Linux on IBM Z and LinuxONE: What's New and what for

Wilhelm Mild
IBM Executive IT Architect
Linux and Container on IBM Z & LinuxONE

Stefan Raspl
Principal Product Manager
Linux & Virtualization on IBM Z and LinuxONE

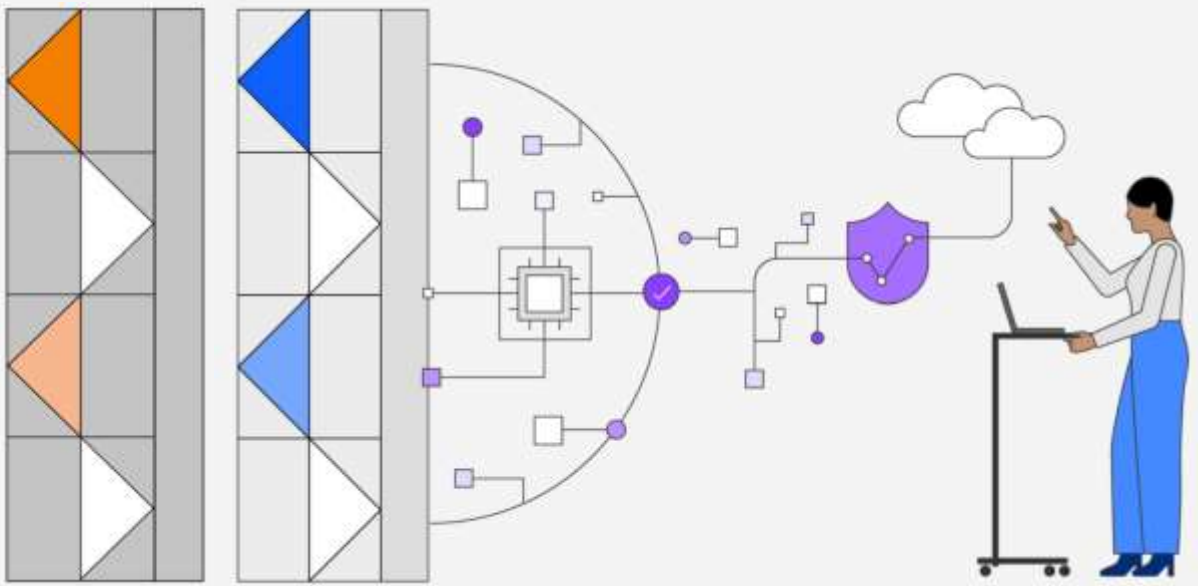

© 2025 IBM Corporation

Contents

IBM Z © 2025 IBM Corporation

Market view of Container workload and Linux on IBM LinuxONE and IBM Z

Installed Linux IFL MIPS at 21% CAGR* - growth!

- Used **in over 60 countries** around the globe
- Used in about **22 industries**

