G08/G09 DPM-Ausblick

Elisabeth Puritscher
elisabeth.puritscher@de.ibm.com
Manfred Gnirss
gnirss@de.ibm.com
Arwed Tschoeke
arwed.tschoeke@de.ibm.com
IBM Client Center Böblingen

GSE GUIDE SHARE EUROPE Deutsche Region

Trademarks

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

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

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.
TEALEAF is a registered trademark of Tealeaf, an IBM Company.
Windows Server and the Windows logo are trademarks of the Microsoft group of countries.
Worklight is a trademark or registered trademark of Worklight, an IBM Company.
UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Notes:

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

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

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

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

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

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

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

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

Notice Regarding Specialty Engines

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

PR/SM and Dynamic Partition Manager

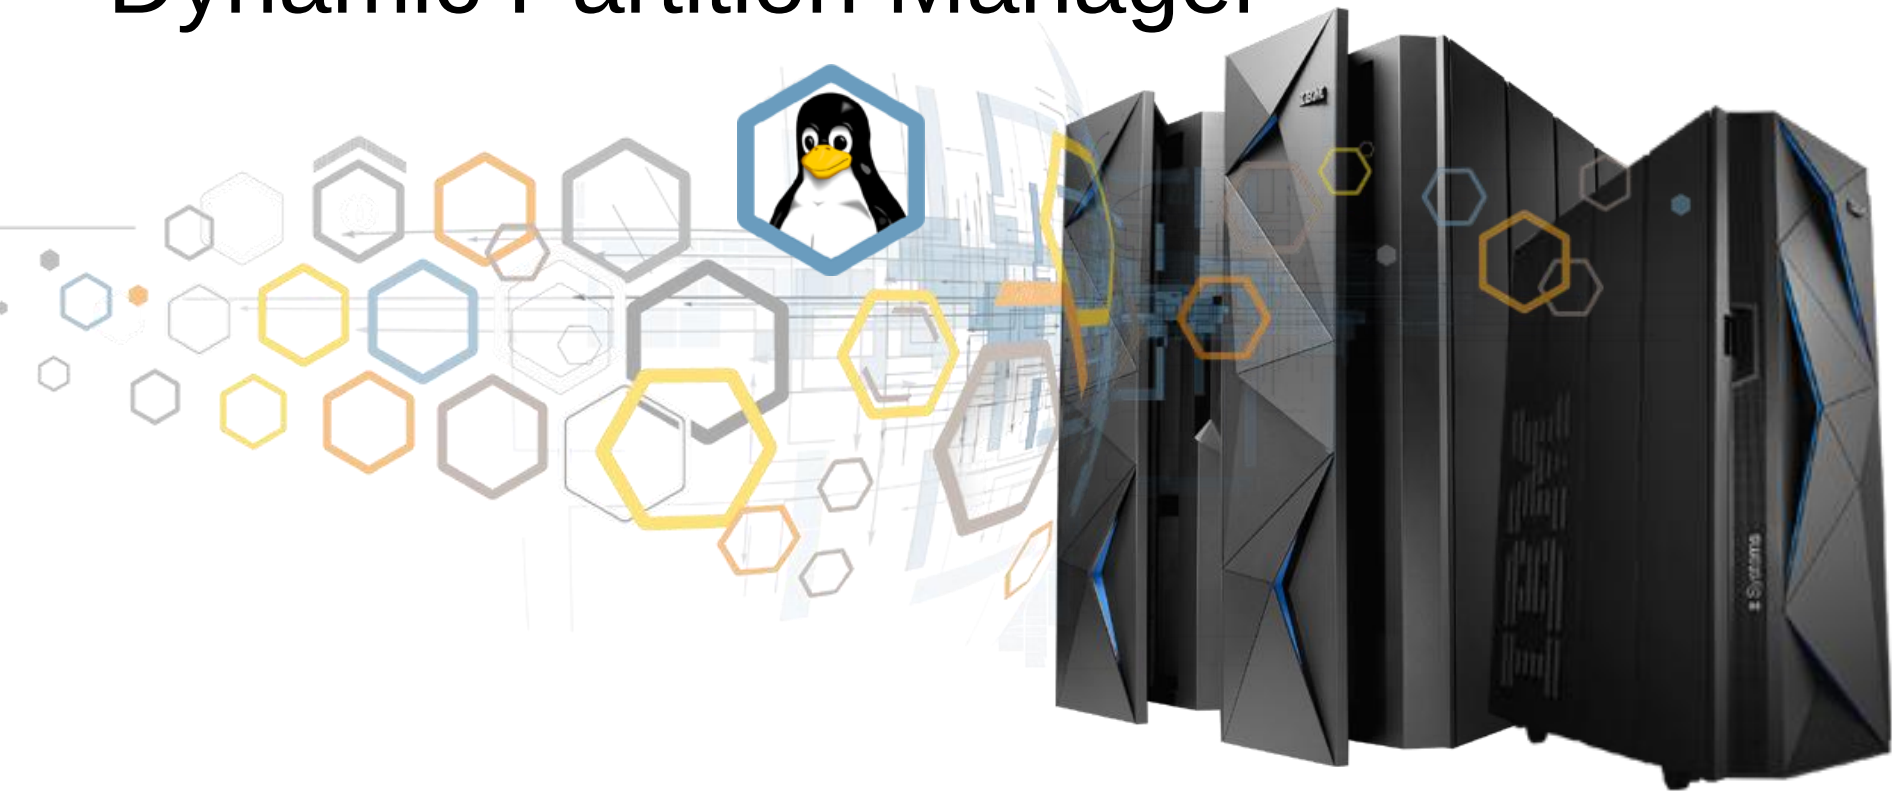

LinuxONE Virtualization Options

- At the hypervisor level (2 IBM options: z/VM® and KVM)
- At the firmware level
  - Processors, memory and I/O resources are divided into 1 or more Logical Partitions or LPARs
  - What runs in each LPAR appears to be running on “bare metal”, and is electronically separate from what runs in any other LPAR, as if they were in physically separate machines
  - LPARs are managed by *one* of two managers:
    - Processor Resource / System Manager (PR/SM™)
    - IBM Dynamic Partition Manager
  - Interaction with the manager is via either:
    - The Support Element (SE) – a dedicated workstation used for monitoring and operating the system
    - The Hardware Management Console (HMC) – a secure Web application accessed through a browser

