

# z/VM Dynamic Memory Management

For the latest version see <https://www.vm.ibm.com/library/presentations/>

**Jeff Gertner**

z/VM Development Lab: Endicott, NY

[gertnerj@us.ibm.com](mailto:gertnerj@us.ibm.com)

© 2018, 2022 IBM Corporation



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

Db2*	FlashCopy*	IBM eserver	OMEGAMON*	XIV*	Z10 BC	zSecure
DirMaint	FlashSystem	IBM (logo)*	PR/SM	z13*	z10EC	zSeries*
DS8000*	GDPS*	IBM Z*	RACF*	z13s	z/Architecture*	z/VM*
ECKD	ibm.com	LinuxONE*	System z10*	z14	zEnterprise*	z Systems*
FICON*	IBM Cloud*	LinuxONE Emperor	System 390*	z15	zPDT	
		LinuxONE Rockhopper	WebSphere*		z/OS*	

\* Registered trademarks of IBM Corporation

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.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

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 other countries.

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.

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.

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

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.

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

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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](http://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.

# Presentation Abstract

The flexibility to reassign (add and remove) system resources is critical to customers. Today's workloads are not static. With Dynamic Memory Downgrade, 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 reIPL of the z/VM image is required to accomplish a change in the memory configuration. This session will describe these new capabilities, give some guidance on use, and walk through some examples.

First

What is Dynamic  
Memory Management,  
and what does it do?

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 increment size** is the smallest amount of storage that can be added (or removed). This is referred to by PR/SM as “storage granularity”; it’s 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.

**Dynamic Memory Downgrade (DMD)** – another term for memory reclamation.



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

# What is Dynamic Memory Management and what does it do?

# Today: z/VM 7.2 – and APARs VM66173 & VM66508

z/VM 7.2 + APAR VM66173, PTF UM35834 (4 TB real memory support)

- `SET STORAGE` command allows a new `RECONFIGURABLE` keyword
  - Added in anticipation of Dynamic Memory Downgrade (DMD) enablement
  - Any specified reconfigurable value will generate an error message
- Also included are enhancements to:
  - Real memory management
  - Paging threshold settings
  - New & improved commands and configuration file statements

z/VM 7.2 + APAR VM66508, PTF UM35859 (VMSES APAR)

- Will automatically handle the situation referenced in the 3/22/21 redalert
  - See <https://www.vm.ibm.com/service/redalert/> for details



# Dynamic Memory Downgrade

- Pre-reqs VM66173
- Reconfigurable storage can be removed from a running z/VM system
- RECONFIGURABLE storage up to 50% of total online storage

## What?

**APAR VM66271,  
PTF UM35561**

## When?

**Available now**

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

# PR/SM's view of your storage

Via the Activation Profile

Central Storage	
Amount in:	Gigabytes (GB) <input type="button" value="v"/>
Initial:	<input type="text" value="8.0"/>
Reserved:	<input type="text" value="2.0"/>

# 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**

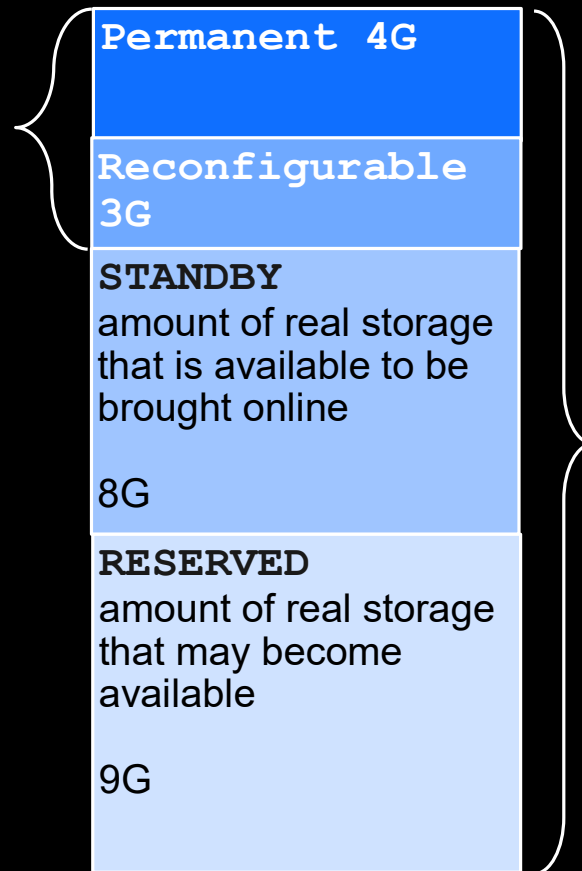
The amount of real storage that can be decreased or increased.

Total reconfigurable storage cannot exceed 50% of configured storage.

#### **Maximum STORAGE**

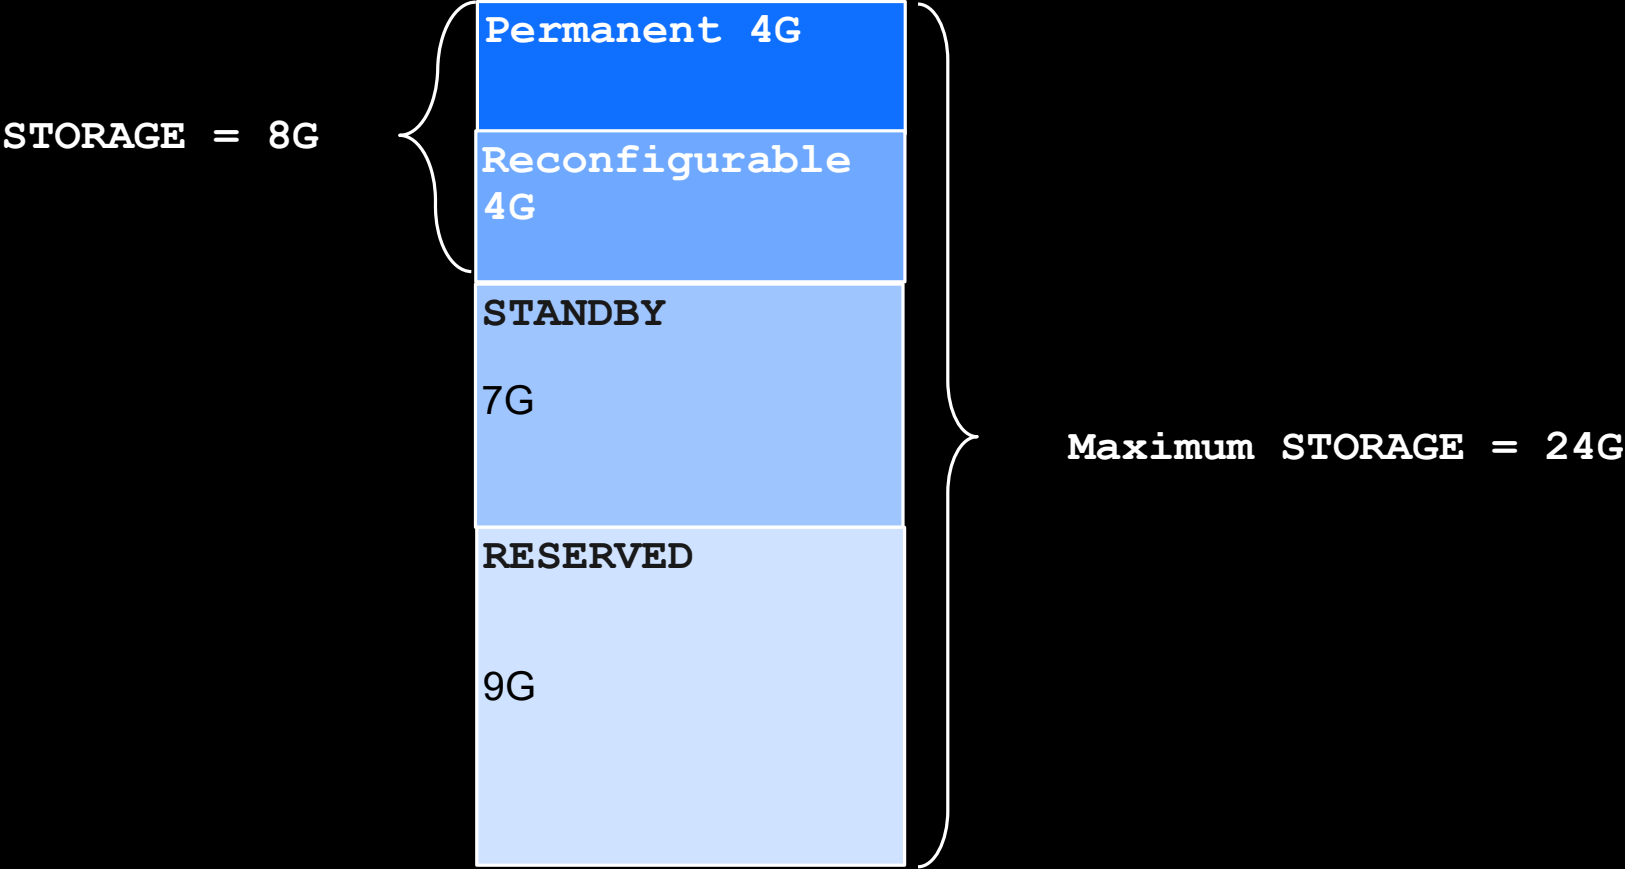
The largest amount of storage that can be brought online to z/VM. Maximum = Initial + Reserved in your LPAR definition, up to 4T.

