45 (Official) Years of Virtualization:

How z/VM has evolved over 50 years to grow with your business needs

Brian Hugenbruch
IBM z/VM Dev Lab Endicott

Wilhelm Mild
IBM R&D Lab Germany

Agenda

- VM – today - THE global virtualization for IBM Z
- VM – beginnings - CP-67 and VM/370
- VM – mascot - the teddy bear history
- Hosting: VM – it's more than virtualization - Profs, Rexx, Pipelines – life as a CMS application developer
- Hosting: VSE – the lovely haven with many happy customer experiences
- Hosting: Linux – where Linux could multiply like rabbits
- Global virtualization – Disk, Virtual Switch, SSI, Clouds in a Box
- Tomorrow's virtualization started 45 years ago and continuous today and tomorrow with your needs

IBM z/VM 6.4

- A release born from customer feedback
  - z Systems Business Leaders Council (zBLC)
  - SHARE dialogues
  - IBM internal T3s (Teach the Teacher)

- Prioritizations set by customers and adjusted by IBM resources and skills

- Two major areas:
  - Technical enhancements that continue to improve TCO and bring direct value
  - Improved quality of life for z/VM system programmers

- New Architecture Level Set (ALS)
  - z196 and z114 or newer
  - Drops z10 EC and BC support

Major Waves of Technology (or, the more things change …)

Back Office Computing

Client Server PC - 1981

World Wide Web and eBusiness

Confluence of Social, Mobile, Cloud, Big Data / Analytics

60’s

80’s

90’s

We are here

Early VM Family

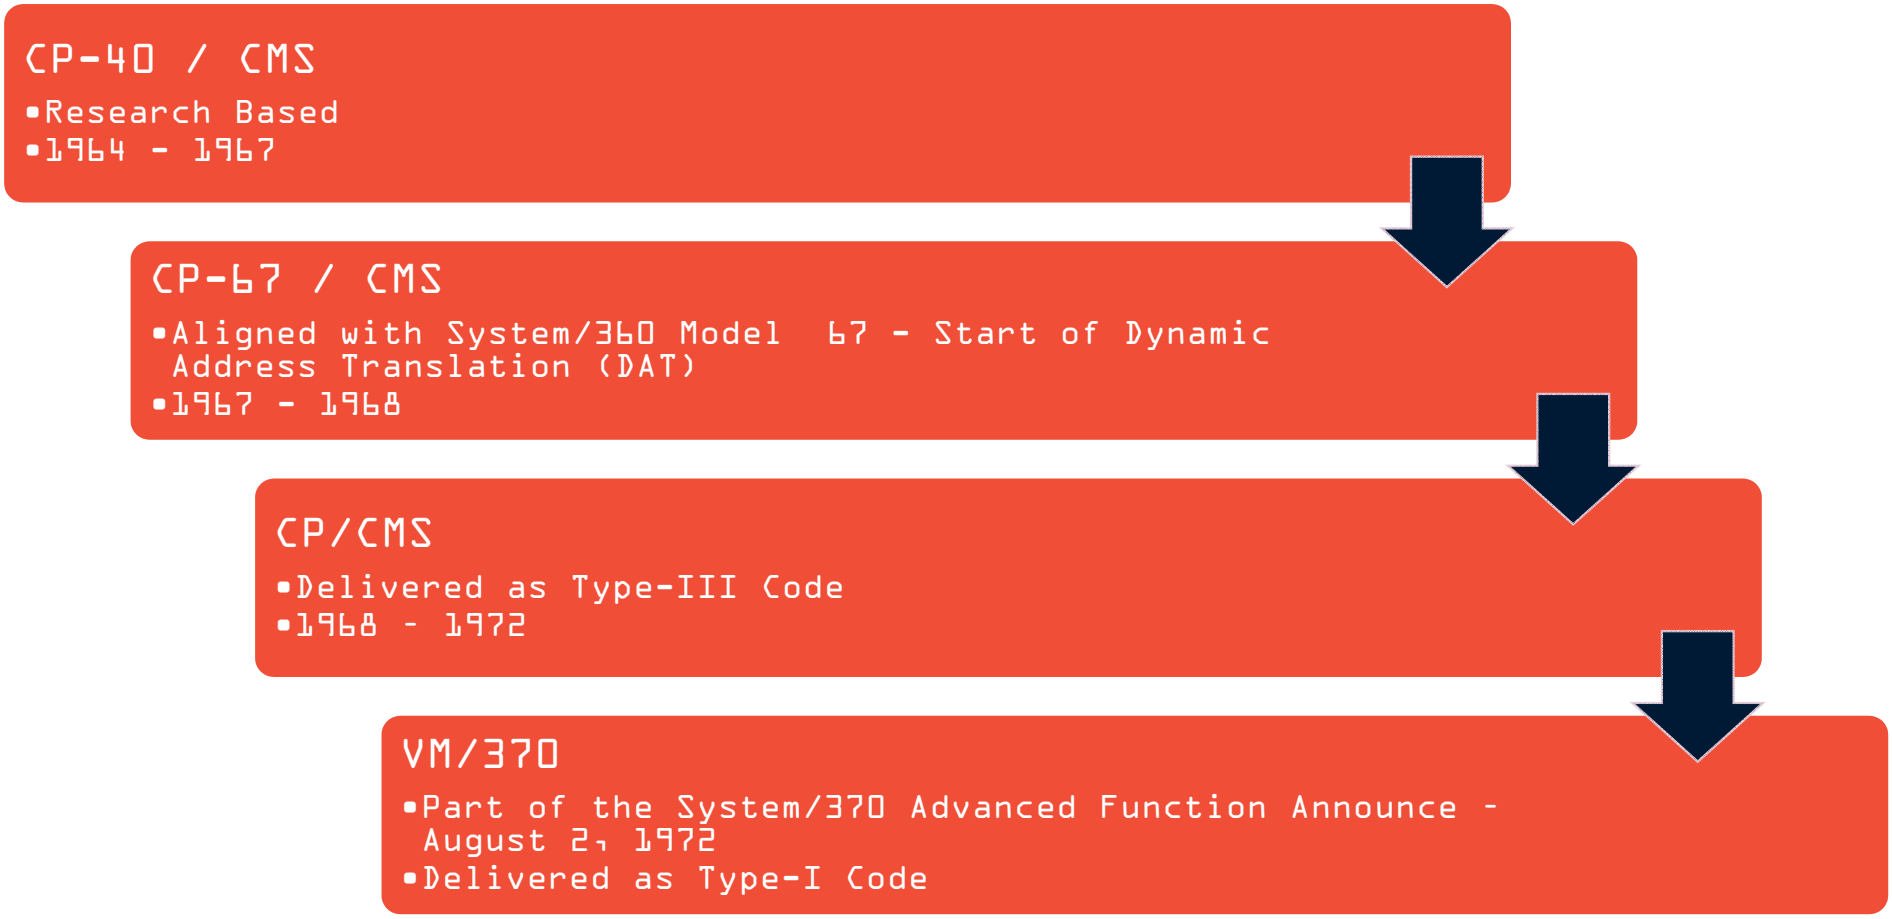

CP-67/CMS releases
1967-1972

- May 1968: Version 1 was released to eight installations
  - It was made available as part of the IBM Type-III Library in June
  - Two time-sharing businesses were launched based on the resale of CP-67/CMS: National CSS and IDC
  - These ventures drew attention to the viability of CP-67/CMS, the S/360-67, and virtual memory
  - As of April 1969 CP-67/CMS had been installed at fifteen sites
- June 1969: Version 2 was released
- November 1971: Version 3.1 was released, capable of supporting sixty CMS users on a S/360-67 (included the 1st "Wheeler Scheduler")
- Early 1972: Version 3.2 was released, a maintenance release with no new functions
  - CP-67 was now running on 44 processors, ¼ of which were inside IBM

CP-67 (remember when things weren't entirely virtual?)

1967-1972

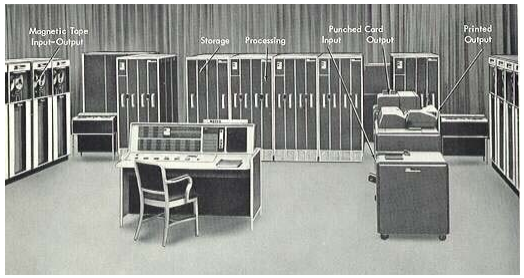

System/360 Model 67

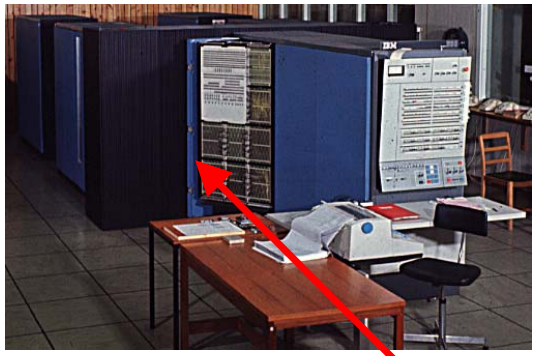

"DAT box"

University of Newcastle Upon Tyne

IBM 1052

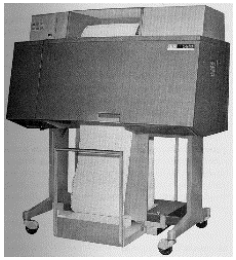

IBM 1403

IBM 2540

IBM 2741

IBM 2314

IBM 2311

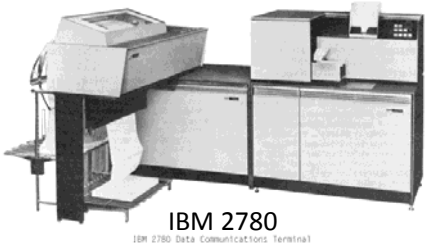

Which brings us to VM/370

The original 3 page announcement letter!

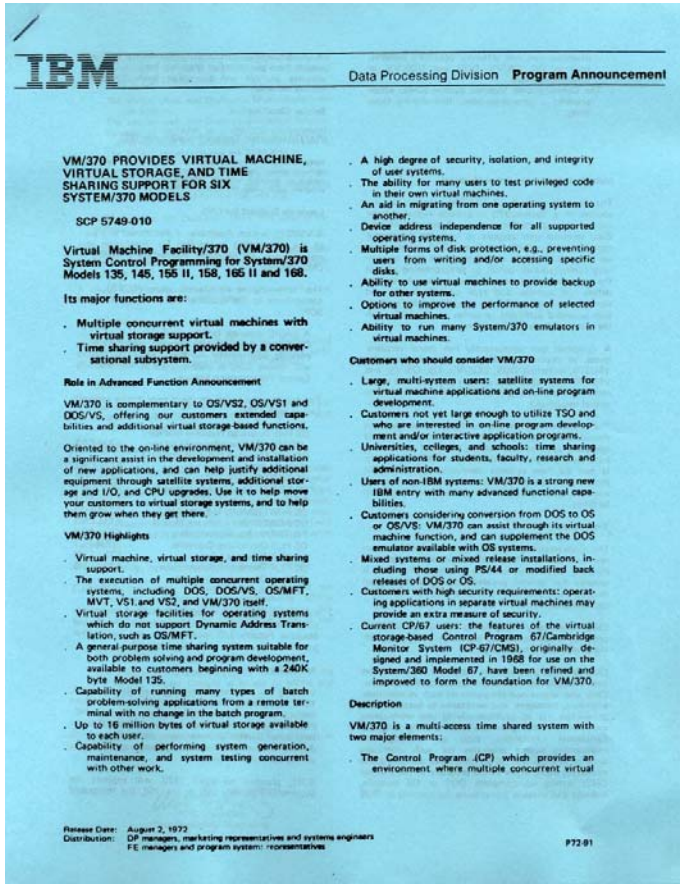

IBM Data Processing Division **Program Announcement**

VM/370 PROVIDES VIRTUAL MACHINE, VIRTUAL STORAGE, AND TIME SHARING SUPPORT FOR SIX SYSTEM/370 MODELS

SCP 5749-010

Virtual Machine Facility/370 (VM/370) is System Control Programming for System/370 Models 135, 145, 155 II, 158, 165 II and 168.

