

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

© 2018, 2021 IBM Corporation



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

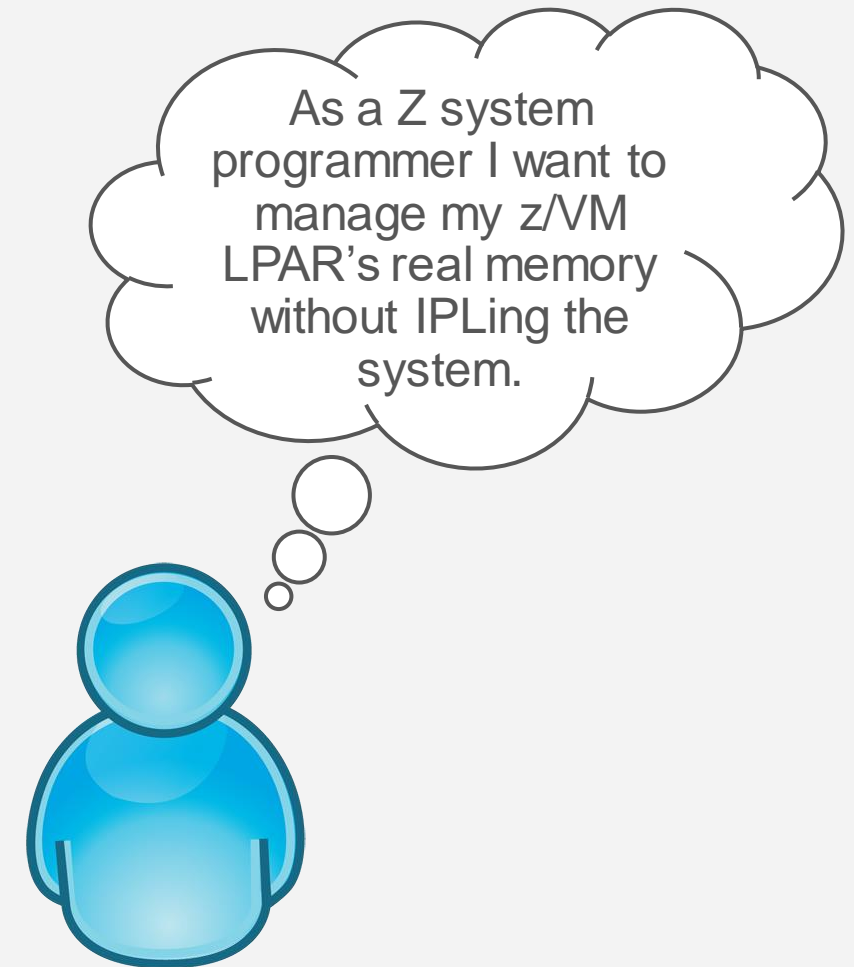
Paging implications,
interactions with other
commands, & conclusion

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](#)

Today: z/VM 7.2 – and APAR VM66173

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

Dynamic Memory Downgrade

- Will pre-req VM66173 (+ VM66508 - <https://www.vm.ibm.com/service/redalert/>)
- Reconfigurable storage can be removed from a running z/VM system
- RECONFIGURABLE storage to be added up to 50% of total online storage

What?

APAR VM66271

When?

3Q21

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. Previously z/VM allowed adding memory, but not removing it.