STORAGE = 7G

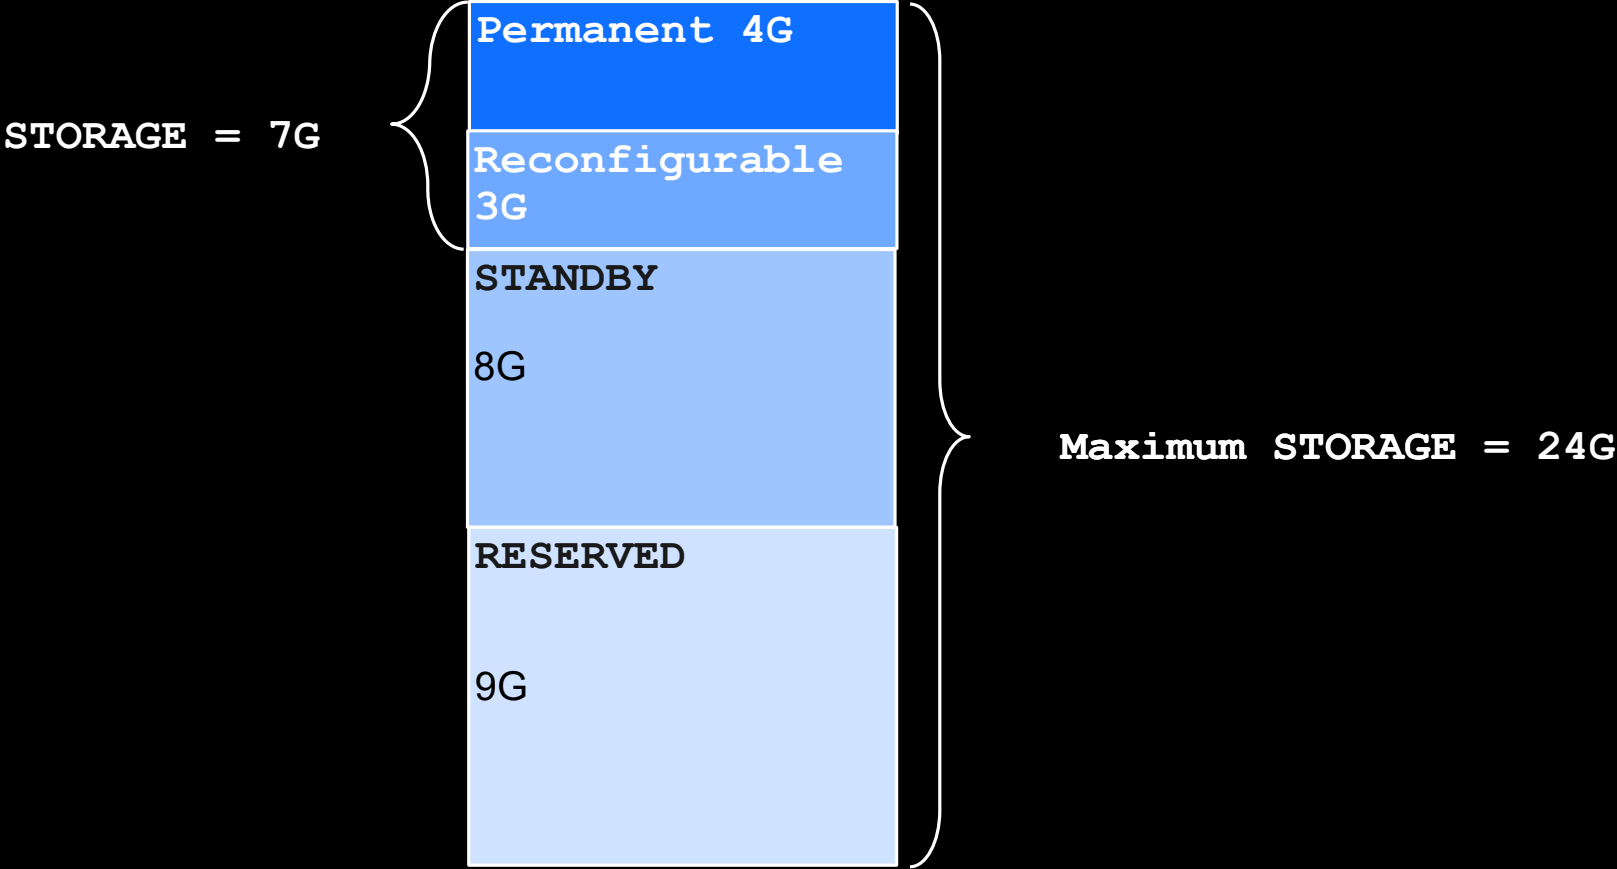


Maximum STORAGE = 24G

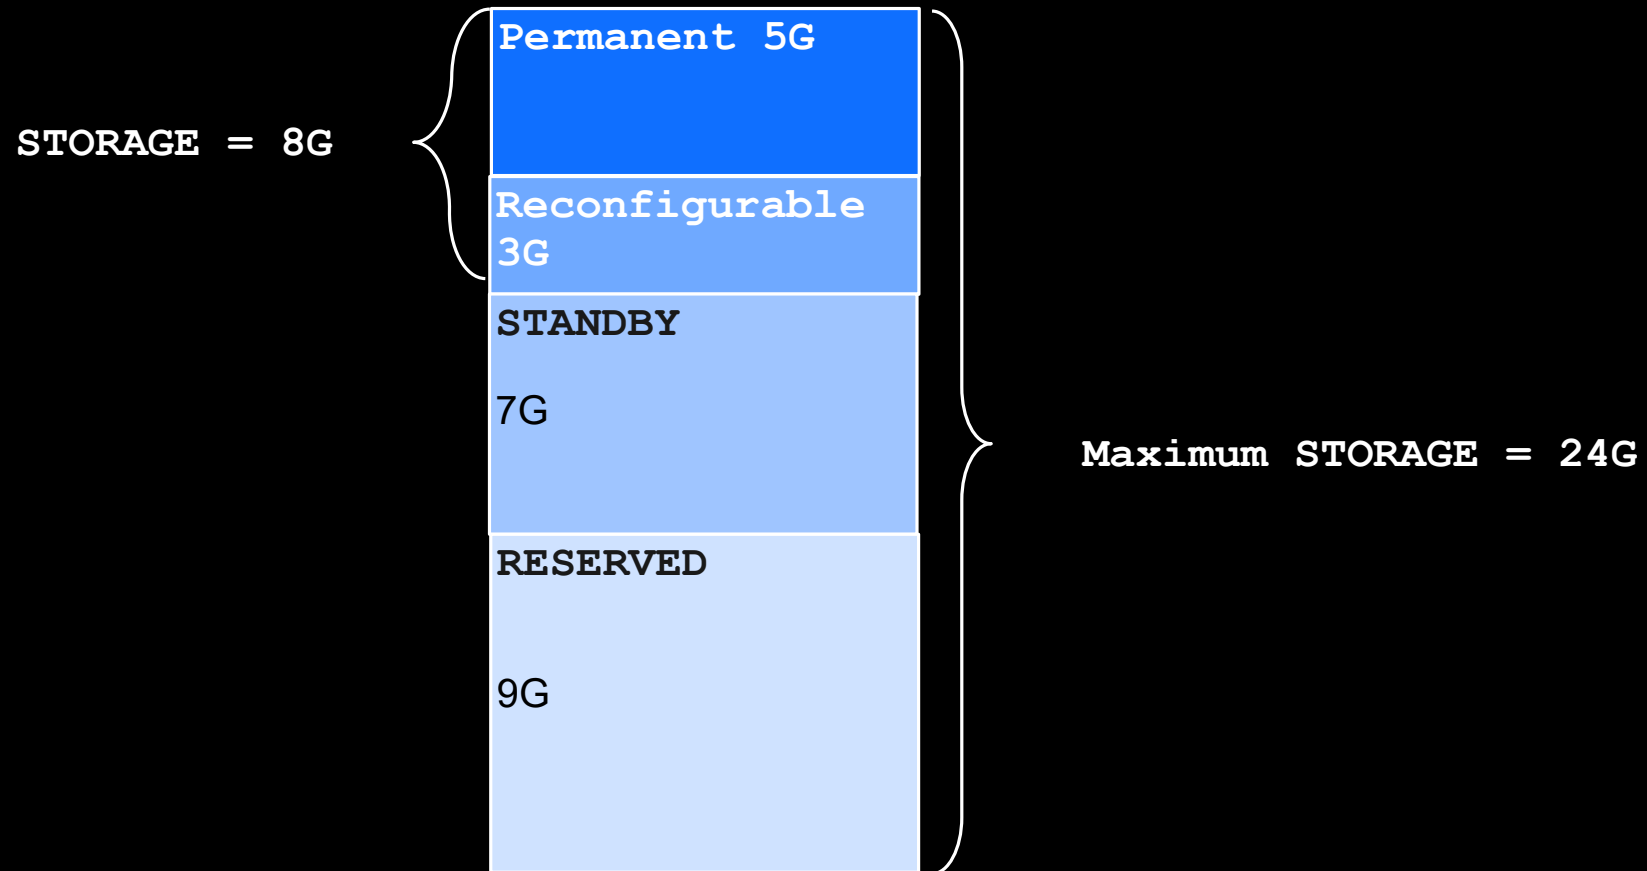
# SET STOR RECONF +1G