Its major functions are:

- **Multiple concurrent virtual machines with virtual storage support.**
- **Time sharing support provided by a conversational subsystem.**

Role in Advanced Function Announcement

VM/370 is complementary to OS/VS2, OS/VS1 and DOS/VS, offering our customers extended capabilities and additional virtual storage-based functions.

Oriented to the on-line environment, VM/370 can be a significant assist in the development and installation of new applications, and can help justify additional equipment through satellite systems, additional storage and I/O, and CPU upgrades. Use it to help move your customers to virtual storage systems, and to help them grow when they get there.

VM/370 Highlights

- Virtual machine, virtual storage, and time sharing support.
- The execution of multiple concurrent operating systems, including DOS, DOS/VS, OS/MFT, MVT, VS1 and VS2, and VM/370 itself.
- Virtual storage facilities for operating systems which do not support Dynamic Address Translation, such as OS/MFT.
- A general-purpose time sharing system suitable for both problem solving and program development, available to customers beginning with a 240K byte Model 135.
- Capability of running many types of batch problem-solving applications from a remote terminal with no change in the batch program.
- Up to 16 million bytes of virtual storage available to each user.
- Capability of performing system generation, maintenance, and system testing concurrent with other work.
- A high degree of security, isolation, and integrity of user systems.
- The ability for many users to test privileged code in their own virtual machines.
- An aid in migrating from one operating system to another.
- Device address independence for all supported operating systems.
- Multiple forms of disk protection, e.g., preventing users from writing and/or accessing specific disks.
- Ability to use virtual machines to provide backup for other systems.
- Options to improve the performance of selected virtual machines.
- Ability to run many System/370 emulators in virtual machines.

Customers who should consider VM/370

- Large, multi-system users: satellite systems for virtual machine applications and on-line program development.
- Customers not yet large enough to utilize TSO and who are interested in on-line program development and/or interactive application programs.
- Universities, colleges, and schools: time sharing applications for students, faculty, research and administration.
- Users of non-IBM systems: VM/370 is a strong new IBM entry with many advanced functional capabilities.
- Customers considering conversion from DOS to OS or OS/VS: VM/370 can assist through its virtual machine function, and can supplement the DOS emulator available with OS systems.
- Mixed systems or mixed release installations, including those using PS/44 or modified back releases of DOS or OS.
- Customers with high security requirements: operating applications in separate virtual machines may provide an extra measure of security.
- Current CP/67 users: the features of the virtual storage-based Control Program 67/Cambridge Monitor System (CP-67/CMS), originally designed and implemented in 1968 for use on the System/360 Model 67, have been refined and improved to form the foundation for VM/370.

Description

VM/370 is a multi-access time shared system with two major elements:

- The Control Program (CP) which provides an environment where multiple concurrent virtual machines can run different operating systems, such as OS, OS/VS, DOS and DOS/VS, in time-shared mode.
- The Conversational Monitor System (CMS) which provides a general-purpose, time-sharing capability.

Multiple Concurrent Virtual Machines

The control program of VM/370 manages the resources of a System/370 to provide virtual storage support through implementation of virtual machines. Each terminal user appears to have the functional capabilities of a dedicated System/370 computer at his disposal. Multiple virtual machines may be running conversational, batch, or teleprocessing jobs at the same time on the same real computer. A user can define the number and type of I/O devices and storage size required for his virtual machine application provided sufficient resources are available with the real machine's configuration.

A customer can concurrently run many versions, levels, or copies of IBM operating systems under VM/370, including DOS, DOS/VS, OS, OS/VS, and VM/370 itself. (See sales manual pages for the major restrictions pertaining to the operation of systems in virtual machines.)

The capability of running multiple virtual machines should assist the customer in scheduling multiple operating systems and various mixes of production jobs, tests, program maintenance, and FE diagnostics. It can aid new systems development, reduce the problems of converting from one operating system to another, and provide more economical backup facilities.

Time Sharing

The Conversational Monitor System (CMS) component of the VM/370 system provides a general-purpose, conversational time sharing facility that is suitable for general problem solving and program development, and can serve as a base for interactive applications.

CMS, specifically designed to run under VM/370, provides broad functional capability while maintaining a relatively simple design.

CMS can help programmers become more productive and efficient by reducing unproductive wait time. CMS also allows non-programmers such as scientists, engineers, managers, and secretaries to become more productive via its problem-solving and work-saving capabilities. CMS gives the user a wide range of functional capabilities, such as; creating and maintaining source programs for such operating systems as DOS and OS on CMS disks; compiling and executing many types of OS programs directly under CMS; setting up complete DOS or OS compile, linkedit and execute job streams for running in DOS or OS virtual machines; and transferring the resultant output from those virtual machines back to CMS for selective analysis and correction from the user's remote terminal.

Service Classification

VM/370 is System Control Programming (SCP).

Note: VM/370 does not alter or affect in any way the current service classification of any IBM operating system, language, program product, or any other type of IBM program while under the control of VM/370.

Language Support for CMS

A VM/370 System Assembler is distributed as a part of the system and is required for installation and maintenance. All necessary macros are provided in CMS libraries.

The following is distributed with VM/370 as a convenience to the customer but is not part of the SCP.

A BASIC language facility consisting of the CALL-OS BASIC (Version 1.1) Compiler and Execution Package adapted for use with CMS. This facility will receive Class A maintenance by the VM/370 Central Programming Service.

The following program products may also be ordered for use with CMS:

OS Full American National Standard COBOL V4 Compiler and Library	5734-CB2
OS Full American National Standard COBOL V4 Library	5734-LM2
OS FORTRAN IV (G1)	5734-F02
OS FORTRAN IV Library Mod I	5734-LM1
OS Code and Go FORTRAN	5734-F01
OS FORTRAN IV H Extended	5734-F03
OS FORTRAN IV Library Mod II	5734-LM3
FORTRAN Interactive Debug	5734-F05
OS PL/I Optimizing Compiler	5734-PL1
OS PL/I Resident Library	5734-LM4
OS PL/I Transient Library	5734-LM5
OS PL/I Optimizing Compiler and Libraries	5734-PL3

Further details on language support and execution-time limitations appear in the manual *IBM Virtual Machine Facility/370: Introduction*, and in the *Program Product* section of the sales manual.

Availability

VM/370 has a planned availability of November 30, 1972, supporting the Dynamic Address Translation facility on the System/370 Models 135 and 145. Planned support for certain advanced VM/370 facilities, other System/370 machines, and additional I/O devices will be via Independent Component Releases on the dates shown below.

ICR1, planned for April 1973, will support the System/370 Models 155 II, the 158, the Integrated File Adapter Feature (4655) for 3330 Model 1 and 3333 Model 1 on the Model 135, and the following additional VM/370 facilities:

- The Virtual=Real and Dedicated Channel performance options.
- The virtual and real Channel-to-Channel Adapter.
- Support of OS/ASP in a VM/370 environment, effective with the availability of ASP Version 3.
- The 3811 Control Unit and the 3211 Printer.

ICR2, planned for August 1973, will support the CMS Batch Facility, the Model 168, and the Integrated Storage Controls (ISCs) for the 158 and 168.

ICR3, planned for December 1973, will support the 165 II.

See the respective program product announcement letters for planned availability of the program products for CMS.

Note: VM/370 requires the system timing facilities (i.e., the Clock Comparator and the CPU Timer).

Maintenance

Maintenance for VM/370 Release 1 will be provided by the VM/370 Central Programming Service until nine months after the next release of VM/370.

Education

See Education Announcement Letter E72-14 for details of VM/370 Introduction (no charge) and additional educational plans.

Publications

IBM Virtual Machine Facility/370: Introduction (GC20-1800), is available from Mechanicsburg. Other manuals to be available at a later date include logic manuals, as well as planning, system generation, command language, system operator, terminal user, and programmer guides. Titles and form numbers will be announced in a future Publications Release Letter (PRL).

Reliability, Availability and Serviceability (RAS)

VM/370 provides facilities which supplement the reliability, availability, and serviceability (RAS) characteristics of the System/370 architecture. See the sales manual or the introduction manual for details.

MINIPERT

VM/370 planning information is available in the MINIPERT Master Library as an aid to selling and installing System/370.

No RPQs will be accepted at this time.

Detailed information on the VM/370 system is in sales manual pages.

W. W. Livingston
Vice President, Marketing

Release Date: August 2, 1972
Distribution: DP managers, marketing representatives and systems engineers
FE managers and program system representatives

P72-91

VM/370 – 5749-010 Release 1 content

- S/370 was announced in June 1970, but these were not announced as being virtual storage capable
- Virtual storage for S/370 was announced on August 2, 1972 with OS/VS, DOS/VS, VM/370
  - VM/370 R1 was available in November 1972 with support for the S/370-135 and S/370-145
  - VM/370 R1 ICR1 (Independent Component Release) was planned for April 1973 with support for the S/370-155 II and S/370-158 and CTCs
  - VM/370 R1 ICR2 was planned for August 1973 with support for the S/370-168 and CMS Batch
  - VM/370 R1 ICR3 was planned for December 1973 with support for the S/370-165 II

- Remote Spooling Communications Subsystem (RSCS)
  - CPREMOTE did not provide a complete inter-system file transfer solution
  - SCNODE was built a replacement using a subsystem supervisor called MSUP and the early network was called SCNET
  - With VM/370, enhancements were made to the spool and hypervisor to add interfaces for a more robust solution
    - The TAG command and interfaces provided routing information in the spool files
  - The RSCS component of VM/370 was released in 1975
  - RSCS was enhanced to support the NJE protocols and was released as the VNET PRPQ in 1976, which later became the RSCS V1 product
  - VNET was the name of the internal network and BITNET was the name of the external academic network, both of which used RSCS

