

z/VM Dynamic Memory Management

Walter Church

z/VM Development Lab: Endicott, NY
wchurch@us.ibm.com

VM Workshop

29 June 2019

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BladeCenter*	FICON*	OMEGAMON*	RACF*	System z9*	zSecure
DB2*	GDPS*	Performance Toolkit for VM	Storwize*	System z10*	z/VM*
DS6000*	HiperSockets	Power*	System Storage*	Tivoli*	z Systems*
DS8000*	HyperSwap	PowerVM	System x*	zEnterprise*	
ECKD	IBM z13*	PR/SM	System z*	z/OS*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs, and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

First

What is Dynamic
Memory Management,
and why do we care?

Then

Planning for Dynamic
Memory Management:
requirements, tools, and
helpful hints.

Followed by

The Nitty-Gritty:
Commands and
Configuration
Statements

And finally

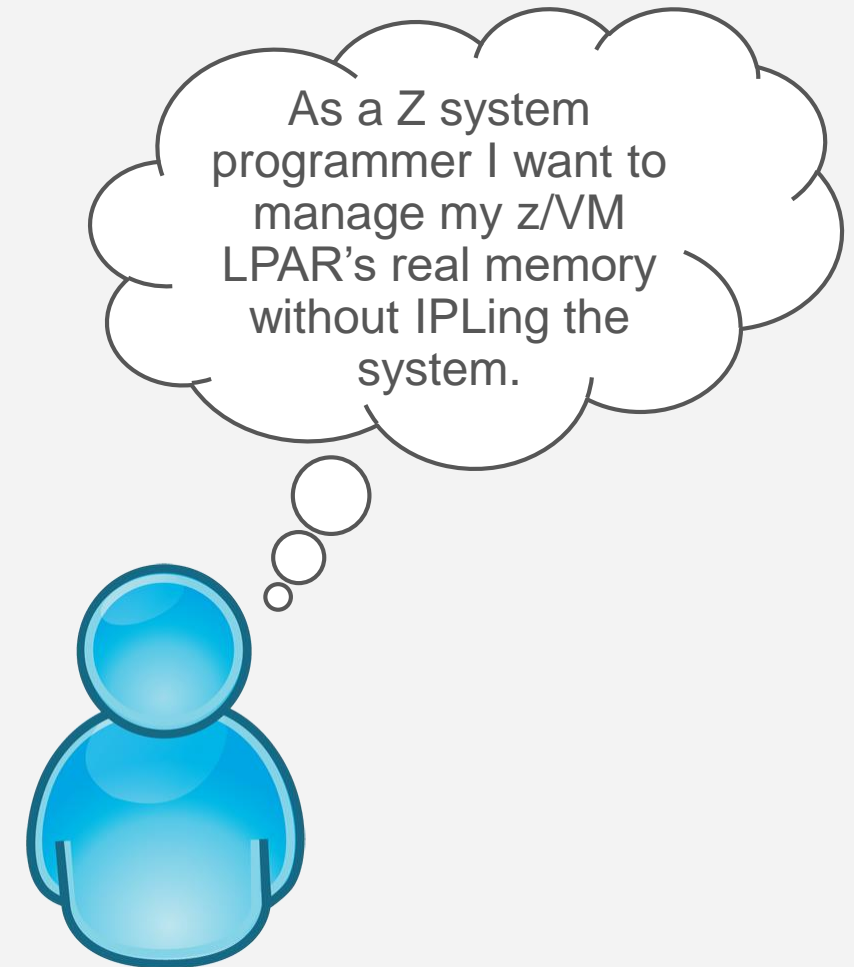
Paging implications,
interactions with other
commands, & conclusion

Problem statement

The **flexibility to reassign** (add and remove) system resources is critical to customers. Today's **workloads are not static**. Having to tolerate a re-IPL to modify the memory configuration is burdensome to customers and contrary to the goal of **continuous operations**.

With Memory Reclamation, a system administrator can **take real memory offline** from a z/VM partition, **making it available** to other partitions on the CPC. The removal will be **dynamic**; no re-IPL of the z/VM image is required to accomplish the change in the memory configuration.

This session will describe these new capabilities, give some guidance on use, and walk through some examples.




[This Photo](#) by Unknown Author is licensed under [CC BY](#)

Definition of Terms

Storage is also called memory on IBM Z hardware.

Storage must be added or removed in multiples of the **storage increment size**. This is determined by the hardware, and mimicked by z/VM second level.

Memory reclamation is the process of removing some storage from your z/VM partition's configured storage and putting it back into standby storage.



z/VM V5.4 Enhancements Announced on August 5, 2008

Dynamic memory upgrade support allows real memory to be added to a running z/VM system. With z/VM V5.4, memory can be added nondisruptively to individual guests that support the dynamic memory reconfiguration architecture. Systems can now be configured to reduce the need to re-IPL z/VM.

Dynamic memory upgrade

System z servers provide a dynamic reconfiguration capability to allow additional main storage available for use in a logical partition (LPAR) while the system is running. The configuration of storage for an LPAR includes an initial amount of main storage and an additional reserved amount, defined in fixed-size increments. The reserved storage can be assigned and accessed when the operating system decides to use it, using the dynamic storage-reconfiguration function.

- The new z/VM dynamic memory upgrade support allows additional main storage to be added to a running z/VM system and also virtualizes the use of this storage. Storage can be added to a running z/VM system without a z/VM system shutdown, LPAR deactivation and reactivation, and z/VM system IPL.

- Another LPAR can be activated when the reserved storage is currently assigned to any active LPAR
- Standby storage is installed

z/VM V5.4 allows a z/VM system to increase the size of its main storage dynamically by bringing online designated amounts of standby storage. In addition, z/VM guests that support dynamic storage reconfiguration can increase and decrease their amounts of main storage dynamically.

z/VM Continuous Delivery News

DISCLAIMER!

<http://www.vm.ibm.com/newfunction/>

Dynamic Memory Downgrade	
Name	Dynamic Memory Downgrade (DMD)
Description	Dynamic Memory Downgrade, AKA Memory Reclamation, would extend the real memory dynamic management characteristics of z/VM to include removing real memory from a running z/VM system. Currently z/VM allows adding memory, but not removing it.
Status	Currently in system test.
Target availability	TBD
Compatibility	Several commands associated with dynamic memory management will have incompatible changes.
Enablement	Applying PTFs will enable code, use of new function will be from commands and configuration file settings. Requires an IPL of the z/VM system to activate the code for this PTF.
Effect	Most mainline paths are unaffected by these changes, but code paths for initialization and adding memory have significant changes.
ISV impact	Some monitor changes are needed. If you have a performance product, please see your vendor.
Linux or hardware interaction	Will require z14 or higher. No known Linux interaction at this time.
Release(s)	z/VM 7.1 SPE
Service details	See below for the IBM service information.
APAR	VM66173
PTF	TBD
RSU	TBD
Sign up	Contact Kerry Wilson - kerryw@us.ibm.com to become a Sponsor User for this SPE.

z/VM 7.1

SET STORAGE command changes

- New PERMANENT keyword
- Remove AS keyword
- No more rounding up to the increment boundary

Increase minimum storage size for a second-level z/VM system

32M → 128M

**z/VM 7.1 available
September 21, 2018**

Dynamic Memory Downgrade

- Improvements to real memory management
- Enhancements to paging threshold settings
- New & improved commands and statements
- Compelling use cases

What?