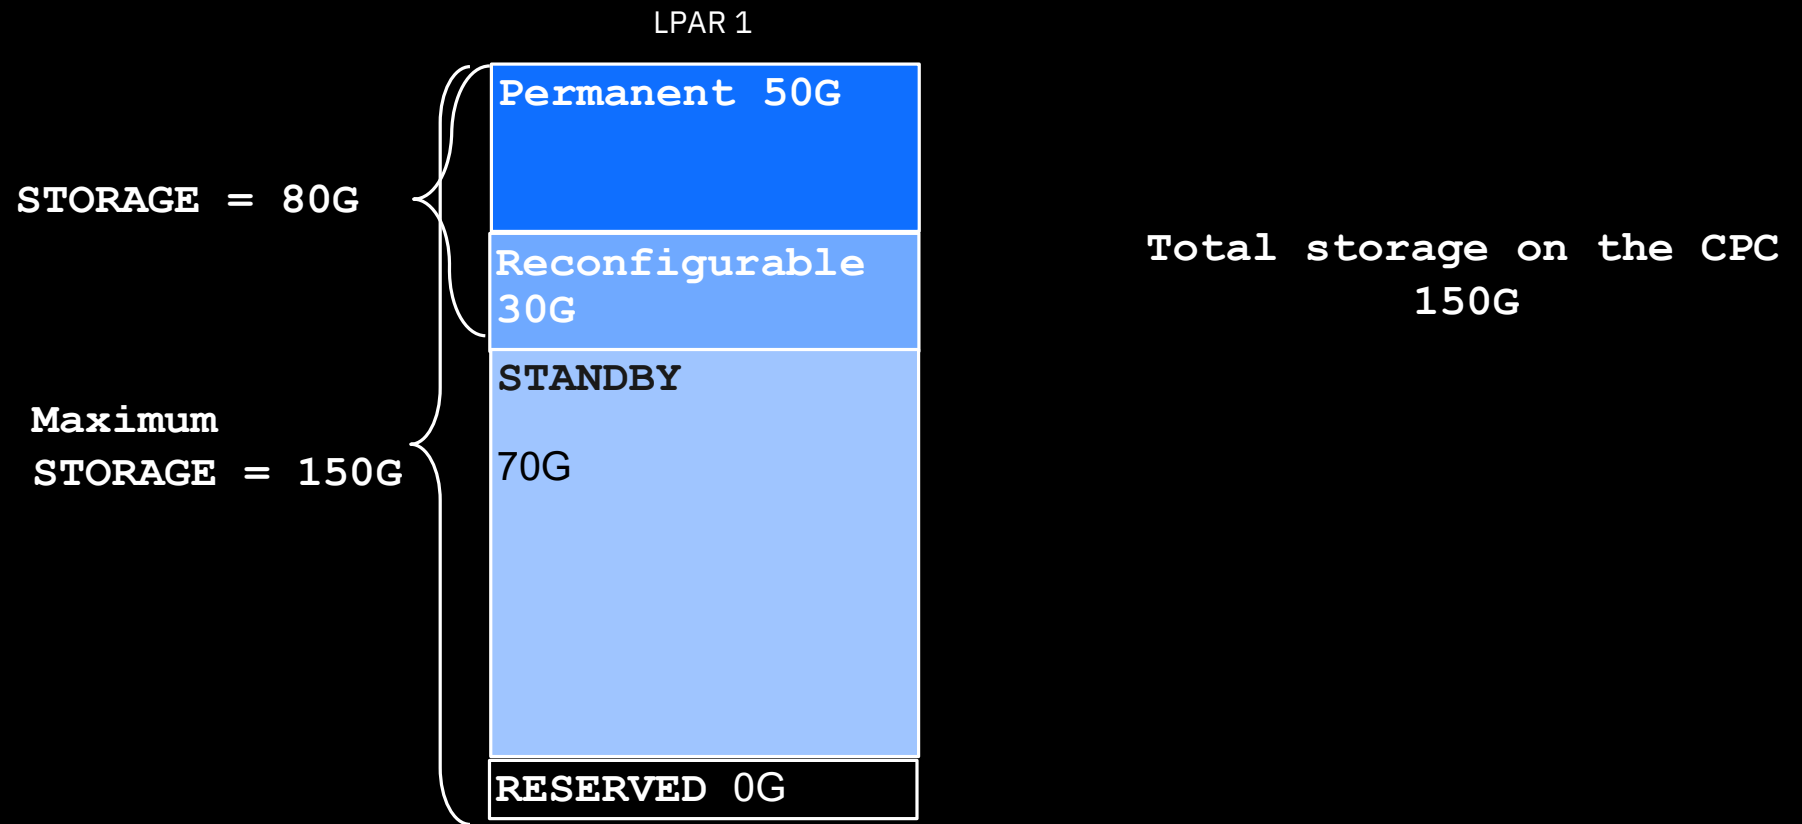
# SET STOR RECONF -1G



# SET STOR PERM +1G

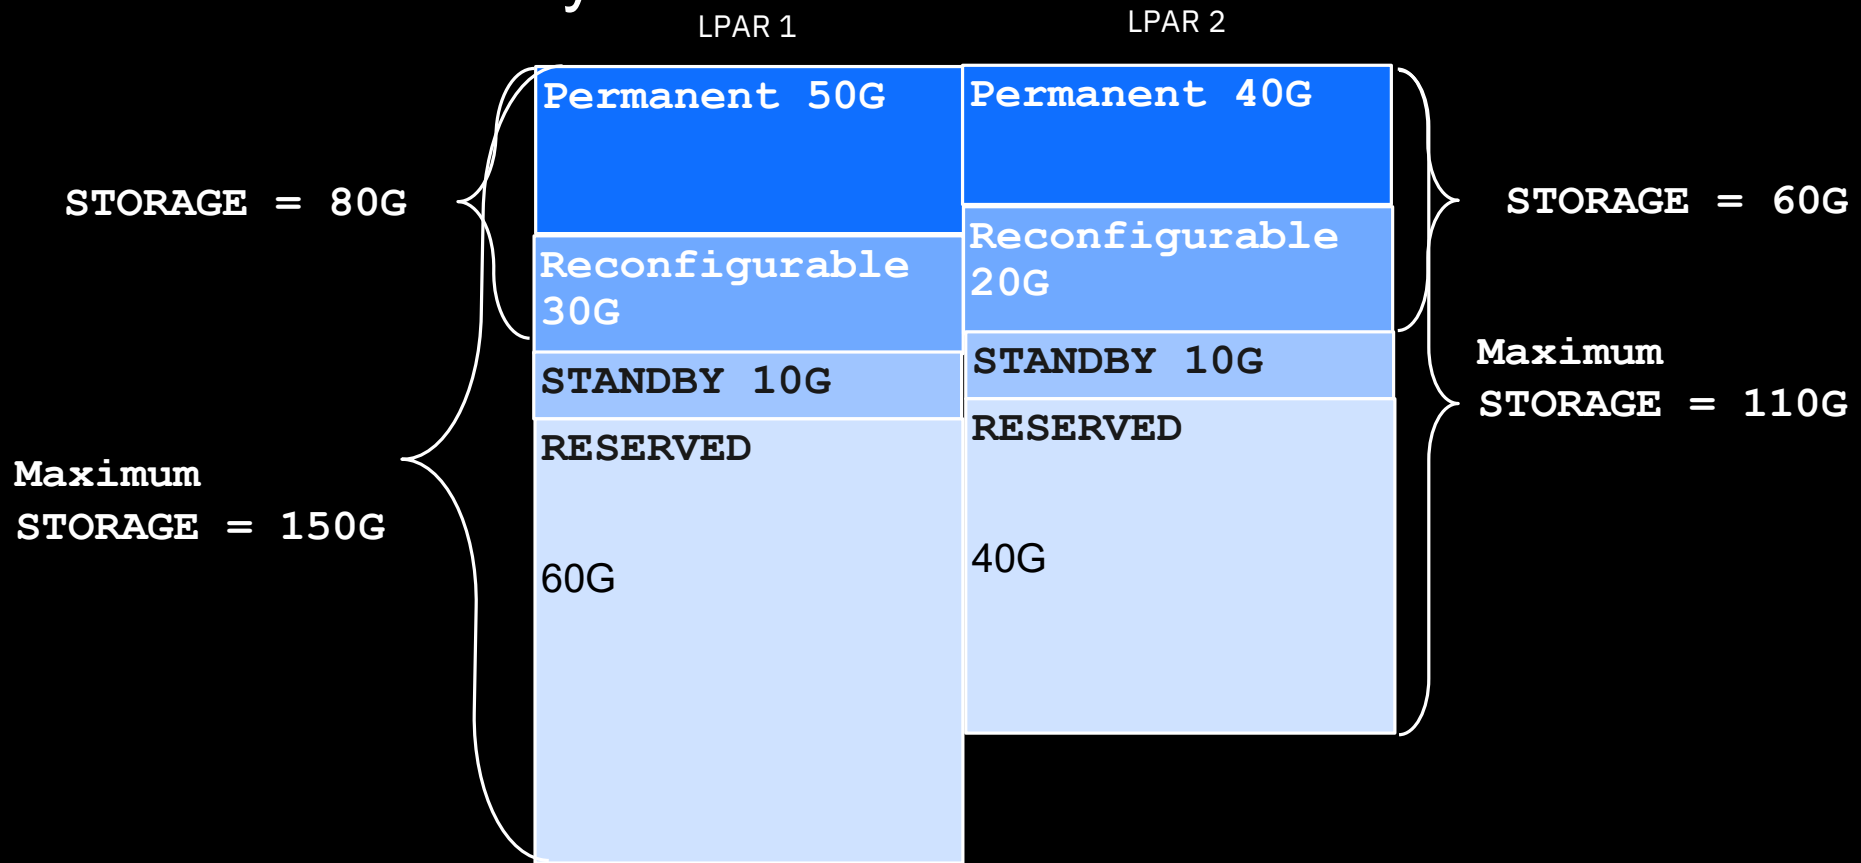


# One LPAR is lonely, but has