z/OS V2 R3 in the Clouds (cumulus nimbus, extratus, cirros) in the CMG imPACt 2017 Conference

08/11/2017

By Alvaro Salla

z/OS 2.3 mantras

Focusing in the critical areas:

- Continuous Deliver
- Simplification (including Performance)
- Security, mainly encryption
- Cloud and other hypers
- z/OS V2.3 is planning to provide a simple, transparent, and consumable approach to enable extensive encryption of data, simplify the overall management of the z/OS ecosystem to increase productivity, and provide a simple, consumable approach for self-service provisioning and rapid delivery of software as a service, while enabling for the API economy.

  (cloud?)

z/OS 2.3 first words

z/OS is designed to support companies most mission-critical work while meeting stringent service levels, illustrated by clients that include the world's leading banks, financial services companies, healthcare enterprises, and governments.

Modernly z/OS needs to face the current hypers. For example, Big Data, Cloud, Analytics, Cognitive are driving higher data and transaction volumes and accelerating the rate of application changes made by enterprises.

In response there is a rapid evolution toward hybrid flexible and scalable IT architectures that rely on combinations of off-premises and on-premises IT resources.

Then, we may say that z/OS 2.3 is able to face the growth in capacity, scale, continuous availability, and throughput required to improve business performance, meet response time objectives, protect sensitive data and transactions, minimize operational risk, and application lifecycle management.

Before z/OS Continuous Delivery

Previously, we had the z/OS Releases upgrades and preventive maintenance trough RSUSs.

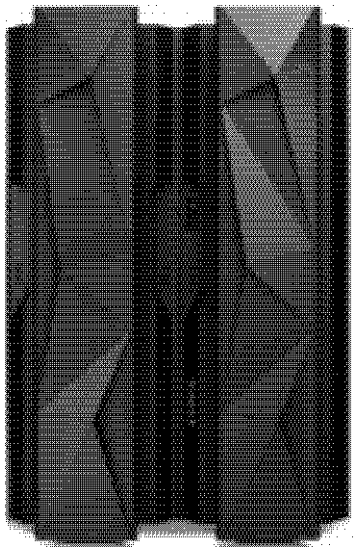

z/OSV1.13

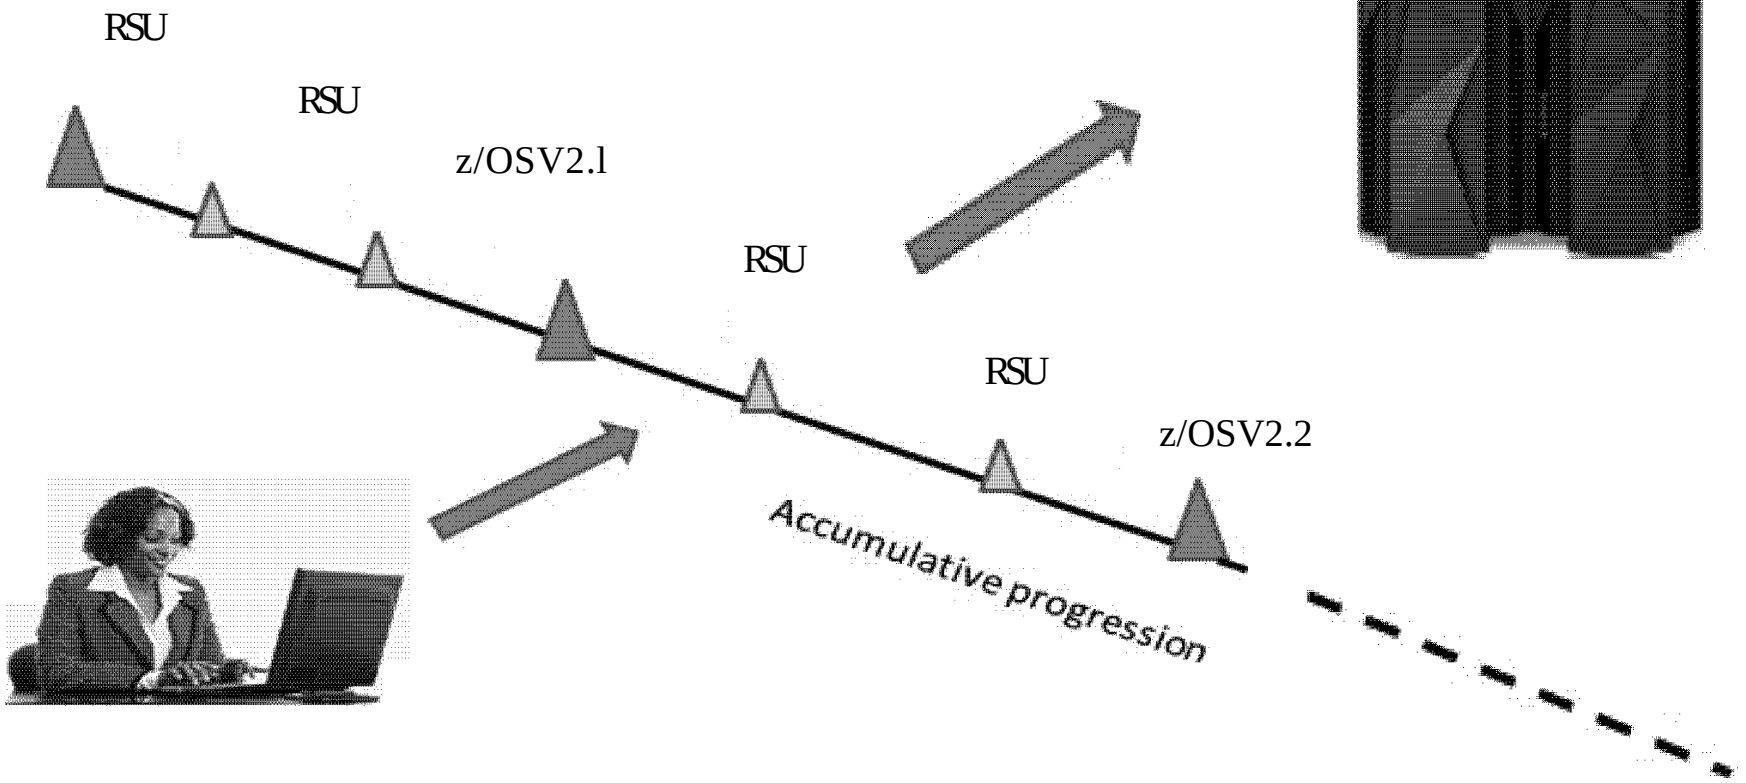

Refreshing Recommended Service Upgrade (RSU)

- When new functions are provided through the z/OS service stream, the Recommended Service Upgrade (RSUs) process is in place to provide a quaUty test in a robust environment for those program temporary fixes (PTFs). Then, RSU provides a consistent, installable and tested preventive maintenance level for z/OS and its subsystems, such as: DB2, CICS, WAS and so on.
- The Small Program Enhancements (SPE) were also inciuded in the RSU stream, so you can wait for the RSU if you want to install more than one SPE in a single instance. However, waitingfor the RSU extends the time to wait for the new functionality to be available on your systems.

z/OS Continuous Delivery?

- z/OS Continuous Delivery is a way of distributing new z/OS software functionalities to customers, as soon they are ready from the labs.
- It addresses the old question: How do you want of being informed of good news? All together (as in a new z/OS release) or drop by drop (in a small program enhancement)?
- The best answer is by the drop, mainly, knowing that:

  - New technologies take time to be digested and accepted by the market , and
  - Customers are usually under staffed in technical personnel and are not able to deal with large amount of modifications alltogether.

- However, z/OS covers a large range of functionality. Some of them are better suited for Continuous Delivery (such as SDSF and z/OSMF). Other areas align more closely with a z/OS release delivery cyde because they are more complex and might be tightly integrated with new microcode and hardware.

z/OS Continuous Delivery with SPEs

Why to wait for a ready and maybe a needed function to gain buisiness and competitiveness

