g. 60% of what they had bought but they still had to pay for the entire machine.

But before I continue with this, I should give you some more terminology. Software pricing is/was mainly divided into two categories: MLC (Monthly License Charge) and OTC (One Time Charge) pricing.

MLC pricing is a monthly charge you are paying in order to use IBM software. Examples of this is the operating system z/OS (OS/390 at the time) and lots of others like e.g. DB2, CICS, Cobol, IMS . . . MLC software is paid in terms of numbers of MSUs (Million Service Units). It's a measurement of the amount of processing work a computer can perform in one hour. This stands in close relationship to the MIPS (Million Instructions Per Second).

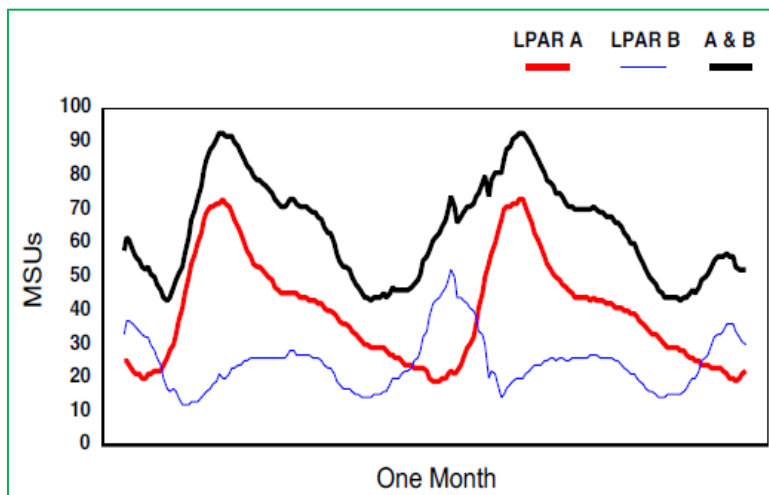
OTC, now better known as IPLA (International Program License Agreement), stands for a One Time Charge. You actually buy the software and you can pay an additional Subscription & Support (S&S) fee that also gives you the right to implement future versions at no extra cost. Examples are the operating system z/VM and lots of 'tools' like TSM, DB2 Utilities ...

Enters Sub-capacity pricing

With VWLC (Variable Workload License Charge) IBM introduced the sub-capacity pricing for MLC. This meant that you closely came to pay what you were really using at that moment. Let me elaborate a bit on this, since this hasn't fundamentally changed since then.

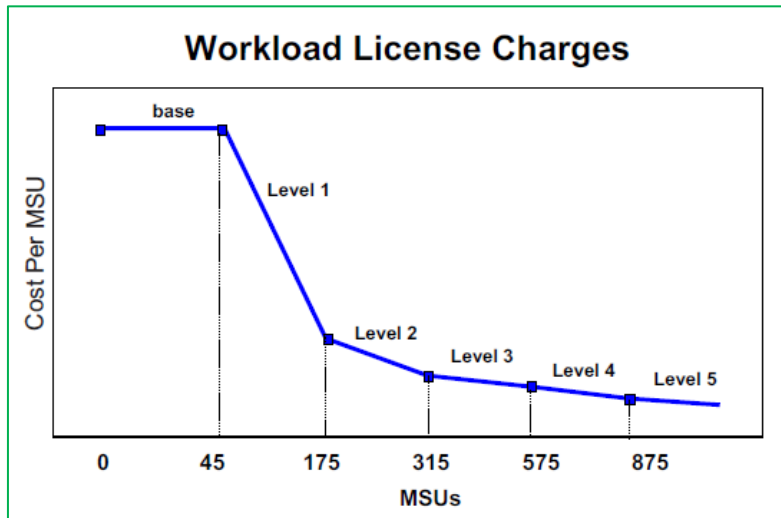
In order to detect what you are actually using, you need a reporting tool. That is SCRT (Sub Capacity Reporting Tool). It registers your usage for an entire month based on SMF records. Some software such as z/OS and DB2 generate their own SMF records for this. Others take on the same level as e.g. z/OS. Out of that it will, per software, determine where your peak is for that software during that particular month. Now, you may argue, one test in acceptance that goes completely out of the roof might cause an enormous peak and then I'm penalized for that one moment during that month. Well, IBM took care of that.