lots of **STANDBY** memory



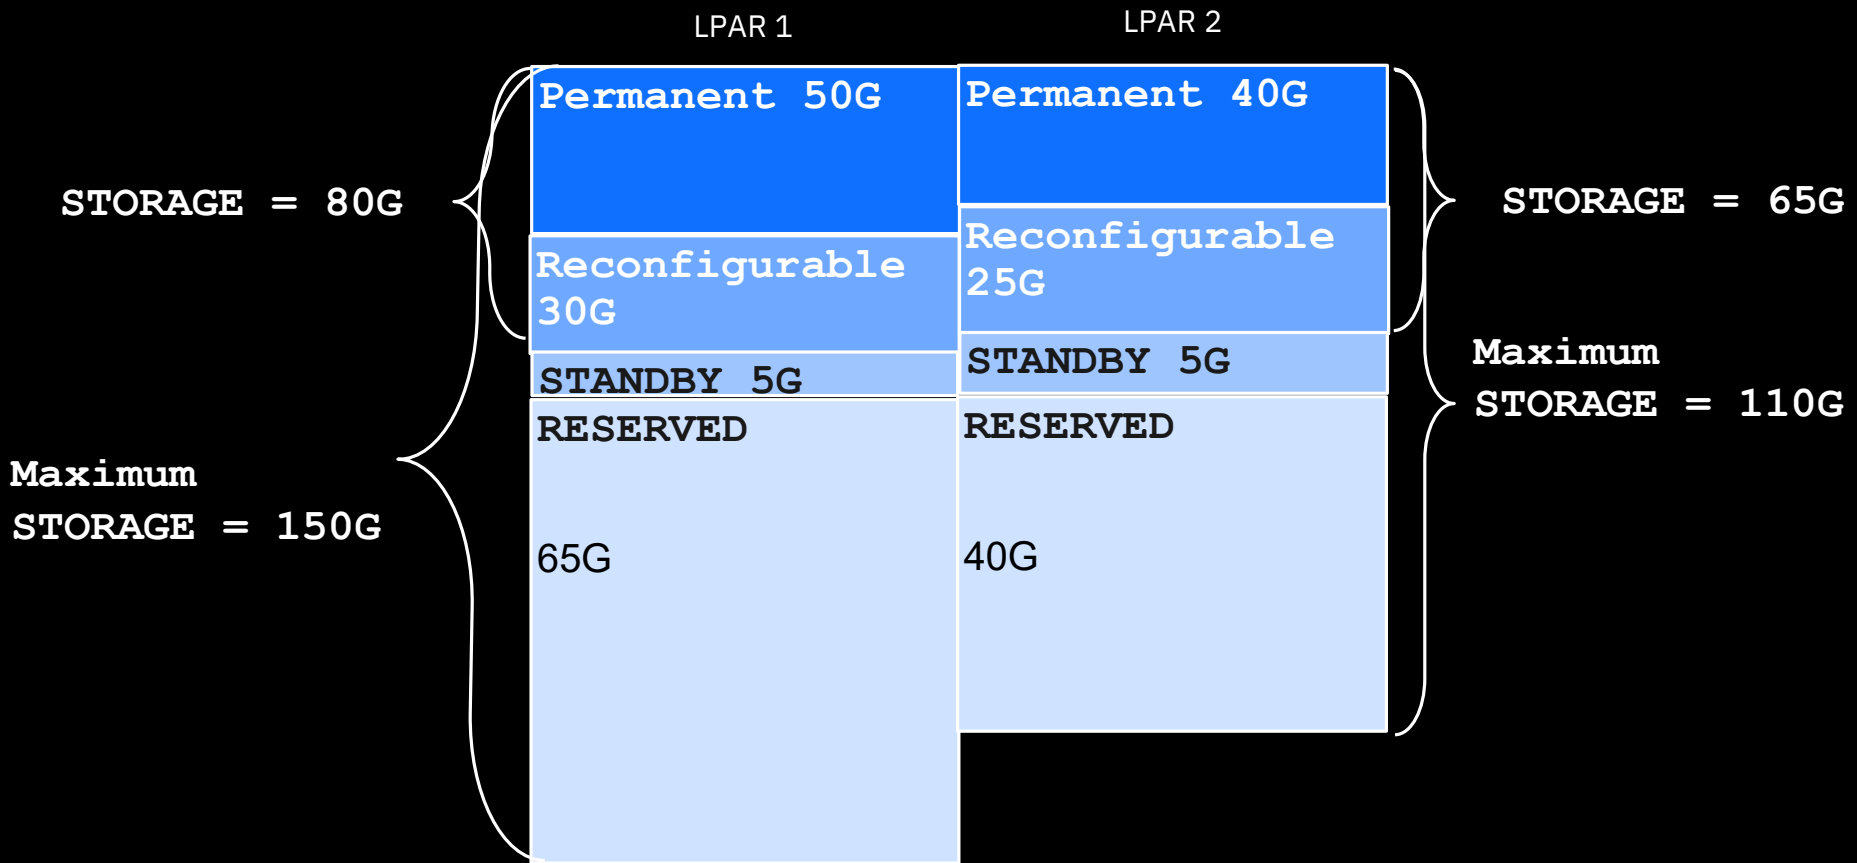


A new partition is activated! I gain a friend but lose some **STANDBY** memory.