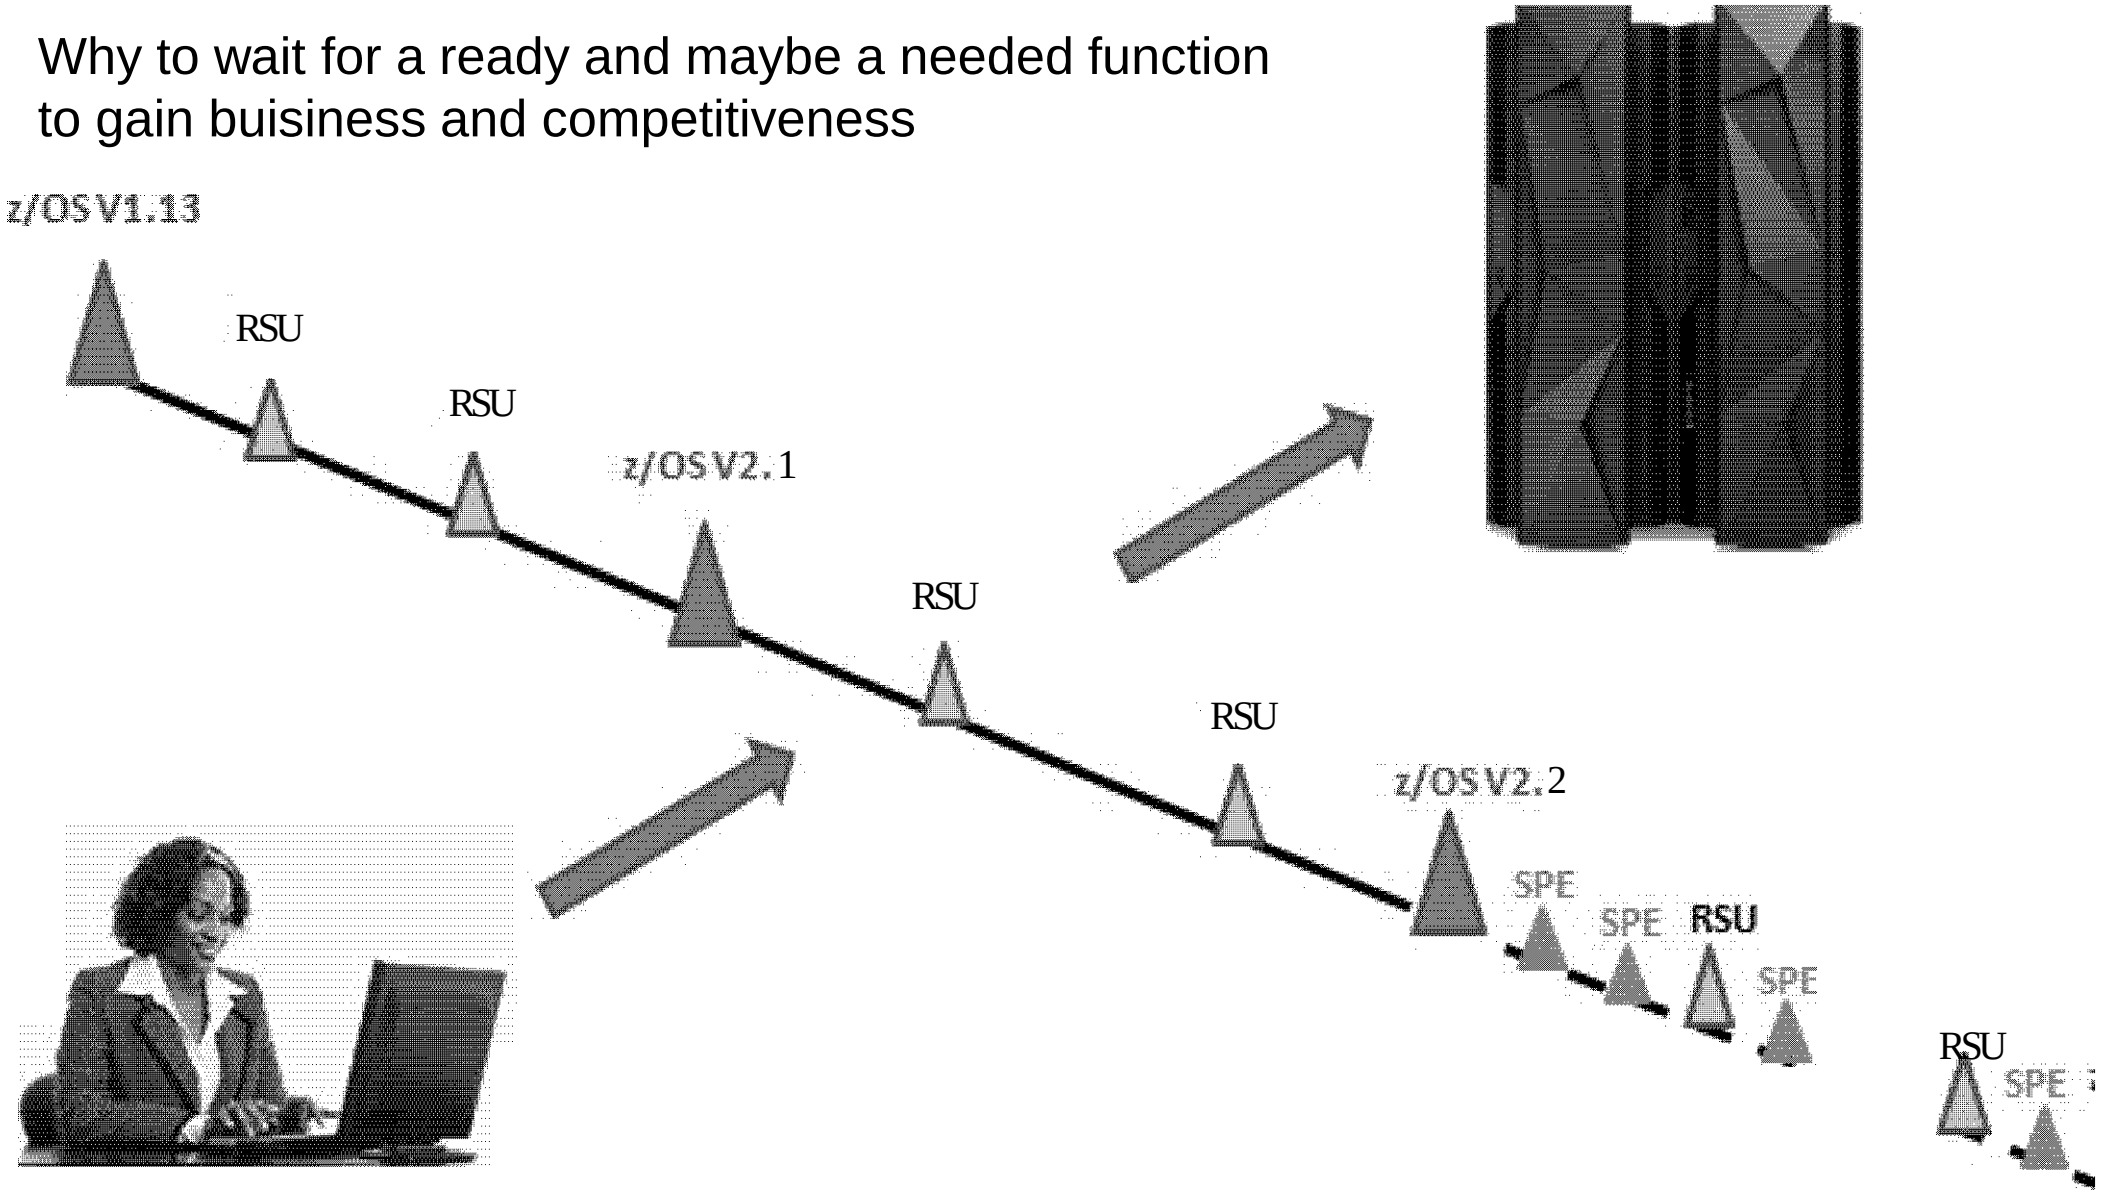

z/OS Continuous Delivery and z/OSMF

- z/OSMF is a z/OS component delivering solutions in a task-oriented, web-browser-based user interface with integrated user assistance. z/OSMF goal is to improve z/OS system programmer productivity, and make functions easier to understand and use. z/OSMF improves skills, as quickly as possible, with least amount of training. Then, you can automate tasks, reduce the learning curve, and improve productivity through a modern, simplified, and intuitive task-based, browser-based interface.

z/OS Continuous Delivery and z/OSMF

At z/OS Continuous Delivery, you might see scenarios in which implementation steps for new functionality are provided in z/OSMF Workflows.

It is expected that providing this information will be more common in the future.

It is recommended that you activate z/OSMF and increase your skills in areas (such as Workflows) to use the functions in a timely manner. The role z/OSMF plays in each SPE varies according to the SPE's content and the necessary implementation tasks. The aim is to ease the implementation effort and and provide a demonstrable consistency and accountability.

z/OS V2 R3 Simplification

z/OS Simplification (I)

z/OS is a very complex operating system only because it solves many complex problems, such as, guarateeing: integrity, security, availability, compatibility, deferentiate performancen to hundretths of concurrent business transactions. However, this complexity never spills to the business end user. Historically, this complexity does not apply to z/OS systems programmers.

z/OS has been delivering simplification efforts, but IT is facing a mixed skills workforce composed of professionals who are new to z/OS and those who are already skilled in z/OS.

Then, z/OS V2.3 is being designed to simplify and modernize the user experience and help make pertinent information readily available and easily accessible.

z/OS Simplification (II)

Among these simplification efforts, we have:

- In z/OS V2.3, z/OSMF component optionally starts late in IPL so that it is available all the time. The logon experience will be improved as well as the initial landing page. Facilities will be added to make administration of z/OSMF easier.

- Improving the time to value for new products with enhancements to installation and configuration through facilities such as portable software packages and guided activities by z/OSMF workflows. Clients will have the tools to start to standardize and simplify the installation and configuration experience.

- Plans are to provide a new z/OSMF plug-in, Sysplex Management, that is designed to provide detailed views of sysplex infrastructure resources such as sysplexes and z/OS systems, CFs and CF structures, CF structure connectors, couple data sets and policies, and coupling links.(zRADAR)

Asynchronous CF Duplexing for Lock Structures

Duplexing and its Types

In the same way that you duplex remotely data sets from a 3390 volume, you can do it with CF structures in two CFs. all Duplex opeartions are synchronous. There are two types of structure duplexing:

- User Managed Duplexing

Performance efficient and the majority of work is executed by the exploiter. Used by DB2 cache structures.

- System Managed Duplexing

Less performance efficient and the majority of work is executed by XES. The CF must be connected by CF links, but only for control information passing. Used by IRLM lock structure. Not very recommended.

Pay attention, it is not the installation that selects the types. It is the exploiters.

System Managed Duplex Considerations

- **The CFRM couple data set must be correctly formatted to support system-managed CF duplexing.**
- **There must be two or more CFs that support system-managed CF structure duplexing, connected to one another with CF-to-CF coupling links.**
- **CFRM policy keyword requirements must be met.**
- **Structures should be evaluated for CF duplexing eligibility (costs such as CF overhead and storage).**
- **Not all structures should exploit CF duplexing.**

Postprocessor Coupling Facility (duplexing)

```
                C O U P L I N G   F A C I L I T Y   A C I V I T Y

       z/OS V1R2              SYSPLEX RMFPLEX2            DAT      E 02/20/2001
                               RPT VERSION V1R2 RMF        TIME     14.29.00
------------------------------------------------------------------      ------------
 COUPLING FACILITY NAME = CF01
 TOTAL SAMPLES(AVG) =  59  (MAX) =  59  (MIN) =  59
------------------------------------------------------------------      ------------
                                 COUPLING FACILITY USAGE SUMM   ARY
------------------------------------------------------------------      ------------
 STRUCTURE SUMMARY
------------------------------------------------------------------      ------------

       STRUCTURE                        ALLOC      CF                    #
TYPE   NAME               STATUS CHG    SIZE       STORAGE               REQ

CACHE  DSNDB1G_GBP0       ACTIVE        21M          0.5      1893
                          PRIM
       DSNDB1G_GBP4       ACTIVE        39M          1.0      6089
                          SEC
       DSNDB1G_GBP5       ACTIVE        39M          1.0      6089
```

Identify duplexed Structures

```
------------------------------------------------------------------      ------------
                                    COUPLING  FACILITY  STRUCTU   RE  ACTIVITY
------------------------------------------------------------------      ------------------------------------

STRUCTURE NAME = DSNDB1G_GBP0      TYPE = CACHE                  STATUS = ACTIVE PRIMARY
SYSTEM     # REQ     --------------  REQUESTS  -------------   --------------  DELA  YED   ESTS  ------------
           TOTAL              #     % OF      -SERV TIME(MIC) -   REASON    #   % OF   REQU  AVG TIME(MIC)  ----
                                                                                          ----
NAME       AVG/SEC            REQ     ALL      AVG    STD_DEV                    REQ     REQ   /DEL    STD_DEV   /ALL
RMF1        8621     SYNC      89      0.3    101.9      34.6      NO        SCH    1     0.1   82.2      111.3     0.1
             4.79    ASYNC   8532     26.3    376.4     355.6              PR WT     4     0.1  285.8      440.2     0.2
                     CHNGD      0      0.0    INCLUDED IN ASYNC          PR CMP  320     0.3   46.5      345.4     0.4
                                                                         DU MP     0     0.0    0.0        0.0     0.0
```

New Delay Reasons:
- peer subchannel wait contention
- waiting for peer completion

Postprocessor CF to CF Activity

```
    z/OS V1R5              SYSPLEX PLEX02            DATE 10/11/2005           INTERVAL 010.00.000
                        CONVERTED TO z/OS V1R5 RMF      TIME 10.00.00            CYCLE 01.000 SECONDS

-------------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = CF05
-------------------------------------------------------------------------------------------------------------------------------
                                                  CF TO CF ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------
          # REQ                              ----------- REQUESTS -----------    -------------- DELAYED REQUESTS ------------
 PEER     TOTAL       -- CF LINKS --                 #    -SERVICE TIME(MIC)-          #     % OF  ------ AVG TIME(MIC) ----
 CF       AVG/SEC     TYPE       USE                 REQ     AVG    STD_DEV           REQ    REQ   /DEL    STD_DEV     /ALL

 CF06      10756       ICP          2        SYNC   10756     0.4       0.0     SYNC    0     0.0    0.0        0.0       0.0
           17.9
```

What a hell?

Asynchronous CF Duplexing for lock structures is planned to be improved that makes duplexing Coupling Facility (CF) Lock structures practical, even at extended distances.

It is planned to deliver a general purpose interface for any CF lock structure exploiters (such as, VSAM RLS) and provides substantial performance advantages for duplexing lock structures.

Because lock structure duplexing is performance wise acceptable, there is no need anymore of Isolation, that is, to keep such structure in an external CF.

On top of that, the exploiter needs to understand that asynch access implies in actions no needed for synch, in order to guarantee the integrity of data in primary structure falure.

Notes:

Asynchronous CF Duplexing is currently available on IBM z13 and IBM z13s systems processors. It requires CFCC Level 21 with service level 02.16, or higher, z/OS V2.2 SPE with PTFs for APAR OA47796 and RMFTM V2.2 reporting support delivered with PTFs for APAR OA49148, CF to CF connectivity via coupling links, and exploitation, for example, DB2 V12.