The SCRT report determines the peak on a 4 hour rolling average. This means that it always takes an average of the past four hours to determine the height of the workload, so extreme peaks are levelled out. This is illustrated on the graph below for two LPARs. And if you want to make sure, you're not going above a certain level you can use a mechanism that is known as capping. You can indicate that a certain LPAR (or group of LPARs) must not go beyond a defined level.



In this illustration, partition A's peak rolling four-hour average is shown to peak at 73 MSUs, during the month. Let's take z/OS as an example. When it would be running solely in partition A it would have its sub-capacity charges based on that 73 MSU value, although the machine capacity is at 100MSU. Likewise, partition B's peak rolling four-hour average is recorded at 52 MSUs. A product running solely in partition B would have its sub-capacity charges based on that 52 MSU value. But since z/OS is running in both LPARs, it will be charged at the combined peak for those LPARs i.e. 98 MSUs. Here we can also illustrate the capping mechanism again: if LPAR B is e.g. a test LPAR you might put up a capping of 45 MSUs and perhaps your combined peak might be lowered to 95 MSUs.

As you can see, the reporting and therefore also the billing is based on MSU's.



The graph above, also shows you the pricing levels for the MSUs. It indicates that you pay a high price for the first MSU's. This is an example for VWLC pricing when it was first introduced. The more MSUs you are reporting, the less you pay per MSU for the higher MSUs. As a matter of fact, there's such a graph for every system and the steps for the smaller systems (z890, z114, zBC9 through z13s) tend to be much smaller. You have expensive 1-3 MSUs and from there on MSUs get less expensive in far smaller steps.

IPLA software also has sub-capacity pricing although there is still some software that is machine based. This software is usually related to the reporting of some MLC software. If you have a machine of 200 MSU and you report only 150 MSU for DB2, then it's sufficient to have bought the equivalent of that 150 MSU for the DB2 Cloning Tools or for the DB2 Utilities.

General price decreases – Technology dividend

So, the stage is set for the following evolutions. The first action taken by IBM is a general one and I already mentioned it before. Customers want to be on a current system, but on the other hand they often postpone buying new systems since there's no real reason for that. So, with the z990, IBM introduced what is now commonly referred to as the technology dividend. In fact it's a very simple maneuver that stimulates every customer to at least make the calculation whether it's beneficial for them to go to the next generation.

When I put 'price decrease' in the title, this is not entirely correct. For a specific generation, there's a correlation between the number of MIPS on that machine and the number of MSUs. This used to be very static, not really changing over the generations. But for the last ten years, we started to call MSUs, software MSUs and the relation with MIPS became less evident. Let me just give you a small table and it will immediately be clear what we mean.

System	Pricing	MIPS	MSU
Z890	EWLC	200	32
Z9 BC	EWLC	200	30
Z10 BC	EWLC	200	25
Z114	AEWLC	200	25
zBC12	AEWLC	200	25
Z13s	AEWLC	200	25

For the first three generations from z890 to z10 BC, for the same amount of processing power, less MSUs were needed to cover it. So, when you had to pay for 32MSUs on a z890, this went down to 25MSs on a z10BC. And I can

assure you, for a mainframe customer, this was a huge benefit on their software cost, well worth the investment in the new hardware. This technology dividend was more or less every time a 5% decrease.

From then on, we see that the correlation between MSUs and MIPS stayed the same. But with the z114 a new pricing was introduced. With AEWLC pricing you payed less per MSU as compared to the EWLC pricing on the z10 BC. Another technology dividend but implemented in a different way. And for the last two generations, we saw the same pricing mechanism, but an extra reduction was given per technology step. So, if you look up the official price for 3 MSU for z/OS, it hasn't changed since the introduction of the z114. But, for the z13s you will have approximately a price reduction on it of about 10%.

General price decreases – New Workloads

The above are pricing decreases every customer could/can enjoy when moving on to the next technology. But, as I already said, IBM is particularly keen on getting new workloads on the mainframe and has gone into great effort doing so.