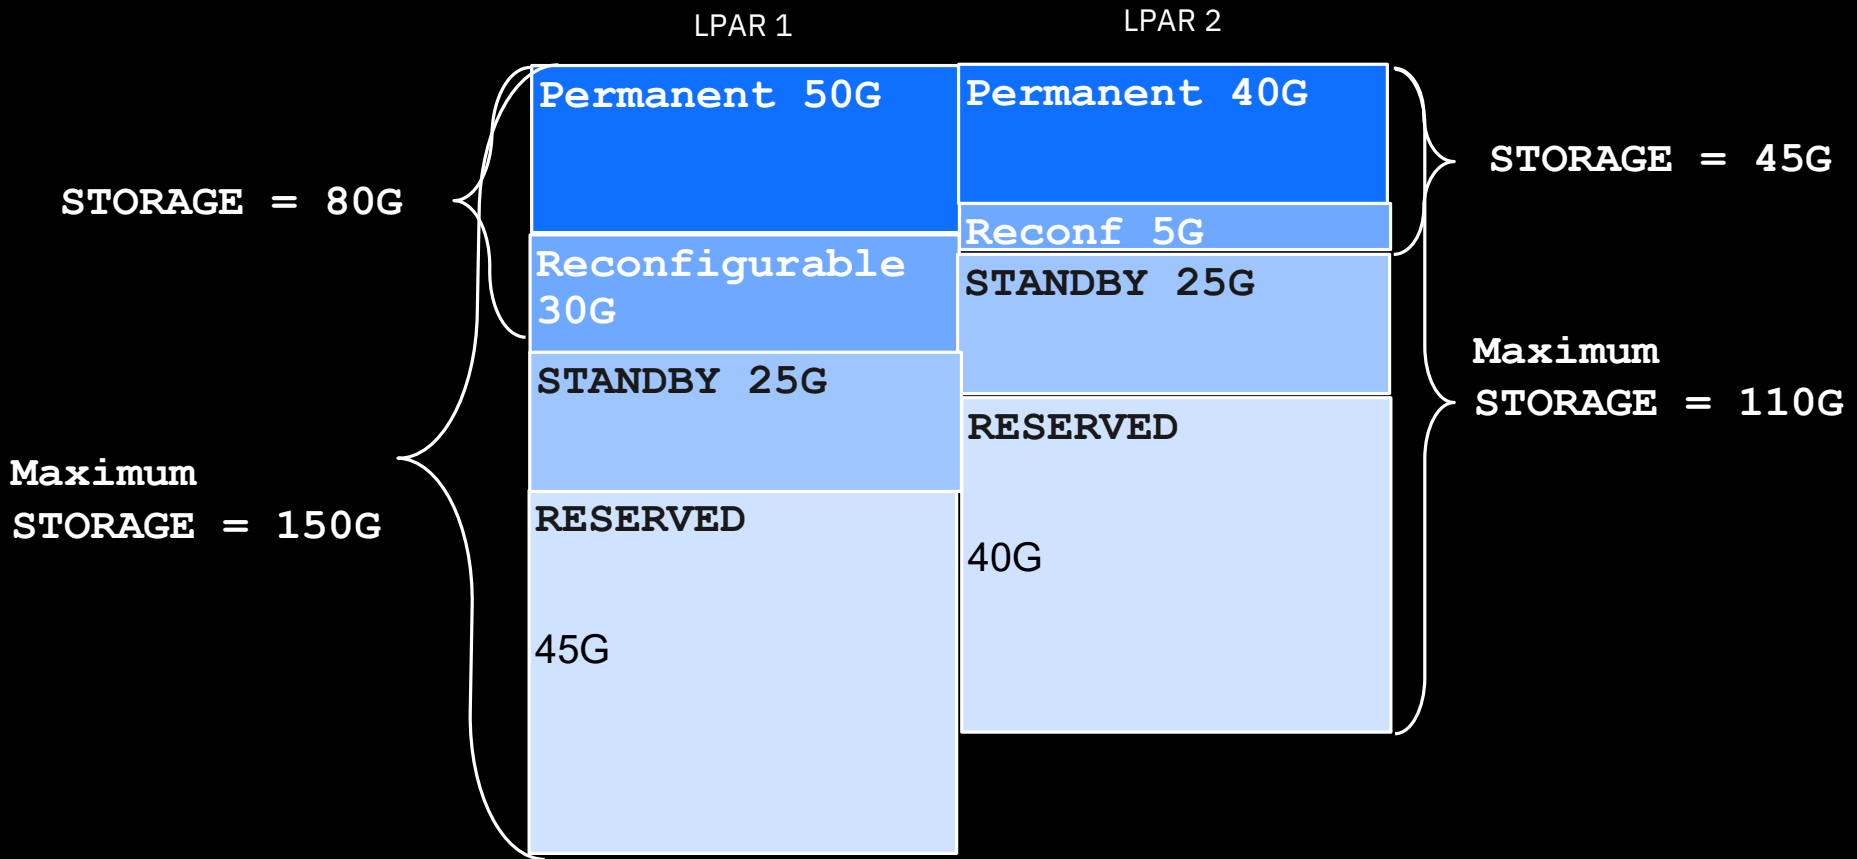
**Total storage on the CPC = 150G = 80G + 60G + 10G (STANDBY)**

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




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

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



**Total storage on the CPC = 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, Rockhopper II (or newer hardware) is necessary for first level real storage reclamations
  - Dependency on z14 firmware enhancements for QDIO and HPMA2
  - For information about recommended hardware service levels, refer to <https://www.vm.ibm.com/memman/dmd.html>
- **Not more than 50% of all online storage may be defined as reconfigurable**
- Reconfigurable storage cannot be added unless there is at least 4G of permanent storage
- Storage additions (and reclamations) must be done in multiples of the storage increment size (QUERY STORAGE will show the increment size).

# When to start planning?



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

# How much permanent storage should I have?

It depends

- Consider what is running today and how it runs (e.g. overcommit, V:R ratio, etc.)
- Consider loads that require permanent storage (e.g. QDIO workloads – large amounts of long term pins)
- Monitor to see if the system performs to expectations

Permanent storage can be added 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 temporary workload. Think about

- 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 temporary workload that must be in permanent storage

- CP control blocks
- I/O and MONITOR pages.



# How much reconfigurable storage should I have?

The following situations lend themselves to reconfigurable storage (temporary workloads)

- Work that happens during special events or at certain periods of time
- Guests that don't always run in the LPAR
- An application that is growing and needs more storage temporarily
- Test LPARs on a storage rich CPC (which would normally be set up to over-commit storage)

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

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

Remember, a decrease in reconfigurable storage will result in a higher V:R ratio

# 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 the highest potential utilization
- 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

```
vir2real
Storage information for VM system GDLLCPX2
CMS users IPL NSSes "CMS GCS" or devices "0190".

Total Virtual storage (only ids not running CMS):           0 MB (  0.0 GB)
Total Virtual storage (only ids running CMS):               9252 MB (  9.0 GB)
Total Virtual storage (all logged on userids):              9252 MB (  9.0 GB)
Total of all Instantiated pages:                            373 MB (  0.4 GB)
Usable real storage (pageable) for this system:            2599481 MB (2538.6 GB)
Total LPAR Real storage:                                    2621440 MB (2560.0 GB)
Permanent storage:                                         2097152 MB (2048.0 GB)
Reconfigurable storage: (Increment size 4G)                524288 MB ( 512.0 GB)
Maximum possible storage:                                   4194304 MB (4096.0 GB)

Total Virtual disk (VDISK) space defined:                   1200 MB (  1.2 GB)
Average Virtual disk size:                                  200 MB

Virtual to (usable) Real storage ratio:                    0.0 : 1
Virtual + VDISK to Real storage ratio:                     0.0 : 1
Virtual to Real ratio (non CMS work only):                 0.0 : 1
Total Instantiated to Real storage ratio:                  0.0 : 1

Paging: 33 volumes active, usable space is:                2244608 MB (2192.0 GB)
Total Paging space in use at 1% utilization:               118 MB (  0.1 GB)
Note: Agelist KeepSlot is DISABLED
```

# 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 a multiple of the increment size
- 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 12G
```

```
Storage reclamation information for VM system GDLLCPX2 on 2021-04-11 at 12:47:46  
CMS users IPL NSSes "CMS GCS" or devices "0190".
```

```
Your system has:
```

```
Total Virtual storage (only ids not running CMS):          0 MB (  0.0 GB)  
Total Virtual storage (only ids running CMS):              9252 MB (  9.0 GB)  
Total Virtual storage (all logged on userids):             9252 MB (  9.0 GB)  
Total of all Instantiated pages:                          372 MB (  0.4 GB)  
Paging usable space:                                     2140625 MB (2090.5 GB)  
Total Paging space in use:                                118 MB (  0.1 GB)  
Paging warning setting:                                   90%
```