APAR VM66173

When?

TBD

Dynamic Memory Downgrade, AKA Memory Reclamation, will extend the real memory dynamic management characteristics of z/VM to include removing real memory from a running z/VM system. Currently z/VM allows adding memory, but not removing it.

PR/SM's view of your storage

Via the Activation Profile

Central Storage

Amount in: ▾

Initial:

Reserved:

z/VM's view of your storage

```
query store
```

```
16:20:24 STORAGE = 8G CONFIGURED = 8G INC = 128M STANDBY = 2G RESERVED = 0
```

```
16:20:24 Permanent = 4G Reconfigurable = 4G Maximum STORAGE = 10G
```

```
Ready;
```

```
query store
```

```
16:20:24 STORAGE = 8G CONFIGURED = 8G INC = 128M STANDBY = 2G RESERVED = 0
```

```
16:20:24 Permanent = 4G Reconfigurable = 4G Maximum STORAGE = 10G
```

```
Ready;
```

What are those new fields in the **QUERY STORAGE** response?

Permanent

The amount of real storage that cannot be decreased, only increased. It contains important CP control structures and long term locked pages.

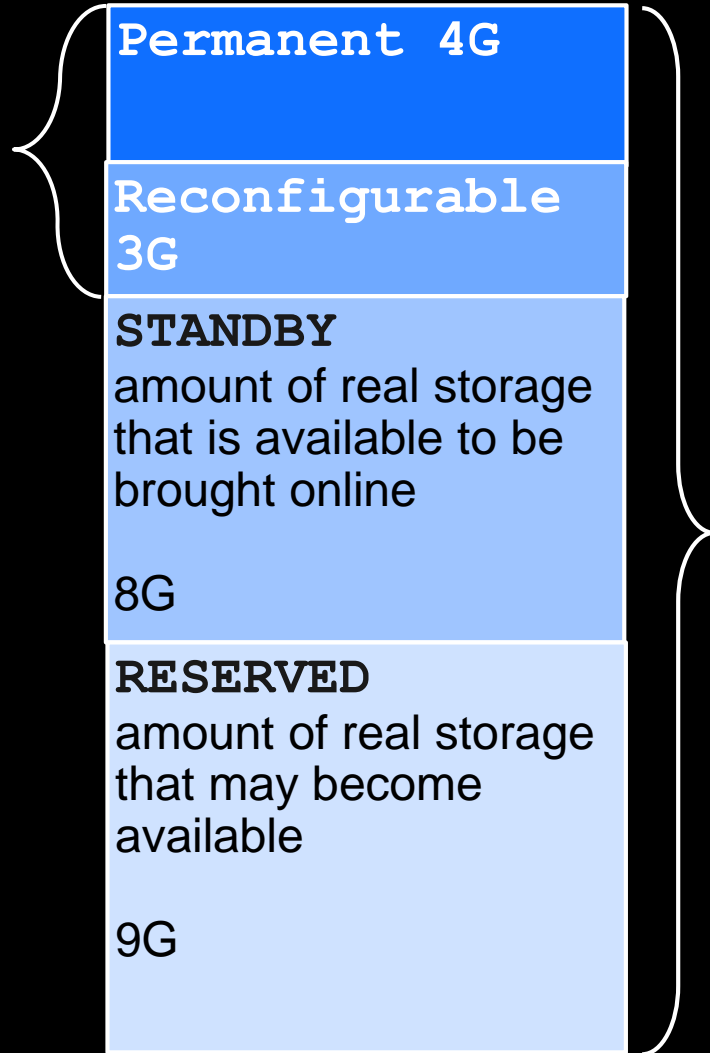
Reconfigurable

The amount of real storage that can be decreased or increased.

Maximum STORAGE

The largest amount of storage that can be brought online to z/VM. Maximum = Initial + Reserved in your LPAR definition, up to 2T.