WLM Sysplex Routing sensitive to upcoming soft capping

Example of using WLM Weights by a Transaction Manager

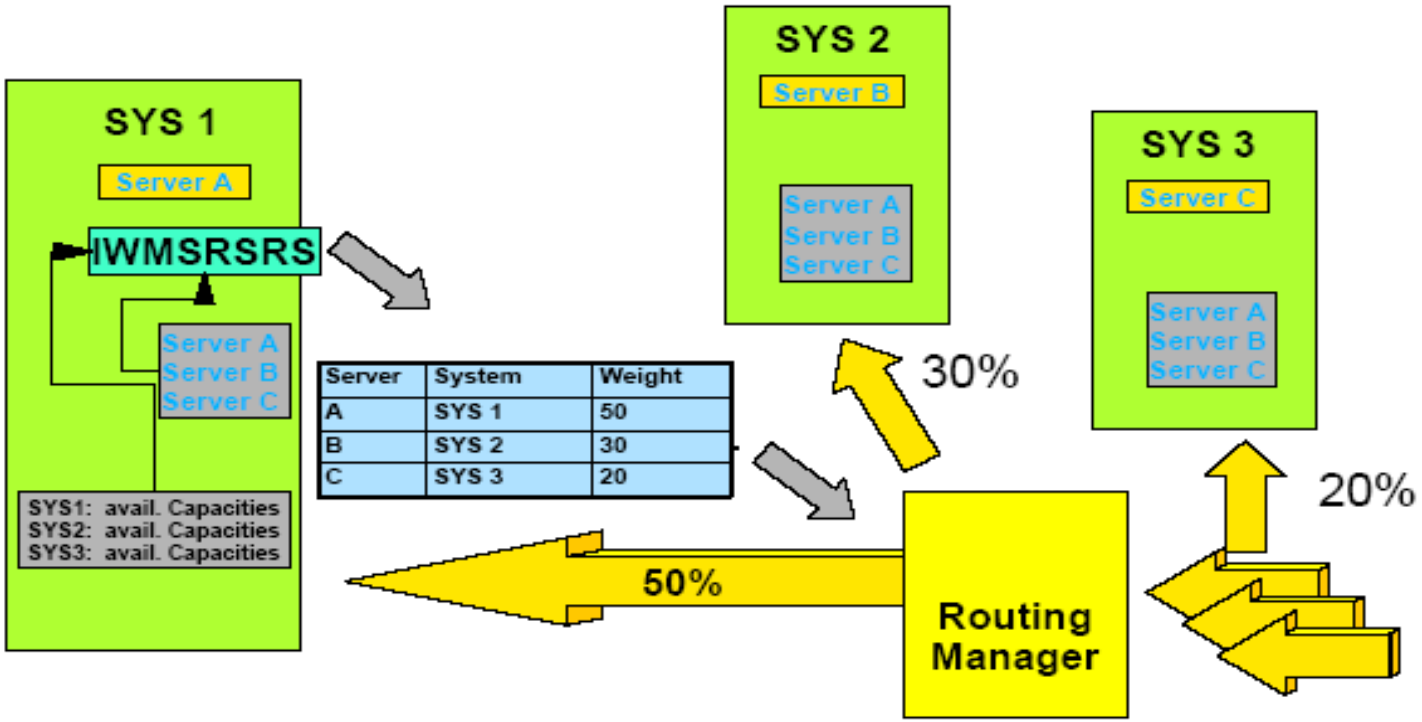

A Transaction Manager uses IWM4SRSC service to request a routing recommendation.
WLM returns a list of servers and weights (numbers between 1 and 64).
The Transaction Manager (exploiter) can use the weights to distribute incoming work.

WLM Available LP Capacity Measurements

In order to implement the Capacity Based model, WLM measures the z/OS (LP) available capacity by:

- **The difference between the actual utilization of this LP and its guaranteed capacity (a result of its LP weight). Obviously, the number and availability of online logical CPs and zIIPs influence the WLM Weights.**

- **Plus some portion of the total unused capacity on the CEC, apportioned based on the relative Weight of this LP compared to the other LPs on the CEC.**

- **WLM also takes into account the relative WLM Importance of the future arriving transaction, and includes capacity used by lower importance workloads in the available capacity value, because that lower importance work could be displaced to cater for the new, higher importance work. Then the same server will get several Weights for every transaction WLM Importance.**

Free LPAR Capacity - Example 1

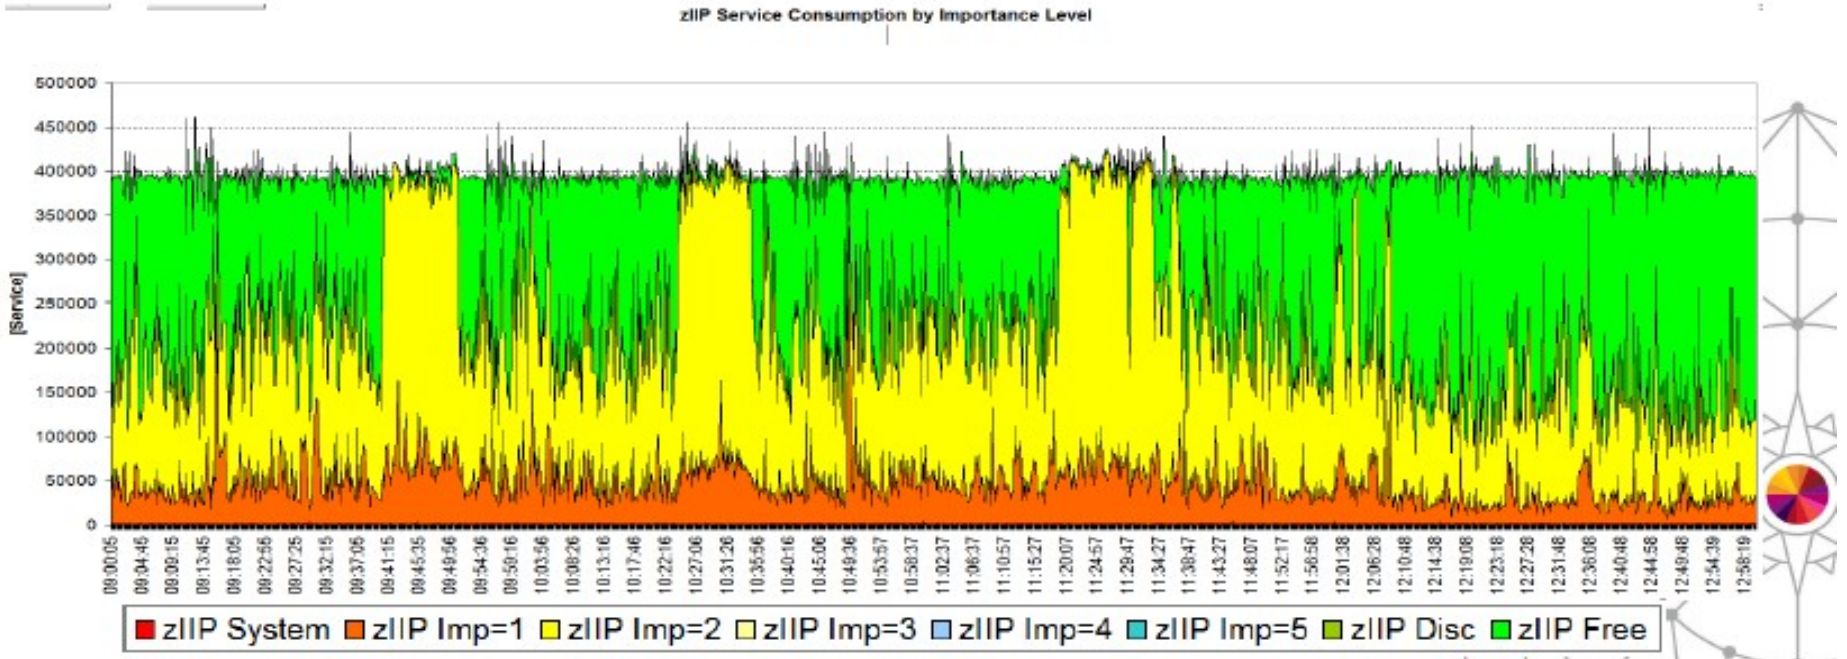

- While an LPAR is running below its weight entitlement and no capping is in effect the total consumed plus free capacity is usually pretty constant.

The Weight entitlement refers to the LP guarantee. BTW, the above phrase is obvious. The Y axis shows Service Units.

WLM Routing Weights Computation

- **Compute capacity-based Weights for servers**
  - Includes adjustment for zIIP processor capacity, crossover importance level weighting
  - Return Weights for each CPU/zIIP processors and combined Weight
  - Frequently scaled to 64

•When multiple servers (AS/enclaves) run on a z/OS divide this z/OS system Weight by the number of servers to derive each server's Weight.

•The IWMSRSRS (with SPECIFIC option) and IWM4SRSC calculate the server Weights based on:

- Capacity Based figures
- Performance index (PI)
- Queue time ratio
- Health indicator

DFSMS V2.r3 news on Simplification

- **DFSMSdfp SAM and VSAM enhancements are planned to provide read-only access to data sets that reside on Peer-to-Peer Remote Copy (PPRC) secondary volumes. Certain applications can take advantage of redundant hardware and avoid interference with production work.**
- **z/OS Allocation and DFSMSdfp enhancements are intended to provide improved performance and scalability for DB2 workloads by allowing the number of concurrent open data sets in a single AS to grow and by improving performance of data set open and close processing.**
- **Enhancements to zFS file systems are designed to allow individual files to be compressed utilizing the zEDC compression card technology. Existing file systems can be compressed while in use. This is not limited to new zFS file systems, and existing file systems can be eligible for compression also.**

More on Simplification

- **XRC (z/GM) is being enhanced to utilize more buffer storage for in-flight updates, making it more resilient to transient events that may otherwise cause suspension or stalls. This is also available on z/OS V2.1 and V2.2 with PTFs for APAR OA49548.**

- **DFSMS provides the ability to encrypt data sets, using either SAF or SMS policies, without changing their application programs. DFSMS uses of the CPACF for extended format version 2 data sets only: BSAM, QSAM, and VSAM. In addition, data set level encryption allows the data to remain encrypted during administrative functions such as backup/restore, migration/recall, and replication.**

- **ITDS implements a new z/OS Health Check that is designed to suggest when the DB2 Reorg (not Defrag) or Runstats utilities should be run for directories in TDBM to help avoid potential performance issues. Runstats updates statistics in the DB2 system catalog about the characteristics of a table, associated indexes, or statistical views.**

WLM and Simplification

- **WLM is planned to be enhanced with an option to cap a system to the MSU value that is specified as the soft cap limit regardless of the four-hour rolling average consumption. An IBM zEC12 (GA2), or higher, server is required. Absolute MSU capping is also available on z/OS V2.2 with PTF UA81256 and on z/OS V2.1 with PTF UA81257 for APAR OA49201.**
- **IFAHONORPRIORITY parameter with Service Class granularity on top of z/OS level (IEAOPTxx). Also WLM resource groups may limit the amount of real storage that may be used by the associated SCs. Both enhancements are designed to provide better control over the execution cost.**

z/OS V2 R3
Security and Crypto

Data Encryption policy for Security

The threat of data breaches in conjunction with compliance mandates are driving the need for clients to adopt extensive use of encryption across their enterprises. z/OS V2.3 replaces application development efforts with transparent, policy-based data set encryption:

•Planning enhanced data protection for z/OS data sets, zFS file systems, and Coupling Facility structures to give users the ability to encrypt data without needing to make costly application program changes.

•Designing new z/OS policy controls to make it possible to use pervasive encryption to protect user data and simplify the task of compliance.

•z/OS Communications Server will be designed to include encryption readiness technology to enable z/OS administrators to determine which TCP and Enterprise Extender traffic patterns to and from their z/OS systems meet approved encryption criteria and which do not.