Evolving VM Product Family

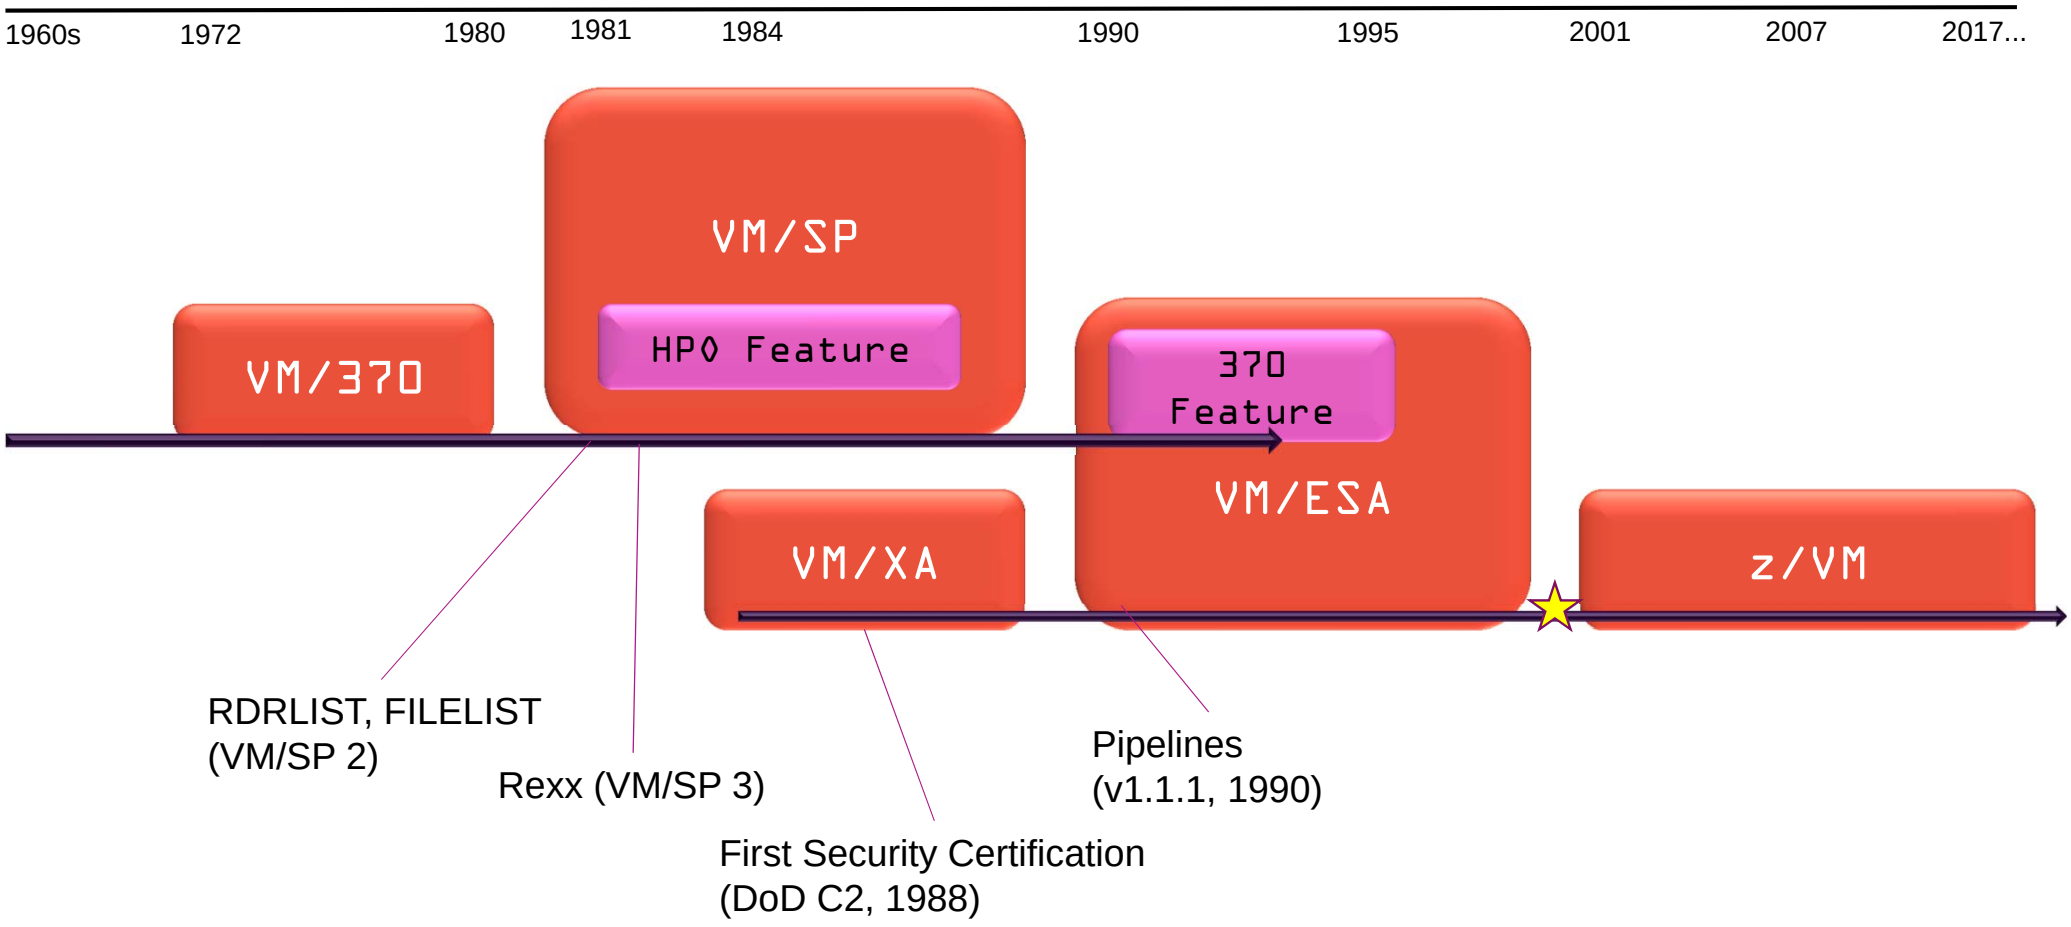

SHARE, VM, and the teddy bear mascot history

- The MVS Group had the turkey as their mascot
  - Changed in the early 1980s to the eagle
- At SHARE 60 in 1983 the VM Group decided to identify newcomers with yellow stickers and old timers with blue stickers, but no one could remember which was which
- Carol Jobusch bought a few hundred teddy bear stickers to identify the “warm, cuddly” old timers, and a mascot was born!

Virtual Machine Technology with z/VM - more than partitioning

- A *Virtual Machine* simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.

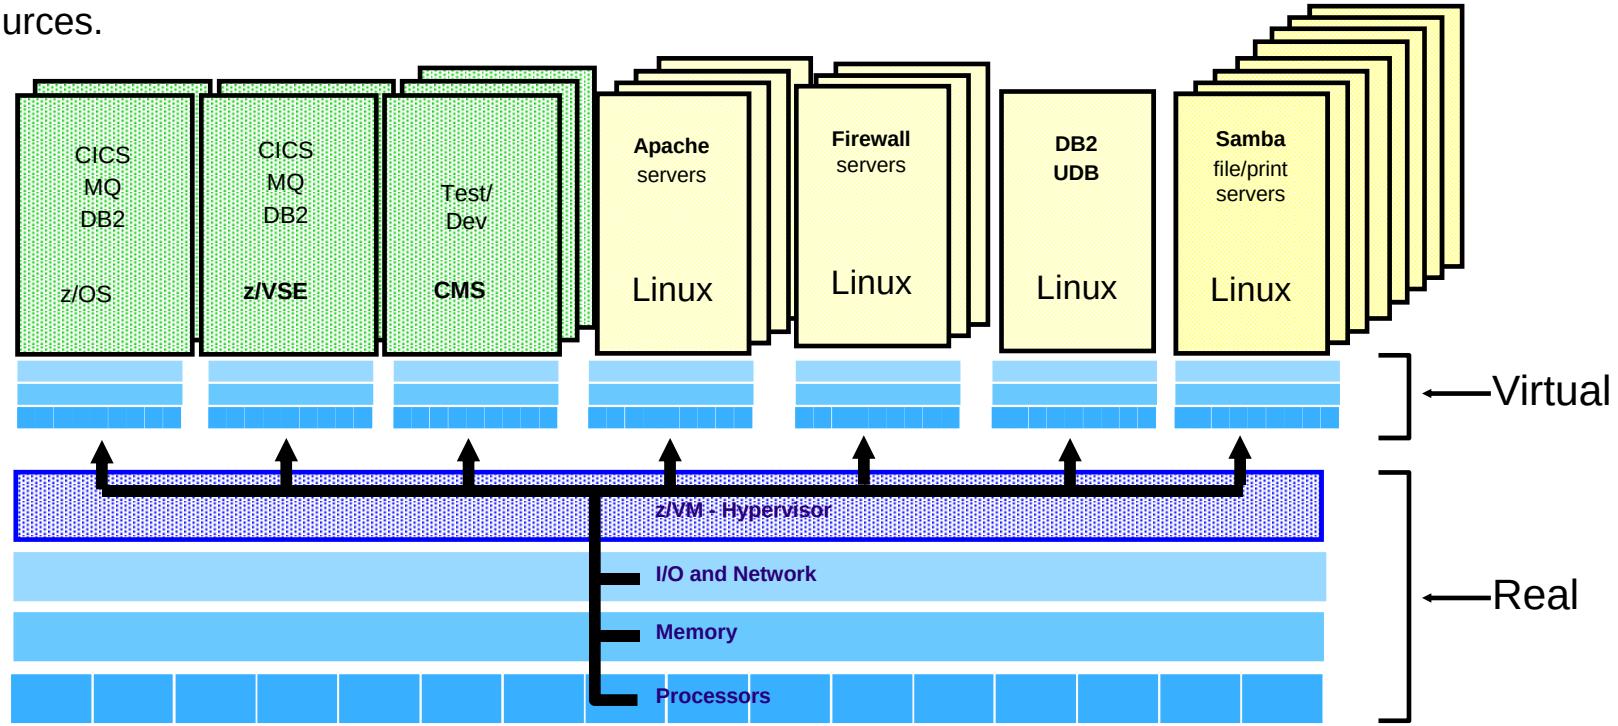

- IBM Z provides the unique capability to run hundreds of Virtual servers on one machine
- Resource sharing and virtualization are key features and provide unmatched flexibility
- Accounting of virtual resources (CPU, I/O, Network,...)

PROFS

- Late in 1981, IBM released the PROFS PRPQ, which had been developed jointly by AMOCO and IBM
- THE collaboration tool that was born within the Virtualization innovations
- Many releases were made available (1983-1997), some of which were:
  - PROFS V1R1 was released in June 1983
  - PROFS V2R1 was released in December 1985
  - PROFS Extended Mail, supporting connections to the Internet, was released in December 1987
  - OfficeVision/VM (aka PROFS V3) was released in October 1989
- By 1987, there were said to be a million PROFS users outside IBM, and IBM itself had become heavily dependent on PROFS
- There are customers using OV/VM today!

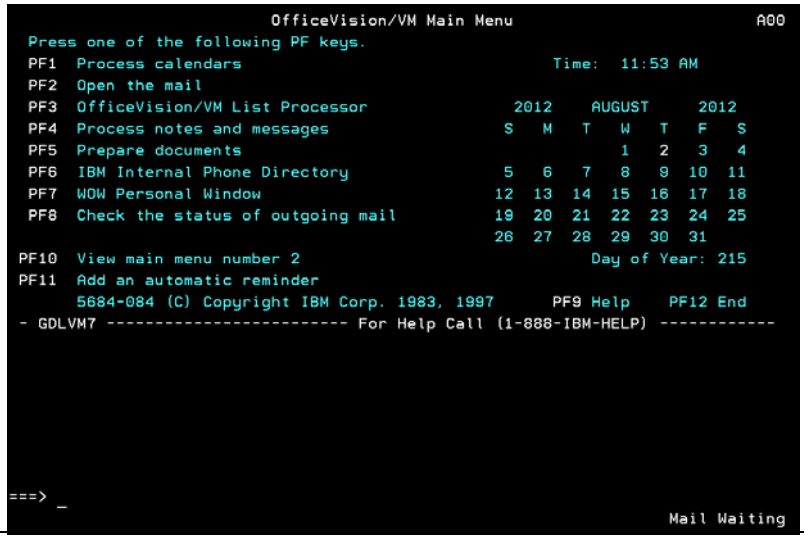

REXX

- REXX (originally REX) was designed and first implemented as an ‘own-time’ project between March 20, 1979 and mid-1982 by Mike Cowlishaw of IBM, originally as a scripting programming language to replace the languages EXEC and EXEC 2
- Distributed internally over VNET, REX was quickly adopted across the internal IBM VM community
- REXX was also intended by its creator to be a simplified and easier to learn version of the PL/I programming language
- It was first described in public at the SHARE 56 conference in Houston, Texas in 1981 where customer reaction, championed by Ted Johnston of SLAC, led to it being shipped in VM/SP R3

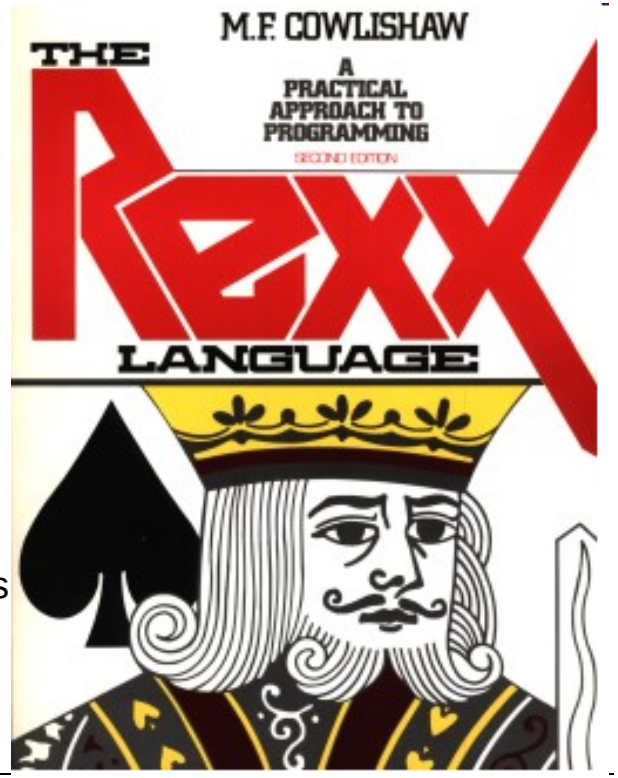

Huge success – innovation with virtualization automation in mind:

- The success led to the port of REXX to multiple platforms and exists today with modern technologies being Object Oriented and automates large virtualized enterprises

z/VM Version 3

- V3.1 announced 2000-10-03, GA 2001-02-23, EOS 2005-12-31
  - –Enabling 64-bit guest operating systems
  - –Real storage constraint relief
  - –Native FlashCopy support for Enterprise Storage Server
  - –Announced as part of the zSeries announcements with the z900, z/OS, and z/VSE
  - –Last MLC version of VM

z/VSE Strategy bases on z/VM Mode LPAR with Linux on z

Hybrid Environment leveraging z/VSE, z/VM, and Linux on z

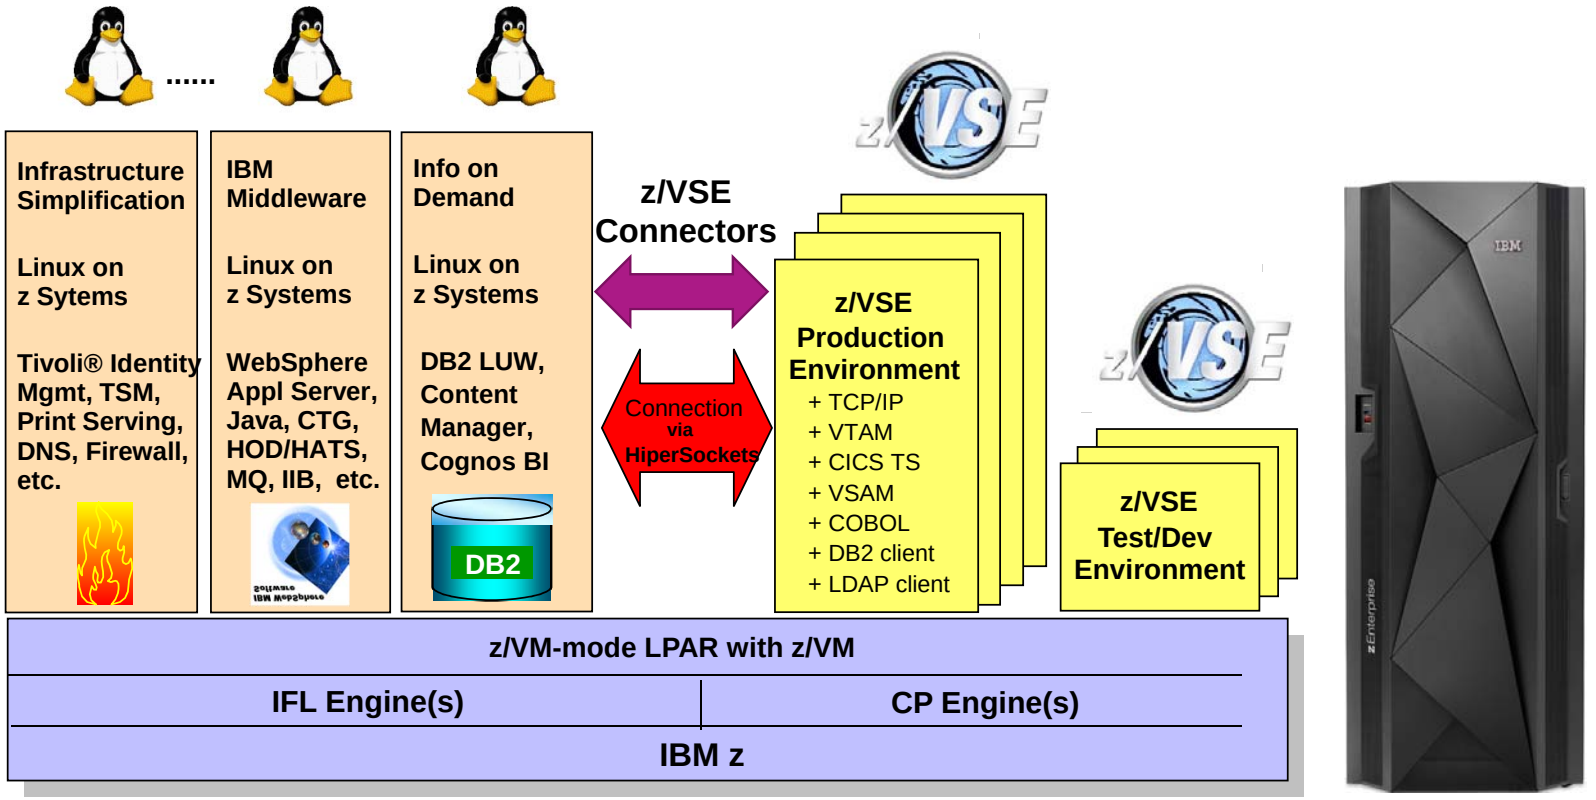

March of the Penguins

To: linas@linas.org
Subject: IBM 390 and Linux
From: Alan Cox <alan@lxorguk.ukuu.org.uk>
Date: Thu, 16 Dec 1999 01:34:47 +0000

They finally delivered code. A decent looking SMP kernel, console and some networking stuff. Glibc, gcc, binutils, gdb patches.

The kernel stuff is in 2.2.14pre14, I'll forward you the other patches if you want.

Alan

Figure 5: Skunkworks No More

Penguins got proud !
…….. And started multiply like Rabbits .

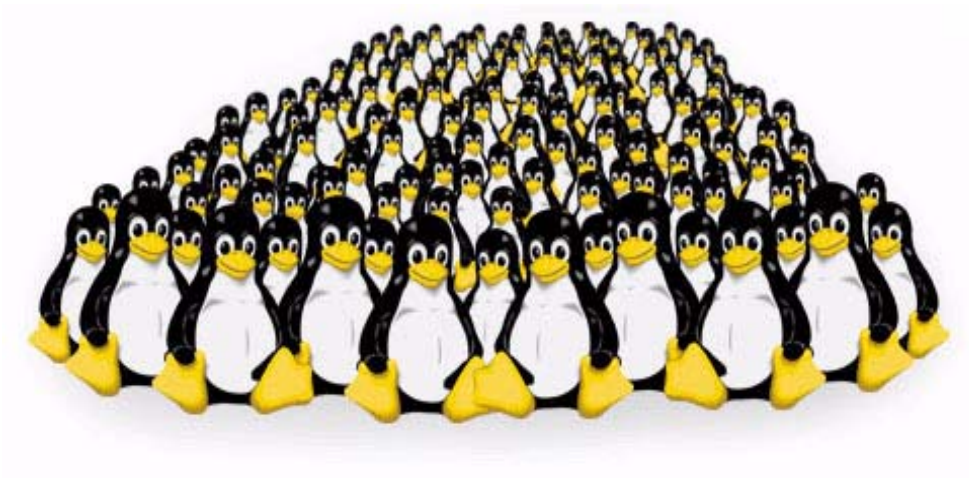

IBM Virtualization with z/VM

- **Vertical virtualization** - Grow workloads without growing number of virtual guest machines
  - one guest can be increased by allocating more resources (CPUs, memory)
- **Horizontal virtualization** – for isolation between servers
  - isolation of guests in a network
  - High availability for applications
- **Dynamically** add, remove and shift physical resources to optimize business results

IBM extreme Virtualization with z/VM on IBM Z

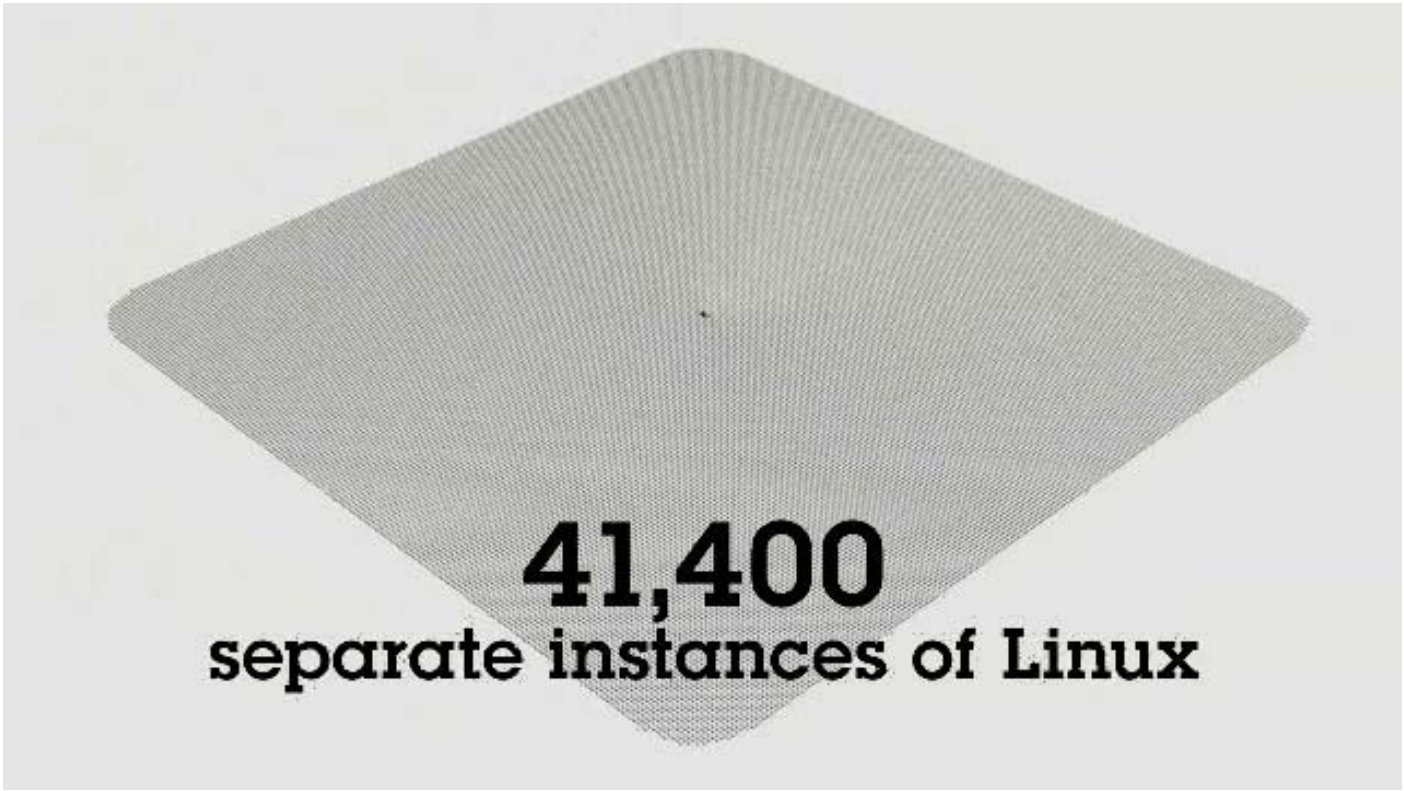

Effective Virtualization with Linux on z and z/VM shared memory

Linux Shared Memory Exploitation for many Virtual machines
z/VM Discontiguous Saved Segments (DCSS)

- **DCSS support is Data-in-Memory technology**
  - Share a single, real memory location among multiple virtual machines
  - Can reduce real memory utilization
- **Use Cases:**
  - **As fast Swap device**
  - **For sharing read only data**
  - **For sharing code (e.g. program executables/libraries)**
- **The large DCSS allows the installation of a full middleware stack in the DCSS (e.g. WebSphere, DB2, etc)**
- **The DCSS becomes a consistent unit of one software level**