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

```
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.

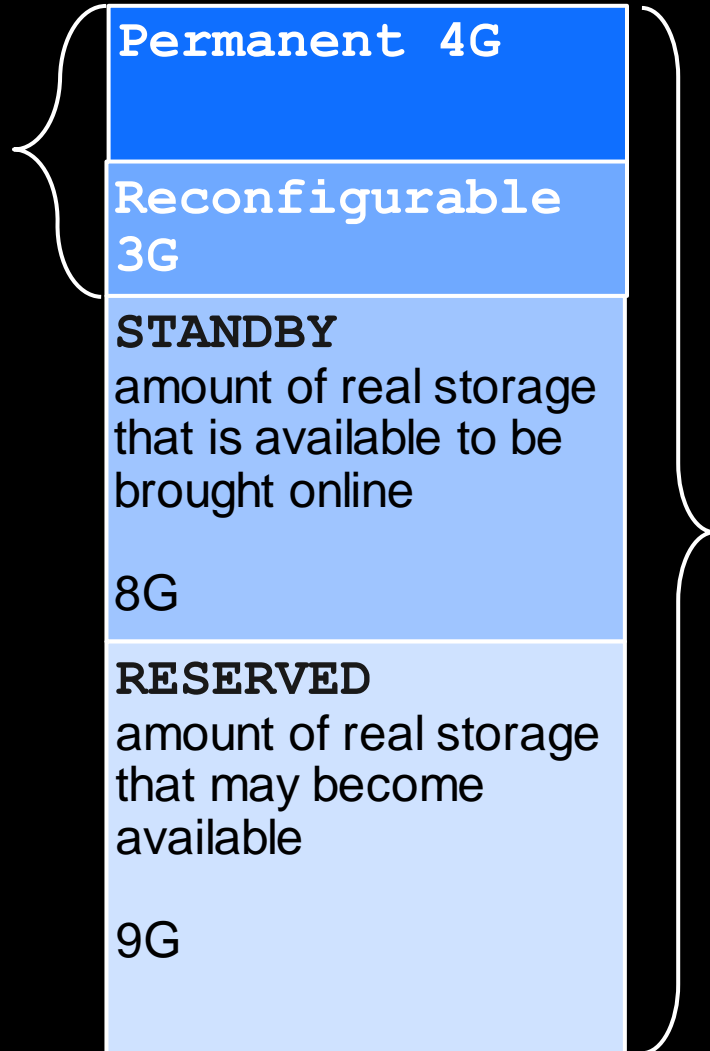
Reconfigurable

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

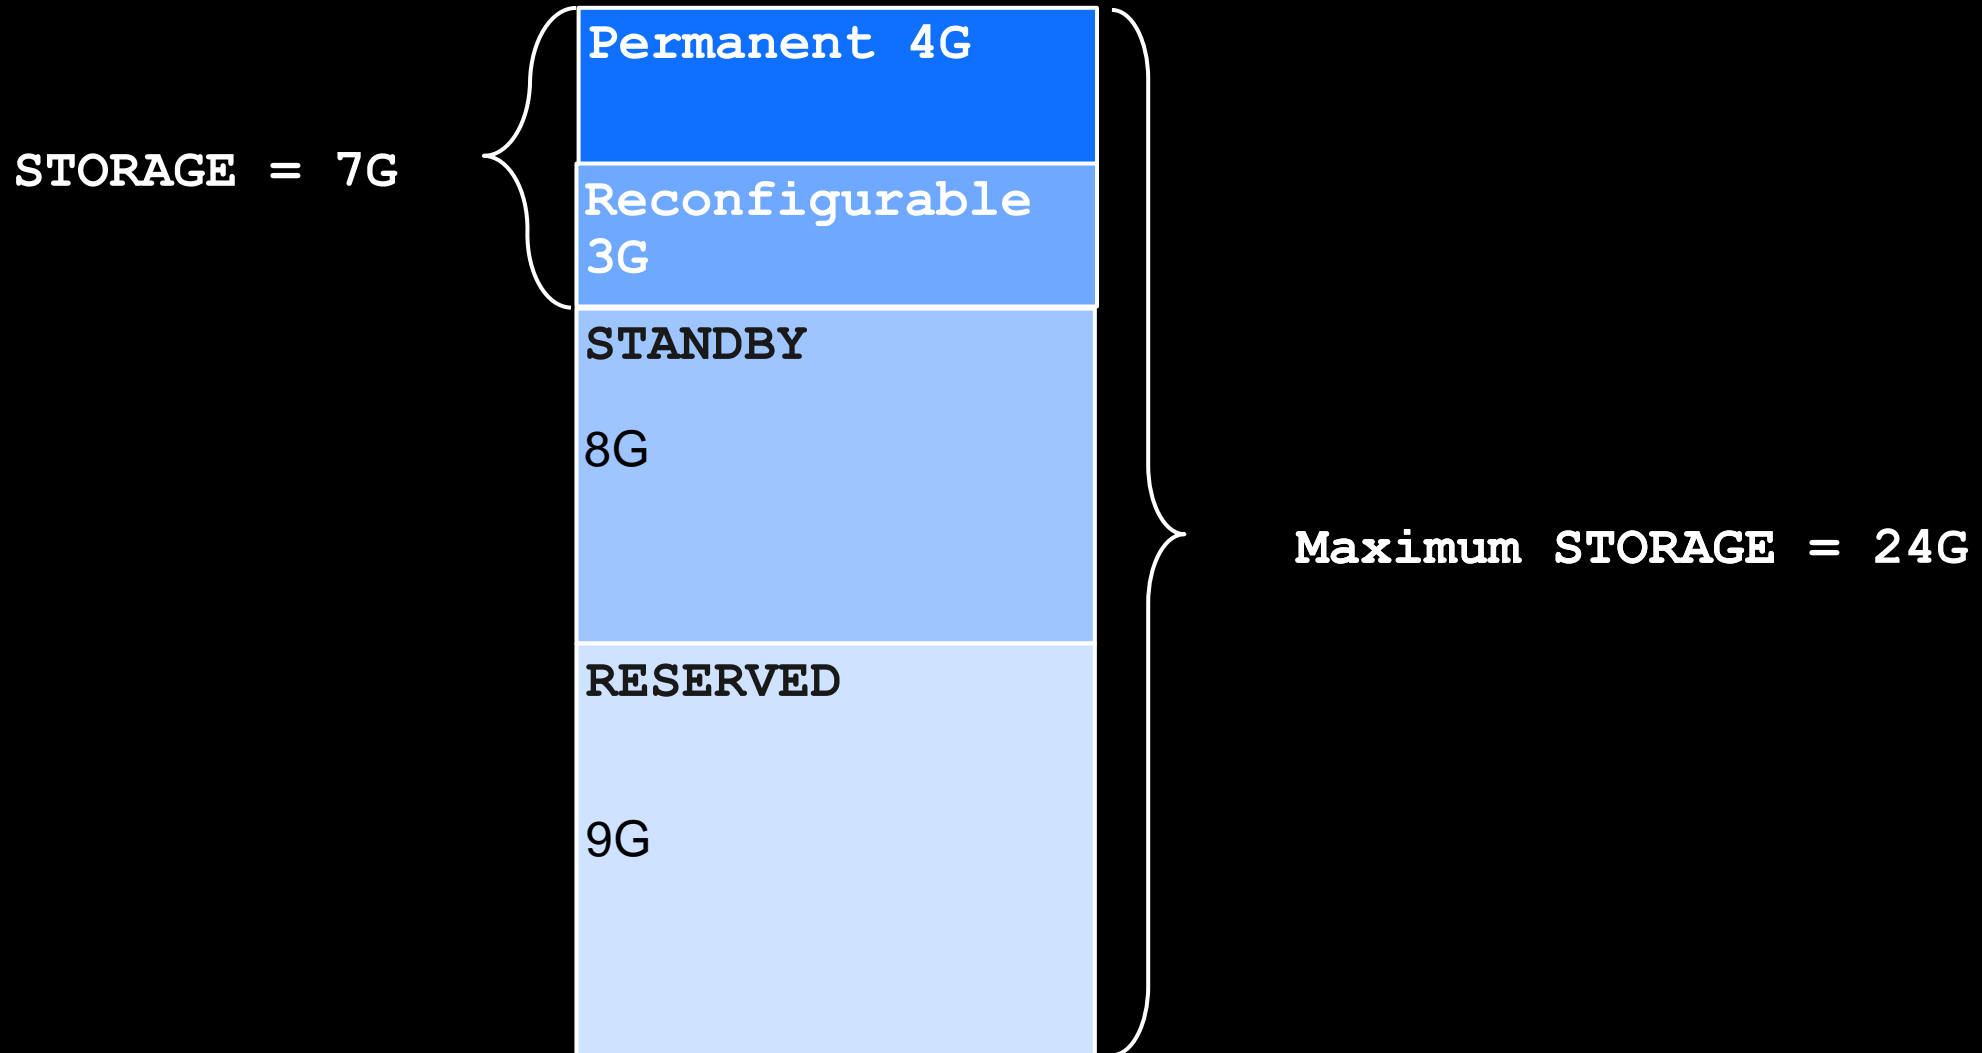
SET STOR RECONF +1G

STORAGE = 8G




Maximum STORAGE = 24G

SET STOR RECONF -1G




SET STOR PERM +1G

