

z/VM – Beyond the Horizon

Frederik Hartmann
z/VM Chief Product Owner
frederik.hartmann@de.ibm.com



Content

3	Why z/VM?
4	z/VM History
5	Initiatives
7	Scalability
14	Usability
18	Focus
20	Discussion

Why z/VM?

Linux on Z

z/VM is the prime hypervisor for data serving or z/OS collocated workload.

Linux on Z and z/VM integration is exceptional and continues to get better.

All Z OSs

z/VM is the only hypervisor on IBM Z supporting all operating systems.

z/OS, z/TPF, z/VM, VSE and naturally Linux on Z are all supported.

Data Serving

Features like FCP passthrough, GDPS support, MVLAG and SSI enable z/VM to be the best place demanding databases and applications.

OCP

z/VM is a great place to host OCP nodes to integrate container-based workloads into an IBM Z environment.

z/VM History

CP-40/CP-67

Released in 1967, CP67 and its research precursor CP-40, were revolutionary time-sharing operating systems for the IBM System/360 Model 67.

It was the first commercially hypervisor in the world.

VM/370

Released in 1972 for the System/370, VM/370 and its different enhancements and versions stayed relevant until 1990.

It introduced features like nested virtualization, a shared file system, and increased supported memory to up to 64 MB.

VM/ESA

Released in 1990 for S/390, VM/ESA was the first release of VM where Linux started to play a major role on the mainframe.


Linux added support for S/390 to kernel 2.2.14 in early 2000.

z/VM

Released in 2000, z/VM is running on on z/Architecture developed and enhanced today.

It supports running 1000s of virtual machines with up to 64 vCPUs and 2TB of memory.

Initiatives



Scalability

The capacity of a system to handle a growing amount of work or to utilize a growing number of resources.




Usability

The extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction



Resiliency

The ability of a system to continue to operate under adverse conditions or stress, even if in a degraded or debilitated state, while maintaining essential operational capabilities.



Hardware Support

The capability of an operating system or application to fully utilize the hardware of the system it is being executed on.

Scalability

The capacity of a system to handle a growing amount of work or to utilize a growing number of resources.

Hardware Resources
from zEC12 in 2014 to
z17 in 2025.

+105%
PUs

Increase from maximum 101 PUs on
the zEC12 to 208 on the z17.

+230%
MIPS

75.000 MIPS on zEC12 to more than
248.000 MIPS on z17.

+2000%
Memory

3 TB on zEC12 to 64TB on z17.

Workload Requirements

Databases

- Db2
- MongoDB
- PostgreSQL
- Oracle

Increase in available memory increases max TPS with constant license costs in many cases.

Languages

- Java
- Python
- Go

Garbage collected languages significantly profit from more memory.

Scalability Targets

Supported Today

Utilization of a full drawer with four z/VM LPARs.

Utilization of full CEC with 16 z/VM LPARs.

Current Hardware Support:

- 40 IFLs with SMT2
- 4 TB Host Memory

Current VM Hardware Support:

- 64 vCPU Support
- 2 TB VM Memory

Minimum Target

Utilization of a full drawer with two z/VM LPARs.

Utilization of full CEC with 8 z/VM LPARs.

Targeted Hardware Support:

- 40 IFLs with SMT2
- 8 TB Host Memory

Targeted VM Hardware Support:

- 64 vCPU Support
- 4 TB VM Memory

Optimum Target

Utilization of a full drawer with one z/VM LPAR.

Utilization of full CEC with 4 z/VM LPARs.

Targeted Hardware Support:

- 60 IFLs with SMT2
- 16 TB Host Memory

Targeted VM Hardware Support:

- 64 vCPU Support
- 4 TB VM Memory

4 KB pages don't cut it

Memory Management based on 4KB pages is limiting memory scalability

- 4TB memory are roughly 1 billion pages
- Every page must be managed individually

Managing these many objects creates a lot of TLB pressure and management overhead

Pushing 4KB pages beyond 4TB will require more and more complex solution for diminishing results.