Security and Crypto

- **z/OS may encrypt CF data, including list and cache structures, controled by CFRM policy using CPACF. The data will be encrypted as it travels on the CF link and will remain encrypted while resident in the CF.**
- **System SSL is compliant with the following RFCs to maintain standards-based security and interoperability:**
  - RFC 6960 X.509
  - RFC 6961
- **PKI Services is planned to support WebSphere Liberty Profile to host the web pages interface. This simplifies installation and exploits the benefits of the smaller footprint of WebSphere Liberty Profile.**

Major Crypto algorithms

- **Cryptography has these major functions:**
  - **Key management**
  - **Data confidentiality and Personal Authentication (PIN)**
  - **Data Integrity (MAC and MDC).**

There are two types of algorithms:

- **Symmetric (encrypt and decrypt keys are the same)**
- **Asymmetric (encrypt and decrypt keys are different)**

Follows the crypto hardware in z13:

- **CP Assist Crypto Function (CPACF) . Executes five instructions: Cipher (KMC and KM), Compute Message Authentication Code (KMAC), Computr Intermediate Message Digest (KIMD) and Compute Last Message Digest (KLMD). Implemented in a inbound synch coprocessor shared by two PUs. Only supports clear Key functions.**
  - **Crypto Express 2 which encompass PCICA and PCIXCC coprocessors.**

Symmetric Cryptography

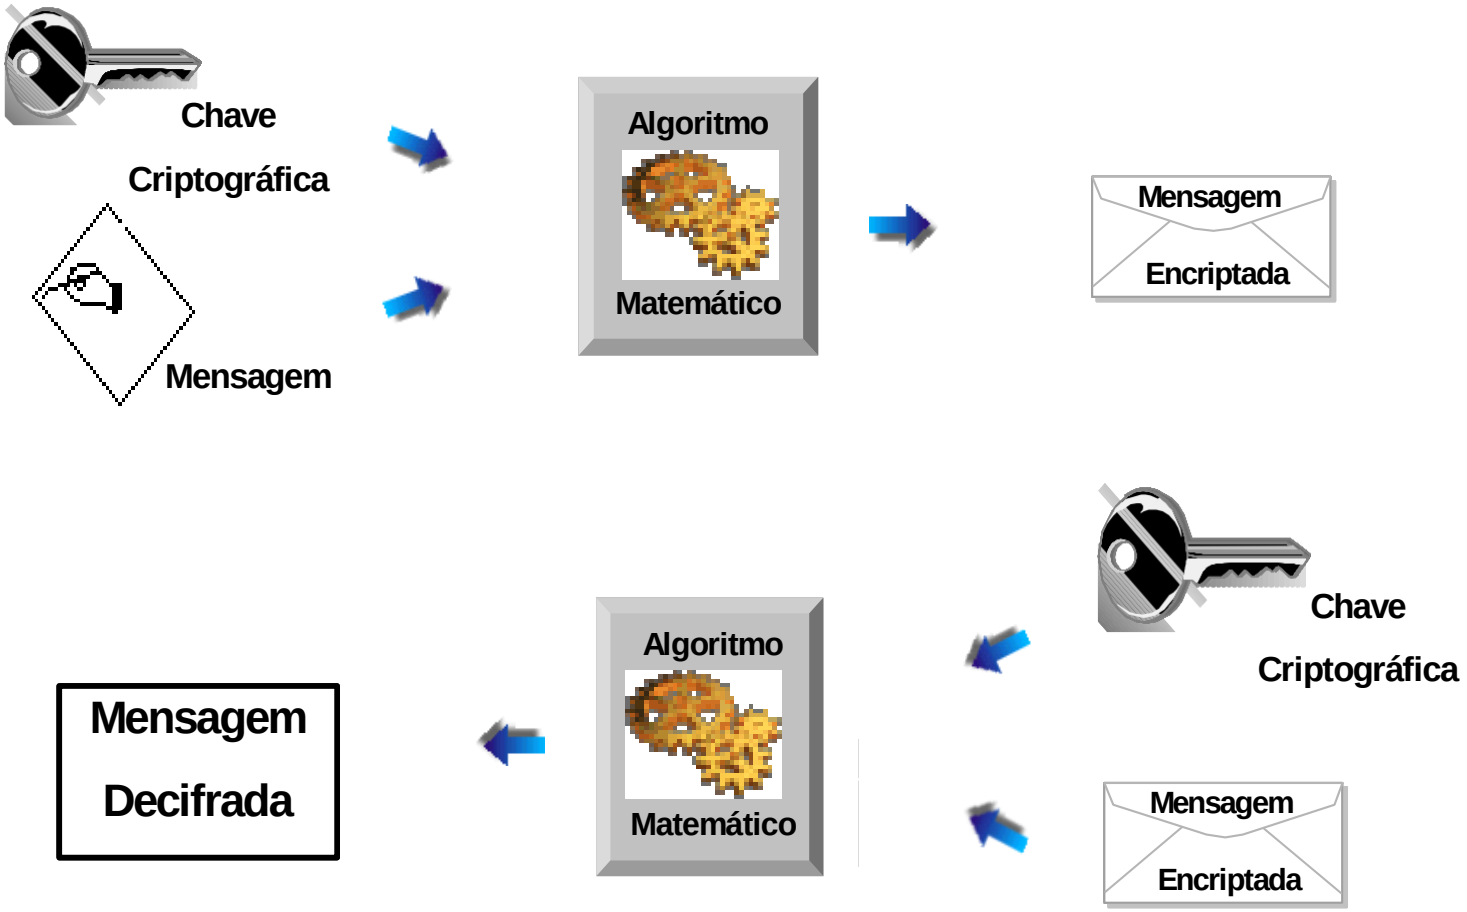

Asymetric Cryptography

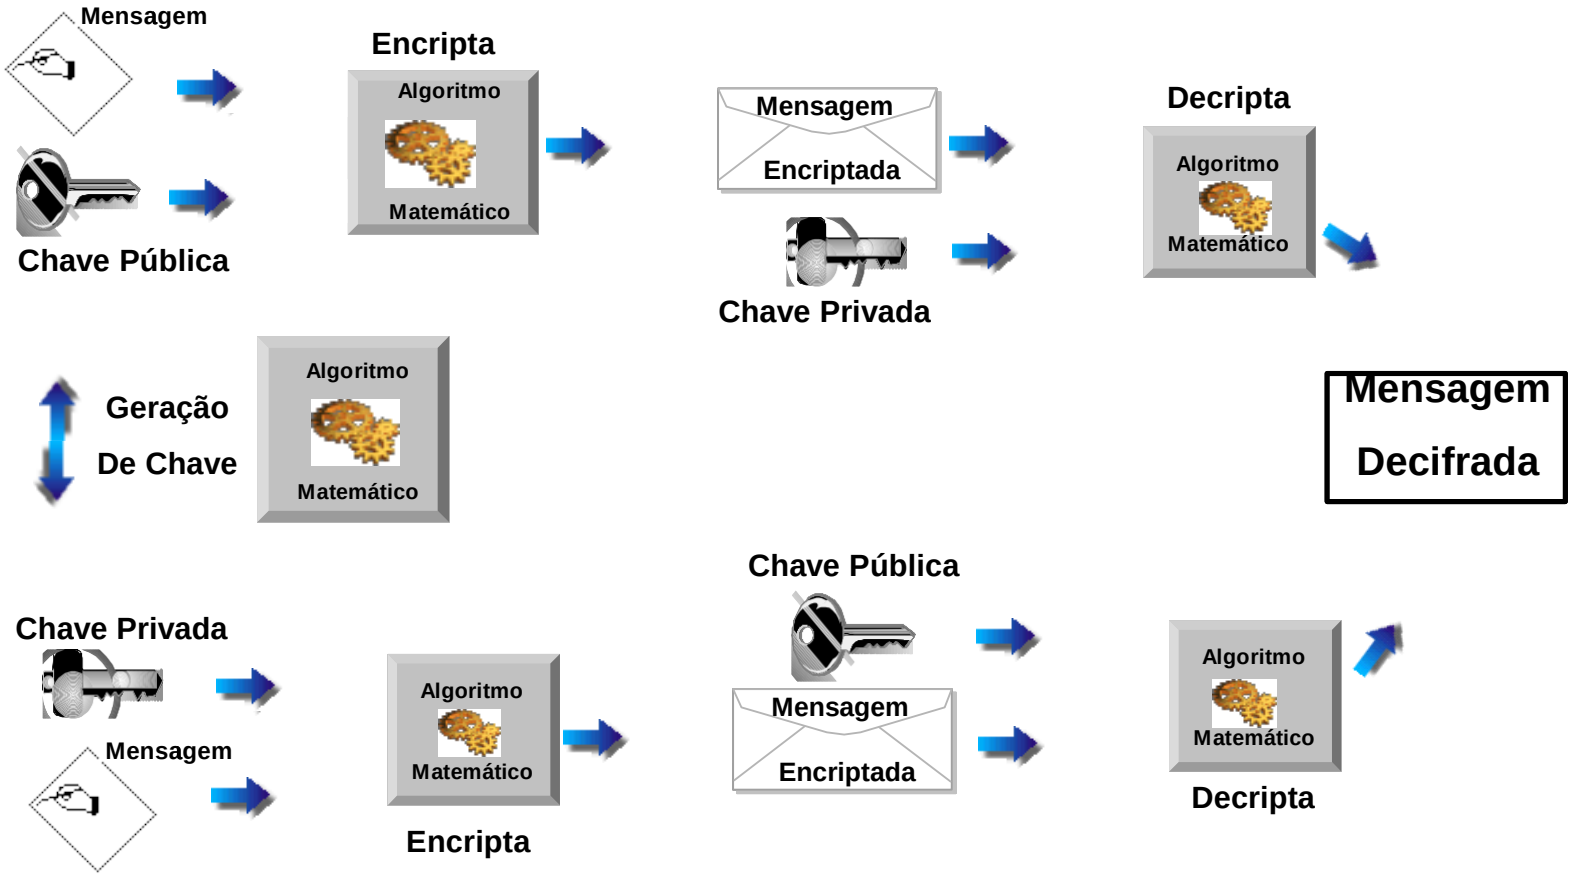

Clear and non-clear key algorithms

Coprocessor with secure keys

tamper proof hardware

(CCF, PCICC, PCIXCC ou CEX2C)

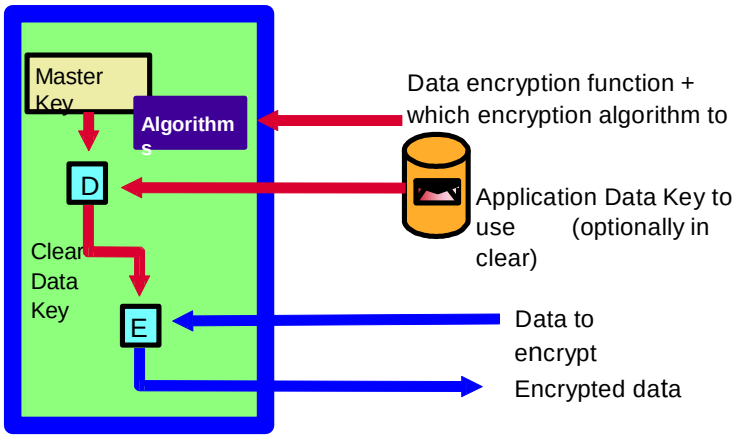

Atributed FIPS 140-1 level 4 (CCF, PCICC e PCIXCC) FIPS 140-2 level 4 (CEX2C) in evaluation

Coprocessor with clear keys

PCICA, CPACF

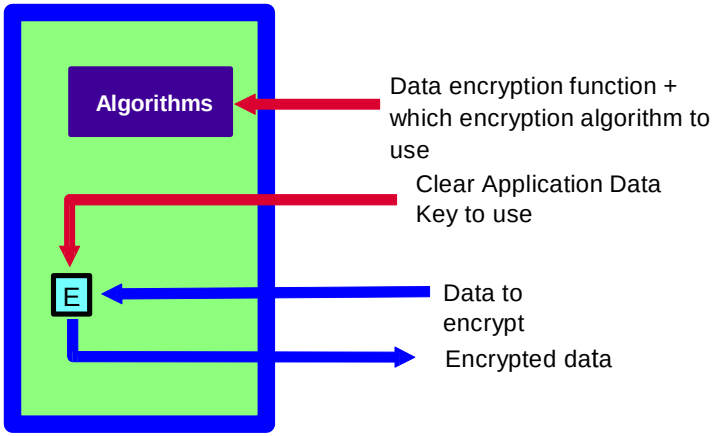

The COP is the CPACF

Each COP is shared by two PUs.