STORAGE = 7G



Maximum STORAGE = 24G

SET STOR RECONF +1G

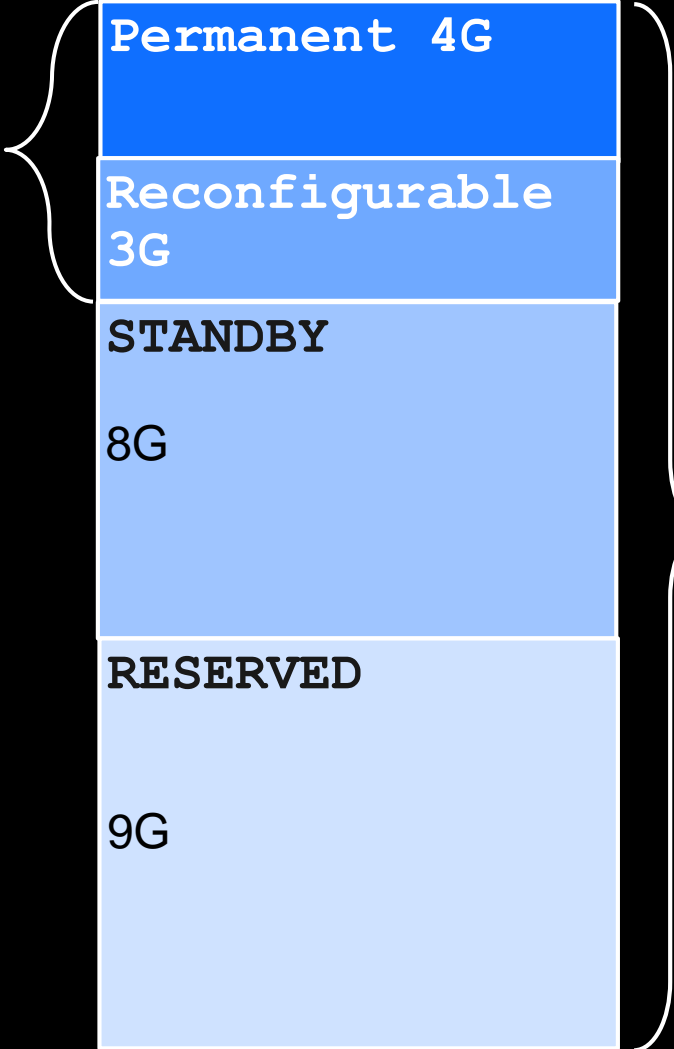
STORAGE = 8G



Maximum STORAGE = 24G

SET STOR RECONF -1G


STORAGE = 7G



Maximum STORAGE = 24G