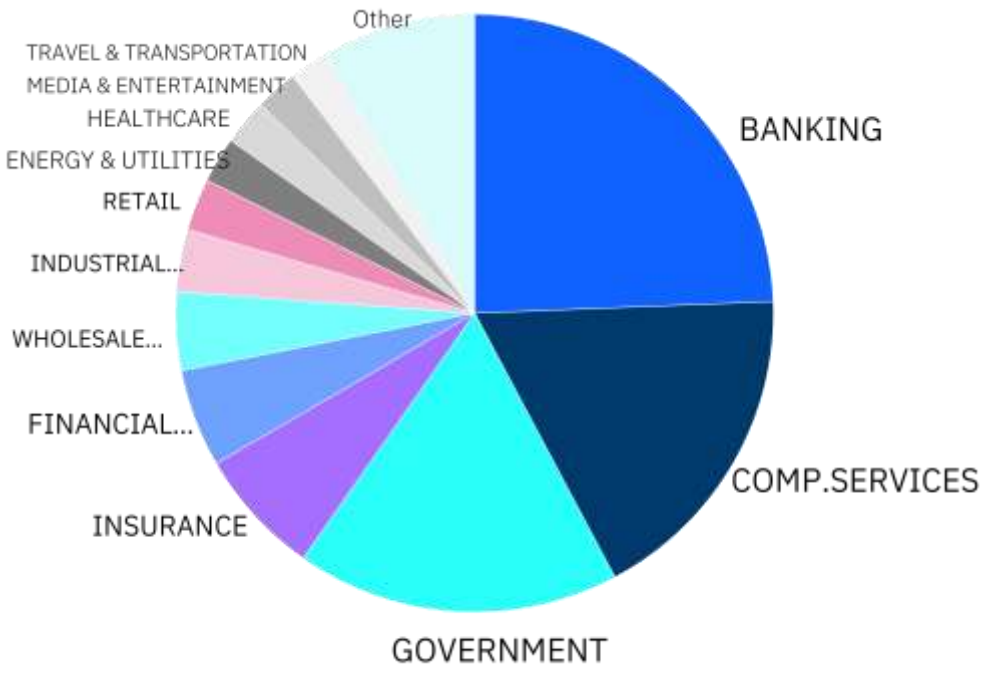

*CAGR - Compound Annual Growth Rate (CAGR)

*based on CustNo 1Q2025

- 96 of the top 100 IBM Z Enterprises are running Linux on Z as of 1Q25*
  - Very large installations with up to hundred and more of IFLs in USA, Japan, Brazil, Asia, Mexico and Germany
  - Small installations with 2 IFLs in all countries and on all IBM Z models
- Many clients run Linux co-located with z/OS® on IBM Z

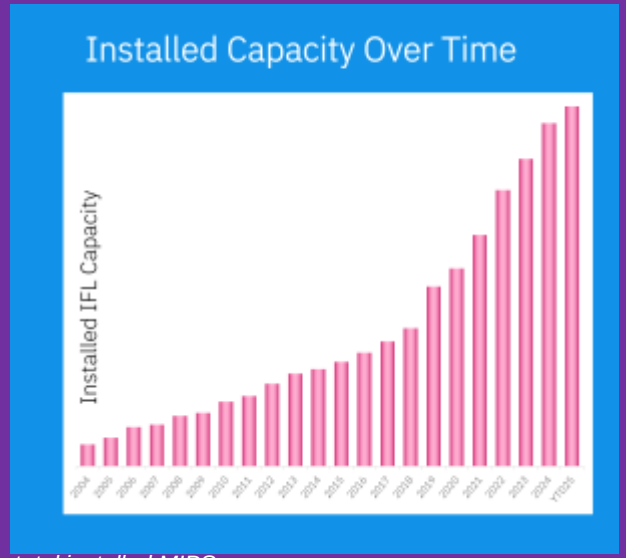

*Top 100 is based on total installed MIPS

© Copyright IBM Corporation 2023

IBM z17 and IBM LinuxONE 5 announced - the twins for different markets

IBM z17

IBM LinuxONE 5

IBM LinuxONE in F1 & the Racing track in Spa Belgium

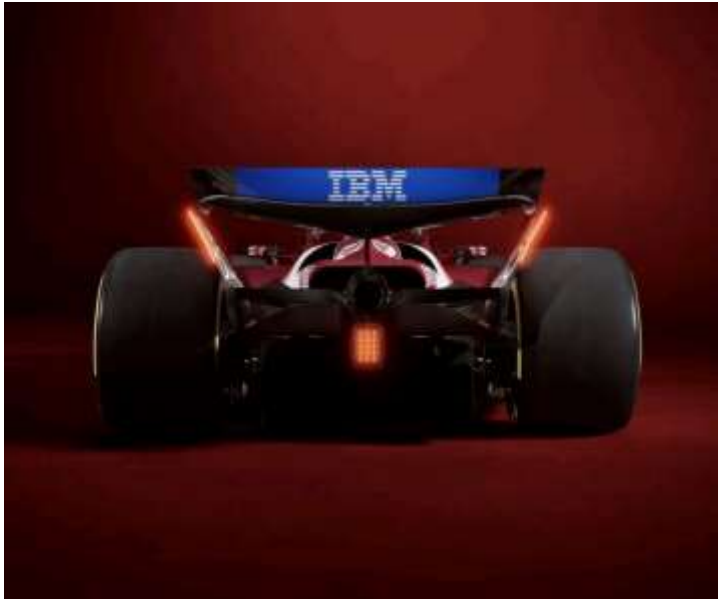

IBM® Linux & LinuxONE timeline

25 Years Linux on IBM Z
10 years of LinuxONE success

2000

Linux® for s390x

- Red Hat® Enterprise Linux®
- SUSE
- Data serving (Oracle, Db2®)