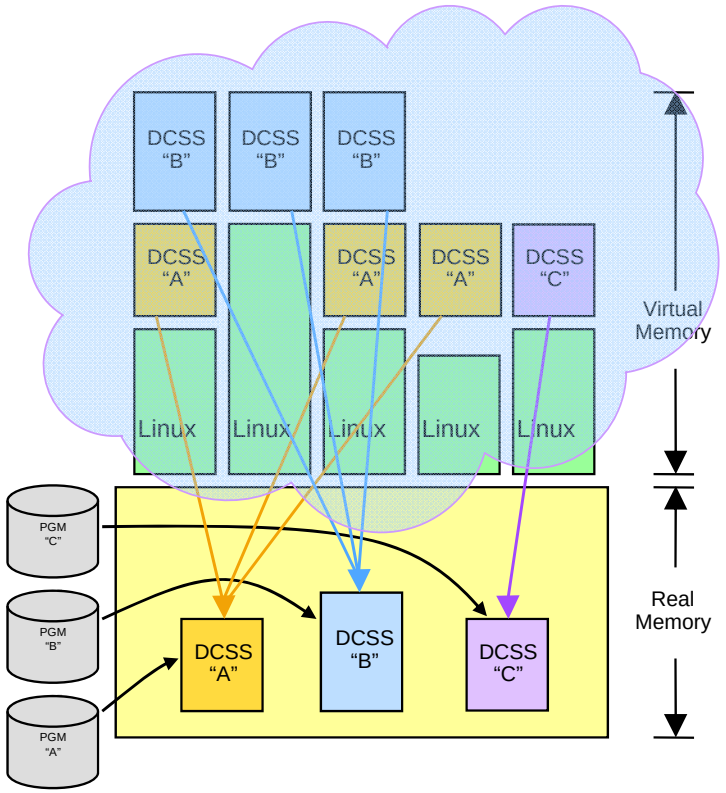

http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03186USEN.PDF

Global Virtualization – with z/VM on IBM Z

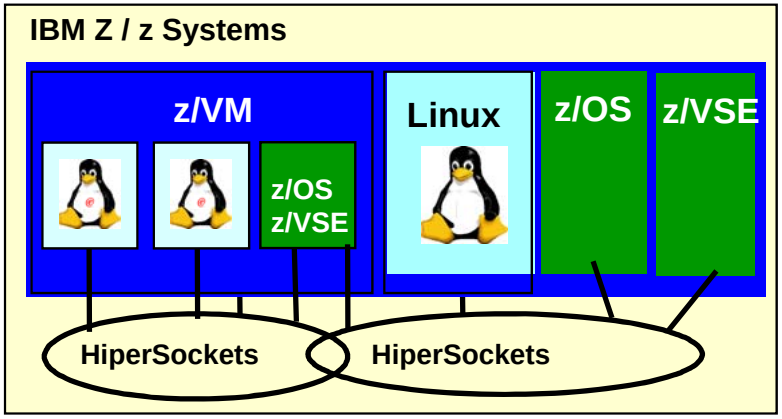

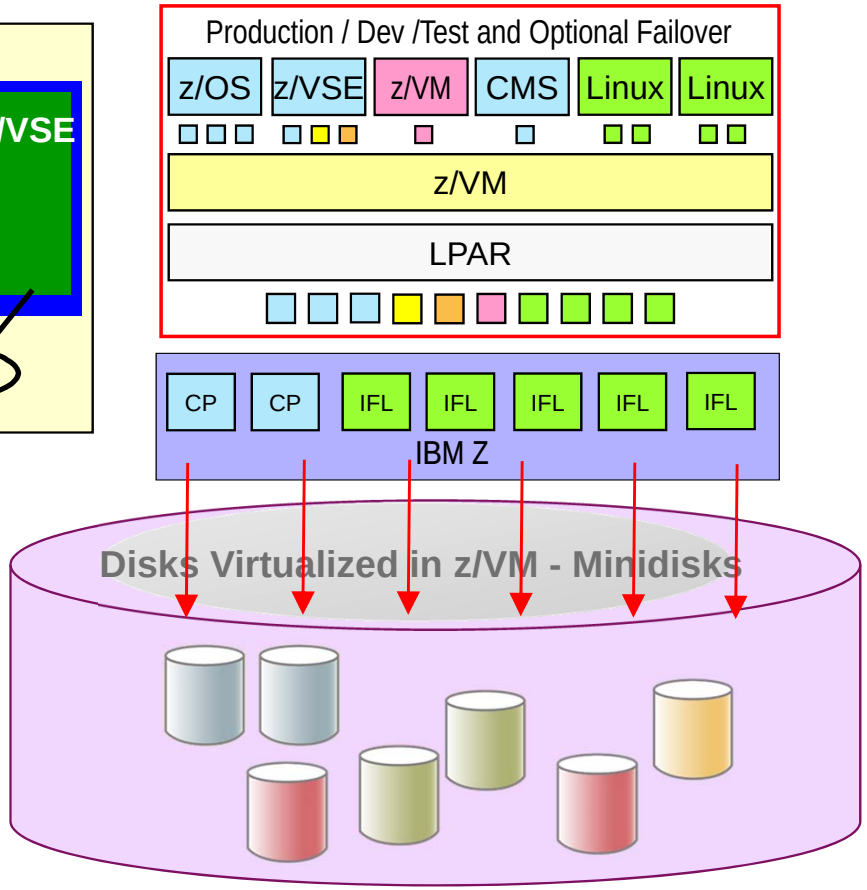

- Network Virtualization
- Memory Virtualization
- Processor Virtualization
- System Virtualization
- Disk Virtualization

22

3D Scalability – non-disruptive with z/VM for operational efficiency

- Multi-dimensional growth and scalability options
  - Dynamically add cores, memory, I/O adapters, devices and network cards
  - Resources may be shared or dedicated
  - Grow **horizontally** (add Linux guests), **vertically** (add to existing Linux guests) and **diagonal** (Mix and Match – Find your scale sweet spot)
  - Grow and scale without disruption to running environment
  - Provision for peak utilization, unused resources automatically reallocated after peak

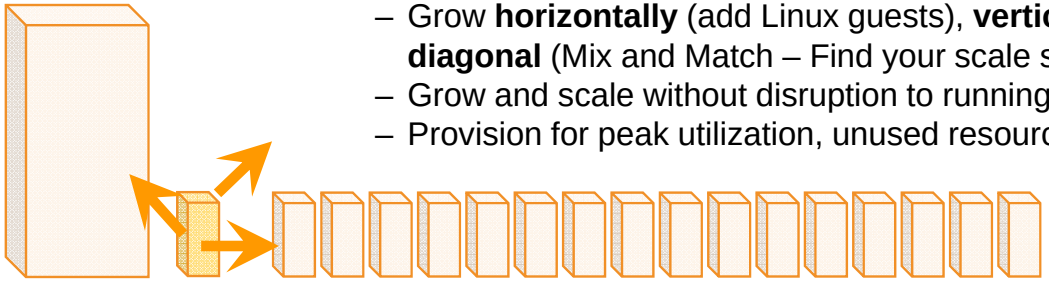

Dynamic Changes Capabilities	z/VM LPAR	Linux Guest
Add CPU	Yes	Yes
Increase weight / share	Yes	Yes
Add memory	Yes	Yes
Add I/O adapter	Yes	Yes
Remove CPU	Yes	Yes
Decrease weight / share	Yes	Yes
Remove memory	No	Yes
Remove I/O adapter	Yes	Yes

- Flexible Resource/Workload Management and High configuration flexibility
  - Advanced workload management enables maximum utilization of the system resources
  - Goal-oriented approach for performance mgmt of a hypervisor
  - Ability to basically do a forklift upgrade to new z Systems server

- Efficiencies of Consolidation
  - Less operational effort based on centralized management, using the same arrangements for administration, security, backup and disaster recovery
  - Less efforts for less IT equipment

And z/VM keeps growing (and re-using pictures for our slides)

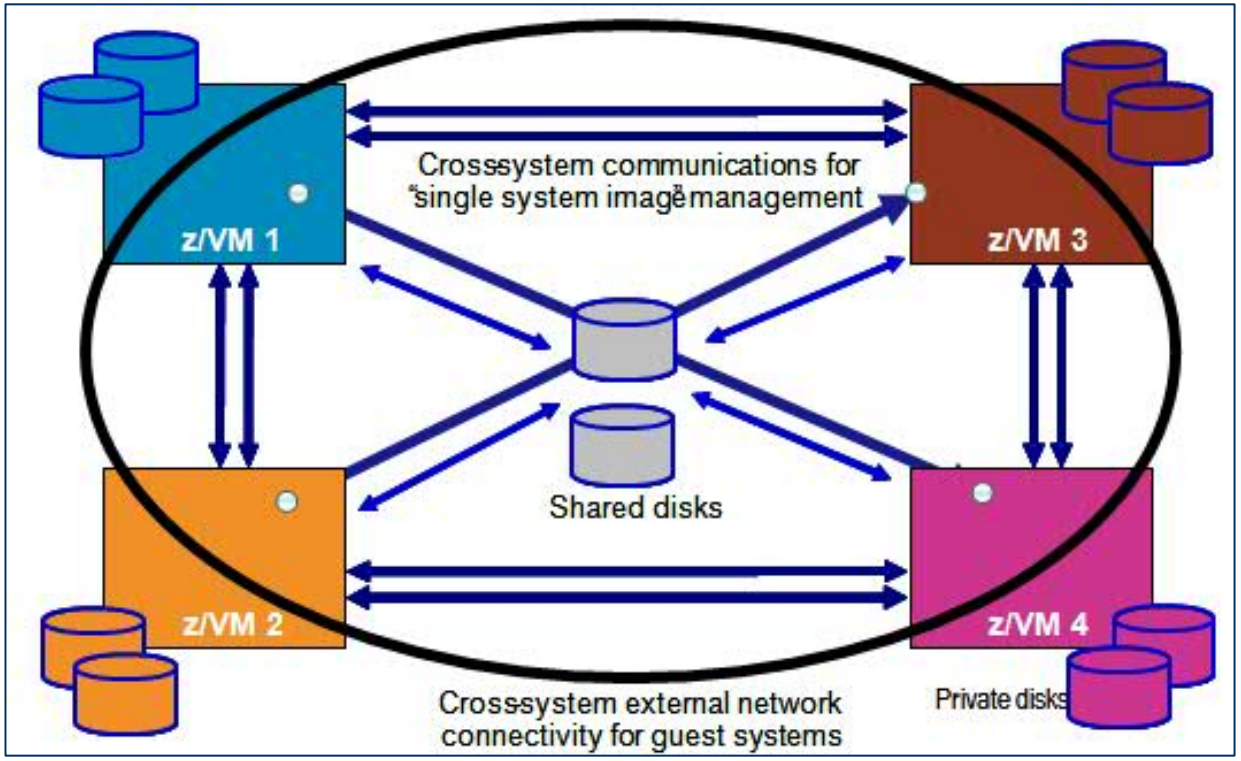

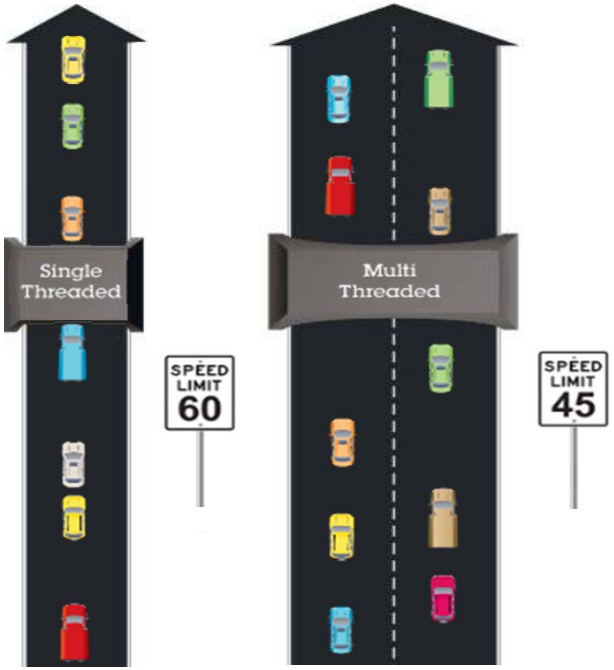

Which approach is designed for the higher volume of traffic? Which road is faster?

**Illustrative numbers only*

A Brief History of Virtualization that IBM is engaged in

IBM has over 45 years of experience in virtualizing our servers. Virtualization was originally developed to make better use of critical hardware. Hardware support for virtualization has been critical to its adoption.