For those who are familiar with the terms, this started off with zOS.e and zNALC. zNALC is the best known of the two and is still in use nowadays. Let's focus on that one. zNALC is a pricing mechanism for New Workloads that run on separate machines or in separate LPARs. The benefit is realized very simply: if you can prove that it's a new workload on your mainframe, answering all the right criteria, your z/OS is only priced at a tenth of the normal pricing. This gives the customer a huge benefit on the software cost since z/OS is one of the most expensive software in MLC pricing.

The major drawback of this, is that it has to be a completely isolated workload in a separate environment. Well, now, that was always a bit the weak point of this solution. Say, you have all your data on the mainframe, as so many customers do, and you develop, let's go modern, an app for your customers who can consult their accounts. What about all the other applications, running in another LPAR, that also have access to that information. This is often very difficult to integrate into your existing environment.

New Workloads – from separated to integrated

That's why, over the years, IBM has made a lot of efforts to give you the best of both worlds: a better pricing for new workloads, but integrated into the existing environment. One of those efforts was e.g. to bring out a new pricing option for certain software solutions such as CICS and DB2. It's called the **Value Unit Edition** and it's basically an equivalent of IPLA software for that product. You buy it once and you pay a Subscription and Support fee for the 'maintenance' and the upgrades of the product. This is something every customer has to investigate whether this is beneficial for them or not. The advantage of this, is that the workload remains integrated into your existing environment.

Another initiative that could have impact on your software cost was the introduction of **specialty engines**. Here we can particularly focus on the zIIP. The zIIP (z Systems Integrated Information Processor) is an additional processor that is added to your system. Specific workloads are offloaded from z/OS to the zIIP. This can be specific DB2 workloads, Java, encryption or XML workloads (you can find a more extensive list over [here](#)). Here's the illustration that is often used to illustrate how the zIIP is working



As you can see, the workload on the general processor is reduced as part of it is now executed on the zIIP. There's an important caveat to be made here: this will only have an influence on your software cost if this happens during your peak period of the month. If so, here's another example of how the software cost can go down due to a hardware investment.

New workloads today

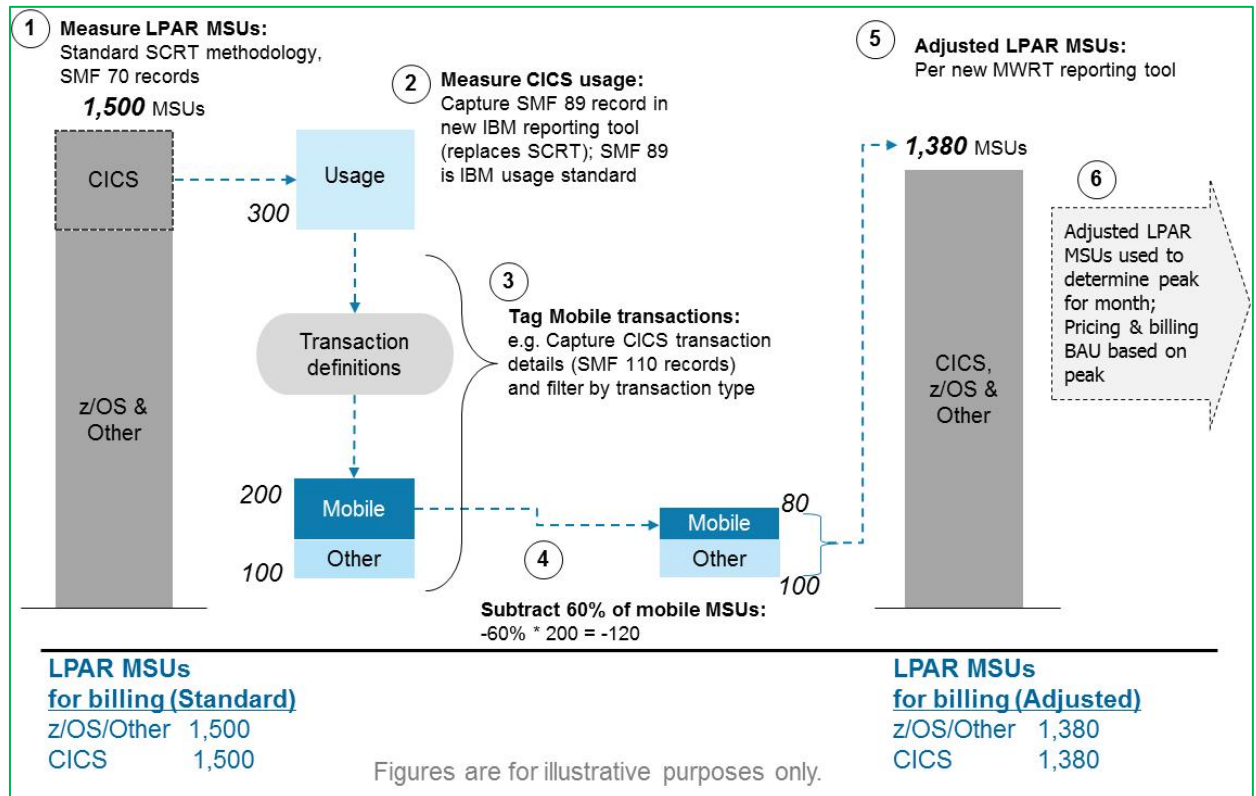
Let's finish off by having a look at the enhancements we recently saw in sub-capacity pricing. The first one was announced about two years ago, the third about two months ago. They have in common that they, again, focus on New Workloads like Mobile or Cloud.