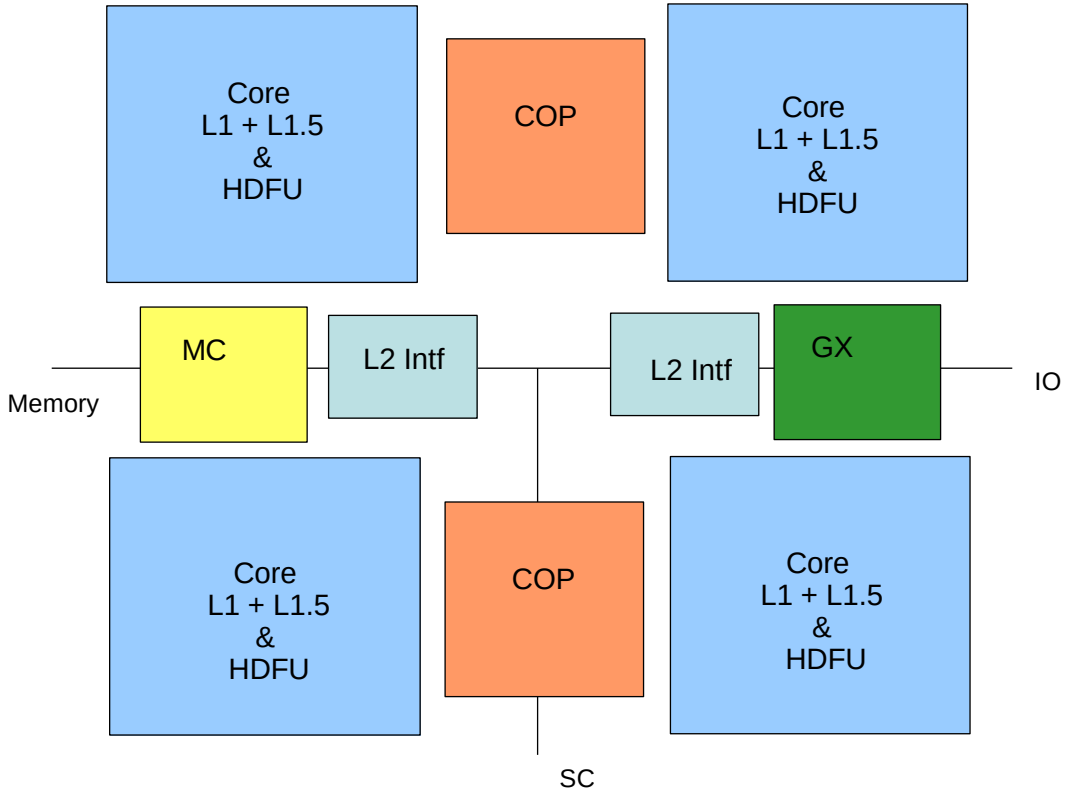

CPACF Algorithms

Data Encryption Standard (DES), which includes:

- Double-key DES (double DES) Triple-key
- DES (triple DES)
- Advanced Encryption Standard (AES) with secure encrypted 128-bit, 192-bit, and 256-bit keys (secure key AES is exclusive to System z10).

Hashing algorithms, such as SHA-1, and SHA-2 support for SHA-224, SHA-256, SHA-384, and SHA-512

- Message authentication code (MAC)
- Single-key MAC Double-key MAC

Pseudorandom number generation (PRNG)

Random number generation long (RNGL) with 8 bytes to 8096 bytes
Random number generation (RNG) with up to 4096-bit key RSA support

More on CPACF algorithms

- Installation of the CP Assist for Cryptographic Functions (CPACF) DES/TDES enablement, feature code 3863. This feature enables the DES and TDES algorithms on the CPACF, which are supported by z/OS, z/VM, and Linux on System z. On the other hand, SHA-1, SHA-256, SHA-384, and SHA-512 are shipped enabled on all servers and do not require this feature.
- Message Autentication Code (MAC) is a value calculated from the message according to a secret shared DES key and sent to the receiver together with the message. The receiver can recalculate the MAC and compare it with the MAC received. If the MAC values are identical, the message has not been altered during transmission.

z13 Crypto Express 5

Crypto Express 5S faster than the previous one, with more algorithms. To crypto the Operlog data at z/OS 2.2

Quem vai decidir onde a API CSNBSYE (Symmetric Key Encipher) ou a API CSNBSYD (Symmetric Key Decipher) irá processar será o ICSF.

A decisão sobre rodar no CPACF ou na Crypto Express será baseada nos parâmetros passados pela API, porém para chaves em claro, 100% das vezes será processada pelo CPACF devido a "performance".

A maneira de "forçarmos" que a API seja executa na placa Crypto Express será definir a chave como segura (encriptada por uma "Master Key"), porém podemos esperar um desempenho pior, pois o processamento será assíncrono enquanto no CPACF o processamento é sincrono.

CPACF não é standard por causa do governo americano.

z13 I/O Structure

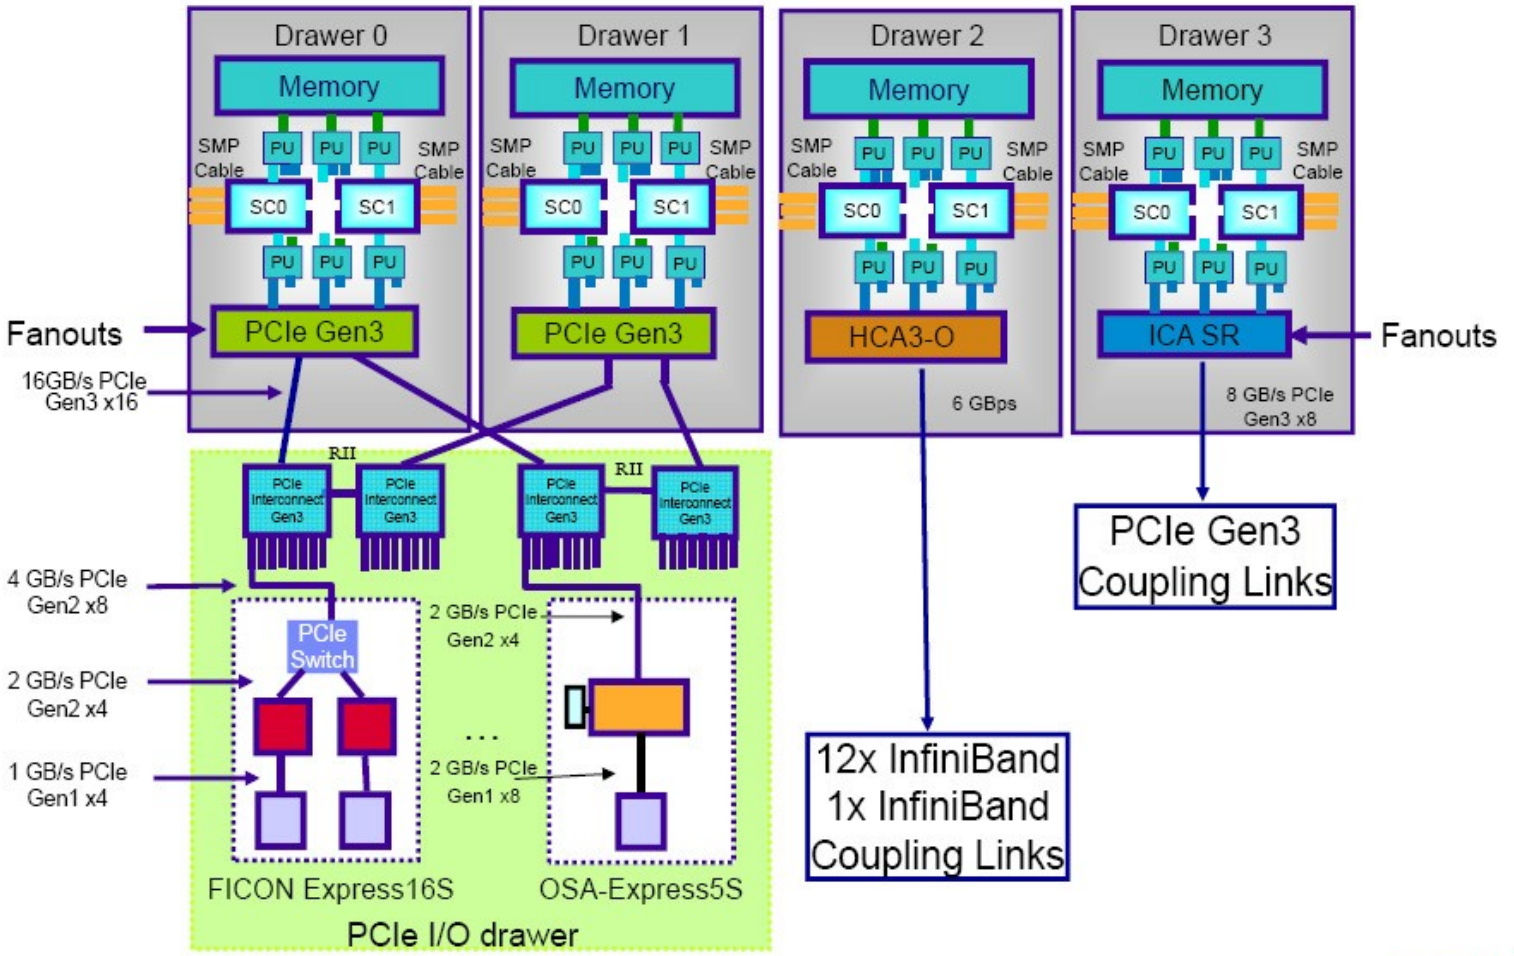

z13 IBM Multi-Factor Authentication

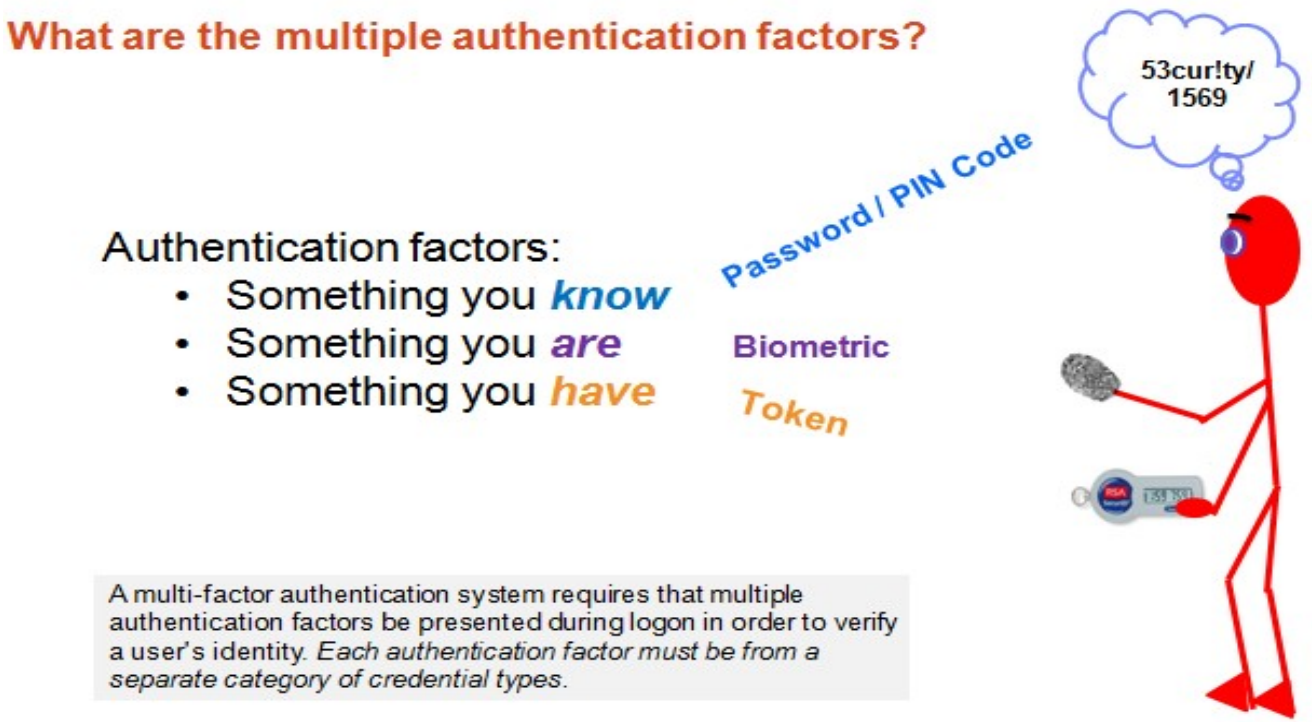

For those of you looking to increase your authentication levels, this presentation introduces you to IBM Multi-factor Authentication (MFA) to give you a basic overview of how you can enrich your authentication levels, to increase the protection of your IT infrastructure and business applications running on z/OS.