STORAGE = 8G



Maximum STORAGE = 24G

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

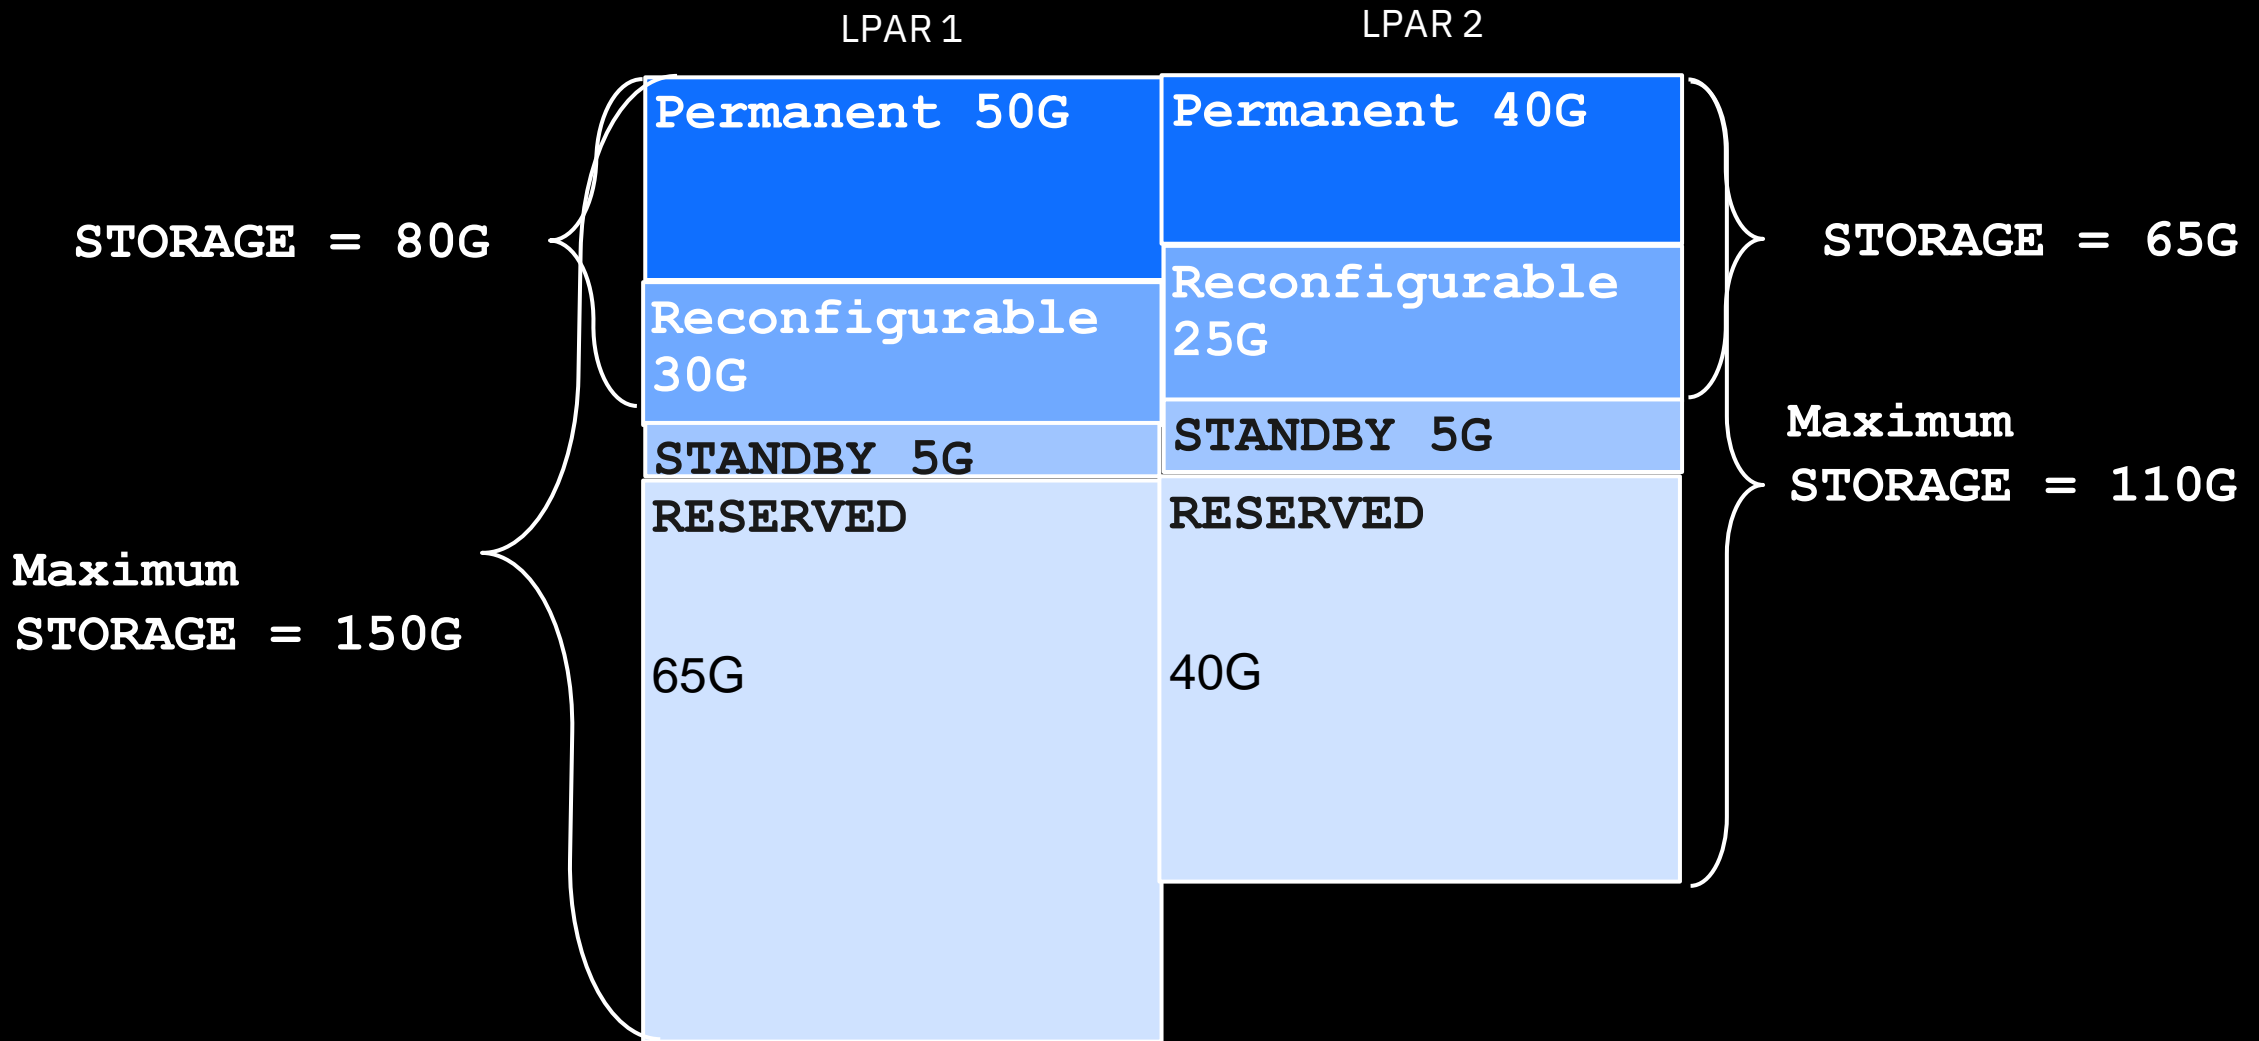


A new partition is activated! I gain a friend but lose some **STANDBY** memory.



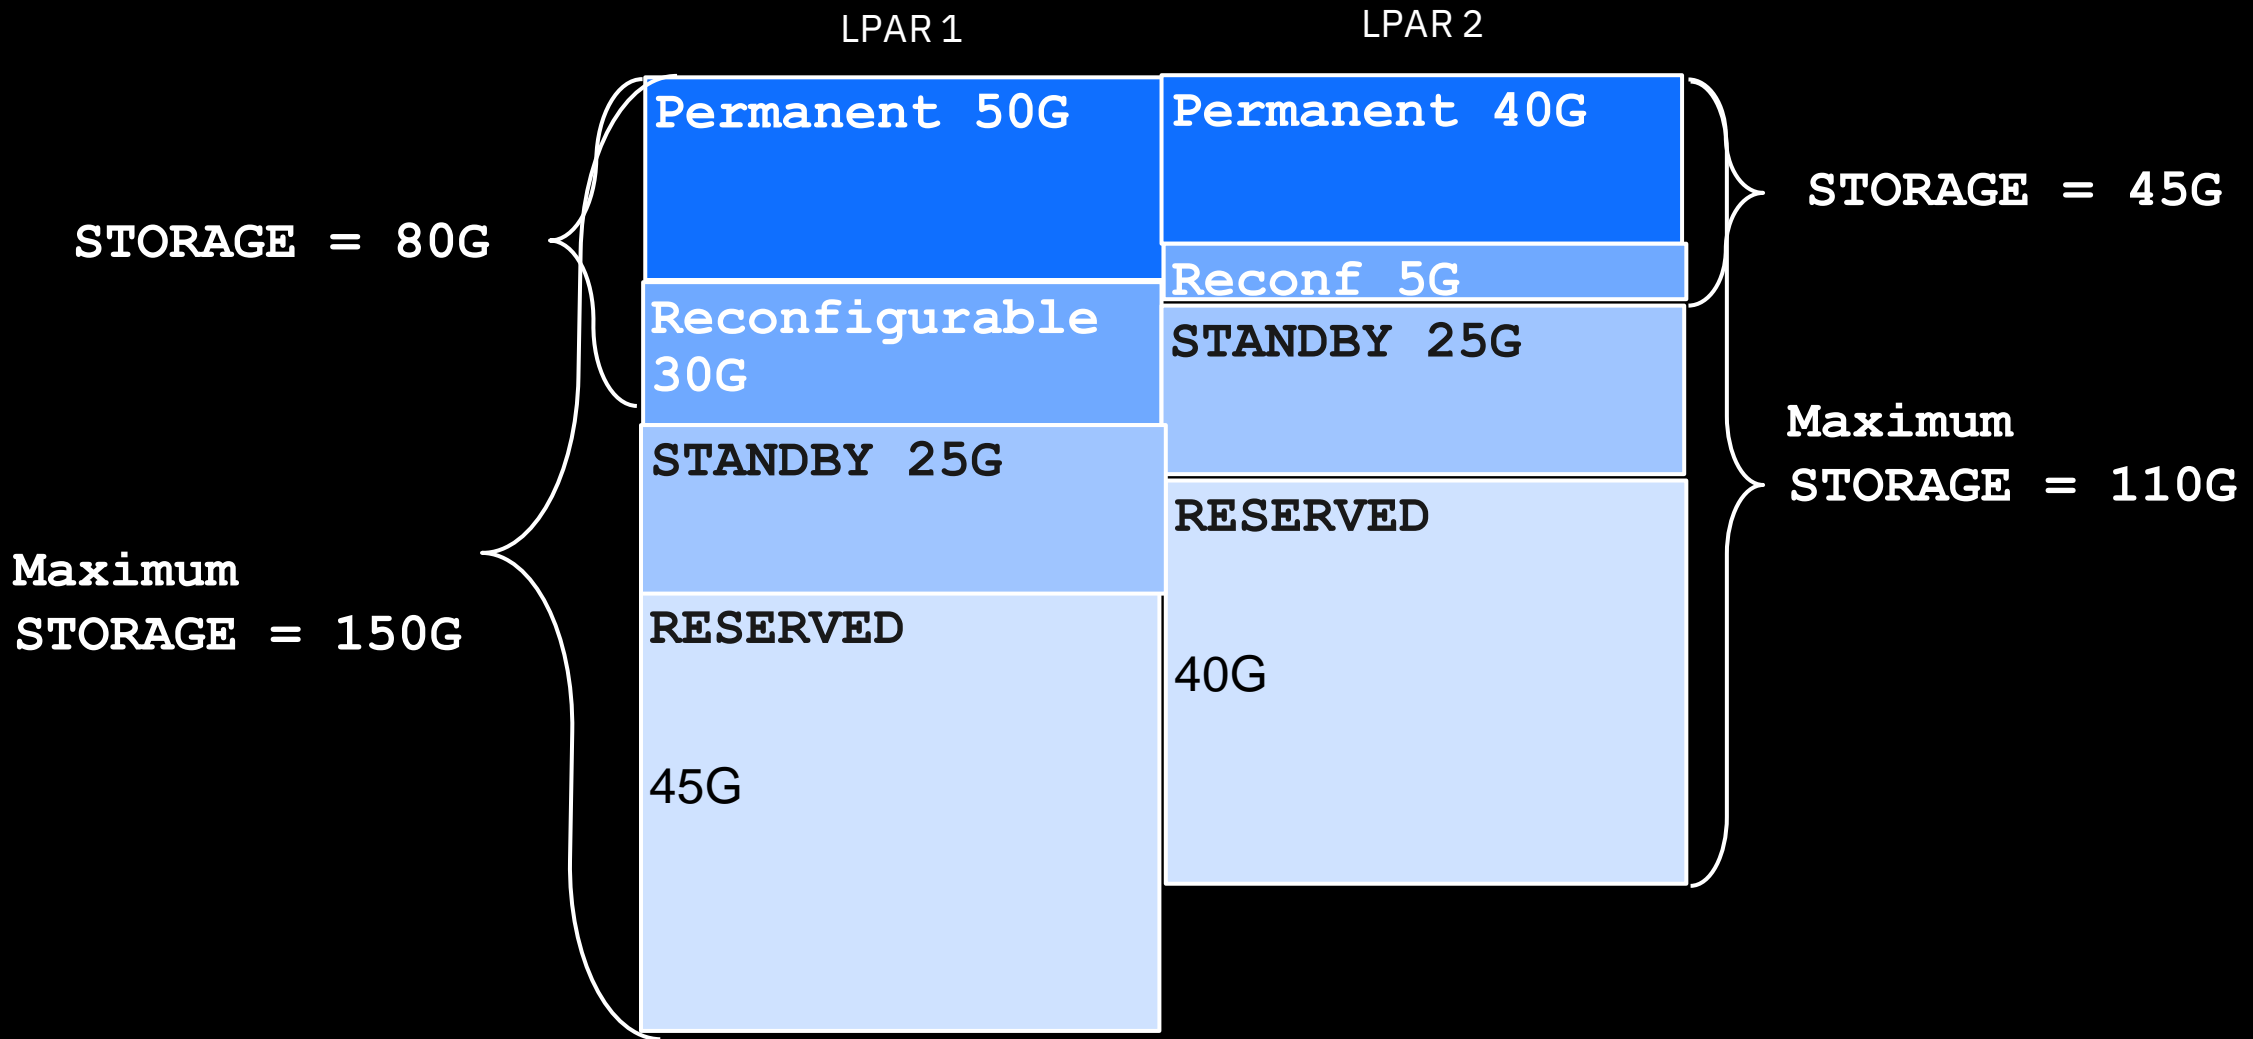
Total storage on the CPC = 150G = 80G + 60G + 10G (STANDBY)