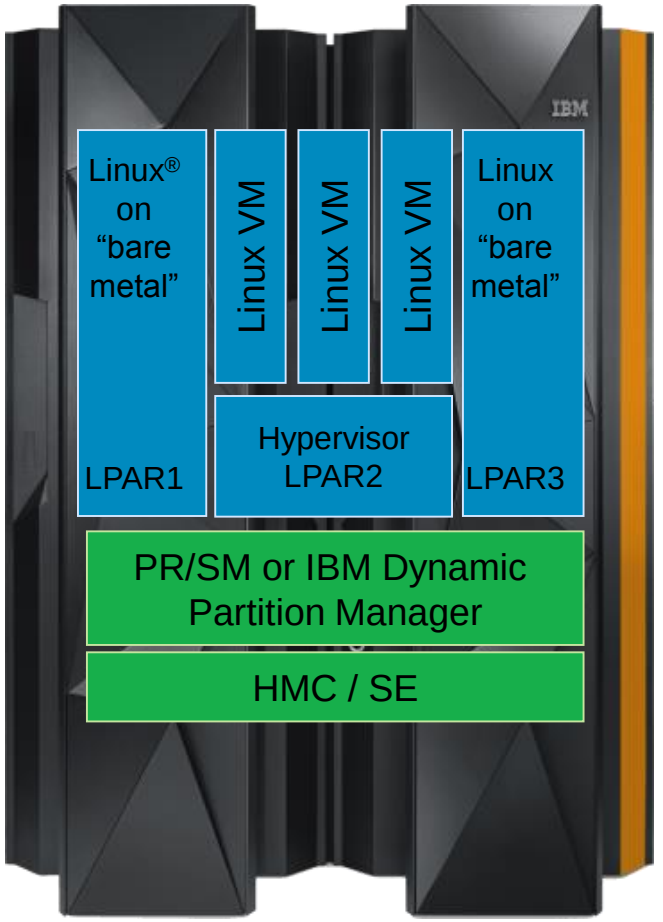

PR/SM or LPAR Hypervisor

- 'Processor Resource/System Manager' (PR/SM) and 'LPAR hypervisor' are commonly used synonymously.
- However the 'LPAR Hypervisor' is the program itself and 'PR/SM' is the facility of the whole
- So PR/SM aka LPAR hypervisor is a Type-1 Hypervisor that manages logical partitions:
  - Each partition owns a defined amount of physical storage
  - Strictly no storage shared across partitions
  - No virtual storage management / paging done by LPAR hypervisor
  - Zone relocation lets each partition start at address 0
  - CPUs may be dedicated to a partition or may be shared by multiple partitions
  - I/O channels may be dedicated to a partition or may be shared by multiple partitions (Multiple image facility, MIF)
  - Each LPAR has its own architecture mode (ESA/390 or z/Architecture)
- PR/SM is shipped with z Systems (considered as part of the firmware)
- PR/SM was initially introduced in 1988 with the IBM 3090 processors
- Beginning with z990, the PR/SM is always loaded (no Basic Mode anymore)
- Separation of logical partitions is considered as good as having each partition on a separate physical machine (Evaluation Assurance Level 5)

How DPM helps in a new Linux Environment

- z Systems and PR/SM require a HW definition
- Dynamic IO – one of the key differentiator of the platform would be nice
- Having the option to have a GUI-based administration
- Overcome the prejudice: z is old school and complicated
- Everything is scripted (that‘s why we need GUIs ;-))

```
*ICP  ICP070I SEARCH FOR '*ICP' TO FIND EACH IOCP MESSAGE
        ID    MSG1='IODF00',MSG2='SYS1.IODF00 - 2014-07-22 16:19',    *
              SYSTEM=(2828,1),LSYSTEM=P008B857,                       *
        RESOURCE PARTITION=((CSS(0),(ZOS1,2),(ZVM1,1),(*,3),(*,4),(*,5*
              ),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),*
              (*,F)),(CSS(1),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7)*
              ,(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)))
        CHPID PATH=(CSS(0),25),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=160,TYPE=FC
        CHPID PATH=(CSS(0),26),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=11C,TYPE=FC
        CHPID PATH=(CSS(0),27),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=161,TYPE=FC
        CHPID PATH=(CSS(0),28),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=11D,TYPE=FC
        CHPID PATH=(CSS(0),F8),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=17C,TYPE=OSD
        CHPID PATH=(CSS(0),F9),SHARED,PARTITION=((ZOS1,ZVM1),(=)),    *
              PCHID=104,TYPE=OSD
        CNTLUNIT CUNUMBR=0030,PATH=((CSS(0),F8)),UNIT=OSA
        IODEVICE ADDRESS=(030,064),UNITADD=00,CUNUMBR=(0030),UNIT=OSA
        CNTLUNIT CUNUMBR=0070,PATH=((CSS(0),F9)),UNIT=OSA
        IODEVICE ADDRESS=(070,064),UNITADD=00,CUNUMBR=(0070),UNIT=OSA
        CNTLUNIT CUNUMBR=0300,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=0,UNIT=2105
        IODEVICE ADDRESS=(300,032),CUNUMBR=(0300),STADET=Y,UNIT=3390
        CNTLUNIT CUNUMBR=0320,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=2,UNIT=2105
        IODEVICE ADDRESS=(320,032),UNITADD=00,CUNUMBR=(0320),STADET=Y,*
              UNIT=3390
        CNTLUNIT CUNUMBR=0340,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=4,UNIT=2105
        IODEVICE ADDRESS=(340,032),UNITADD=00,CUNUMBR=(0340),STADET=Y,*
              UNIT=3390
        CNTLUNIT CUNUMBR=0360,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=6,UNIT=2105
        IODEVICE ADDRESS=(360,032),UNITADD=00,CUNUMBR=(0360),STADET=Y,*
              UNIT=3390
        CNTLUNIT CUNUMBR=0380,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=8,UNIT=2105
        IODEVICE ADDRESS=(380,032),UNITADD=00,CUNUMBR=(0380),STADET=Y,*
              UNIT=3390
        CNTLUNIT CUNUMBR=0400,PATH=((CSS(0),25,26)),                  *
              UNITADD=((00,032)),CUADD=A,UNIT=2105
        IODEVICE ADDRESS=(400,032),CUNUMBR=(0400),STADET=Y,UNIT=3390
```