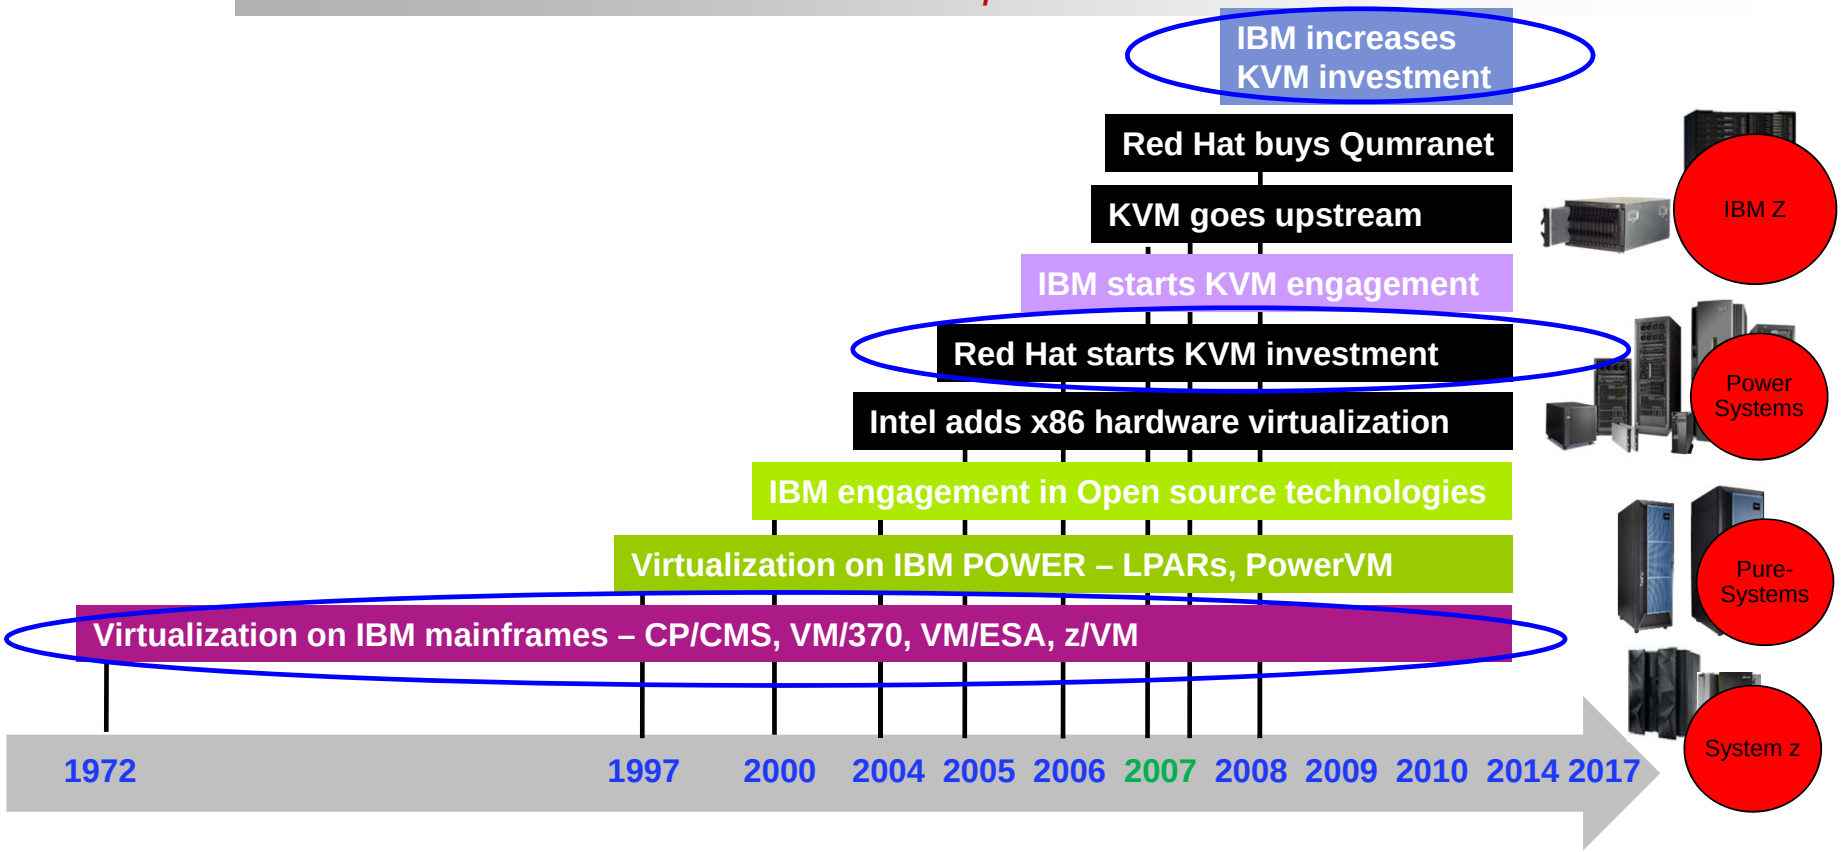

IBM Z Virtualization

Build-in and Shared Everything Architecture -towards a SDD or HCI

IBM Z & LinuxONE Systems

- Provisioning of virtual servers in seconds
- High granularity of resource sharing (<1%)
- Upgrade of physical resources without taking the system down
- Scalability of up to 1000's of virtual servers
- More with less: more virtual servers per core, sharing of physical resources
- Extensive life-cycle management
- HW-supported isolation, highly secure (EAL5+ or EAL4+ certified)

Distributed platforms

- Limited virtual server scalability per core
- Scaling requires additional physical servers
- Operational complexity increases with growth of virtual server images
- VMware, Xen, Hyper-V focus on x86, no HW management across multiple platforms

IBM Z & LinuxONE Systems Virtualization Options

IBM LinuxONE Systems now has three strategic virtualization platforms

1. IBM Dynamic Partition Manager (DPM) or Processor Resource /System Manager (PR/SM)
2. IBM z/VM
3. KVM

NEW

KVM — **KVM provides an open source choice for IBM LinuxONE Systems virtualization for Linux workloads. Best for clients that are not familiar with z/VM and are Linux centric admins.**

z/VM — **IBM Proprietary Server Virtualization that is completely integrated into the full stack. Complete hardware awareness. Supported on IBM LinuxONE Systems. z/VM will continue to be enhanced to support Linux Workloads.**

DPM or PR/SM — **Divide one physical LinuxONE System into up to 85 logical partitions (LPAR) running isolated and secured in parallel. Share resources across LPARs or dedicated to a particular LPAR. Control the LPAR virtualization via the new Dynamic Partition Manager (DPM) or with PR/SM.**

What is Dynamic Partition Manager?

DPM graphical interface

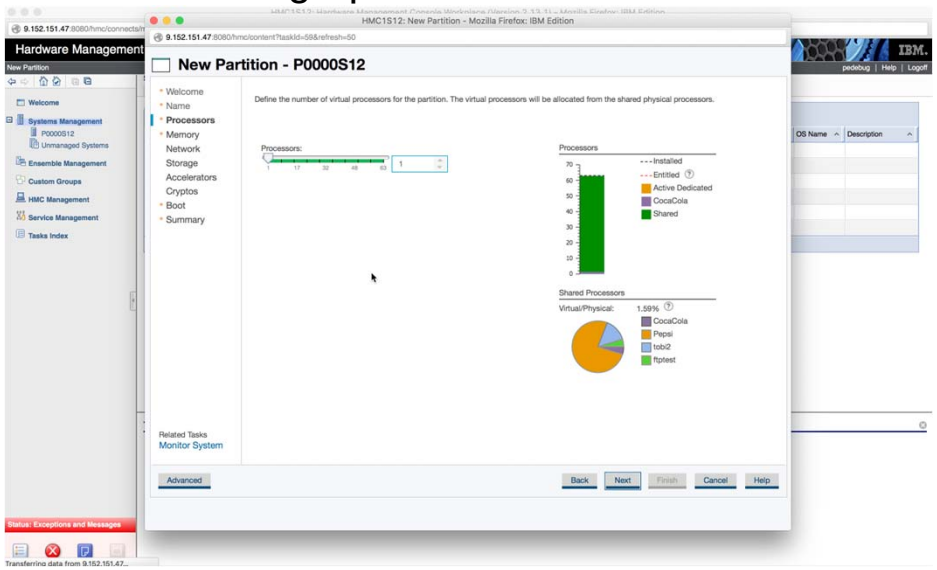

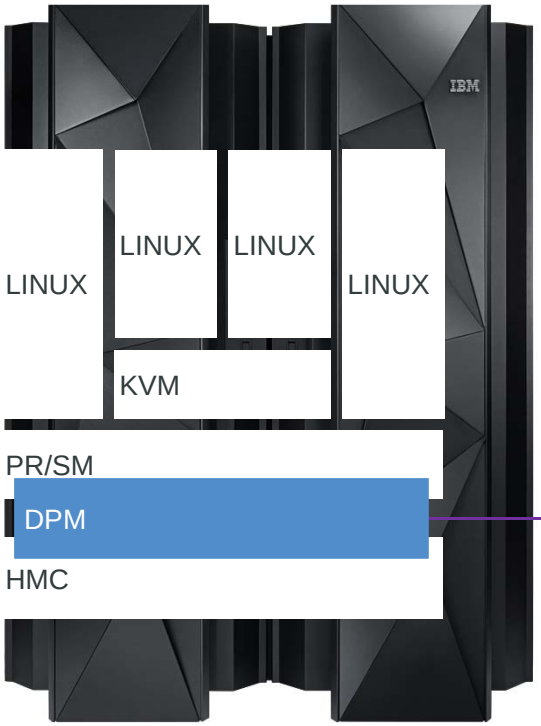

- Provide simplified, consumable, enhanced Partition life-cycle and integrated dynamic I/O management capabilities.
- Provides the technology foundation that enables APIS for IaaS and secure, private Clouds.