A picture of a Crypto card

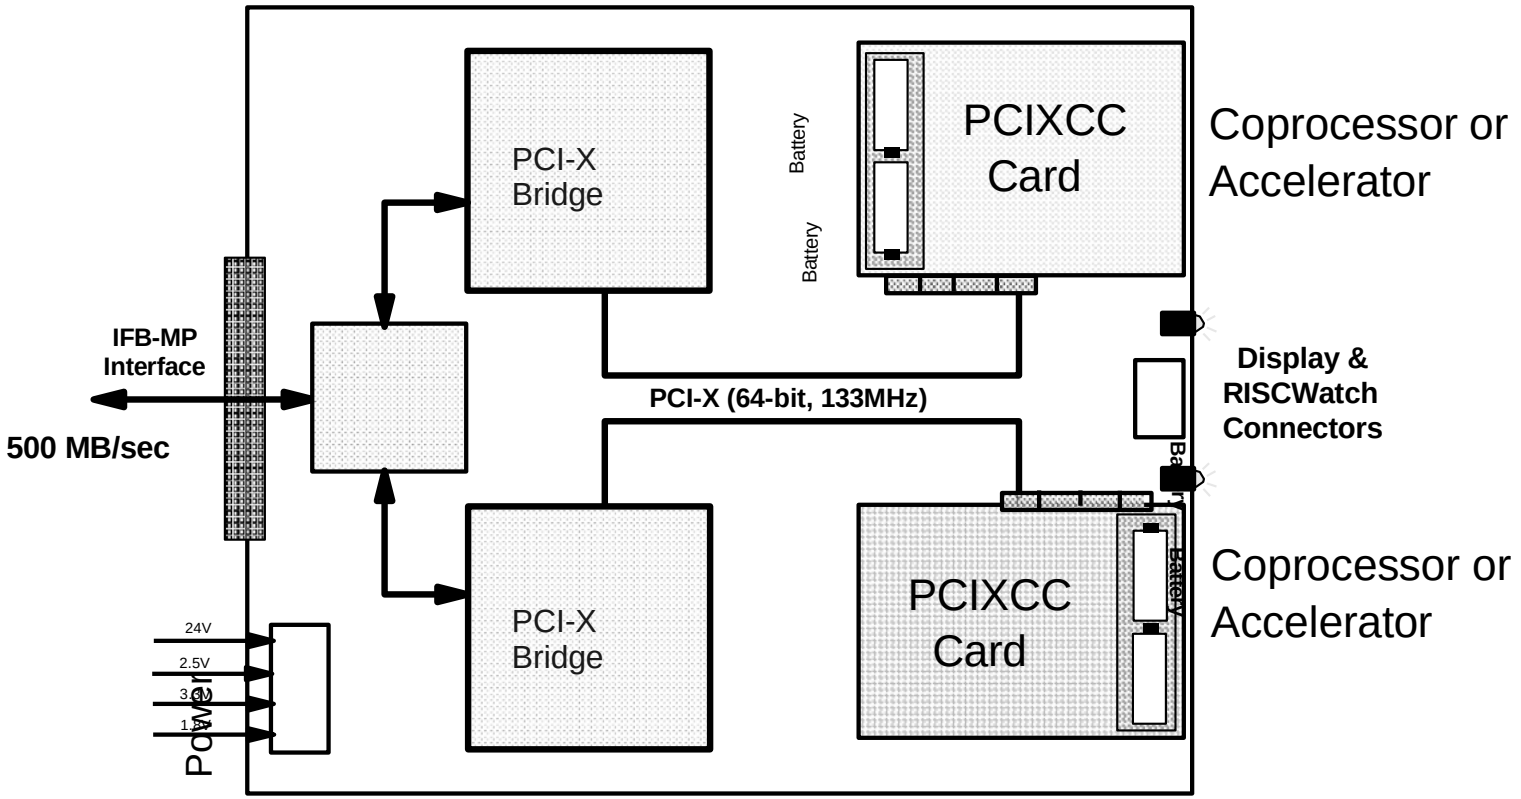

Crypto card Asynchronous Algorithms

Data Encryption Standard (DES):

Double-key DES (double DES)

Triple-key DES (triple DES)

DES key generation and distribution

•PIN generation, verification, and translation functions

Note: Keys must be provided in clear form only.

• Pseudorandom number generator (PRNG)

• Public key algorithm (PKA) facility and RSA Clear key

More on Crypto cards

- Each Crypto card feature contains two PCI-X adapters. Each adapter can be configured as a cryptographic coprocessor or accelerator. During the feature installation both PCI-X adapters are configured by default as coprocessors.
- Crypto card feature does not use CHPIDs, but each feature is assigned two PCHIDs, one per PCI-X adapter. They must have a candidate list describing the LPs that can access them.

Integrated Cryptographic Service Facility (ICSF)

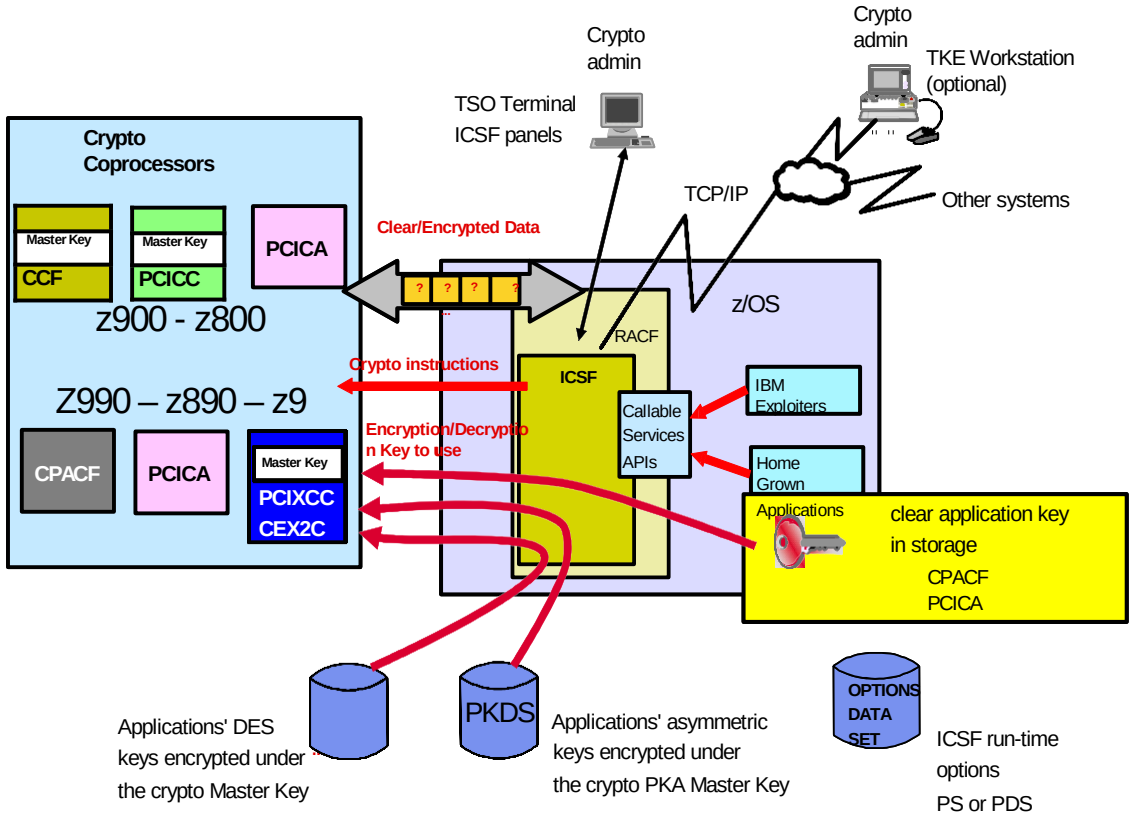

Notes:

ICSF is the z/OS component in charge of crypto. It has all the APIs to be used for the routines requiring crypto. Based in the type of the request, ICSF chooses one hardware or assist feature to execute it (synchronous CPACF or assynchronous Crypto Express2 adapter - accelerator or coprocessor).

Modernly ICSF manages crypto keys allowing for example, tape subsystem encryption. IBM announced crypto capability to DASD but in the controller.

Postprocessor Crypto RMF

C R Y P T O H A R D W A R E A C T I V I T Y

PAGE 6

z/OS V1R8 SYSTEM ID SYS1 DATE 11/28/2006 INTERVAL 14.59.946
RPT VERSION V1R8 RMF TIME 16.30.00 CYCLE 1.000 SECONDS

-------- CRYPTOGRAPHIC COPROCESSOR --------

		TOTAL			KEY-GEN
TYPE	ID	RATE	EXEC TIME	UTIL%	RATE
PCIXCC	0	0.00	0.0	0.0	0.00
	1	0.01	3205	32.1	0.01
	2	83.04	1.1	8.8	0
	3	0.00	0.0	0.0	0.00
CEX2C	4	210.8	4.4	93.3	1.91
	5	186.4	4.8	89.6	1.85

---------- CRYPTOGRAPHIC ACCELERATOR ----------

		TOTAL			ME(1024)			ME(2048)			CRT(1024)			CRT(2048)		
TYPE	ID	RATE	EXEC TIME	UTIL%	RATE	EXEC TIME	UTIL%	RATE	EXEC TIME	UTIL%	RATE	EXEC TIME	UTIL%	RATE	EXEC TIME	UTIL%
PCICA	6	165.2	1.3	21.5	107.1	1.1	11.8	0	0	0	58.1	1.7	9.7	0	0	0
	7	892.3	3.6	64.3	350.1	4.1	28.6	0.00	0.0	0.0	512.6	2.4	24.7	29.65	18.5	11.0
	8	684.8	3.5	47.8	260.4	4.0	21.0	0.00	0.0	0.0	402.4	2.3	18.6	22.02	18.5	8.1

------- ICSF SERVICES -------

	DES ENCRYPTION		DES DECRYPTION		MAC		HASH		PIN	
	SINGLE	TRIPLE	SINGLE	TRIPLE	GENERATE	VERIFY	SHA-1	SHA-256	TRANSLATE	VERIFY
RATE	4975K	497.5	12438	1244K	12438	4975K	497.5	0.0	1244K	1244
SIZE	0.75	100K	10.00	0.01	10.00	0.01	10000	0		

Explaining the previous report

This report provides performance measurements on selected ICSF activities (at ICSF Services) running on CPACF or Crypto feature. It shows the rate and the size of the data.

–Using Data Encryption Standard (DES), which is probably best known and scrutinized encryption algorithm, to encipher and decipher data.

–Generating and verifying message authentication codes (MAC). The MAC is a value calculated from the message according to a secret shared DES key and sent to the receiver together with the message. Receiver can recalculate MAC and compare it with the MAC received. If MAC values are identical, the message has not been altered during transmission.

–Using public hash functions. A hash is calculated from the transmission data according to a public key or function in cases where it is impossible to share a secret key. If the recalculated hash is identical to the one calculated before transmission, data integrity is ensured.

–Translating and verifying PINs.

Report pictured applies to z9 or z10 because there is CEX2C data.

Details on RMF Crypto report (II)

Sessions Criptographic Coprocessor and Accelerator Describes the Crypto feature performance such as average response time. Follows commeon such:

- **The two CEX2C are overloaded with 93.3% and 89.6% respectively, if these high figures are often presented try to shift work or buy more coprocessors. Just one of the PCIXCC are really in use. This fact needs to be investigated.**
- **There is an unbalanced activity in the outboard coprocessor PCICAs, see that PCCA 7 utilization is already above 60%. ME(1024) means operations in 1024-bit-ME format.**
- **There is not such data for CPACF, the synchronous coprocessor**

RMF Workload Activity AS/enclave States

POLICY=WLMPOL1 WORKLOAD=GRI SERVICE CLASS=GRI2 RESOURCE GROUP=NONE PERIOD=1 IMPORT=1

GOAL	EX VEL 75.0%	PERF INDX	AVG ADRSP	-USING%-- CPU	I/O	- EXEC TOTAL	DELAYS % -- I/O	CPU	---DLY%-- - UNKN	IDLE	CRYPTO%- USG	DLY	% QUIE
ACTUALS *ALL	49.4%	1.5	15.0	23.3	21.1	47.6	45.7	1.9	0.0	5.9	0.2	0.0	0.0
SYSB	65.8%	1.1	2.0	2.1	14.4	8.5	8.4	0.1	0.0	75.0	0.0	0.0	0.0
SYSG	79.5%	0.9	2.0	32.0	21.2	13.7	10.5	3.1	0.0	33.2	0.0	0.0	0.0
SYSII	52.3%	1.4	3.0	29.3	21.4	46.2	45.3	0.9	0.0	3.2	0.3	0.0	0.0
SYSJ	41.7%	1.8	2.0	40.7	12.2	56.4	54.4	2.0	0.0	3.3	0.3	0.0	0.0
SYSK	54.2%	1.4	3.0	23.0	31.2	45.7	43.3	2.4	0.0	0.1	0.3	0.0	0.0
SYSL	45.3%	1.7	3.0	18.5	26.8	54.7	52.6	2.1	0.0	0.0	0.2	0.0	0.0

Postprocessor Workload Activity (crypto)

All WLM Using and Delay counters shown in this report should add to roughly 100%. Any deviation is because of the possibility of overlap between Using and Delay states.