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 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
- ❑ **Not more than 50% of all online storage may be defined as reconfigurable**
- ❑ 4G of permanent storage is an enforced *minimum* and thereafter no more than 50% of storage can be reconfigurable.
- ❑ Storage additions (and reclamations) must be done in multiples of the storage increment size (QUERY STORAGE will show the increment size).

How much reconfigurable storage should I have?

The following situations lend themselves to reconfigurable storage (variable 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 variable 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

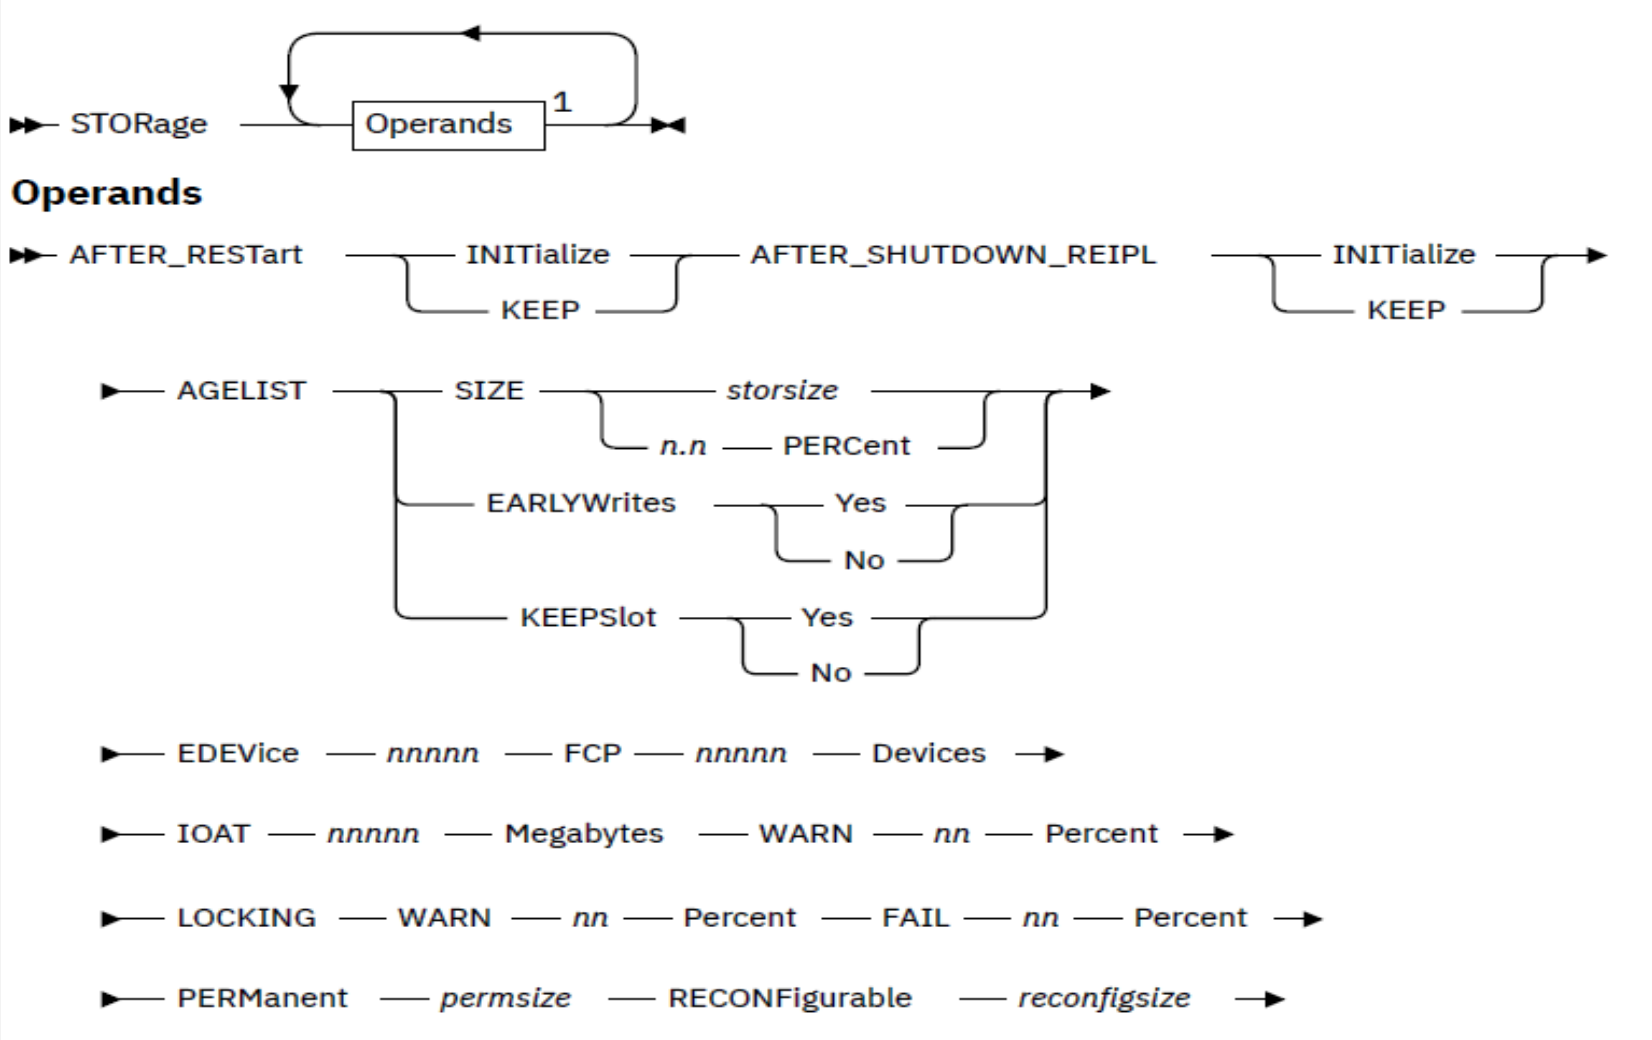