DPM

Powerful and easy

Secure Service Container Framework Overview

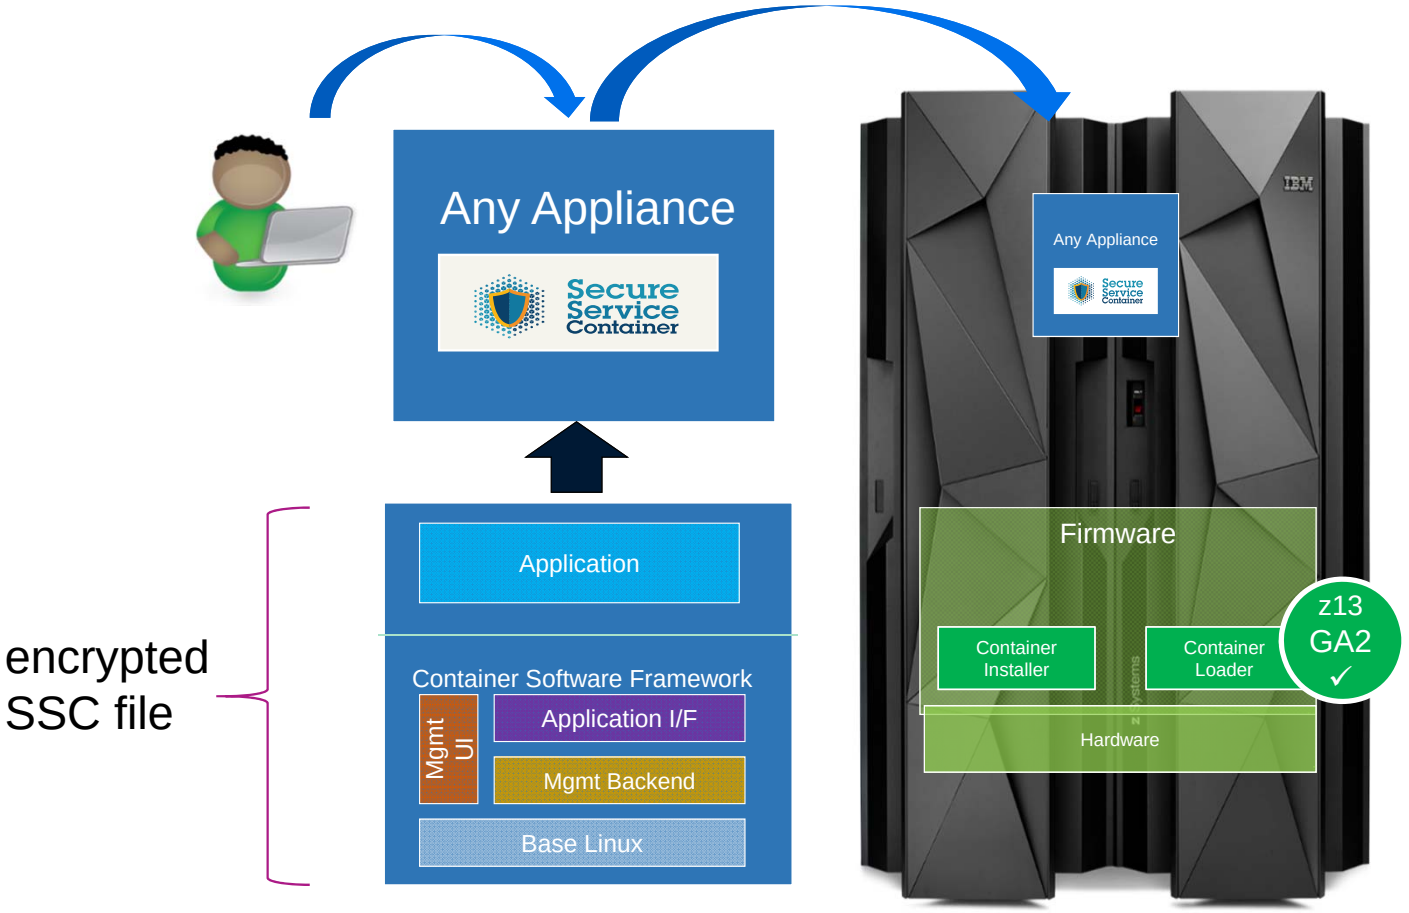

Linux virtual server vs. Docker containers overhead

- The virtualization approach usually provides a high level of isolation and security as all communication between the guest and host is through the hypervisor.
- It is also usually slower and incurs some overhead due to the infrastructure emulation.
- To reduce this overhead, another level of virtualization called "**container virtualization**" was introduced which allows to run multiple isolated user space instances on the same kernel.
- Containers layered approach share common files and use copy-on-write filesystems

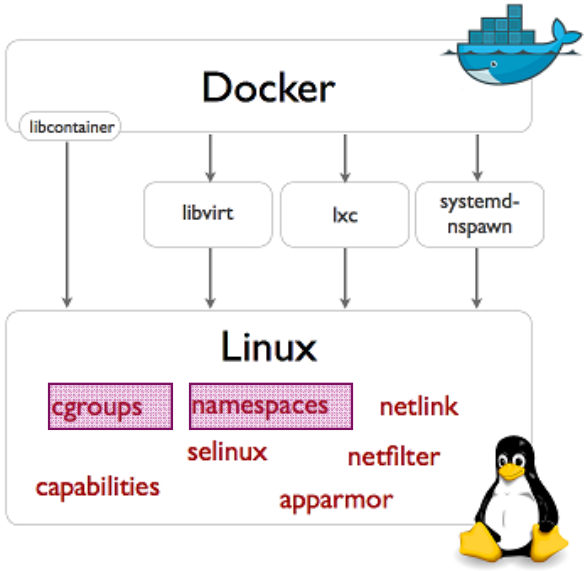

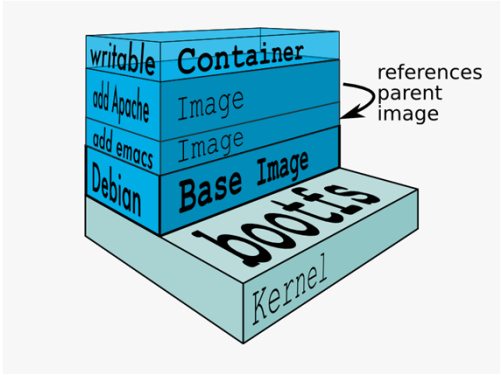

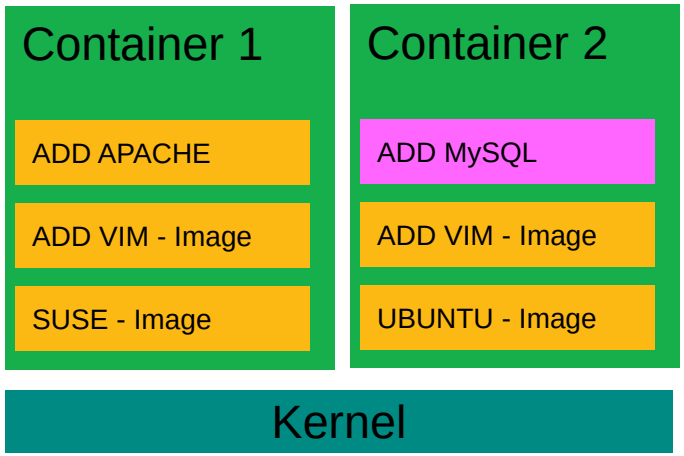

IBM Z Virtualization

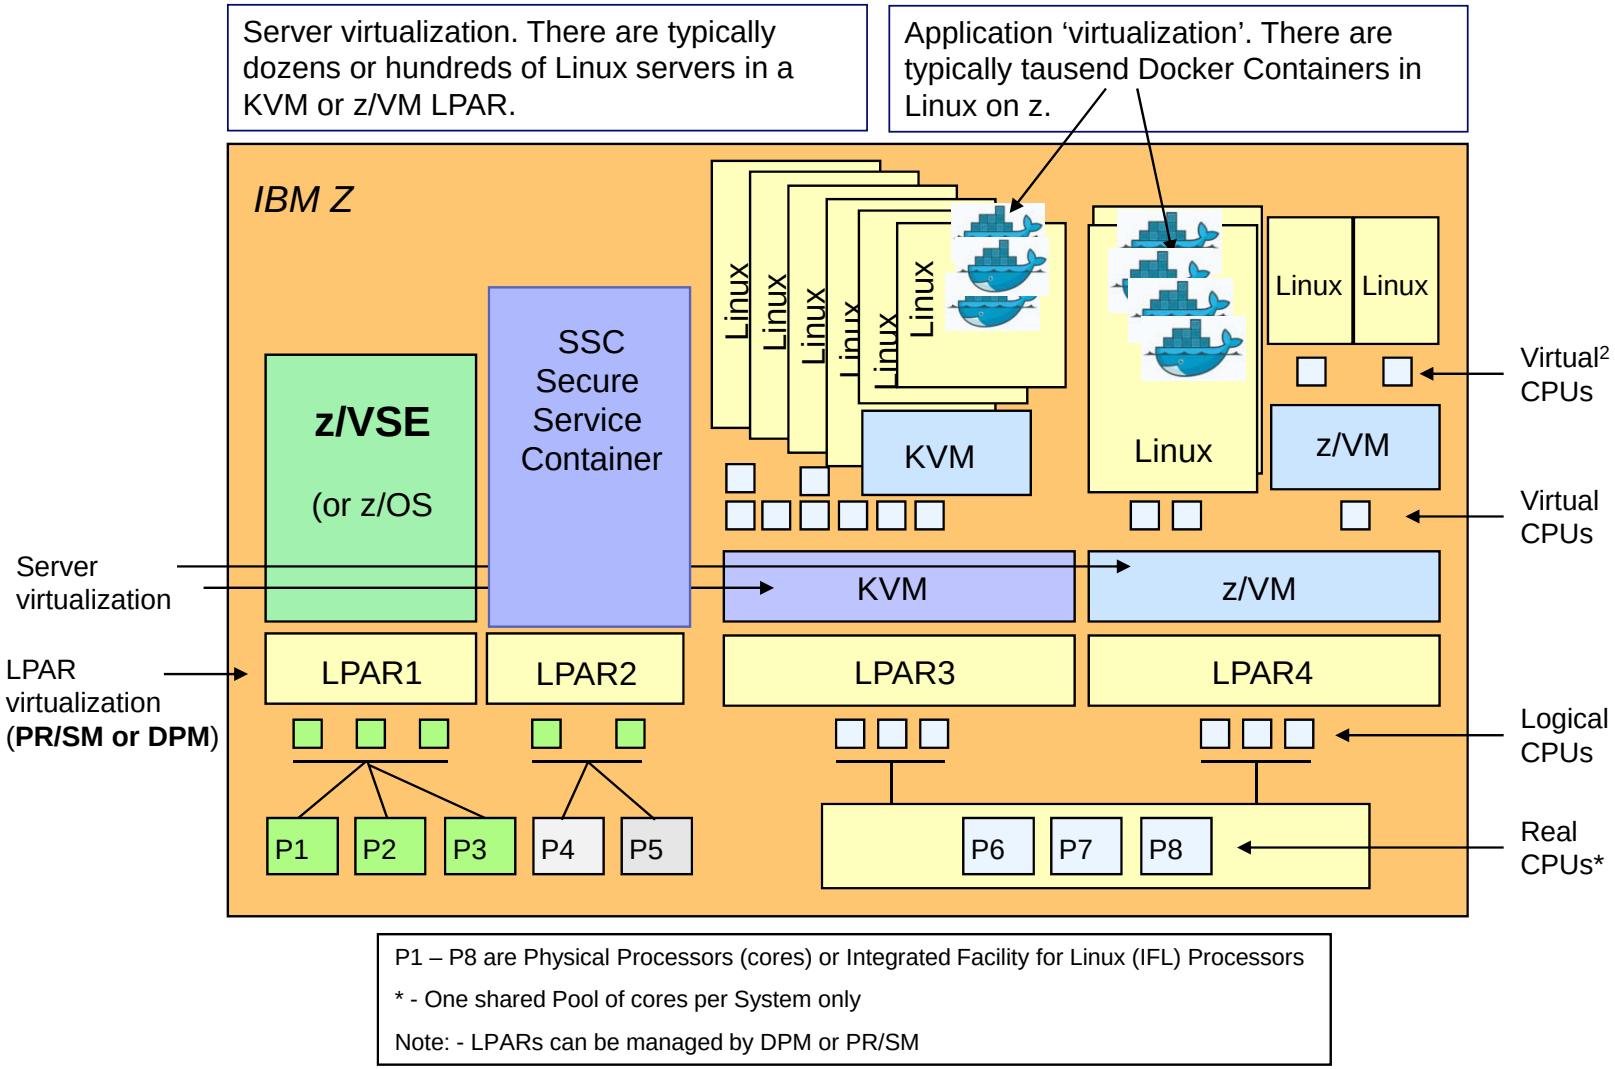

IBM LinuxONE Portfolio ™

Linux Your Way

Linux without Limits

Linux without Risk

A cloud in a single footprint.