2015

IBM® LinuxONE

- Ubuntu
- Data serving (MongoDB)

2018

IBM® LinuxONE II

- IBM Db2® Analytics Accelerat
- Co

xONE 4

- antum-safe encryption
- Sustainability
- Red Hat OpenShift Ansible® Automation Platform
- Data serving (Fujitsu, EDB)
- Core banking (Finacle)

2025

IBM® LinuxONE 5

- Cost efficiency
- Scalable AI
- Confidential Containers (Red Hat OpenShift CoCo
- Red Hat OpenShift Virtualization-technology preview
- Red Hat OpenShift AI-technology preview
- HashiCorp integration

© 2025 IBM Corporation

IBM® LinuxONE Emperor 5: Unlock potential

System structure

- 208 customer cores
- Memory- Up to 64TB; 60% increase

I/O

- Simplified system IO architecture
- Double networking port density per IO drawer

Frequency

@5.5 GHz*

Cache Growth

40% increase*

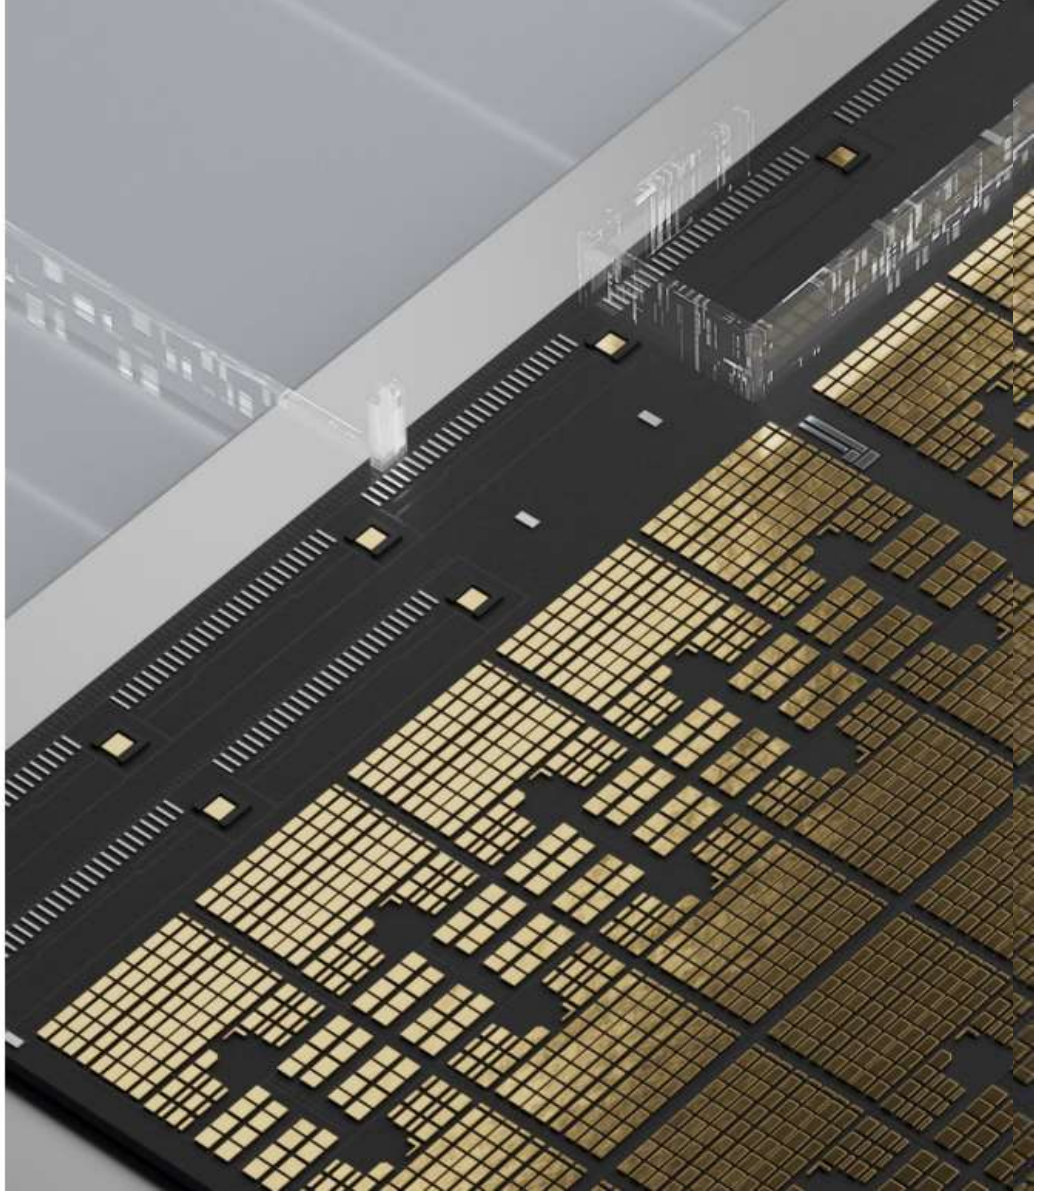

© 2025 IBM Corporation

*Compared to the IBM® LinuxONE Emperor 4

IBM z17 & IBM® LinuxONE 5 Hardware and System Structure

Total System Capacity Growth
- ~20% drawer capacity growth for standard models
- ~12% drawer capacity growth for max config model

I/O
- Double density with new FICON Express 32G (4-port)
- Converged network adapter (RoCE and OSA)

Memory
- Up to 64TB RAIM memory
- 60% more memory capacity per drawer over IBM z16

Sysplex
- Improved connectivity and configuration flexibility with up to 512 coupling CHPIDs per CP (increased from 384)
- Up to 128 LR coupling links per CPC (increased from 64)
- Up to 8 logical connection per physical link (CHPIDs/port) (increased from 4)