- **Mobile Workload Pricing (MWP)**
Used when IBM programs such as CICS, DB2, IMS, MQ, or WebSphere Application Server are processing mobile transactions from phones or tablets,
- **z Systems Collocated Application Pricing (zCAP)**
Used when net new instances of IBM programs such as CICS, DB2, IMS, MQ, or WebSphere Application Server are added in support of a new application not previously running on a mainframe server
- **z Systems Workload Pricing for Cloud (zWPC)**
Used when IBM programs such as CICS, DB2, IMS, MQ, or WebSphere Application Server are processing transactions from net new public cloud applications

The three have in common that they may reduce the cost of growth for these target applications by potentially reducing the reported peak capacity values for sub-capacity charges. This remains a constant throughout: it must have an impact on year peak reporting, otherwise you'll see no benefit for your software cost.

Integration of new workloads into your existing environment is one thing, distinguishing which of the workload is new (e.g. mobile "coming from phones or tablets") is another thing. So, just in case you started to think that I exaggerated at the beginning and you think software pricing isn't all that complicated, let me explain how Mobile Workload Pricing can influence your reported peak and how the workload is recognized as 'mobile' workload. This gives me the opportunity to illustrate all the mechanisms I talked about earlier.

Here's an illustration on MWP



- Let's assume that we are talking about the 4hr rolling average peak for this LPAR. Normally your SCRT tool would report a peak 4hr rolling average of 1.500 MSU. That is, 1.500 for z/OS and 300 for CICS. Maybe this customer is using a zIIP which could already have an influence on the reported peak that would otherwise perhaps have been 1.800 MSUs.
- For CICS, as we already indicated, 300MSU is reported. That means that for z/OS you will pay for 1.500 MSUs and for CICS you will pay 300 MSUs. Other software like DB2 have their own reporting but let's assume they have the same MSUs as z/OS.
- Now you have to determine which part of CICS is used for Mobile Transactions and which part is not. This can be a particularly difficult one, since the same transaction might be executed many times, but it can originate from a mobile app or perhaps also from a plain and simple terminal. This is something a customer has to agree upon with IBM and it's based upon e.g. specific fields of SMF records
- Based on that, you come to the conclusion that 200 out of the 300 MSUs that are reported by CICS actually have a mobile origin.
Well, you get a reduction of 60% on that part of the CICS workload reducing that part of 200 MSUs to 80 MSUs.
- The good news is that for that LPAR, all the software in that LPAR, including z/OS and DB2 get the same reduction. This means that the reported 1.500 MSUs is lowered to 1.380 MSUs.
- The rest is BAU (Business As Usual), you can use this to determine the peak for that month. It is e.g. possible that at another moment your z/OS reports 1.400 MSUs at a moment that no mobile workload enters the system. This would mean that for that month your peak will be at 1.400 MSUs.