```
Projected real and virtual storage ratios:
```

	Current	Post-Reclaim
Total LPAR Real storage:	2560.0 GB	2548.0 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:	0.0 : 1	0.0 : 1
Virtual + VDISK to Real storage ratio:	0.0 : 1	0.0 : 1

# CHKRECLM EXEC

Estimated paging space utilization after reclamation:

	Low	High
Percent of paging in use:	0.6 %	0.6 %
Paging in use:	12.2 GB	12.2 GB
Paging usable space:	2090.5 GB	2090.5 GB

Note: AGELIST KEEPSLOT NO is in effect. Therefore, more paging will result in more page space usage. Consult the CP Planning and Admin book chapter on Allocating DASD space for more info.

# VIR2REAL EXEC and CHKRECLM EXEC

- **VIR2REAL** is an existing EXEC available on the z/VM downloads page:  
<https://www.vm.ibm.com/download/packages/descript.cgi?VIR2REAL>
  - The updated EXEC now uses CP commands/responses available with 4 TB real memory support (APAR VM66713) to display the amount of Permanent, Reconfigurable, and Maximum memory sizes
  - Updates contained in version updated on 2021-03-23 or later
- **CHKRECLM** is a new EXEC that is now available on the downloads page:  
<https://www.vm.ibm.com/download/packages/descript.cgi?CHKRECLM>

# Performance Toolkit can help too!

- Performance Toolkit for z/VM reports can also help with determining your workload's size
- FCX103, Storage Utilization – 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



# Performance Toolkit can help too! (continued)

- FCX292, User Page Utilization Data – UPGUTL – gives you information about how much storage your guests are using

```

FCX292      CPU nnnn  SER nnnnn  Interval HH:MM:SS - HH:MM:SS  Perf. Monitor
-----
          <----- Storage ----->
          <----- Resident ----->
          <----- Invalid But Resident ----->
Data      Spaces      <---- Total ----> <--Locked--> <--- UFO ---> <--- PNR ---> <--AgeList-->      Base
Userid    Owned    WSS  Inst Resvd T_All  T<2G  T>2G  L<2G  L>2G  U<2G  U>2G  P<2G  P>2G  A<2G  A>2G  XSTOR  AUX  Space  Nr of
>>Mean>> 10.6 5645M 6113M 256M 5391M  34M 5356M 585.1 41472  .0 2743K 292.6  46M  73.1 6008K  .0 3305M 6341M 56
User Class Data:
PERF      .0  16M  35M  .0  16M  .0  16M  .0  .0  .0  512K  .0  15K  .0  204576  .0  34M  3072M  1
Service   .0 1972K 5196K 2048K 2048K 4096 2044K .0 77824 .0 1204K .0  .0  .0  376K  .0 3640K 32768K 1

User Data:
APCSCIF   0  952K  972K  0   4K   0   4K   0   0   0   0   0   0   0   0   0  968K  64M
BK200001  0 3318M 3074M  0 3072M  0 3072M  0   0   0   0   0   0   0   0  2160K 3328M
BK200007  0 3318M 3074M  0 3072M 12M 3060M  0   0   0   0   0   0   0   0  3074M 3328M
BK200008  0 3318M 3074M  0 3072M  76K 3072M  0   0   0   0   0   0   0   0  3074M 3328M
BK300012  0 4194M  16G  0 6814M 164K 6814M  0   0   0 8728K  4K 265M  4K 273M  0  16G 16640M
BK400015  0 100G  100G  0   86G  760M  85G  0   0   0  47M  0 1085M  0  50M  0  60G  100G