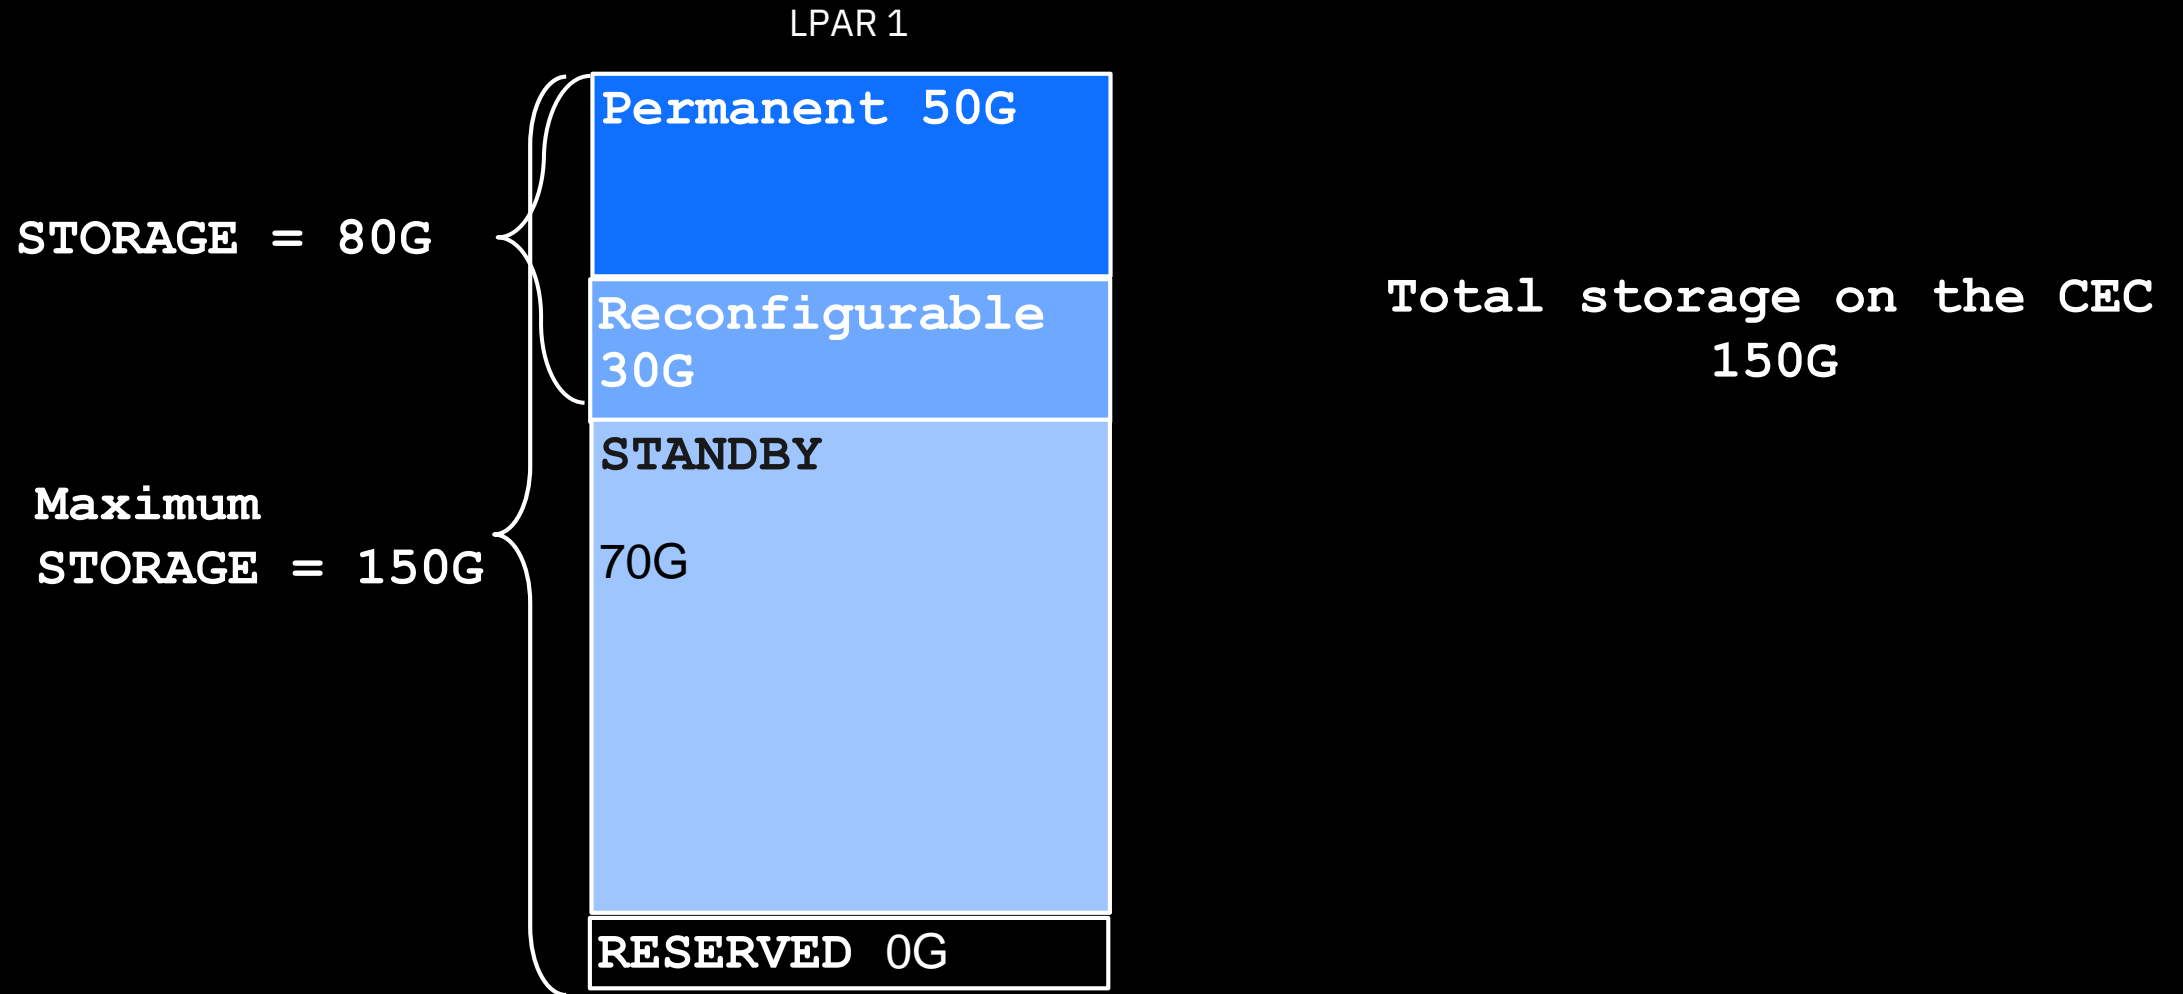
SET STOR PERM +1G

STORAGE = 8G



Maximum STORAGE = 24G

One LPAR is lonely, but has lots of **STANDBY** memory

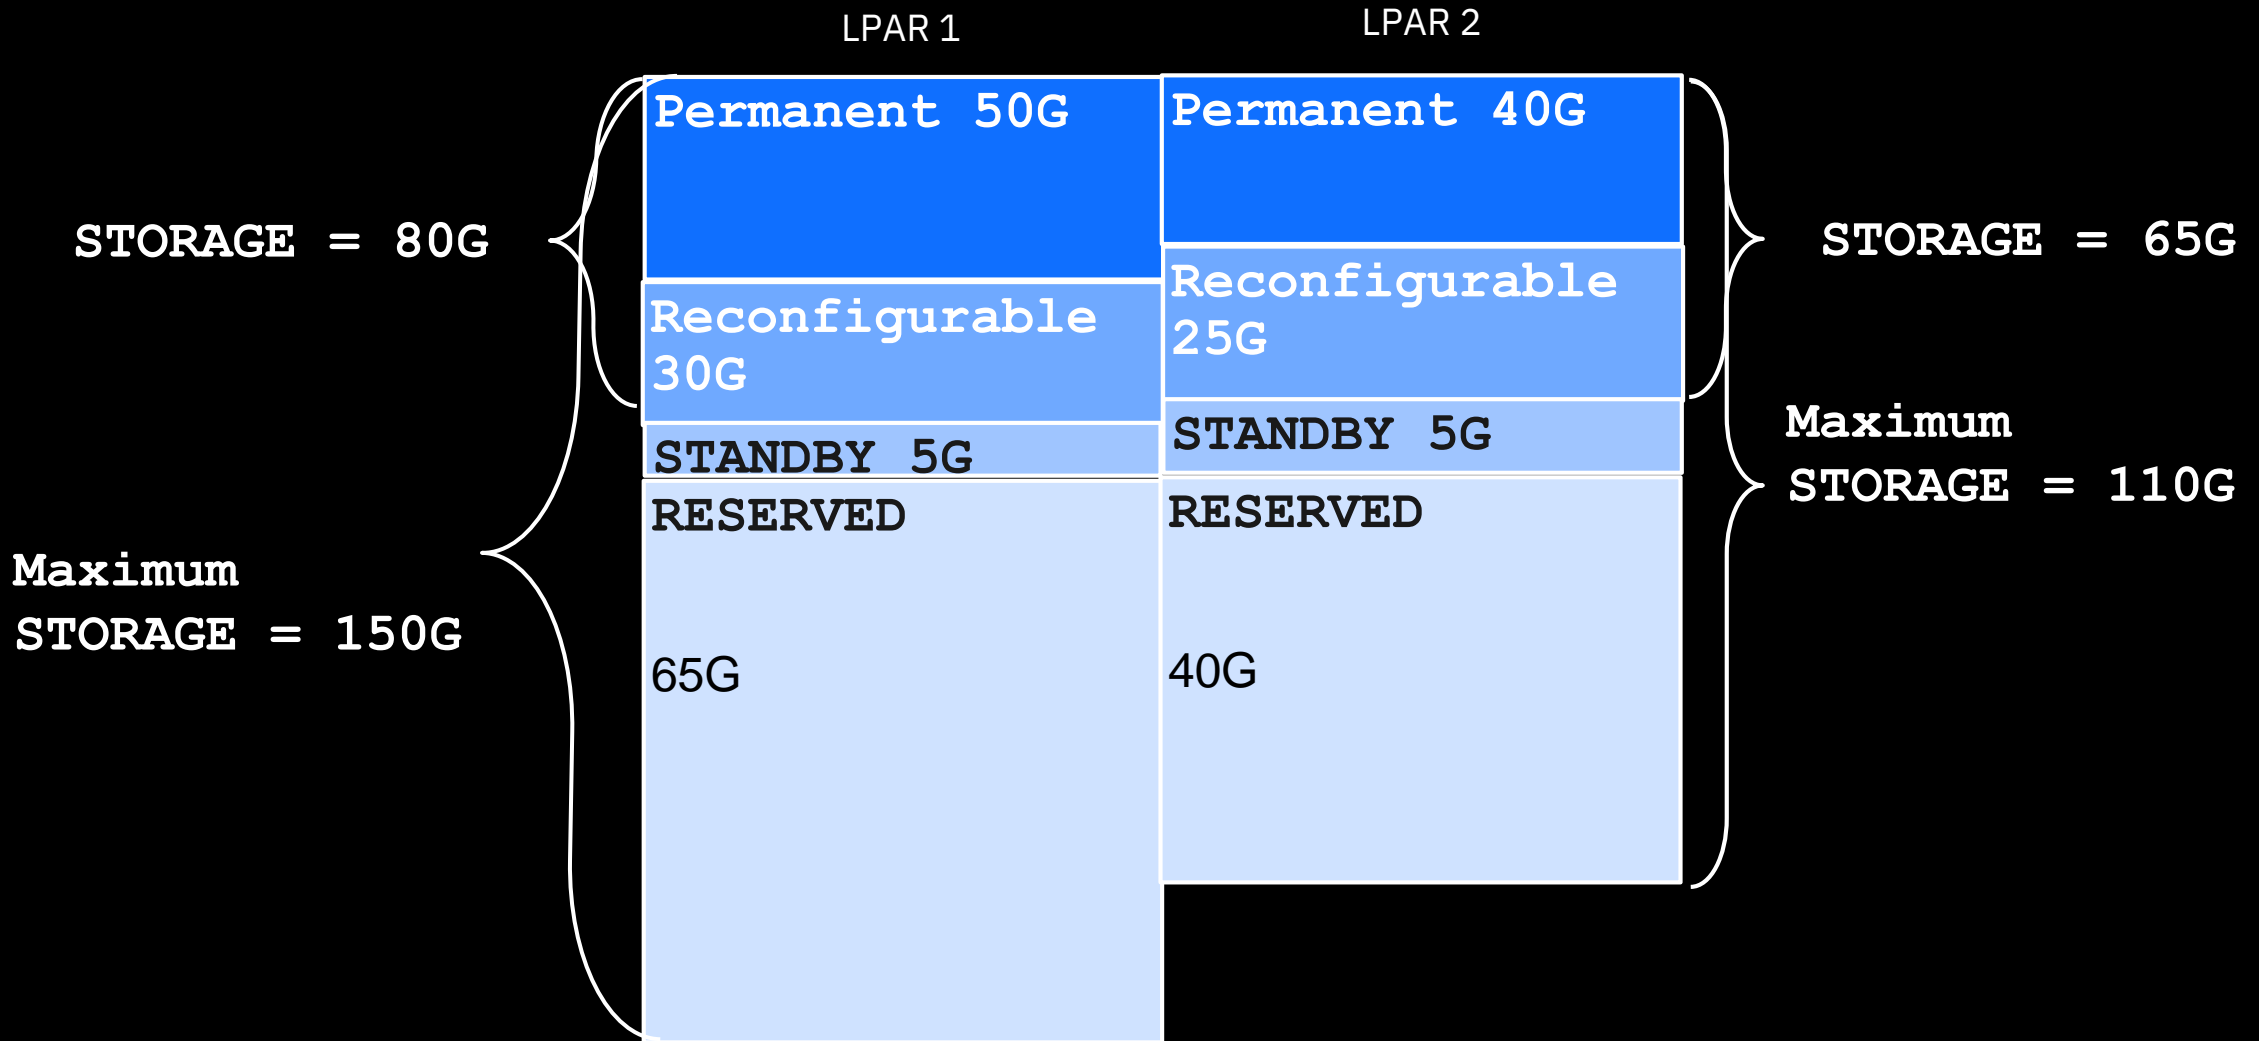