Well, this little example concludes our short walk through Software Pricing history. I might just add one source of information. There's a very elaborate IBM internet page that explains you utterly everything about software pricing you ever wanted or, perhaps, didn't want to know. It's all [there](#) for you.

Announcements

IBM z/OS Version 2 Release 2 enhancements and statements of direction

This announcement ([ZP16-0504](#)) “describes new capabilities designed to support:

- Asynchronous CF duplexing for lock structures
- z/OS platform software installation improvements
- IBM Cloud Provisioning and Management for z/OS
- ICSF enhancements for Crypto Express5S updates
- z/OS Client Web Enablement Toolkit adds REXX support
- Real-time SMF Analytics infrastructure support”

And there are also a couple of Statements of Direction of which one caught my particular interest: “IBM intends to support new capability in z/OS for metering and capping workloads over CPU and memory consumption. This capability will be delivered in stages with the initial focus on workloads that run only on specialty engines.” I’m curious what that will be about.

If you want more information on z/OS 2.2, take a look at the z/OS V2R2 Knowledge Center home [page](#).

Price Changes effective January 1, 2017

It’s a traditional, I know, but we still have to mention it, the annual price [changes](#) and the ‘change’ is going in the direction you already suspected: the only way is up !

“Today, IBM® announces increases in monthly license charges (**MLC**) on select operating system and select middleware software programs and their features. For the affected programs, the price changes will only apply to the following monthly license charge (MLC) software pricing metrics: VWLC, AWLC, CMLC, EWLC, AEWLC, PSLC, FWLC and TWLC. **The price increase for a given program will be approximately 4%** depending on the features selected.

I think that not many MLC products escape the price increase: all commonly used software is there: CICS, DB2, IMS, Cobol, PL/I, z/VM, z/OS, z/VSE... Have a look at the list for your particular software. And from the calculations I’ve made so far, 4% seems to be pretty accurate.

Hardware withdrawal: IBM TS3500 Tape Library models L23 and L53 and select features

Do you remember when these were first announced. No? I also had to look it up: the L23 and L53 were first announced in May 2006 and now their End of Marketing date was announced ([ZG16-0129](#)). Effective October 7, 2017 you will no longer be able to order the L23 (base frame for 3592) and L53 (base frame for LTO) frames.

As a replacement we have the TS4500 L25 and TS4500 L55 frames.

Since the D23 and D53 frames are not mentioned, I guess that you can still order additional expansion frames on the TS3500 tape library.

DB2 12 for z/OS and DB2 Tools

Along with the Tools ([ZP16-0415](#)) DB2 12 was announced ([ZP16-0492](#)) on October 4, 2016 and is generally available since October 21, 2016.

There's a lot that can be discussed about this new version, so just let us focus on a couple of focal points

- **Faster analytical insights.** DB2 12 is the leading enterprise data server for business critical transactions and analytics. It delivers highly concurrent queries 100x faster providing deeper insights.
 - o Analytics workload improvements including optimized performance and improved zIIP offload.
 - o Deliver insights faster with up to 50% elapsed query time improvement.
 - o New fast index traversal block (FTB's)
 - o Improved DB2 Analytics Accelerator performance
- **Scale and speed for mobile era.** Enhanced support for cloud and mobile workloads with dramatic scalability and performance improvements Increasing maximum table size from 16 TB to 4 PB and utilizing RESTful API technology.
 - o Support for SQL as a service (SQLaaS) through RESTful connectivity
 - o Enhanced support for cloud and mobile workloads with dramatic scalability improvements for tables, increasing from 16 TB to 4 PB
 - o Enhanced support for the next generation of mobile applications
 - o More [info](#) on how to Build a 'DB2 for z/OS mobile application using IBM MobileFirst'
- **Reduced cost and enhanced performance.** 23% lower CPU cost through advanced in-memory techniques and enhanced continuous availability, scalability, and security for business-critical information.
 - o All parallel child task processing is now zIIP eligible.
 - o Up to 30% CPU improvement for query workloads and even higher CPU improvements in select query workloads exploiting DB2 12 query optimization. Up to 10% online transaction processing (OLTP) CPU savings with larger memory and activation of memory exploitation features.
 - o High-volume insert performance equaling two times throughput increase for concurrent sequential insert without clustering.

z/VM V6R4 announcement

You might say that z/VM was somewhat bleeding to death, left in a corner when the rollercoaster called 'Linux on z' came along So z/VM is more thriving than ever on the mainframe system and it's clearly positioned as the hypervisor for Linux Workloads in this announcement of z/VM 6.4 'IBM z/VM V6.4 delivers industry-proven advanced virtualization capabilities to support the increasing demands of a hybrid cloud environment ([Z16-0111](#))'. Whether you have a lot or just some limited knowledge of z Systems.