Command ==> _
F1=Help F4=Top F5=Bot F8=Fwd F12=Return
    
```

Shows number of instantiated pages, add up the numbers in this column for User Data: to get the sum of all instantiated pages over all users

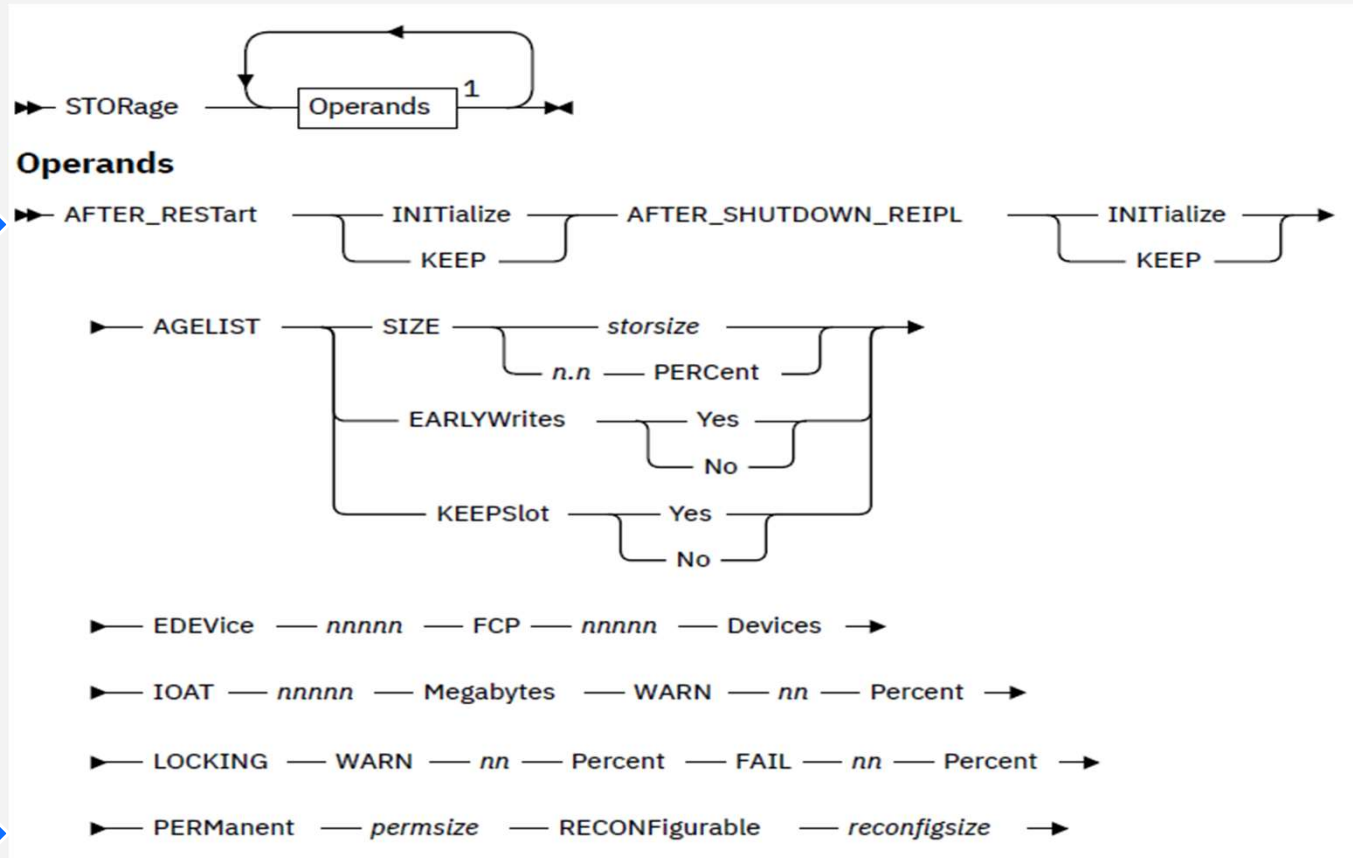
# The Nitty-Gritty: Configuration Statements and Commands

# CP Variables

- z/VM New Function Variable List page:  
<https://www.vm.ibm.com/newfunction/varlist.html>
- Issue QUERY VARIABLE to find out what function your z/VM system has

APAR	CP Variable	Value	Meaning
VM66173	CP.FUNCTION.MEMORY.RECLAIM	1	Value indicates that dynamic removal of reconfigurable memory is supported.
VM66173	CP.LIMIT.MEMORY.RECONFIG	0	Value indicates that 0% of a system's real memory can be designated as reconfigurable. However, all the functions needed for DMD are available and functional (like system configuration options for PERM and RECONFIG).
VM66173	CP.LIMIT.MEMORY.REAL_TOTAL	4	Value indicates that 4 TB of real memory is supported.
VM66271	CP.LIMIT.MEMORY.RECONFIG	50	Value indicates that 50% of a system's real memory can be designated as reconfigurable.

# STORAGE System Configuration Statements



# STORAGE System Configuration Statements (continued)

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

Central Storage	
Amount in:	Gigabytes (GB) ▾
Initial:	8.0
Reserved:	2.0

You can specify the storage for z/VM to IPL with using the system configuration **STORAGE** statements:

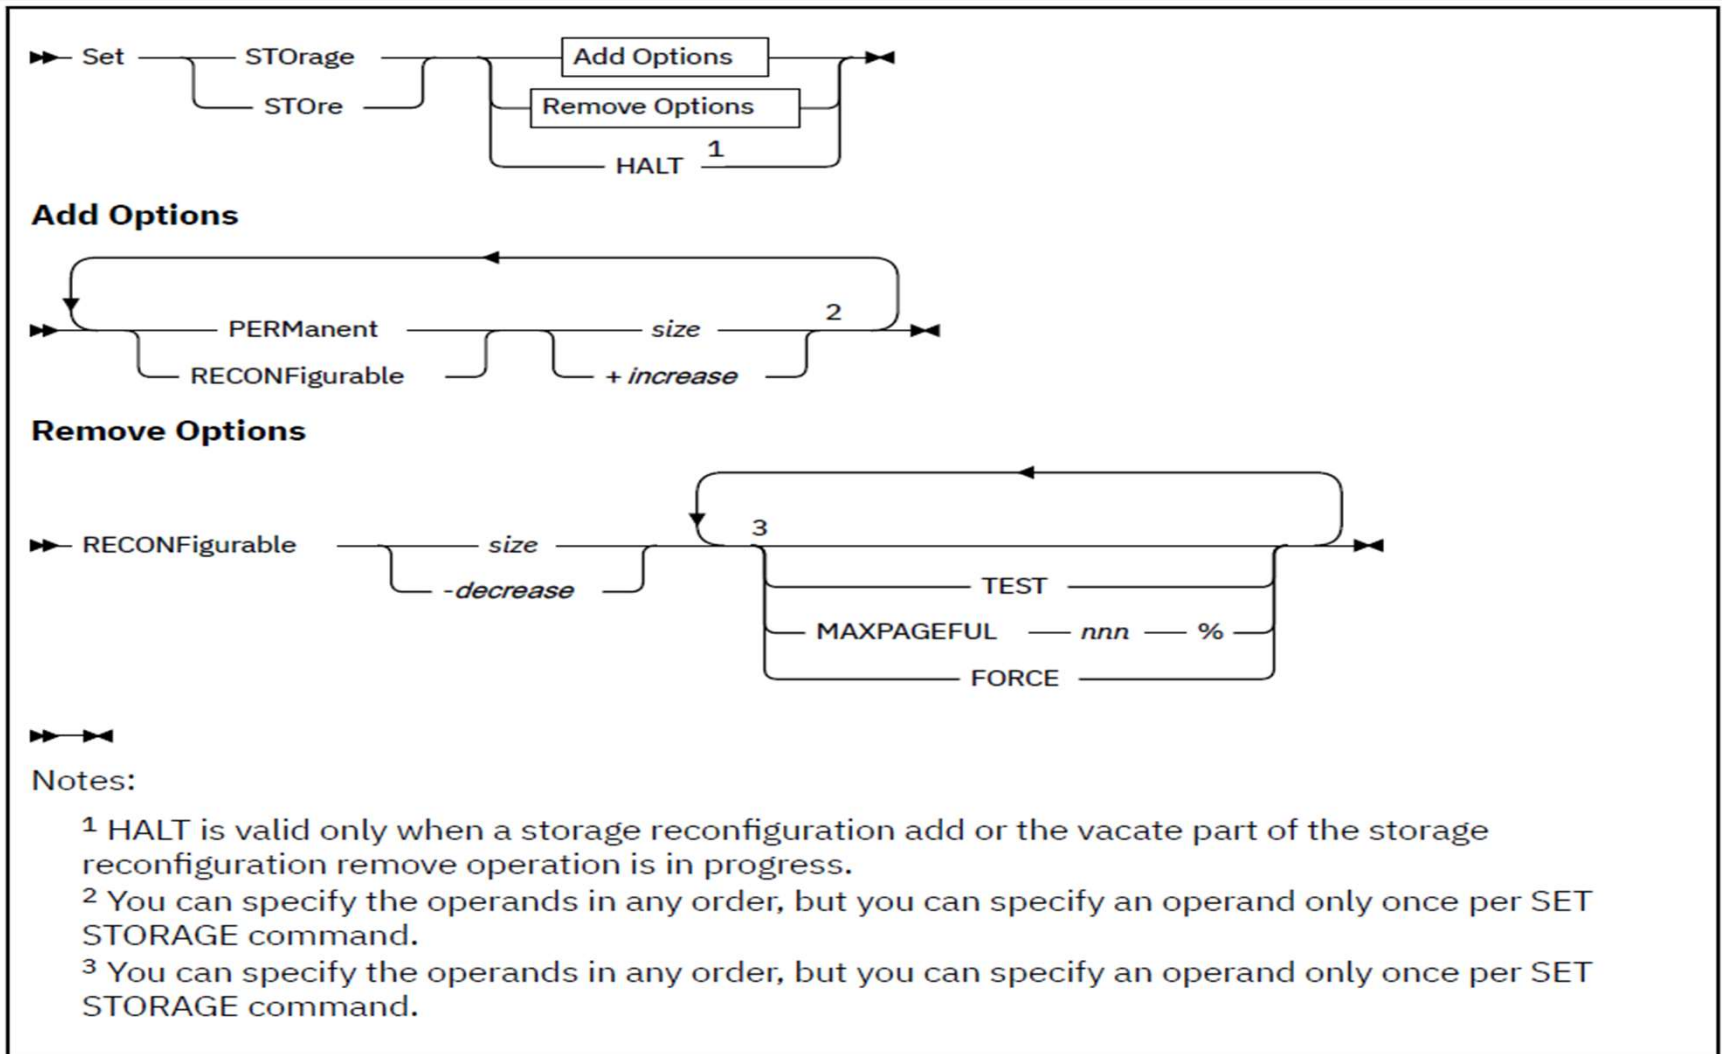
```
STORAGE PERManent 8G
```

```
STORAGE RECONFigurable 2G
```

... Or maybe

```
STORAGE PERM 4G RECONF 4G
```

# Dynamic Storage Reconfiguration



# Dynamically Adding Storage

Add PERManent and RECONFigurable storage in one command

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

- When adding both PERManent and RECONFigurable storage in one command, PERManent storage is always processed first

Storage initialization changes made in APAR VM66173 (4 TB real memory support)

- Storage initialization may still be taking place after a SET STORE add command has completed. Another SET STORE command may be issued as soon as this process has completed, even if frame initialization is ongoing
- QUERY FRAMES will show a non-zero value for NotInitialized= if frames are still being initialized

SET STORE commands will be rejected if a storage reconfiguration is currently in progress

QUERY STORE commands will indicate if a storage reconfiguration is in progress.

# Dynamically Removing Storage

## **SET STOR RECONF -1G**

HCPPCC2593I Storage reclamation viability test passed with MAXPAGEFULL value of 90%. 19% of paging space could be required for the current workload.

HCPPCC2581I Storage reconfiguration to remove 1G of reconfigurable initiated by OPERATOR.

HCPPCC2582I Storage reconfiguration by OPERATOR is complete. Permanent = 4G Reconfigurable = 1G

Use `TEST` to determine the viability:

## **SET STOR RECONF -200G TEST**

HCPPCC2650E Storage reclamation viability test failed with MAXPAGEFULL value of 90%. 150% of paging space could be required for the current workload.



## Dynamically Removing Storage (continued)

Use `MAXPAGEFULL` to set maximum paging percent:

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

```
HCP2650E 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 = 5G CONFIGURED = 5G INC = 128M STANDBY = 9G RESERVED = 0
Permanent = 4G Reconfigurable = 1G Maximum STORAGE = 14G
Storage increase in progress. Elapsed time = 00:00:01
Target: Permanent = 13312M
Total to add: Permanent = 9216M
Remainder to add: Permanent = 9216M
Storage reconfiguration is active.
```

## Monitoring a Storage Reconfiguration (continued)

Use new RECONFiguration option of Q STOR to check status of an in-progress storage reclamation:

### Query Store RECONFiguration