A new partition is activated! I gain a friend but lose some **STANDBY** memory.



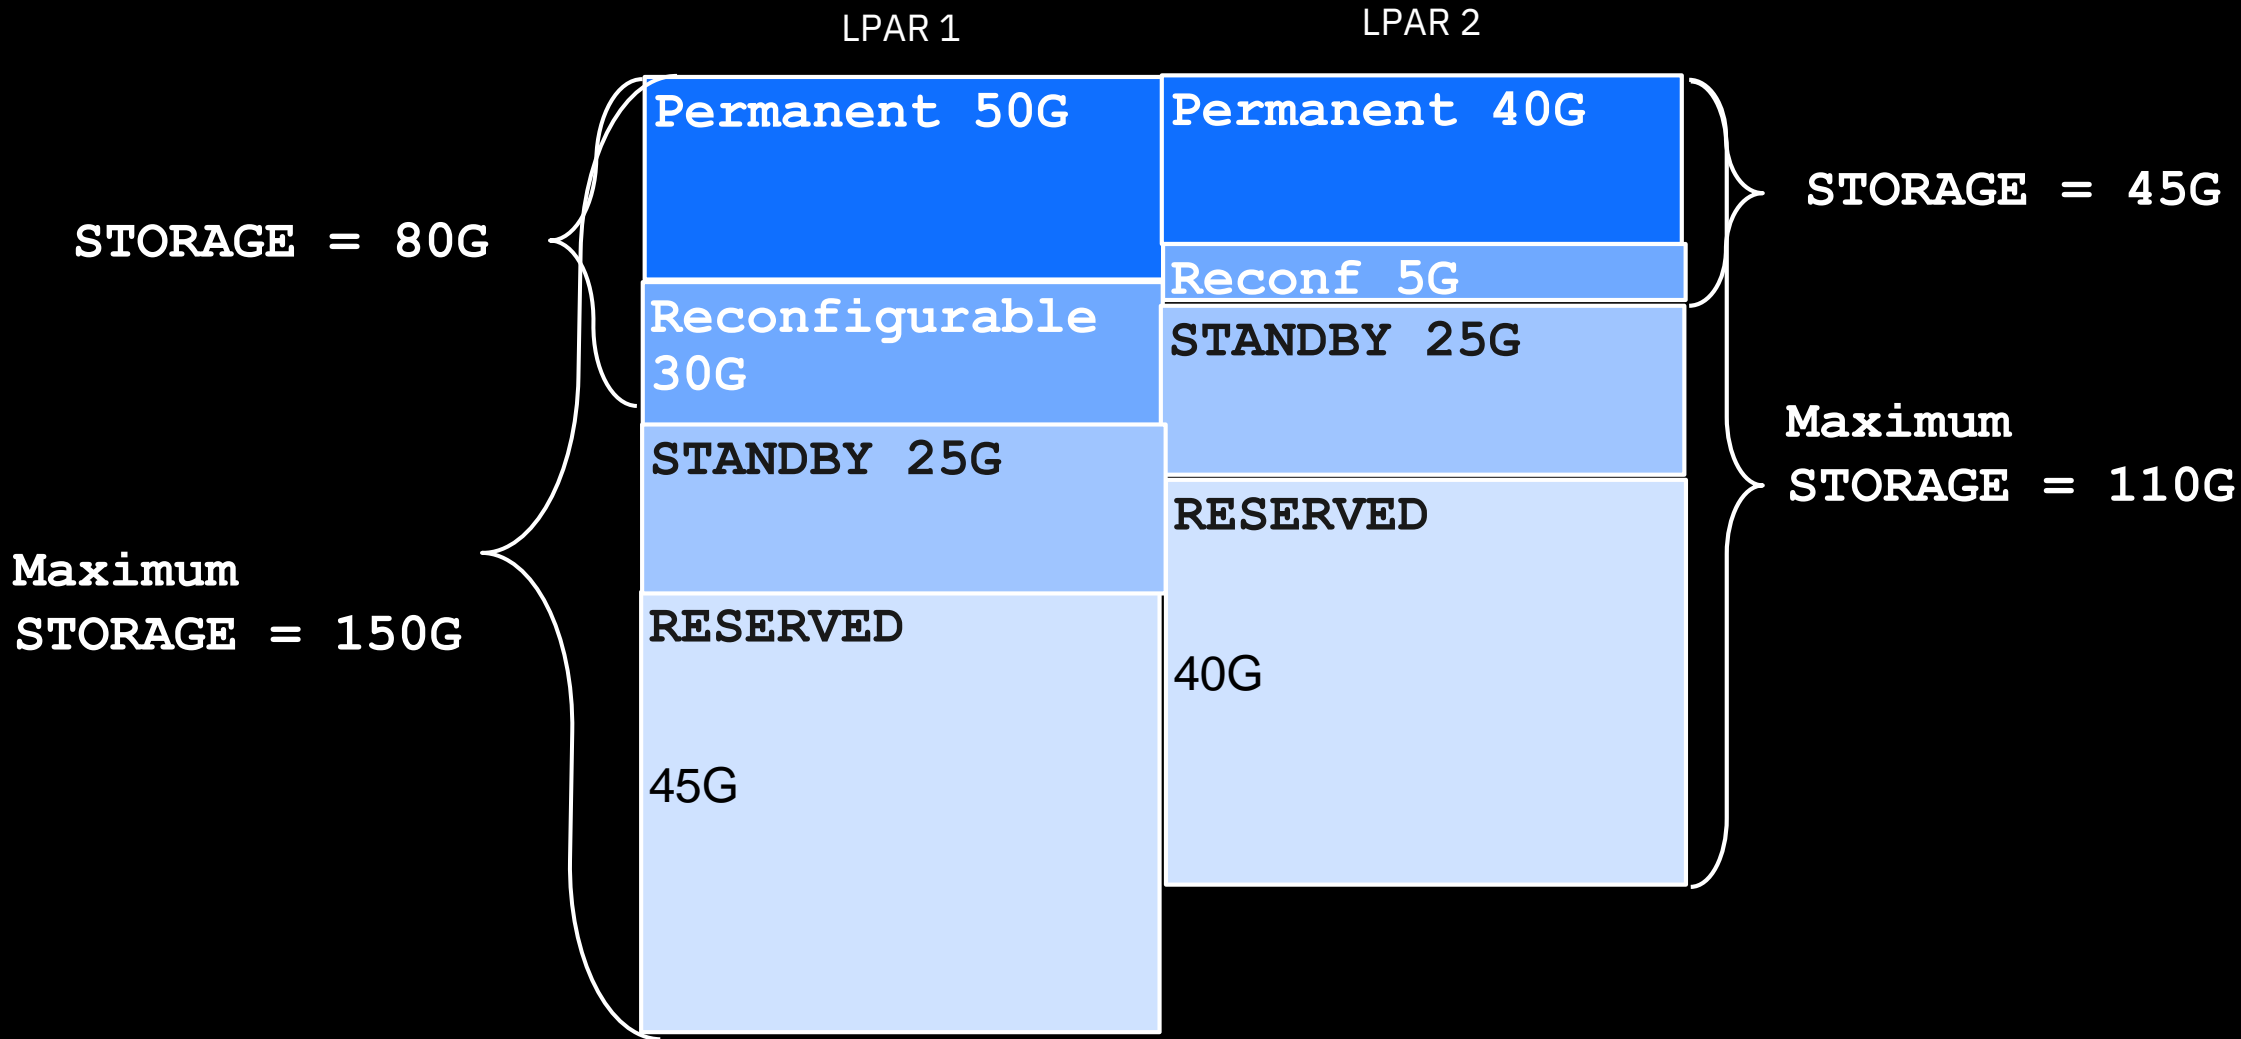
Total storage on the CEC = 150G = 80G + 60G + 10G (STANDBY)