"z/VM V6.4 virtualization technology is designed to run hundreds to thousands of Linux servers on a single IBM z Systems or LinuxONE server (...) The ability of z/VM to support multiple machine images and architectures provides a highly flexible production and test environment for z Systems and LinuxONE operating systems (...)

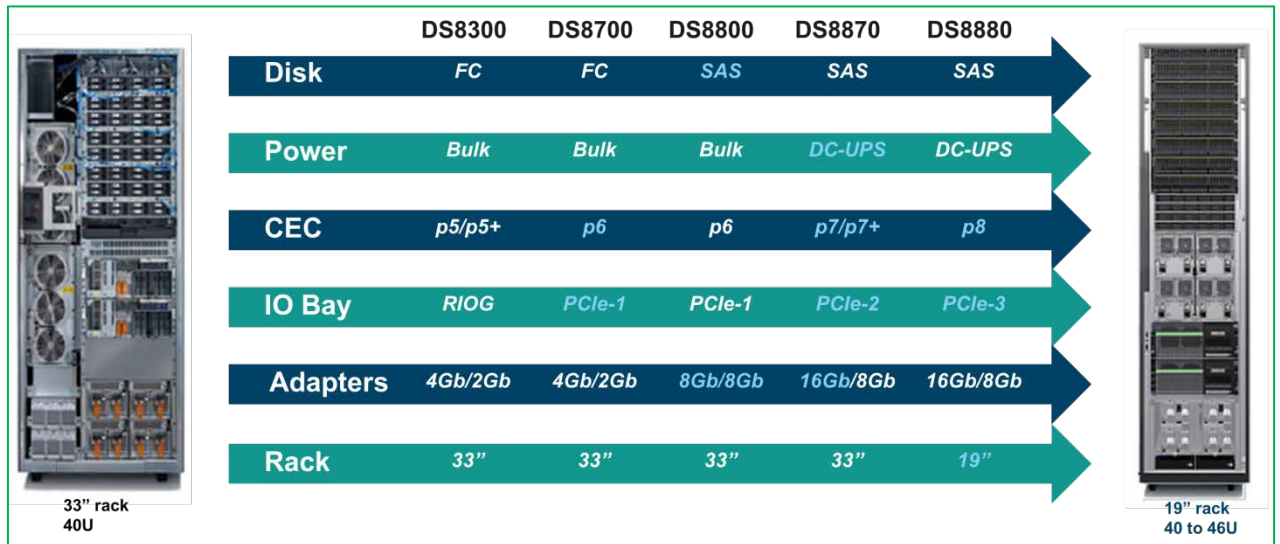
z/VM V6.4 provides support for the IBM z13™, IBM z13s, IBM LinuxONE Rockhopper™, IBM LinuxONE Emperor™, and zEnterprise® (196, 114, EC12, BC12) servers, as well as Red Hat, SUSE, and Ubuntu Linux distributions.

Improvements (...) allow low-end devices such as IBM Storwize® V7000, V840, and V9000 to be attached to a z/VM host, removing the need for a SAN Volume Controller.

z/VM V6.4 is a supported environment using IBM Dynamic Partition Manager for Linux-only systems with SCSI storage. This simplifies system administration tasks for a more positive experience by those with limited mainframe skills".




DS8888 – the new all flash DS8xxx system

This is an illustration that returns at every DS8000 announcement and I still like it a lot, as it gives you at a glance what's changed over the years and where the focus is nowadays. It pretty much explains itself, doesn't it?



On May 31, IBM announced ([ZG16-0067](#)) the All Flash Array storage system in the DS8000 series: the DS8888. "This model delivers up to 2.5 million IOPS random I/O workload environments. DS8888 supports 24- and 48-core processors, with 1 and 2 TB system memory respectively, 128 host adapter ports, and 16 high-performance flash enclosures, up to 480 flash cards".