To help an installation monitor crypto performance, WLM started to sample the Using and Delay states for Crypto, while still considering these states as Unknown because WLM cannot do anything to help. WLM delays apply only to resources that WLM may act upon.

The RMF Workload Activity report, shows crypto data coming from WLM sampling of ASs/enclaves in relation to crypto hardware. These samples are consolidated in a SC period basis and have the following meaning:

- **Crypto% USG - a TCB or SRB was found be using a crypto asynchronous coprocessors (Crypto Express2)**
- **Crypto% DLY - a TCB or SRB was found to be in the queue for the asynchronous coprocessors (Crypto Express2)**

Then, the Delay (UNKN) field includes some using time, as for using crypto asynchronously.

In the report, we see that the DUs of such SC period are mostly in an unknown state, 98.1% (field DLY% UNKN), where 49.9 was using asynchronously the crypto coprocessors. As a general recommendation the delay per crypto should not be greater than 10 to 20%.

Improving Crypto Performance (ICSF Parameters)

- **Crypto and SAF requests:**
  - KEYAUTH(NO) - doubles number of crypto operations when using a key label. Up to 30% boost with (NO) option
  - CHECKAUTH(NO) - no SAF calls for authorized programs

Compartilhando CKDS antes do HCR7730

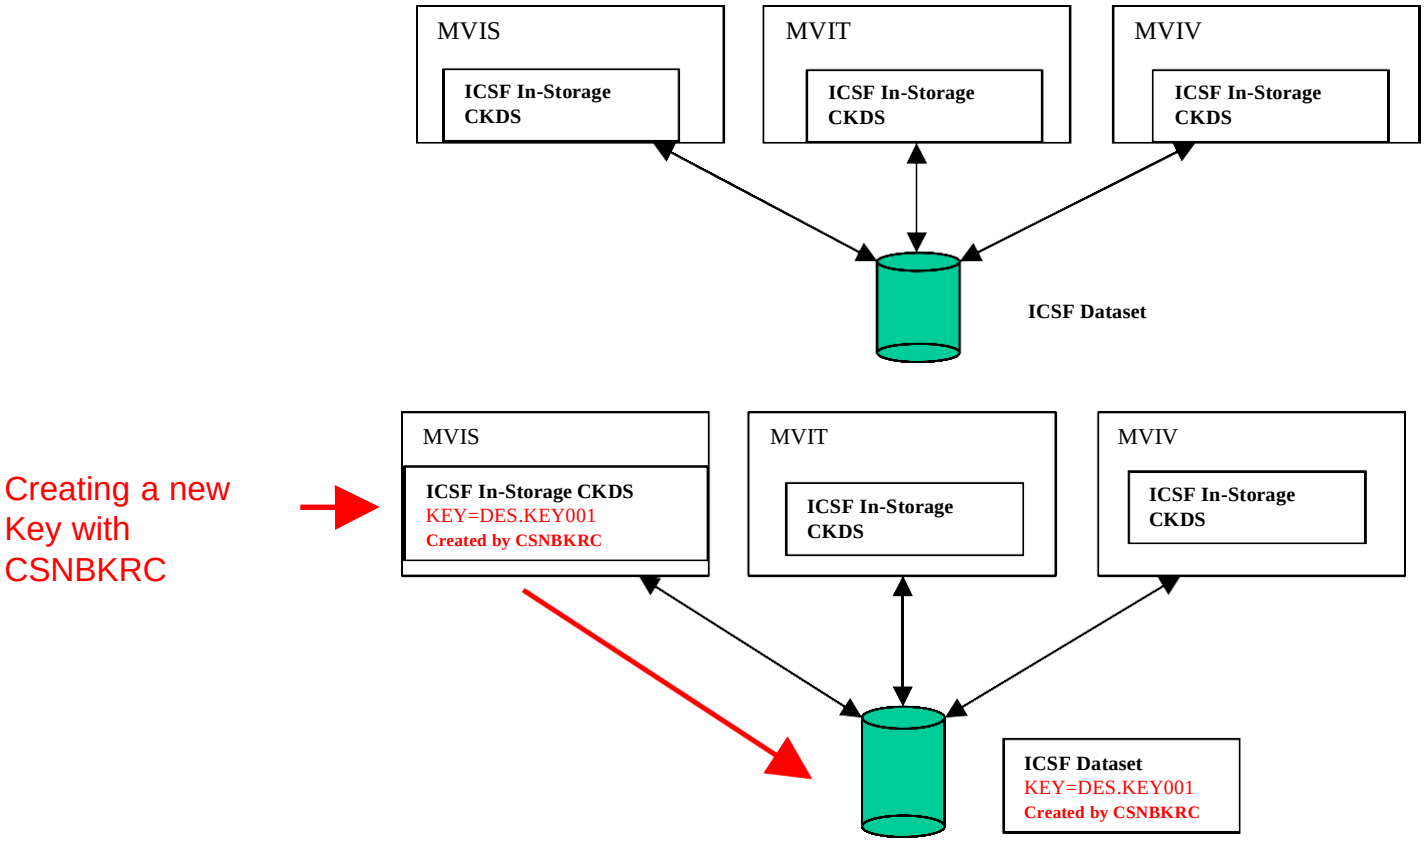

CKDS Sysplex wide cache coherence

- SYSPLEXCKDS(YES,FAIL(xxx))
  - A sysplex broadcast message informing sysplex members of the CKDS update and requesting them to update their in-storage CKDS copy
- SYSPLEXCKDS(NO,FAIL(xxx))
  - No sysplex broadcast of the update
- SYSPLEXCKDS(YES,FAIL(YES/NO))
  - ICSF initialization will end abnormally if the request to join the ICSF sysplex group fails during ICSF initialization. FAIL only applies when the first is YES. FAIL(NO) indicates that the join proceeds but the Sysplex sharing protocol is not used.

Crypto Domains

Each cryptographic coprocessor has 16 physical sets of registers or queue registers, each set belonging to a domain, as follows:

A cryptographic domain index, from 0 to 15, is allocated to a logical partition via the definition of the partition in its image profile; the same domain must also be allocated to the ICSF instance running in the logical partition via the Options Data Set.

Each ICSF instance accesses only the Master Keys or queue registers corresponding to the domain number specified in the logical partition image profile at the Support Element and in its Options Data Set. Each ICSF instance is seeing a logical crypto coprocessor consisting of the physical cryptographic engine and the unique set of registers (the domain) allocated to this logical partition.

Defining Crypto Resources at LPAR Profile

- Usage domain index
- Control Domain Index
- PCI Cryptographic Coprocessor Candidate List
- PCI Cryptographic Coprocessor Online List

This is accomplished through the Customize/Delete Activation Profile task.After this operation, any changes to the image profile require a DEACTIVATE and ACTIVATE of the logical partition for the change take effect, so the cryptographic definition is disruptive to a running system. Howevre, through the use of Change LPAR Cryptographic Controls task the changes are not disruptive.

Customize Image Profile Crypto

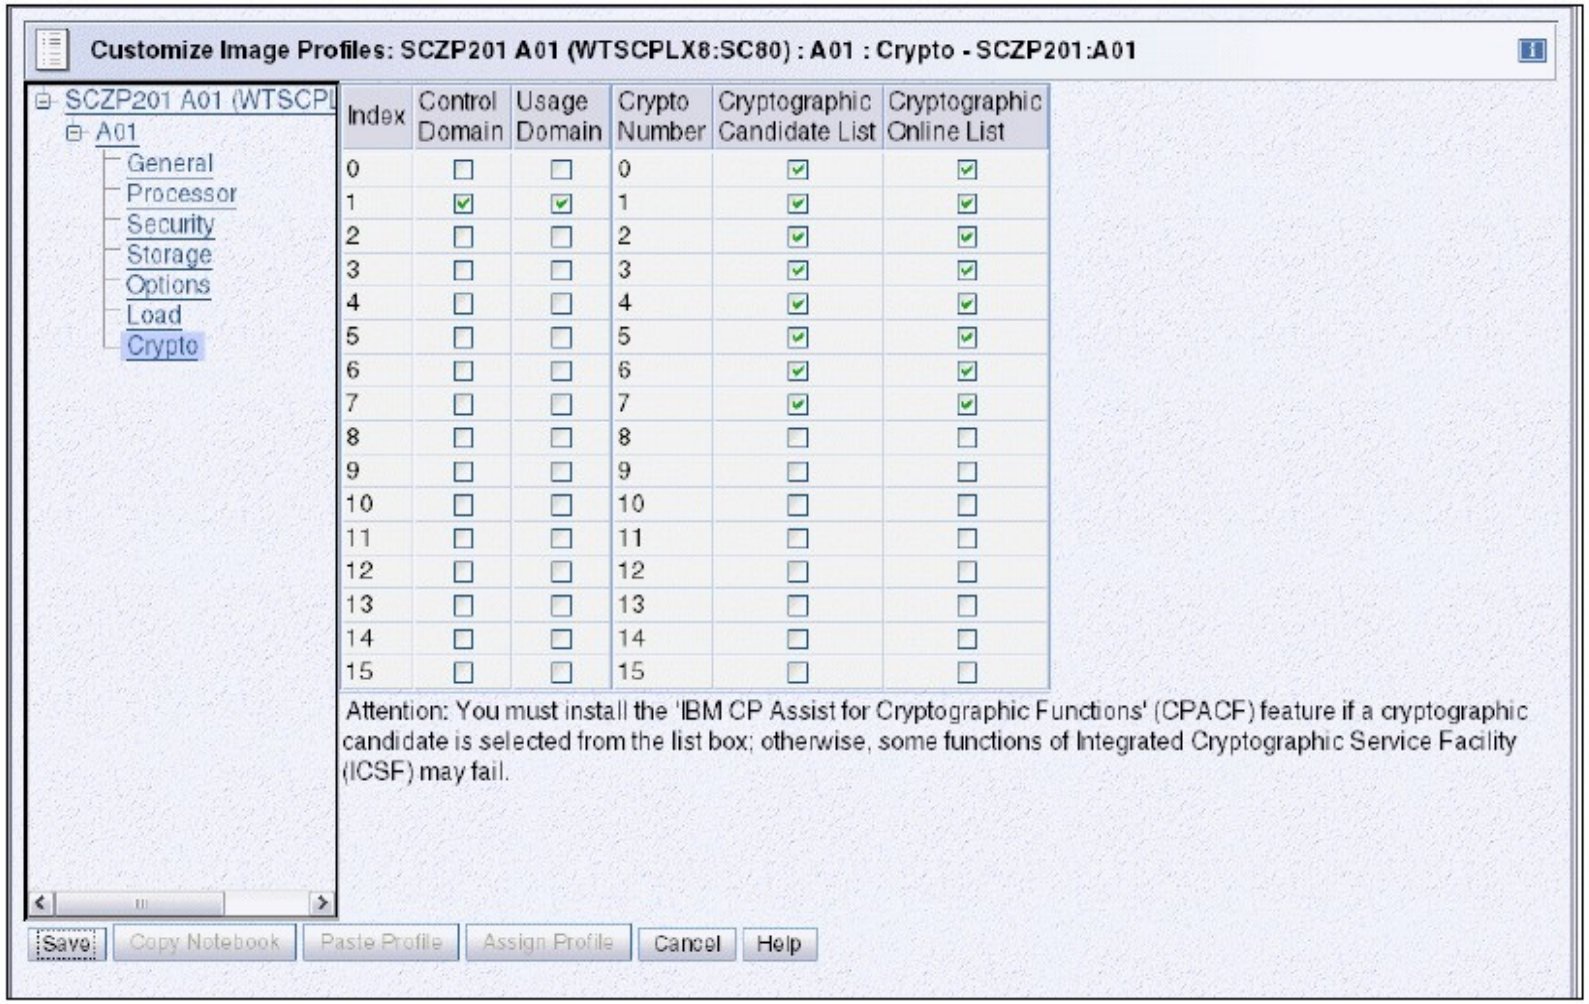

Domain Indexes

- Usage domain index

Identifies cypto coprocessor domains assigned to LP for all crypto coprocessors that are configured on the LP. The numbers selected should match domain numbers entered in Options data set when starting this zOS instance of ICSF. The same usage domain index can be used by multiple LPs regardless of which CSS they are defined to, but the combination PCI-X adapter number and usage domain index number must be unique across all active LPs.