Solution:

“Persistent” Virtual Machines

- Backed by 2GB regions from a new 2GB region memory pool.
- 2GB regions significantly reduce memory management overhead, roughly by a factor of 2M.
- Improved performance by reduced TLB usage
- New VM type will be limited in functionality to Linux on Z requirements.
- VM relocation not supported
- Mixing allocation pools not supported
- Sufficient memory must be available for VM start

-50% TLB usage
+400% max LPAR size
+400% max VM size

Solution:

2GB region memory pool

- Dividing physical memory into a pool managed in 4KB pages and a pool managed in 2GB regions.
- Size of pools can be configured within reasonable constraints.
- Reclaiming 2GB regions requires the termination of the virtual machine it is allocated to.
- z/VM will always allocate at least one 2GB region for its own use.
- It will not be possible to unassign memory from the 2GB region pool from the LPAR.

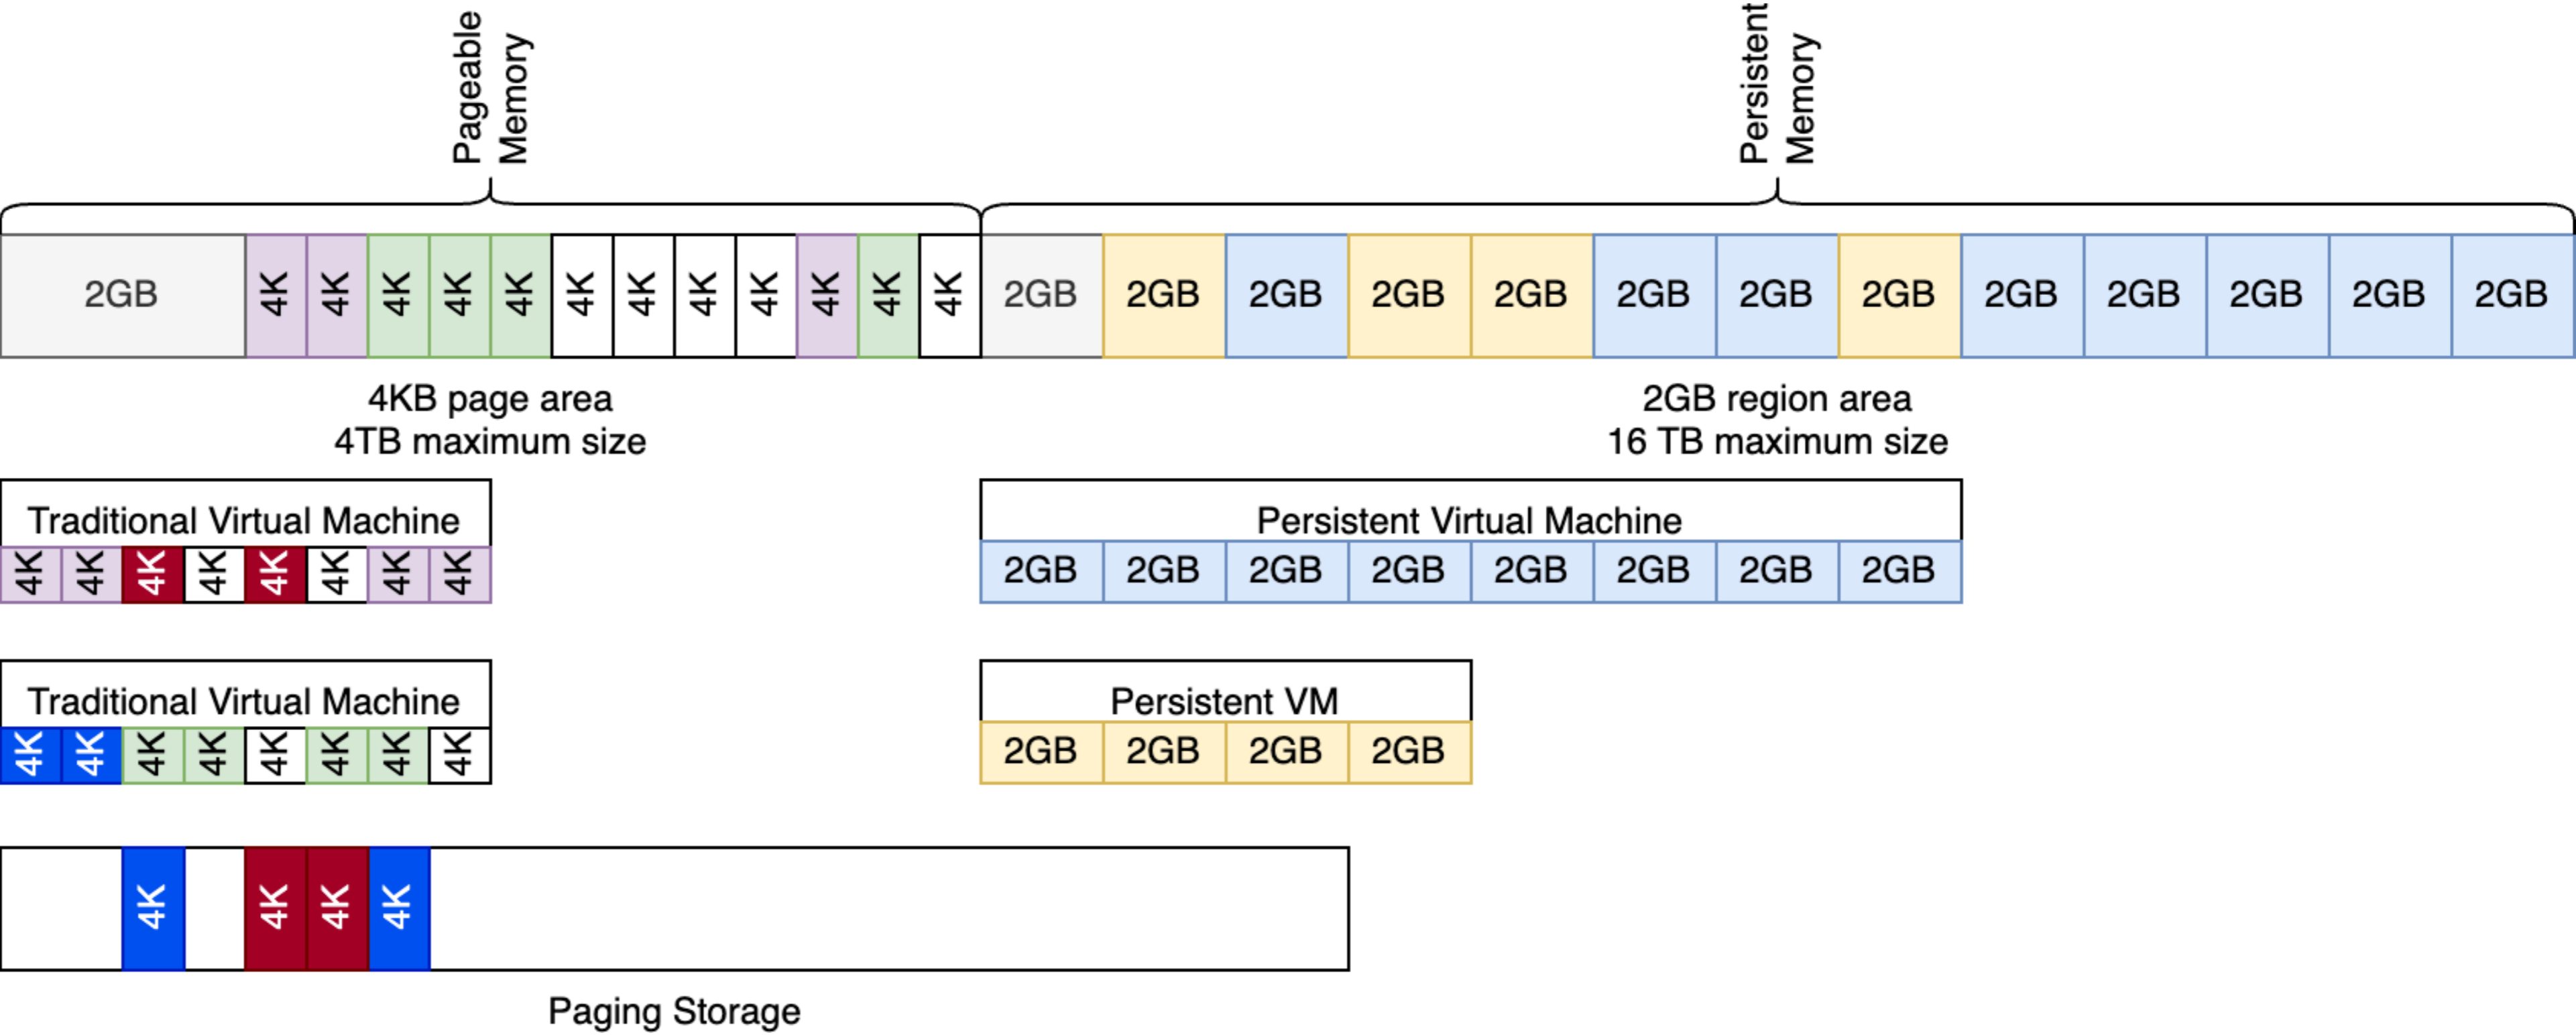
-50% TLB usage
+400% max LPAR size
+400% max VM size