So here's an overview of the DS8000 family as it is marketed today.

	Business Class	Enterprise Class	Analytic Class
			
	DS8884	DS8886	DS8888
	Affordable hybrid-flash block storage solution for midrange enterprises	Faster hybrid-flash block storage for large enterprises designed to support a wide variety of application workloads	Fastest all-flash block storage with superior performance designed to address the most demanding business workloads
Model	984	985 (Single Phase) 986 (Three Phase)	982
Cache	256GB	2TB	2TB
FC/FICON ports	64	128	128

Media	768 HDD/SDD 96 flash cards	1536 HDD/SDD 192 flash cards	480 Flash cards
Max raw capacity	2.6 PB	5.2 PB	384 TB

In addition to that I just want to draw your attention to a couple of new characteristics that were introduced as well.

“RAID 10 support for flash

RAID 10 will be supported for flash in addition to the already supported RAID 5. No intermixing is allowed for different capacities or different RAID types within a single flash enclosure. To support RAID 10, the flash enclosure must be fully populated with flash cards.

Three-phase power support

Three-phase power will be available for DS8886 and DS8888. Three-phase power reduces DS8886 individual rack capacity by two drive enclosures, however maximum system capacity remains at 1536 drives.

Thin provisioning for CKD

Support for the use of thin provisioning with extent space efficient (ESE) volumes is now available for CKD. In combination with the use of small extents, this replaces the previous track space efficient volume function, delivering a higher performance capability for sequential write workloads.

DS8884 and DS8886 support the ability to configure FlashCopy target volumes as thin provisioned.

SuperPAV

SuperPAV extends upon the previous HyperPAV capabilities by enabling alias devices to be used across multiple logical subsystems (LSS). With this release, an alias device can be used for any base device on the same DS8000 server and in the same path group on the server. SuperPAV will enable performance to be sustained with fewer overall alias devices and offer greater parallelism for individual larger volumes, improving scalability and performance.”

Hints and Tips

Hot Topics #30

This one is always worth mentioning: last August IBM z/OS [Hot Topics \(issue 29\)](#) was published. The newsletter has a rich content:

- IBM's strategy on Cloud Computing with an article by Frank De Gilio, IBM's Chief Architect for Cloud entitled 'Becoming a Hybrid Cloud hero in three easy steps'.
- z/OSMF (several articles)
- Some articles on z/OS migration
- z/OS reporting enhancements for Mobile Workload Pricing
- New LPAR Capping options
- ...

Last call for MES upgrades on your z12 system(s)

Let me freshen up your memory, just in case you wanted to have a MES upgrade on your z12 system(s) like disruptively adding memory, adding books or I/O cards, you have time till the end of the year. Or e.g. you plan on acquiring an IDAA next year, but if you stay on your zEC12 or zBC12, you will have to order the necessary 10GbE cards while it's still possible. Afterwards it will be a no go. This was announced half of February this year, shortly after the announcement of the z13s ([ZG16-0021](#)).

During 2017 'upgrades' that are delivered through a modification to the machine's Licensed Internal Code (LIC) will still be possible. E.g. adding additional CBU records or activating memory that's already in the machine. Here's an overview I made a couple of months ago:

Type	Family	GA	WDFM HW	Yr/Mt from GA	WDFM LIC	Yr/Mt from GA
2965	z13s	mrt/16				
2964	z13	mrt/15				
2828	zBC12	sep/13	dec/16 (*)	3.3	dec/17	4.3
2827	zEC12	sep/12	dec/16 (*)	4.3	dec/17	5.3
2818	z114	sep/11	jun/14	2.9	jun/15	3.9
2817	z196	sep/10	jun/14	3.9	jun/15	4.9
2098	z10 BC	okt/08	jun/12	3.8	jun/13	4.8
2097	z10 EC	feb/08	jun/12	4.4	jun/13	5.4

(*) zEC12/zBC12 - ROHS countries - New Models and upgrades to zEC12/zBC12 WDFM jun/16
 (*) zEC12/zBC12 - Other countries + Hardware upgrades within zEC12/zBC12 WDFM dec/16

EOS dates Operating Systems

In this topic, we give you an overview of the EOS dates of the operating systems. If you have questions regarding this topic or End of Support dates of other software you can always contact us at Realdolmen and we will gladly help you out.