Partition 2 adds 5G of memory, I now have less STANDBY




Total storage on the CEC = 150G = 80G + 65G + 5G (STANDBY)

Partition 2 subtracts 20G of memory, I now have more STANDBY



Total storage on the CEC = 150G = 80G + 45G + 25G (STANDBY)

The other LPAR is deactivated. More **STANDBY** for me!



Planning for Dynamic Memory Management on z/VM

Requirements and Restrictions

- ❑ z/VM LPAR on z14, Emperor II, or Rockhopper II is necessary for first level real storage reclamations.
 - Dependency on z14 firmware enhancements in QDIO and HPMA2
- ❑ Other improvements are available on all supported hardware.
 - New **STORAGE** configuration statement to preserve or reset the storage configuration across IPLs
 - Ability to set a paging warning threshold customized to your system
- ❑ New restrictions included in the z/VM 7.1 base release (regardless of hardware)
 - Minimum second level z/VM system storage is now 128M (old minimum was 32M).
 - Storage additions (and reclamations) must be done in multiples of the storage increment size.

When to start planning?



- Before the system is IPLed!
- You should plan how much memory to have in reconfigurable vs permanent storage.

Can't I just make all my storage reconfigurable?



[This Photo](#) by Unknown Author is licensed under [CC BY](#)

- No!
- At least 4G of permanent storage is required before any reconfigurable storage can be added.
- Permanent storage is necessary for CP control structures and pages that might be locked long term.
 - We wouldn't want to remove the storage that contains the table that describes storage!
 - We wouldn't want a reclamation to get stuck on a locked page while a **SPXTAPE** operation is ongoing!
- 4G of permanent is an enforced *minimum* but for larger systems more permanent storage is recommended.

How much permanent storage should I have?

Think about the normal workload on your LPAR

- How much storage is that?
- Do you expect any overcommit?

That storage should be permanent

Permanent storage can be added to later, but it cannot be removed dynamically!

If possible, use **VIR2REAL EXEC** or **MONITOR** data to look at your system workload before you add the variable workload

- Storage originally defined for the LPAR
- The virtual to real ratio for your core workload
- Storage instantiated on a typical day

Add a little extra to your permanent storage to cover the parts of your variable workload that must be in permanent.

- CP control blocks
- I/O and MONITOR pages

How much reconfigurable storage should I have?

The following situations lend themselves to reconfigurable storage

- Work that happens during special events or at certain periods of time
- Guests that do not always run in the LPAR
- Applications that are growing and might need its own LPAR in the future

These workloads are perfect for reconfigurable, you can have storage when you need it and give storage back to another LPAR when you don't

Think about the amount of storage you would use for that workload and the new virtual to real ratio for your system after the new workload and storage are added.

VIR2REAL EXEC

- Gathers information about the system and the users currently running on it
- Differentiates between users running CMS (by looking at which NSSes or devices they IPL) and those not running CMS
- Looks at both their total virtual storage and instantiated storage
 - Shows what the system looks like at this moment in time
 - Total virtual shows you a highest utilization case
- Compares these numbers to the real storage available to the system
- Gives information on the storage available to the system as well as the increments in which it may be increased
- Looks at the paging space available, current utilization and what the paging space utilization would be **if the guests used all their available virtual storage**

VIR2REAL EXEC

```
Storage information for VM system PTCVMD01
CMS users IPL NSSes "CMS GCS" or devices "0190 0490".

Total Virtual storage (only ids not running CMS): 205312 MB (200.5 GB)
Total Virtual storage (only ids running CMS):      1104 MB ( 1.1 GB)
Total Virtual storage (all logged on userids):    206416 MB (201.6 GB)
Usable real storage (pageable) for this system:  81127 MB ( 79.2 GB)
Total LPAR Real storage:                          81920 MB ( 80.0 GB)

Total Virtual disk (VDISK) space defined:        43256 MB ( 42.2 GB)
Average Virtual disk size:                       158 MB

Virtual to (usable) Real storage ratio:          2.5 : 1
Virtual + VDISK to Real storage ratio:           3.1 : 1
Virtual to Real ratio (non CMS work only):        2.5 : 1

Paging warning setting: 90%
Amount of permanent storage available to system: 256G
Amount of reconfigurable storage available to system: 20G
Maximum amount of storage that can be brought online to this system: 2T
Storage increment size: 4G
```

VIR2REAL EXEC

Percent of paging space needed for these virtual storage totals:

Virtual storage for all logged on guests: 113% 9

Virtual storage for non-CMS guests only: 112% 10

Virtual+VDISK for all logged on guests: 136% 11

Paging: 78 volumes active, usable space is: 183072 MB (178.8 GB)

Total Paging space in use, 25% utilization: 46188 MB (45.1 GB)

Note: AGELIST KEEPSLOT YES is in effect. Therefore, more paging may not result in more page space usage.