- Control domain index

Identifies the crypto coprocessor domain indexes that can be administered from this LP if it is being set up as the TCP/IP host for the TKE. This index must include the usage domain index specified for the LP.

PCI Cryptographic Lists

- PCI Cryptographic Coprocessor Candidate list

Identifies the crypto coprocessor numbers that are eligible to be accessed by this LP. Select the coprocessor numbers, from 0 to 15, that identify the PCI-X adapters to be accessed by this LP.

You can also add the crypto coprocessor candidate list after activating the LP non disruptively using the Change LPAR Cryptographic Controls task from the SE.

If the crypto coprocessor number and usage domain index combination for the coprocessor selected in the LP Online list is already in use by another active LP, activation of the LP controls task fails;

- PCI Cryptographic Coprocessor Online list

Identifies the crypto coprocessor numbers that are automatically brought online during LP activation. The numbers selected in the Online list must also be part of the Candidate list.

Add Crypto feature to a logical partition

You can preplan the addition of Crypto features to a logical partition on the Crypto page in the image profile by defining the Cryptographic Candidate List, Cryptographic Online List and usage, and Control Domain Indices in advance of installation. By using the Change LPAR Cryptographic Controls task, adding Crypto dynamically to a logical partition without an outage of the logical partition is possible. Also, dynamic deletion or moving of these features

no longer requires preplanning. Support is provided in z/OS, z/VM for guests, z/VSE, and Linux on System z.

Attention: Cryptographic coprocessors are not tied to partition numbers or MIF IDs. They are set up with AP numbers and domain indices. These are assigned to a partition profile of a given name. The customer assigns these lanes to the partitions and continues to have the responsibility to clear them out when their users change.

Although PCI-X cryptographic adapters have no CHPID type and are not identified as external channels, all logical partitions in all channel subsystems have access to the adapter (up to 16 logical partitions per adapter). Having access to the adapter requires setup in the image profile for the partition. The adapter must be in the candidate list. For details about setting up the image profile, see IBM System z10 Enterprise Class Configuration Setup, SG24-7571.

Cloud and Hipers

Most IT Hypers

Since the early dawn times of IT, the hypers were introduced. Usually ar old ideas with new appearance and new wording...Hypers keep consultants, like me, alive, although some of them are for real... Currently, we have four of them:

- Cloud (or Clouding) Big Data
- Analytics
- Cognitivity
- Internet of Things (IOT)
- Three Dimension Printers (tele transport)

The first four are suported by Mainframe hardware and software, mainly Clouds.

Initial chat

Hi guys, Before we start , I want to share some comments and feelings with you about cloud in the Mainframe platform:

- I thought that cloud was just another fashionable buzzword (like many others)that wold vanish somewhere in the future. Maybe, I am wrong....
- Friends told me that Cloud is just another name for the tools that already existed in the Mainframe for many moons. Maybe they are partially right ....
- Foes said that the Mainframe (IBM included) is totaly out of the Cloud model of computing. They are totaly wrong.....
- Right now to me, Cloud was designed by someone unfamiliar with the Mainframe, in order to solve many problems of the distributed platform.

Beacuse of this in a sense Cloud looks like a Mainframe. Unhapilly in the act of to solving such problems, other were introduced, with no solution at this very moment. I am totaly right....

Cloud as perceived by a Mainframer

Cloud is maybe the collision of these old concepts:

- Outsourcing
- Virtualizaton
- Thin Client
- SMS

Cloud Definition

Then let us compare the Clouds features with the ones already in the Mainframe,
For such, we need to find a reliable Cloud definition (not that easy).

- Before we go on, please give your own idea about what is a Cloud.
- I found one that looks serious, from the National Institute of Standards and Technology (NIST):

“Cloud computing is a model for enabling ubiquitous, convinient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, service, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This Cloud model is composed of five essential characteristics, three service models, and four deployment models.” (uaohhhhh!)

Cloud Essential Characteristics (I)

As the huge majority of definitions, this one instead os defining “what is” the entity being defined, itjust describes its properties (ubiquitous, convenient, on-demand and so on). Then, let us start by comparing the Cloud essential characteristics with the Mainframe counterparts:

On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automaticall) without requiring human interaction witheach service provider. I a m n o t s u r e . . . . h o l d y o u r h o r s e s

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations). Yes, we do!

Cloud Essential Characteristics (II)

Resource pooling. The provider's computing resources are pooled to serve multipie consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that customer generally has no control or knowledge over the exact location of provided resources but may be able to specify location ata higher levei of abstraction (country, state, or datacenter). Resources: storage, processing, memory, and network bandwidth. Yes, we do! And we invented such, we use to calling it Virtualization

Cloud Essential Characteristics (III)

- Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time. Yes, we can do! And we invent such. Examples: WLM policy, WLM Capacity provisioning, capping, CUoD, flexible Memory, etc.
- Measured service.Cloud sistems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled and reported, providing transparency for both the provider and the consumer of the utilized service. Yes, we do! We measure optionally every resource and every user.

Unsolved Cloud Problems

- How about Capacity planning? The resources may look infinite, but they are not.
- How about the maintanaceand use of pre-production environment?
- Security? Who does that? If anyone is doing....
- How about Crypto Clear keys? Government agencies do not allow a logical partition dealing with crypto keys, share memory with other logical partitions.
- How about Buisness Continuance? Maybe using the Cloud is a disaster site?

z/OS Cloud

By Cloud on z/OS, we meas what can a MF customer do in order to improve his current MF workload processing using the Cloud model of computing in their MFs.
Thinking of that, we have two types of z/OS Clouds:

- Provisioning Cloud that is the implementation of the On Demand Self Service applied to z/OS Transaction Managers, such as, CICS, IMS, WAS, MQ....Then, it is a service model SaaS, that is called Middlewareas a service (MWAaS).
- Storage Cloud

- PS: For the moment Cloud os z/OS is not tackling the problem of a company using the MF to be a Cloud provider for other companies.

z/OS Provisioning Cloud

In order to understand the goodies of z/OS Cloud Provisioning, let us see how a new application is deployd nowadays.

- Firstly the infrastructure people should be invoked in order to create the Middleware environment in pre-existent CICS (for example), with DB2, RACF profiles and so on. Usually this is a long time consuming task.

- Secondly, you are increasing the load on the CICS Middleware tasks and address spaces and by the Law of the Diminishing Returns (LDR) increasing the service time of those transactions....

z/OS Provisioning Cloud

- Provisioning Cloud is the Service Model “Software as a Service (SaaS)”, where SW is MF Middleware products, such as: CICS, MQ, IMS, WAS, DB2.
- By this reason , this service Model is called Middleware as a Service (MWaaS)
- Then MWaaWS is a collection of Middleware (CICS, DB2, IMS....) software resources automatically provisioned to your requirements (for example, for a new application). By the way , this is na old CICS function called CICS Deploymen Manager.
- By provisioning, I mean to create for example a new CICS Instance associated with the new application users (called Tenants) with:

- New libraries for load modules.
- New options.
- New address spaces (TOR, AOR, FOR, TSOR)
- New RACF rules

z/OS Provisioning Cloud

- Then, a consequence of MWaaS is the increase in the CICS address spaces population.
- A Tenant is a collection of one or more users who collectively own MW software resources. It could be a department, a function (payroll, credit) or a phase (test, development, pre-production).
- The new CICS instance is created automatically by actions and commands discribed in Cloud structures called workloads through the use of z/OSMF.
- These workloads are activated by on Demand Self Service clicking:

  –Via a sample portal
  –Via published APIs
  –Via bluemix integration, openstack, RYO
  –Communication via email when appropriate
  –Hou about z/OSMF?

Cloud at z/OS 2.3 announcement

As the API economy develops, customers are incented to move their IT operations from a cost center model to a revenue-generating profit center model. The z/OS platform, known for its outstanding vertical scalability and speed, coupled with leading-edge security and reliability, provides the foundational capabilities that are ideal for private cloud service delivery:

- A z/OSMF enhancement is planned to support workflow extensions for IBM Cloud Provisioning and Management for z/OS.
- z/OS V2.3 is planned to deliver Real-Time SMF Analytics infrastructure support, which will enable faster processing for high-volume SMF data, providing the response time required for real-time analysis of SMF data in analytics and cloud application.
- Enabling the z/OS platform with these cloud capabilities will deliver innovations not only in certain infrastructure elements and components of the z/OS, but also in selected levels of various z/OS software subsystems such as CICS(R) Transaction Server for z/OS, IBM IMSTM for z/OS, IBM DB2(R) for z/OS, IBM MQ for z/OS, and IBM WebSphere(R) Application Server for z/OS.

API economy

Today's leading enterprises are transforming digitally, jumping head first into the API economy. Transformations are being driven by connected devices and consumers' thirst for compelling brand experiences—all generating a vast and ever-growing amount of data.

In the API economy, application programming interfaces (APIs) act as the digital glue that links services, applications and systems. This allows businesses to make the most of their data to create compelling customer experiences and open new revenue channels.

More on such announcement

Companies need to support the IBM United States Software Announcement 217-085 IBM is a registered trademark of International Business Machines Corporation 3 growth of current business workloads while also supporting emerging cognitive and cloud workloads. They need to keep costs under control and manage their skill sets.

The expectations for IT are high:

- Grow existing workloads while adding new ones, all delivered with improved service at reduced cost.
- Analyze data at the point of need for improved business insight. Deliver services through new delivery models, such as a private or hybrid cloud.
- Simplify management and reduce demands on skills.
- Manage risk, security, and compliance to meet evolving regulatory requirements.

Summary

Today's economy requires organizations to quickly consume, manipulate, and deliver vast amounts of information, extracting business insight while tapping into the capabilities of cloud services. The information must be securely managed, processed, and delivered across the globe. Such a fundamental shift away from traditional processing needs calls for a highly responsive and reliable platform that can support new workloads without impacting service levels of mission-critical work.

The enhancements planned to be delivered in z/OS V2.3 provide the platform to smoothly transition businesses for this new IT environment. z/OS V2.3 is designed to deliver innovations to build the next-generation infrastructure.

Big Data Definition

Big data is a term for data sets that are so large or complex that traditional data processing applications are inadequate". Major reasons are IOT and smart phones.

What is OpenStack?

OpenStack is a set of software tools for building and managing cloud computing platforms for public and private clouds. Backed by some of the biggest companies in software development and hosting, as well as thousands of individual community members, many think that OpenStack is the future of cloud computing. OpenStack is managed by the OpenStack Foundation, a non-profit that oversees both development and community-building around the project.

OpenStack lets users deploy virtual machines and other instances that handle different tasks for managing a cloud environment on the fly. It makes horizontal scaling easy, which means that tasks that benefit from running concurrently can easily serve more or fewer users on the fly by just spinning up more instances. For example, a mobile application that needs to communicate with a remote server might be able to divide the work of communicating with each user across many different instances, all communicating with one another but scaling quickly and easily as the application gains more users.

More on OpenStack

•And most importantly, OpenStack is open source software, which means that anyone who chooses to can access the source code, make any changes or modifications they need, and freely share these changes back out to the community at large. It also means that OpenStack has the benefit of thousands of developers all over the world working in tandem to develop the strongest, most robust, and most secure product that they can.

SMF Analytics Application

Just one customer generates 2.5 TB per day. On top of data due to CICSPlex the SMF data must be consolidated by system and by crossing CICS. Then, the basic question is:

How to extract good information from SMF data in order to improve my system?

To help customers the following functions were introduced:

- **zEDC for SMF data**
- **Security SMF records signing at z/OS 2.2**
- **SMF Streaming support providing real time access to SMF data**
- **Implementation of z/OS Apache Spark platform**