“Persistent” Memory Architecture

4 KB backed traditional virtual machines get their memory from the 4KB area, their memory can be paged out or not instantiated.

2GB backed virtual machines get their memory from the 2GB area, their memory is allocated at their start and will remain fixed.

It will not be possible to start 2GB backed virtual machines if insufficient memory is available.



2GB Region Pool Configuration

The size of the 4KB pageable memory will be customer configurable within the following constraints:

- The minimum size of the 4KB pool cannot be less than 2GB.
- The maximum size of the 4KB pool cannot exceed 4TB.
- At least one 2 GB region remains for the CP owned 2GB region.

-50% TLB usage
+400% max LPAR size
+400% max VM size

Challenge:

Managing memory on 4KB pages is essential for paging & relocation but significantly inhibits memory scalability.

How important is paging & relocation for virtual machines larger than 245-512 GB for you?

Usability

How well a specific user in a specific context can use a product to achieve a defined goal effectively, efficiently and satisfactorily.

Linux Integration

Many required administration utilities are currently only available in CMS and z/CMS.

We plan to make interfaces to essential system operation and administration utilities available through native Linux interfaces.

ICIC Integration

We see ICIC as the strategic solution for z/VM operation and plan to further enhance the integration of z/VM into ICIC.

Additionally, we're looking at ways to simplify the installation and configuration of the z/VM ICIC integration.

Terminology

We plan to better align z/VM with the industry terminology to reduce the entry barrier and confusion.

We plan to gradually phase out the term central storage in favor of real memory to better align with industry terminology.

We plan to gradually phase out the usage of guest in favor of virtual machine to better describe their purpose.

Service Changes

z/VM 7.4 introduced linear service per component, where every released APAR has a prerequisite against the previously released APAR.

We're looking at ways to further improve the service experience in the future.

Challenge:

Modernization without breaking changes has limited effectiveness.

What "legacy" features(370ACCOM, APPC, TSAF, ...) do you rely on?

Question:

What kind of performance management & observability solution do you need?

Do you value simplicity over depth?

Do you value usability over features?

Does your company have a strategic observability solution?

Focus

The most important and most strategic usage of z/VM today is as hypervisor for Linux on Z, z/OS and other Z operating systems.

Focus

Sometimes it's more about what we won't do than what we are doing or plan to do.

z/VM will focus its investment on scalability and usability in the core of its mission:

To be the best Z Hypervisor possible.

Unfortunately, this means that some things that are part of the product today will not be in the future: EREP, LDAP, NFS, ...

We believe that integration into industry standard tooling is of more value than to focus on z/VM specific solutions.

Discussion

Are we on the right track?

Do you believe that Usability/Simplification and Scalability are the most important issues of today?

What do you plan to use z/VM for in the future?

Where would you place your investments into resiliency?

Application or Infrastructure?

How much memory will your biggest virtual machines need?

Do you agree that virtual machine relocation will become less relevant in the future?