CHKRECLM EXEC

- Collects all the same information as **VIR2REAL** and calculates the same ratios, but then models taking away the specified amount of storage from your system and recalculates the same information
- Allows you to compare how your system would look after a reclamation
- Will warn you if a reclamation isn't possible because of lack of reconfigurable storage, or if the value specified is not on an increment boundary
- Virtual to real overcommit ratio is a guideline many customers use
- Shows a range of paging space utilization, because there could be multiple copies of a page

CHKRECLM EXEC

CHKRECLM 20G

Storage reclamation information for VM system GDLMCT1 on 2017-11-13 at 09:26:24
CMS users IPL NSSes "ZCMS CMS GCS" or devices "190 990".

Your system has:

Total Virtual storage (only ids not running CMS):	6144 MB (6.0 GB)
Total Virtual storage (only ids running CMS):	6518 MB (6.4 GB)
Total Virtual storage (all logged on userids):	12662 MB (12.4 GB)
Total of all Instantiated pages:	25591 MB (25.0 GB)
Paging usable space:	1662976 MB (1624.0 GB)
Total Paging space in use:	166297 MB (162.4 GB)
Paging warning setting:		90%

Projected real and virtual storage ratios:

	Current	Post-Reclaim
Total LPAR Real storage:	276.0 GB	256.0 GB
Virtual to (usable) Real storage ratio:	0.0 : 1	0.0 : 1
Virtual to Real ratio (non CMS work only):	0.0 : 1	0.0 : 1
Total Instantiated to Real storage ratio:	0.1 : 1	0.1 : 1

CHKRECLM EXEC

Estimated paging space utilization after reclamation:

	Low	High
Percent of paging in use:	10.0 %	11.2%
Paging in use:	162.4 GB	182.6 GB
Paging usable space:	1624.0 GB	1624.0 GB

Note: AGELIST KEEPSLOT YES is in effect. Therefore, more paging may not result in more page space usage.

Consult the CP Planning and Admin book chapter on Allocating DASD space for more information.

VIR2REAL EXEC and CHKRECLM EXEC

- **VIR2REAL** is an existing EXEC available on the z /VM downloads page:
<http://www.vm.ibm.com/download/packages/descript.cgi?VIR2REAL>
 - There are planned updates to this EXEC to go along with the DMD APAR
 - Updated EXEC will be available around the time the APAR is available
- **CHKRECLM** is a new EXEC that will be made available on the downloads page after the DMD APAR is available

PerfKit can help too!

- IBM Performance Toolkit screens can also help with determining your workload's size
- FCX103, Storage Utilization Screen – STORAGE – provides information about the real storage available on the system

```
FCX103      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor

Main storage utilization:
Total real storage      2'048GB
Total available        2'048GB
Offline storage frames      0
SYSGEN storage size     2'048GB
Shared storage         11'228KB
FREE stor. subpools     5'540KB
Subpool stor. utilization 67%
Total DPA size         2'032GB
Locked pages           44217
Reserved user storage   4'155MB
Set reserved SYSMAX     0KB
Trace table            9'648KB
Pageable               2'031GB
Storage utilization     107%
Tasks waiting for a frame 0
Tasks waiting for a page 4743/s
Standby real stor. size  0KB
Reservd real stor. size  0KB

XSTORE utilization:
Total available          0KB
Att. to virt. machines  0KB
Size of CP partition    0kB
CP XSTORE utilization    ...%
Low threshold for migr.  ....kB
XSTORE allocation rate   ....s
Average age of XSTORE blks ...s
Average age at migration ...s

MDCACHE utilization:
Min. size in XSTORE     0KB
Max. size in XSTORE     0KB
Ideal size in XSTORE    0KB
Act. size in XSTORE     0KB
Bias for XSTORE         .00
Min. size in main stor.  0KB
```



SYSGEN storage size shows the amount of storage online to the z/VM system

The Nitty-Gritty: Commands and Configuration Statements

How do I define reconfigurable storage?

- To bring reconfigurable storage online you must have
 - 4G of permanent storage
 - **QUERY STORE** shows some standby storage
 - A z/VM partition on a z14, Emperor II, or Rockhopper II if you are using it 1st level
- Define it via the new operands on the **STORAGE** system configuration statement
 - **STORAGE PERManent 4G**
 - **STORAGE RECONFigurable 4G**
- To determine the storage configuration at IPL
 - **QUERY STORE IPL**
 - Also great for building automation based on how your storage is configured
 - Operator's console for messages
- Re-define it dynamically using the **SET STORAGE** command
- IPL parameter **STORE=** is still supported, but only defines permanent storage

STORAGE System Configuration Statements

- Let's say your LPAR activation profile has **INITIAL** = 8G and **RESERVED** = 2G

Central Storage


Amount in: ▼

Initial:

Reserved:


- You can manipulate the storage available to your LPAR with new **STORAGE** statements
 - **STORAGE PERManent 8G**
 - **STORAGE RECONFIgurable 2G**

Dynamic Storage Reconfiguration




NOW ASYNCHRONOUS!

Add:



Remove:



Notes:

- 1 HALT is valid only when a storage reconfiguration add or the vacate part of the storage reconfiguration remove operation is in progress.
- 2 You can specify the operands in any order, but you can specify an operand only once per SET STORAGE command.

Dynamically Adding Storage

- Add PERManent and RECONFigurable storage in one command

```
SET STOR PERM +32G RECONF +32G
```

- No more AS operand

```
SET STOR AS 100G  
SET STOR PERM 100G
```

- Storage initialization at start up may take some time. A SET STORE command will be rejected if storage initialization is still active.

Dynamically Removing Storage

SET STOR RECONF -2G

HCP2593I Storage reclamation viability test passed with MAXPAGEFULL value of 90%. 50% of paging space could be required for the current workload.

HCP2581I Storage reconfiguration to remove 2G of reconfigurable initiated by OPERATOR.

HCP2582I Storage reconfiguration by OPERATOR is complete. Permanent = 40G Reconfigurable = 200G

Use TEST to determine the viability

SET STOR RECONF -200G TEST

HCP2593E Storage reclamation viability test failed with MAXPAGEFULL value of 90%. 150% of paging space could be required for the current workload.