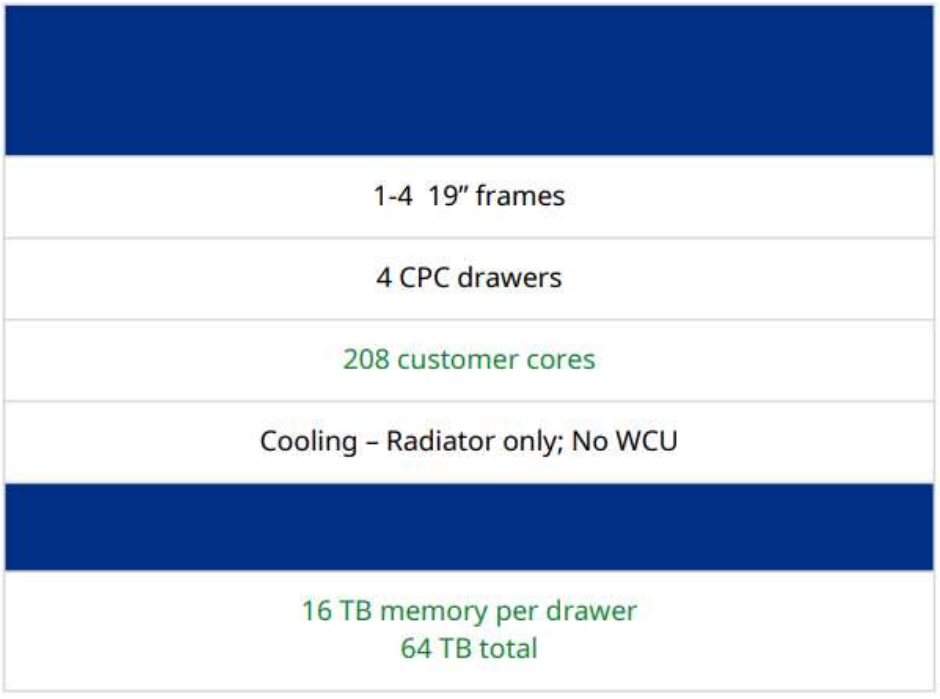

IBM Telum™ II Processor

- 5nm technology (Samsung), 5.5GHz
- 8 cores with 20% area reduction and 15% power reduction
- 40% more cache per core
- 24.1 Miles of wire per chip (5.3 more miles than z16)
- 43.0 Billion transistors (20.5B more than z16)
- **NEW:** On-chip Data Processing Unit (DPU): Increased I/O performance, RAS, and reduced latency
- 2nd-gen AI Accelerator for high-speed inferencing
- 8x dedicated AI processing per core

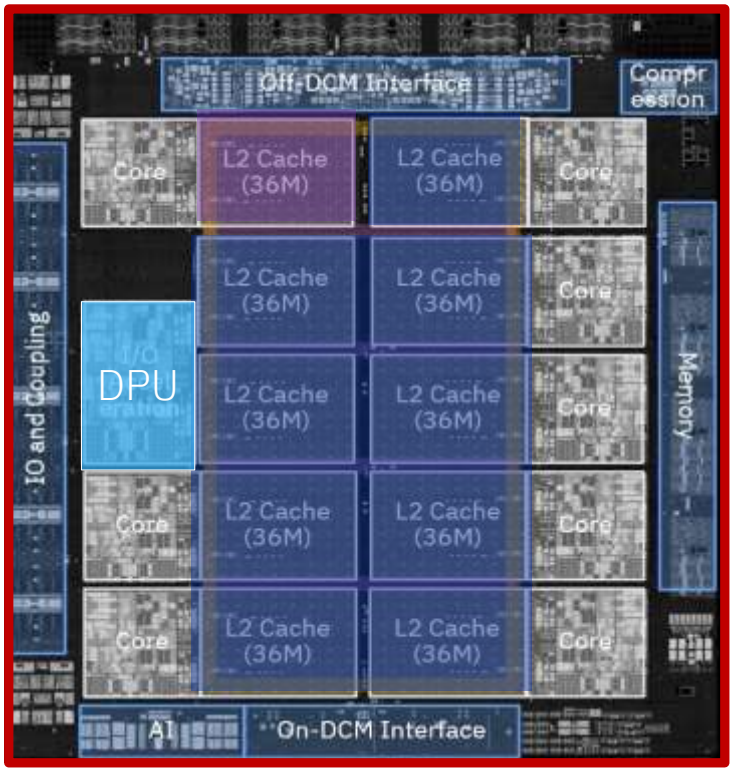

Why is IBM Z & LinuxONE different?

- Performance, I/O
- On-chip accelerators for compression, encryption, AI
- Continuous Availability
- Security
- Scalability
- Manageability
- Total Cost of Ownership

2025 IBM Corporation

Linux on IBM Z & LinuxONE Distributions

IBM Z © 2025 IBM Corporation

Linux on IBM Z & LinuxONE Distributions: Red Hat

- **Red Hat Enterprise Linux 10**
  - 05/2025 RHEL 10 GA: Kernel 6.12, GCC 14.2.1
  - EOS: 31 May 2035; ELS: 31 May 2038
- **Red Hat Enterprise Linux 9**
  - 05/2022 RHEL 9 GA: Kernel 5.14, GCC 11.2.1
  - 05/2025 RHEL 9.6
  - EOS: 31 May 2032; ELS: 31 May 2035

But don't forget:
Keep it current, apply z-Stream updates!!!