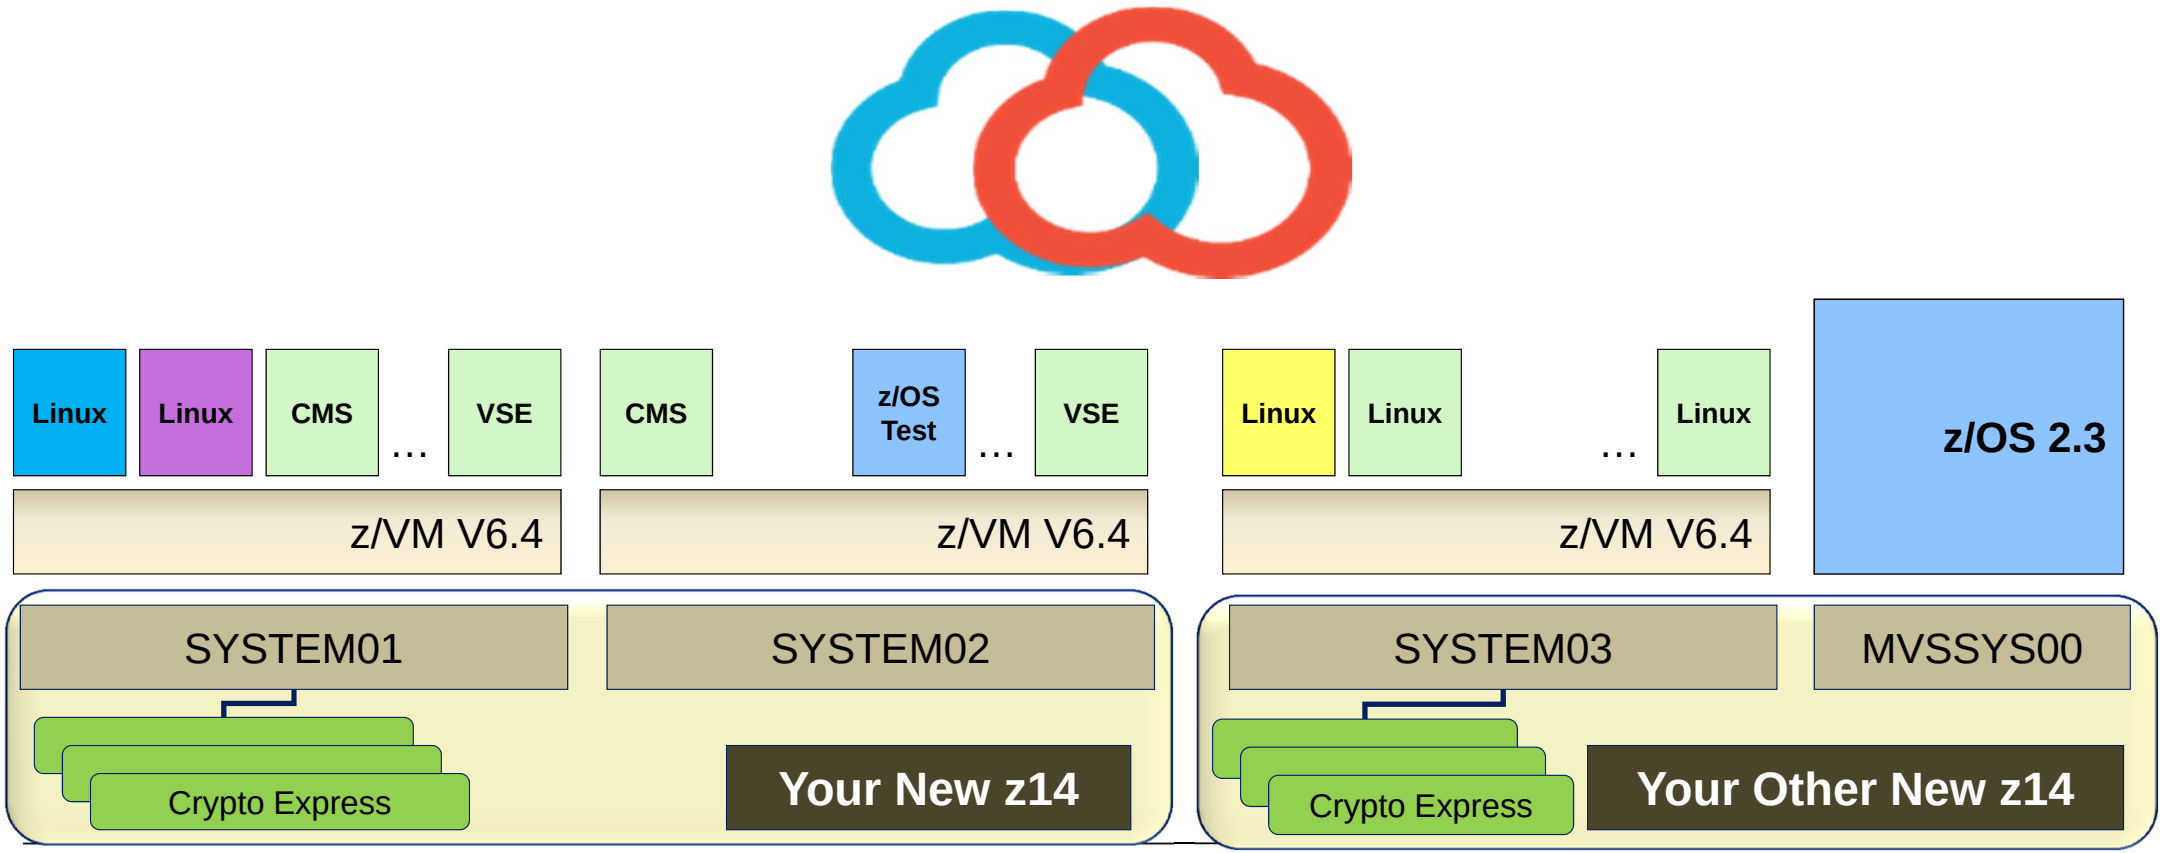

z/VM – a comprehensive and sophisticated suite of virtualization function

- Over 50 years of continuous innovation in virtualization
  Refined to support modern business requirements
  Exploiting hardware technology for economical growth

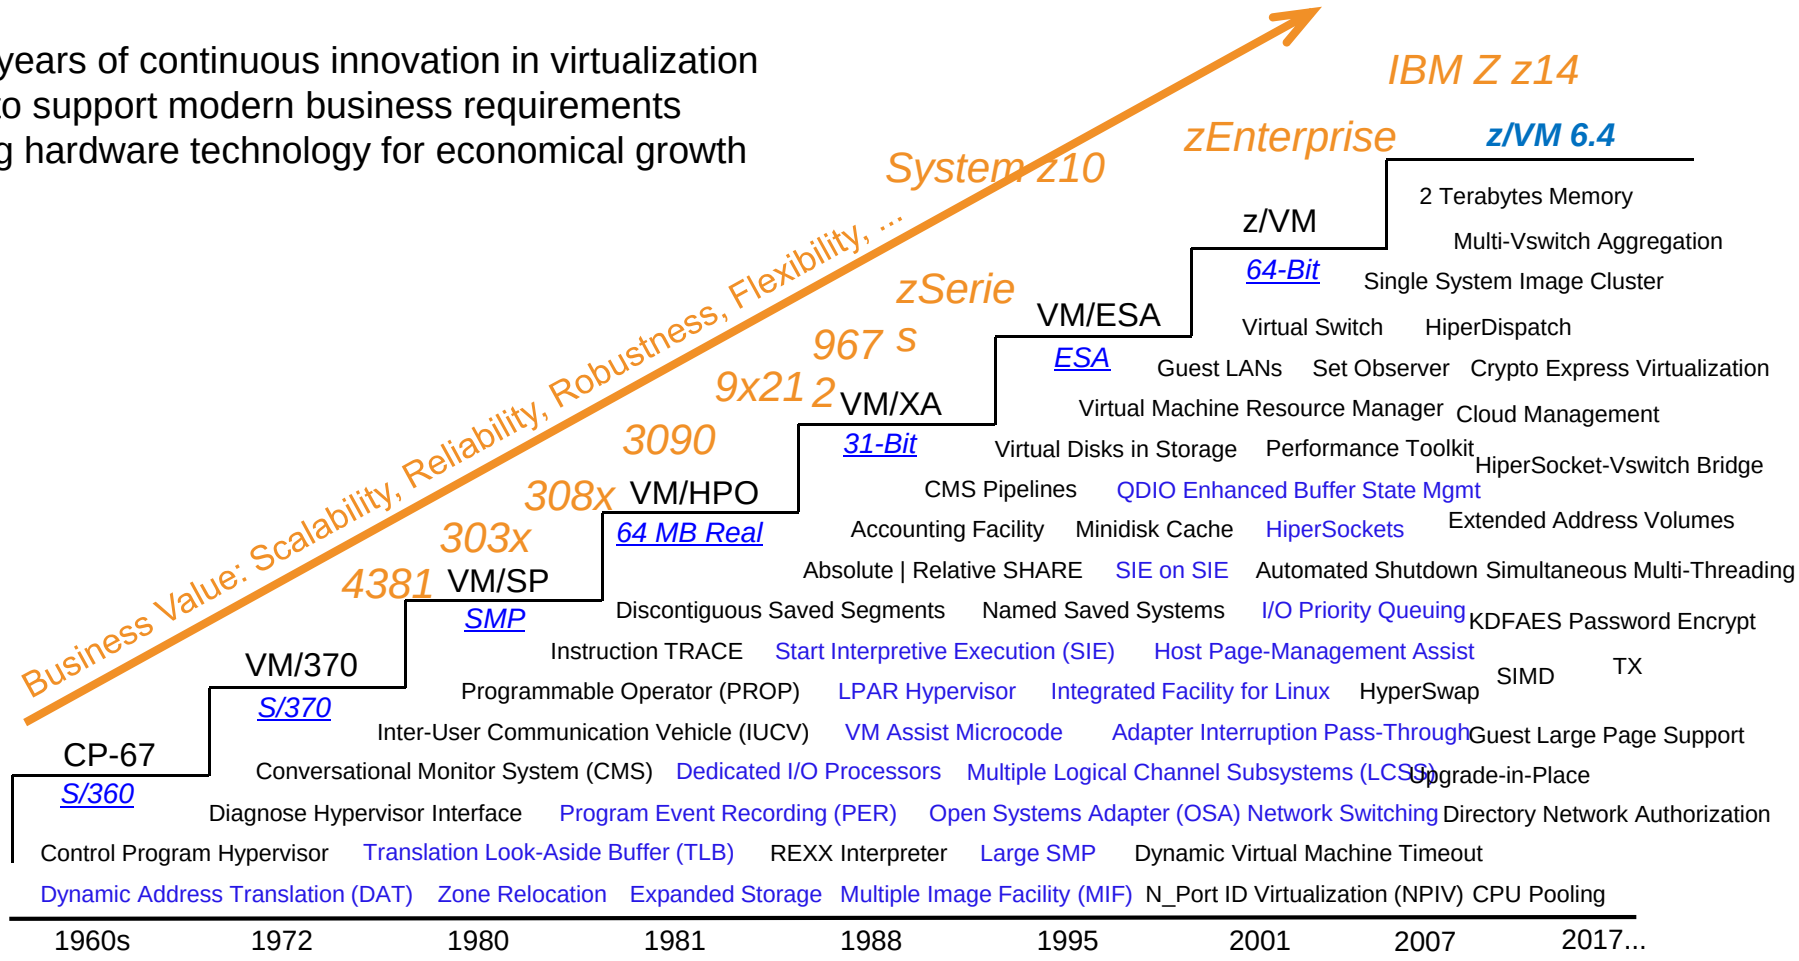

Key VM Attributes

© 2017 IBM Corporation

VM Mobile Museum (VMMM) – privately held

Happy triple on IBM Z since many years

- From the beginning sympathy on all levels
- Solid Virtualization
- High Reliability
- Scaling on demand with highest flexibility

Bibliography

- Melinda Varian – VM history papers
  - http://www.leeandmelindavarian.com/Melinda/
- Chuck Boyer – “The 360 Revolution”
  - ftp://ftp.software.ibm.com/s390/misc/bookoffer/download/360revolution_040704.pdf
- IBM Archives: Valuable resources on IBM’s history
  - http://www.ibm.com/ibm/history/
- Bill Bitner – “40 Reflections for 40 Years of z/VM” blog
  - http://zvm40for40.blogspot.com/
- Neale Ferguson – “The Last 30 Years of VM”, z/Journal, May 2012
  - http://www.mainframezone.com/article/the-last-30-years-of-vm
- Bill Bitner and Susan Greenlee – “z/VM – A Brief Review of Its 40 Year History”
  - http://www.vm.ibm.com/vm40bday.html

Bibliography

- IBM Systems Journal
  - “A virtual machine time-sharing system”
    R. A. Meyer and L. H. Seawright, Volume 9, Number 3, 1970
  - “Evolution of a virtual machine subsystem”
    E. C. Hendricks and T. C. Hartmann, Volume 18, Number 1, 1979
- IBM Journal of Research and Development
  - “The Origin of the VM/370 Time-Sharing System”
    R. J. Creasy, Volume 25, Number 5, 1981
  - “System/360 and Beyond”
    A. Padegs, Volume 25, Number 5, 1981
- IBM’s 360 and Early 370 Systems
  - E. W. Pugh, L. R. Johnson, and J. H. Palmer, MIT Press, 1991

Trademarks

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

BladeCenter*	GDPS*	IBM z13*	PR/SM	System z9*	zSecure
DB2*	HiperSockets	IBM z14	RACF*	System z10*	z/VM*
DS6000*	HyperSwap	OMEGAMON*	Storwize*	Tivoli*	z Systems*
DS8000*	IBM LinuxONE Emperor	Performance Toolkit for VM	System Storage*	zEnterprise*	
ECKD	IBM LinuxONE Rockhopper	Power*	System x*	z/OS*	
FICON*	IBM Z*	PowerVM	System z*		

* 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.