Product-Version-Release	Date Availability	Date Withdrawn from Marketing	Date End Of Support (EOS)	Date End of Extended Support
z/OS v1.10	09/2008	09/2009	09/2011	
z/OS v1.11 (*)	09/2009	09/2010	09/2012	
z/OS V1.12 (**)	09/2010	10/2011	09/2014	09/2017
z/OS V1.13	09/2011	01/2014	09/2016	09/2019
z/OS V2.1 (***)	09/2013	01/2016	09/2018	09/2021
z/OS V2.2	09/2015	09/2017	09/2020	09/2023
z/VM V5R3	06/2007	09/2008	09/2010	
z/VM V5R4 (****)	09/2008	03/2012	12/2017 (no longer tied to the EOS for z9)	
z/VM V6R1	10/2009	11/2011	04/2013	
z/VM V6R2	12/2011	07/2013	06/2017	
z/VM V6R3	07/2013		12/2017	
z/VM V6R4	11/2016			
z/VSE V4R1	03/2007	10/2008	04/2011	
z/VSE V4R2	10/2008	11/2010	10/2012	
z/VSE V4R3	11/2010	06/2012	10/2014	
z/VSE V5R1	11/2011	05/2014	06/2016	
z/VSE V5R2	04/2014	03/2017		
z/VSE V6R1	11/2015			

Green: active versions
Blue: available versions
Red: announced versions

(*) Last version to offer Lifecycle Extension support

(**) First version with a 3-year extended support after End of Support. Fee-based corrective service.

(***) A recent [announcement](#) set the End of Marketing date for z/OS 2.1 on 25/01/2016

(****) No longer supported by the z13 and z13s

Agenda

Our agenda informs you of events that may be of interest to you: webcasts, Proofs of Technology, events in the Brussels IBM Forum and GSE Working Group meetings. Since it's the end of the year, not many new 2016 sessions have been announced so far.

We also have an agenda in the right hand column of our [blog](#) entitled 'Upcoming events Belgium'. The agenda on the blog is updated regularly so you don't have to miss any events.

Datum	Meeting	Agenda	Information and registration
13-17/11/2016	DB2 IDUG 2016 EMEA Tech Conference at Crown Plaza Brussels	<p>Agenda can be found here</p> <p>In short:</p> <ul style="list-style-type: none"> • Five days of education sessions • Half and full-day workshops • More than 100 one-hour technical sessions • Three expert panels on z/OS, LUW & Application Development 	Information
15/11/2016	GSE Young professionals Working Group meeting at Realdolmen headquarters in Huizingen	<ul style="list-style-type: none"> • Innovation in Payments @ KBC [Stefaan Vanneste – KBC] • Testimonial of a Young Professional Part I [Elona Dervishi – Euroclear] • Cognitive Era [Laurent Boes - IBM] • Testimonial of a Young Professional Part II [Gregoire Del Lisse – Euroclear] 	Information
29/11/2016	Webcast IBM Systems Magazine on 'Where's My Bottleneck Now? Staying Ahead of Your 4HRA' (content by MVS Solutions Inc.)	<p>How to</p> <ul style="list-style-type: none"> • Stay ahead of your 4HRA • Use new software pricing to your advantage • Understand how your 4HRA is actually calculated and which peak drives your machines vs. which peak drives your software bill (hint: they're usually not the same!) 	Information and registration
14/12/2016	GSE z/SO Working Group Meeting at Realdolmen headquarters in Huizingen	To be determined	Information



Small print

We would like to point out that all texts in this newsletter are based upon the interpretation by Realdolmen of the information that is at our current disposal. Therefore Realdolmen cannot be held responsible for an incomplete interpretation of the data at hand.

www.realdolmen.com

Industriezone Zenneveld A. Vaucampsaan 42 1654 Huizingen
Tel +32 2 801 55 55 Fax +32 2 801 55 99 BTW/TVA be 0429.037.235
rpr BRUSSEL/RPM BRUXELLES