Not more than 50% of all online storage may be defined as reconfigurable.

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
- **CHKRECLM** is a new EXEC that will be made available on the downloads page after the DMD APAR is available

The Nitty-Gritty: Configuration Statements and Commands

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: ▾

Initial:

Reserved:

- You can specify the storage available to your LPAR with the **STORAGE** statements

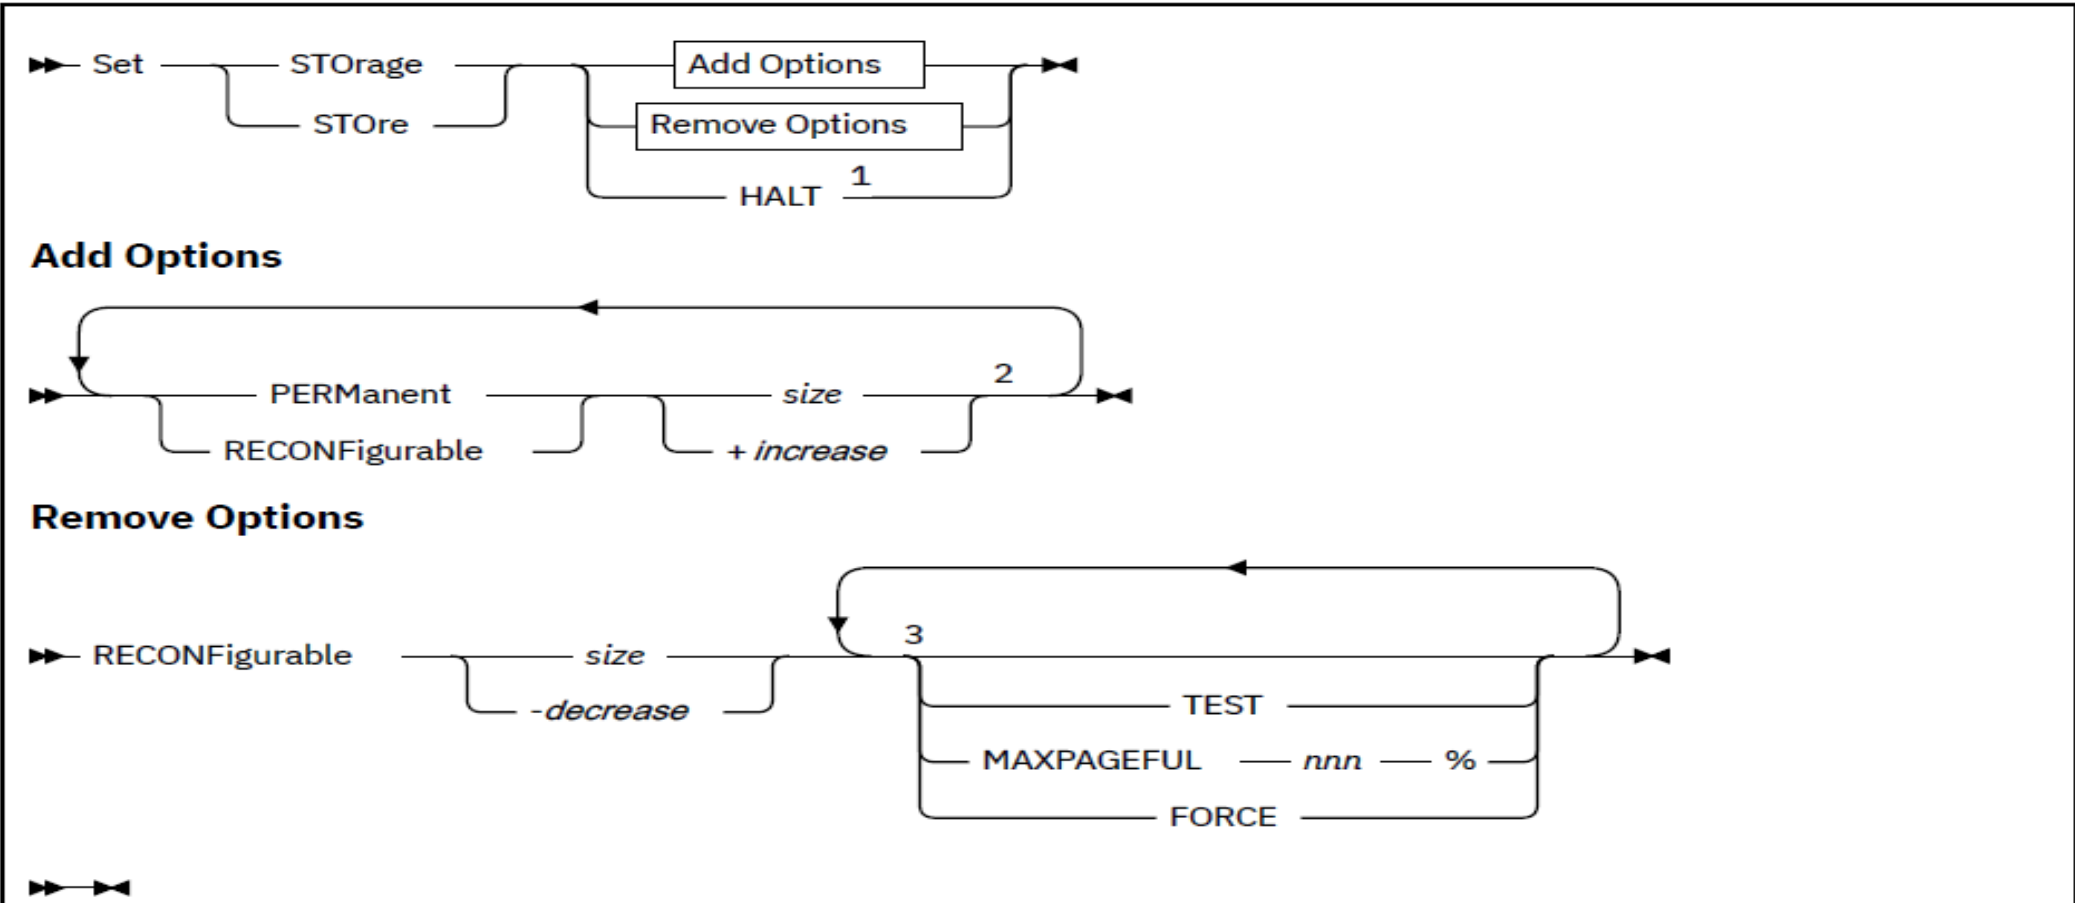
– **STORAGE PERManent 8G**

– **STORAGE RECONFigurable 2G**

... Or maybe

– **STORAGE PERM 4G RECONF 4G**

Dynamic Storage Reconfiguration



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.
- 3 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
```