short version: No Texteditor required to get started, dynamic IO available

Partition and I/O device management at the HMC

- Linux partitions or KVM partitions correspond to LPARs under standard PR/SM
- Supports only Linux and Linux based hypervisors
- Creation of I/O Configuration Data Set (IOCDS) is “under the covers”
  - Supports dynamic updates of I/O
- Hardware and operating system message displays are unchanged
- Problem determination and maintenance continues to exist on the System Element (SE)
- On/Off Capacity on Demand (OOCoD) and Customer Initiated Upgrade (CIU) supported for Linux

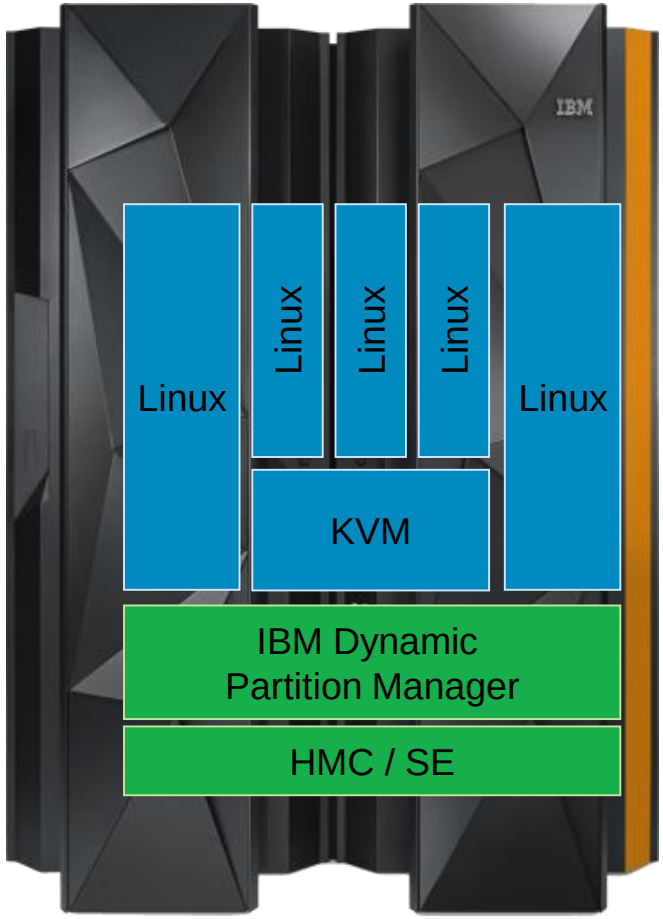

System Requirements

- IBM LinuxONE Rockhopper™ or IBM LinuxONE Emperor™
  - IBM Dynamic Partition Manager feature code 0016
  - Two dedicated OSA-Express5S 1000BASE-T Ethernet #0417 features
  - Server can be in standard PR/SM mode or Dynamic Partition Manager mode – set at initial IML
- Only supports Fibre Communication Protocol (FCP)
- IBM KVM and/or Linux without a hypervisor
- Does not support
  - Any hypervisor other than IBM KVM
  - ECKD™ disk
  - IBM zAware
  - GDPS® Virtual Appliance

[1]Requires firmware level 27 (GA2)

Additional technical information

- Separate HMCs are NOT required; HMC networks with both IBM Dynamic Partition Manager servers and standard PR/SM servers
- Only SERVICE or SYSPROG userids can enable/disable IBM Dynamic Partition Manager
- All alarms exist as before, plus there are 3 or 4 new types of alarms, primarily around utilization being added
- Improved support over today's ASCII Console support – ability to open a console window
- HMC REST API support – everything in the UI can be done via the REST API

Herzlichen Dank

- Demo