Dynamically Removing Storage

- Use `MAXPAGEFULL` to set maximum paging percent

```
SET STOR RECONF 0 MAXPAGEFULL 95%
```

```
HCP2593E Storage reclamation viability test failed with MAXPAGEFULL value of  
95%. 150% of paging space could be required for the current workload.
```

- Use `FORCE` to skip any viability checks (use at your own risk!)

```
SET STOR RECONF 0 FORCE
```

```
HCP2581I Storage reconfiguration to remove 2G of reconfigurable initiated by  
OPERATOR.
```

```
HCP2582I Storage reconfiguration by OPERATOR is complete. Permanent = 4G  
Reconfigurable = 0
```

**Using FORCE is very risky. It can cause a PGT004
abend if the system runs out of paging space!**

Monitoring a Storage Reconfiguration

- Use new `RECONFiGuration` option of `Q STOR` to check status of an in-progress storage addition

Query `STorage RECONFiGuration`

```
STORAGE = nu CONFIGURED = nu INC = nu STANDBY = nu RESERVED = nu
```

```
Permanent = nu Reconfigurable = nu Maximum STORAGE = nu
```

```
Storage increase in progress. Elapsed time = hh:mm:ss
```

```
Target:           {Permanent = nu} {Reconfigurable = nu}
```

```
Total to add:    {Permanent = nu} {Reconfigurable = nu}
```

```
Remainder to add: {Permanent = nu} {Reconfigurable = nu}
```

Monitoring a Storage Reconfiguration

- Use new `RECONFiGuration` option of `Q STOR` to check status of an in-progress storage reclamation

Query STore RECONFiGuration

`STORAGE = nu CONFIGURED = nu INC = nu STANDBY = nu RESERVED = nu`

`Permanent = nu Reconfigurable = nu Maximum STORAGE = nu`

`Storage decrease in progress. Elapsed time = hh:mm:ss`

`Target Reconfigurable = nu`

`Total to remove = nu`

`Remainder to remove = nu`

`MAXPAGEFULL = nnn% | Forced to no limit`

Halting a Storage Reconfiguration in Progress

```
SET STOR HALT
```

- Terminate a storage reconfiguration in progress
- Any portion of the reconfiguration already completed will not be undone

```
HCP2593E Storage reclamation viability test failed...
```

- Halt by the system if viability test fails while storage reclamation is in progress

Memory Reclamation on Guest Systems

- z/VM second-level will support memory reclamation on any level of hardware
- `DEFINE STORAGE` has a new `INCRement` operand
 - Useful for testing out reclamations and understanding how increment size affects adding and removing storage
 - When your second level z/VM system does not have any `STANDBY` or `RESERVED` storage defined the increment size will be 1M

Storage and Initialization

- If you leave your system how it is today (with no changes to the SYSTEM CONFIG) your storage will come up as `Permanent`.
 - If you have `STANDBY` or `RESERVED` storage available, you may be able to add `RECONFIGURABLE` storage later.
- `STORE= IPL` parameter overrides any other storage setting
 - The amount of storage specified will be `Permanent`
- If you have less than 4G of storage, it will always be initialized as `Permanent` storage. `Reconfigurable` is only available after you have 4G of permanent.
- New system config statements change storage handling after a restart.

Reminder: Always use CPSYNTAX to check your configuration file for errors! Find problems BEFORE your system restarts!

Storage After System Restart

There are new configuration statements to help you keep or ignore your dynamic storage changes

– Specify behavior after a CP Abend or PSW RESTART

- `STORage AFTER_REStart INITialize | KEEP`

– Specify behavior after a SHUTDOWN REIPL:

- `STORage AFTER_SHUTDOWN_REIPL INITialize | KEEP`

Want to ignore any dynamic storage changes made? Use `INITialize`

– `INITIALIZE` uses the values specified on the `STORAGE` statement in the system configuration file

Want to keep storage the way it was before the restart? Use `KEEP`

– `KEEP` uses the amount of permanent and reconfigurable storage online at system termination.

– As long as the CP nucleus being IPLed has the same version and release level as that of the terminating system.

Storage After a Restart

If your storage does not match what you would expect, check `QUERY STORE IPL`

```
q stor ipl
10:53:05 STORAGE = 10G CONFIGURED = 10G INC = 128M STANDBY = 10G RESERVED = 0
10:53:05 Permanent = 4G Reconfigurable = 6G Maximum STORAGE = 20G
10:53:05 IPL actual:      Permanent = 4G Reconfigurable = 6G
10:53:05 IPL requested: Permanent = 2G Reconfigurable = 6G
10:53:05 IPL requested data source: STORAGE system configuration statement(s)
10:53:05 Last start was a system IPL
```

Keep in mind...

- Time it takes to complete a reclamation depends on size and system workload
- Try to avoid simultaneous storage reconfigurations on other partitions
- Specified storage values must be a multiple of the increment size (`INCRement`)
- Limit reconfigurable memory to amount expected to reclaim
- CP manages reconfigurable and permanent storage differently
- Ensure workload is stabilized before initiating a reclamation
- A storage reconfiguration may halt, but will not be automatically undone
- Update your system configuration file to match dynamic storage changes

Paging implications, interactions with other commands, and conclusion

How much paging is too much?

- Think about your over commitment ratio on your system and how much paging space you normally expect to use
- IBM has always warned you when you used 90% of your paging space
 - We will always warn you at this level (it's dangerous!)
 - Now you can set your own warning level as well!
- Set your own warning level
 - **SYSTEM CONFIG** statement
`PAGING WARNING nnn%`
 - **Dynamic command**
`SET PAGING WARNING nnn%`
- An informational message is issued to alert the system operator whenever the system pages over the percent specified as well as when you go over 90%

Storage reconfiguration and paging

- Removing storage can cause more paging
 - We will have less main storage to work with!
 - The act of vacating and removing storage uses some storage temporarily
- Is it okay to exceed your normal paging space utilization on a storage reclamation?
 - Maybe, if you're removing storage so you can move workload to a new system
 - Maybe not, if the storage being removed should be excess then you wouldn't want to see a large paging increase
- CP will not do a storage reclamation if it will cause paging to go above the warning threshold you've set
- Using the `SET PAGING` command to reset the warning threshold will not affect any current reclamations running.
- You can override the paging warning threshold for a particular memory reclamation
 - Using the `MAXPAGEFULL` operand on `SET STORAGE`
 - Decimal integer in the range of 0 to 100

QUERY command updates

QUERY STORE IPL

- Storage configuration specification at the last start of the system
- Could be used for new automated bring-up procedures

QUERY FRAMES

- Now shows permanent and reconfigurable frames
- Also shows vacating frames: the number of frames being taken offline by the SET STORAGE remove operation

QUERY PAGING

- Shows the new WARNING nnn%

QUERY SXSSTORE

- Frame Table Active Range
 - Shows you the range of addressable frames for the System eXecution Space, some of the frames within that range might be offline.
 - Does not show the actual size of the SXS.


Conclusion

An exciting new capability, memory reclamation, is coming soon to z/VM 7.1 systems near you.

Memory reclamation is only available for the z14 or equivalent hardware.

It introduces a new type of memory, reconfigurable, which requires some planning to use

It includes other goodies like the ability to set another paging warning threshold and keeping or tossing out dynamic storage changes on restart



Thank you!

Contact Information:

Walter Church

z/VM Storage Management Development

wchurch@us.ibm.com