```
STORAGE = 12544M CONFIGURED = 12544M INC = 128M STANDBY = 1792M RESERVED = 0
Permanent = 7G Reconfigurable = 5376M Maximum STORAGE = 14G
Storage decrease in progress. Elapsed time = 00:00:01
Target Reconfigurable =      2048M
Total to remove      =      5120M
Remainder to remove  =      3328M
MAXPAGEFULL          = Forced to no limit
Storage reconfiguration is active.
```

# 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

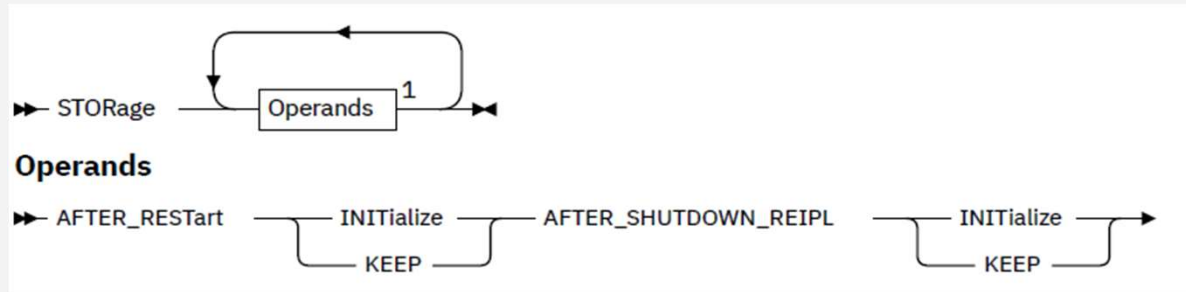
```
HCP2650E Storage reclamation viability test failed...
```

- As a reclamation progresses, the system will periodically check viability so...
- The system will halt an in-progress reclamation if the viability test fails.

# Memory Reclamation on Guest Systems

- z/VM second-level will support memory reclamation on z13
- Useful for testing out reclamations
  - Use `INCRement` operand of `DEFINE STORAGE` command to see how increment size affects adding and removing storage
  - Second level guest must specify `DEFINE STORAGE` command with `STANDBY` and/or `RESERVED` storage
- When your second level z/VM system has no `STANDBY` or `RESERVED` storage defined, the increment size will be 1M and you cannot use `SET STORAGE` to add or remove storage.

# System Configuration Restart Options



Use new `AFTER_REStart` and `AFTER_SHUTDOWN_REIPL` options to specify how storage is initialized:

- `AFTER_REStart` specifies how to initialize storage after an abend
- `AFTER_SHUTDOWN_REIPL` specifies how to initialize storage after a `SHUTDOWN REIPL` command

In both cases:

- `INITIalize` will initialize storage in the “usual” way using system config options (if available)
- `KEEP` will initialize storage based on the storage configuration at the time of the restart/reIPL.

# System Configuration Restart Options (continued)

Want to ignore any dynamic storage changes made prior to reIPLing? Use `INITIALIZE`

- Specify behavior after a SHUTDOWN REIPL:

```
STORage AFTER_SHUTDOWN_REIPL INITIALIZE
```

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

- Specify behavior after a CP Abend or PSW RESTART

```
STORage AFTER_RESTART KEEP
```

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

- As long as both the CP nucleus being IPLed and the CP nucleus being SHUTDOWN are DMD capable.  
*This behavior can only be changed in the system configuration file.*

# Storage After a Restart

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

```
q store ipl
STORAGE = 256G CONFIGURED = 256G INC = 1G STANDBY = 0 RESERVED = 0
Permanent = 200G Reconfigurable = 56G Maximum STORAGE = 256G
IPL actual:    Permanent = 200G Reconfigurable = 56G
IPL requested: Permanent = 200G Reconfigurable = 200G
IPL requested data source: STORAGE system configuration statement(s)
Last start was a system restart from SHUTDOWN REIPL
```



## Keep in mind...

- Ensure workload is stabilized before initiating a reclamation
- Specified storage values must be a multiple of the increment size (`INCRement`)
- Limit reconfigurable memory to the amount expected to be reclaimed
  - There is some overhead associated with managing reconfigurable storage
- Time it takes to complete a reclamation depends on size and system workload
- Try to avoid simultaneous storage reconfigurations on other partitions
- A storage reconfiguration may halt, but will not be automatically undone
- Update your system configuration file to match dynamic storage changes if you want to keep them going forward.

# Paging & Pinning implications, and interactions with other commands

# Paging considerations

- How much paging space do I need?
  - Consider the total amount of potential memory your guests will use
  - How much is instantiated in aggregate on a given day?
  - Consult *CP Planning and Administration* Chapter 23, *Paging Space* for guidance
- Be prepared for spikes in virtual storage use – this translates to real storage use
  - z/VM will warn you when paging space is 90% full (always)
- Set your own warning level!
  - **SYSTEM CONFIG** statement  
`PAGING WARNING nnn%`
  - **Dynamic command**  
`SET PAGING WARNING nnn%`
  - This value will be used to send an alert to the system operator whenever the paging space exceeds the specified threshold (as well as when you go over 90%)
- If set, the `PAGING WARNING` percentage will be the default `MAXPAGEFULL` value used in reclamation viability checks.

# Storage reclamation and paging

- Removing storage can cause more paging
  - We will have less main storage to work with!
  - The act of removing storage uses some storage temporarily
- Is it okay to exceed your normal paging space utilization during a storage reclamation?
  - Maybe, if you're removing storage so you can move workload to a new system
  - Maybe not, the storage being removed should be excess so you wouldn't expect to see a large increase in utilization
- CP will not do a storage reclamation if it will cause paging utilization 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.
- As previously mentioned, 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.

# Pinning considerations

- Pinned page characterization
  - All pins by z/VM are characterized as short or long term
  - Short term pins (not a problem)
- A pinned page cannot be reclaimed until unpinned
  - Long term pins can hold up a reclamation
  - z/VM will revisit pinned pages
- Although z/VM is designed to ensure that reclamation completes in the most common scenarios, there is the possibility a DMD may stall
  - Some of these stall scenarios are addressed with firmware updates reference later.
  - Guidance and tools for how to determine if a DMD has stalled and how to recover found at <https://www.vm.ibm.com/memman/dmd.html>
- DMDQPIN EXEC
  - Aids in determining if a “stalled” DMD is due to pinned guest pages
  - **DMDQPIN** is a new EXEC that is now available on the downloads page : <https://www.vm.ibm.com/download/packages/descript.cgi?DMDQPIN>

# 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, reconfigurable, and vacating frames (number of frames being taken offline):

```
Permanent=52428800 Reconfigurable=14680064 Vacating=0
```

- Also shows frames not yet initialized:

```
NotInitialized=0
```

## QUERY PAGING

- Shows the new WARNING nnn%

## QUERY SXSTORE

- 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

- Always shows the host logical storage address ranges, which do not necessarily correspond with actual sizes.

# Summary

An exciting new capability, memory reclamation, is now available on z/VM 7.2 with APAR VM66271

Memory reclamation is available only for the IBM z14 (or newer) family of servers (or equivalent)

For information about recommended hardware service levels, refer to <https://www.vm.ibm.com/memman/dmd.html>

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

Other enhancements like the ability to monitor reconfigurations and keep or reset dynamic storage changes on restart.

# Thank you!

Contact Information:

**Jeff Gertner**

**z/VM Storage Management Development**

**[gertnerj@us.ibm.com](mailto:gertnerj@us.ibm.com)**



# Presentation Contributors

- Bill Bitner
- Walter Church
- Emily Hugenbruch
- Steve Jones
- Stephanie Rivero
- Brian Thompson
- Jeff Gertner

# Suggested Use Cases

Use `VIR2REAL` to figure out what the ideal paging percentage is for your system

- Set `PAGING WARNING` near your ideal paging percentage
- You might set the `PAGING WARNING` a bit lower than your average or ideal paging percentage in order to be alerted before paging starts ramping up

Use Dynamic Memory Upgrade to add memory to an LPAR and move workload there during a planned outage, then move the workload and memory back after its original LPAR is available

Use DMU and DMD in conjunction with Live Guest Relocation to rebalance workload in an SSI.

# What about VMRELOCATE?

Best practice: Do not reconfigure storage while actively adjusting a partition's workload.

Suggested approach:

1. Add storage to target partition
2. Relocate guest to target partition
3. Remove storage from source partition

Can you relocate a guest while reconfiguring the host storage?

- Yes, but we don't recommend it
- Be careful! Both relocation and reconfiguration can temporarily cause more storage usage

Do storage reconfigurations proactively, not reactively!

# Storage and Initialization

- If you leave your system how it is today (with no changes to the SYSTEM CONFIG) whatever storage that is online to the LPAR will come up as `Permanent`
  - If you have `STANDBY` or `RESERVED` storage available, you may be able to add `RECONFIGURABLE` storage later
- If you have 4G storage or less, it will always be initialized as `Permanent` storage
- Use new system configuration statements to change how storage is initialized:
  - Statements for `Permanent` and `Reconfigurable` storage
  - Statements for storage configuration after a restart or re-IPL
  - If there are multiple `STORAGE` entries, the last one “wins”
  - **Reminder: Use `CPSYNTAX` to check your configuration file for errors!**
- The IPL parameter `STORE=` overrides `SYSTEM CONFIG` specifications for `Permanent` and `Reconfigurable` storage (all storage initializes as `Permanent` storage)
- Unused storage available to the LPAR is release to the CPC for use by other LPARs