- 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
- A SET STORE command will be rejected if storage reconfiguration is currently in progress
- A QUERY STORE command 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

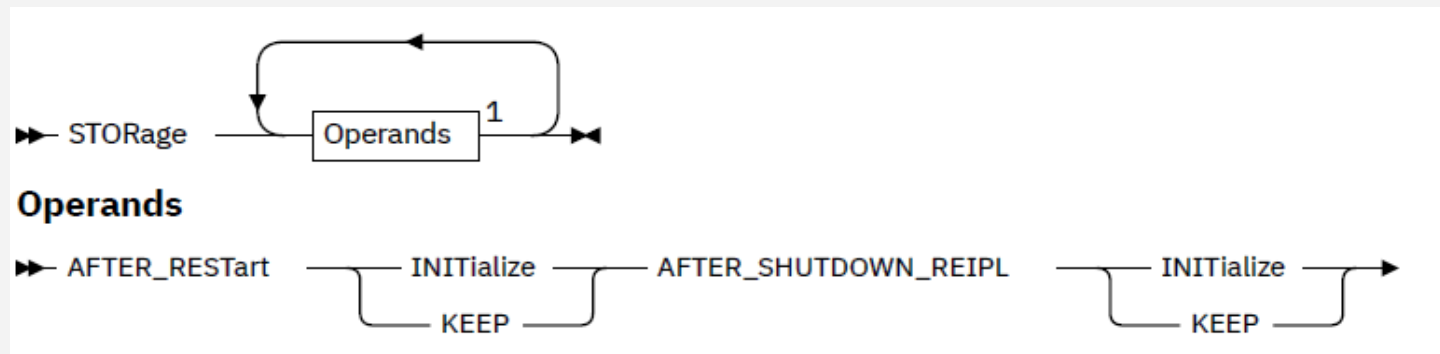
```
SET STOR HALT
```

- 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

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

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 the amount expected to be reclaimed (there is some overhead associated with managing reconfigurable storage)
- 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 if you want to keep them going forward

Paging implications, interactions with other commands, and conclusion

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

Conclusion

An exciting new capability, memory reclamation, is coming soon to z/VM 7.2 systems near you

Available via APAR VM66271

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

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

It includes other enhancements like the ability to set another paging warning threshold and the ability to keep or reset dynamic storage changes on restart

Thank you!

Contact Information:

Jeff Gertner

z/VM Storage Management Development

gertnerj@us.ibm.com