- **Red Hat Enterprise Linux 8**
  - 05/2019 RHEL 8 GA: Kernel 4.18, GCC 8.2.1
  - 05/2024 RHEL 8.10
  - EOS: 31 May 2029; ELS: 31 May 2031
- **Red Hat Enterprise Linux 7**
  - 06/2014 RHEL 7 GA: Kernel 3.10, GCC 4.8
  - 09/2020 RHEL 7.9
  - ~~**EOS 30 June 2024**~~; **ELS: 30 June ~~2026~~ 2028** (see here for details)
- **Lifecycle Details**
  https://access.redhat.com/support/policy/updates/errata

Tag Legend

x.y for RHEL <x> Update <y>, e.g. 7.4 for RHEL7.4

IBM Z © 2025 IBM Corporation

Linux on IBM Z & LinuxONE Distributions: SLES

- **SUSE Linux Enterprise Server 15**
  - 07/2018 SLES 15 GA: Kernel 4.12, GCC 7.1 / 7.3
  - 06/2024 SLES 15 SP6: Kernel 6.4. GCC 7.5 / 13.2
  - EOS 31 July 2031; LTSS: 31 July 2034

- **SUSE Linux Enterprise Server 12**
  - 10/2014 SLES 12 GA: Kernel 3.12, GCC 4.8
  - 12/2019 SLES 12 SP5: Kernel 4.12, GCC 4.8
  - ~~**EOS 31 Oct. 2024**~~; LTSS: 31 Oct 2027

- **Lifecycle Details**
  https://www.suse.com/lifecycle/

But don't forget:
Keep it current, apply maintweb updates!!!

Tag Legend

x.y for SUSE SLES <X> SP <Y>, e.g. 12 SP3 for SLES12 SP3

IBM Z © 2025 IBM Corporation

Linux on IBM Z & LinuxONE Distributions: Canonical

- **Ubuntu 24.04 (Noble Numbat)**
  - 04/2024 GA: Kernel 6.8, GCC 13.2.0, LTS-Release
  - EOS: April 2029; ESM: Apr 2034
- **Ubuntu 22.04 (Jammy Jellyfish)**
  - 04/2022 GA: Kernel 5.15, GCC 11.2.0, LTS-Release
  - EOS: April 2027; ESM: Apr 2032
- **Ubuntu 25.04 (Plucky Puffin)**
  - 04/2025 GA: Kernel 6.14, GCC 15
  - EOS: January 2026

But don't forget:
Keep it current, apply SRU updates!!!

- **Ubuntu 20.04 (Focal Fossa)**
  - 04/2020 GA: Kernel 5.4, GCC 9.3.0, LTS-Release
  - EOS: April 2025; ESM: Apr 2030
- **Ubuntu 18.04 (Bionic Beaver)**
  - 04/2018 GA: Kernel 4.15, GCC 7.2.0, LTS-Release
  - ~~EOS: April 2023~~; ESM: April 2028
- **Ubuntu 16.04 (Xenial Xerus)**
  - 04/2016 GA: Kernel 4.4, GCC 5.3.0+, LTS-Release
  - ~~EOS: April 2021~~; ESM: April 2026
- **Lifecycle**
  - Regular releases every 6 mths, supported for 9
  - LTS releases every 2 years, supported for 5
  - LTS enablement stack provides newer kernels
  - http://www.ubuntu.com/info/release-end-of-life

Tag Legend

x.y for Ubuntu x.y, e.g. 16.04 for Ubuntu 16.04 LTS

IBM Z © 2025 IBM Corporation

Linux Distributions and Hardware Certifications

	z17™	z16®*	z15®*	z14®*
	IBM LinuxONE 5	IBM LinuxONE 4 *	IBM LinuxONE III *	IBM LinuxONE II *
RHEL 10	●	●	●	
RHEL 9	●	●	●	●
RHEL 8	●	●	●	●
RHEL 7		●	●	●
SLES 15	●	●	●	●
SLES 12	●	●	●	●
SLES 11				●
Ubuntu 24.04	●	●	●	●
Ubuntu 22.04	●	●	●	●
Ubuntu 20.04		●	●	●
Ubuntu 18.04			●	●
Ubuntu 16.04			●	●

* entire product portfolio

Last update: June 05, 2025

See here for latest updates and details, including ***certified Linux distributions by machine***.

IBM Z © 2025 IBM Corporation

Latest on IBM Z & LinuxONE Features and Packages

IBM Z © 2022 IBM Corporation

Linux / Networking

Networking Adapter Convergence

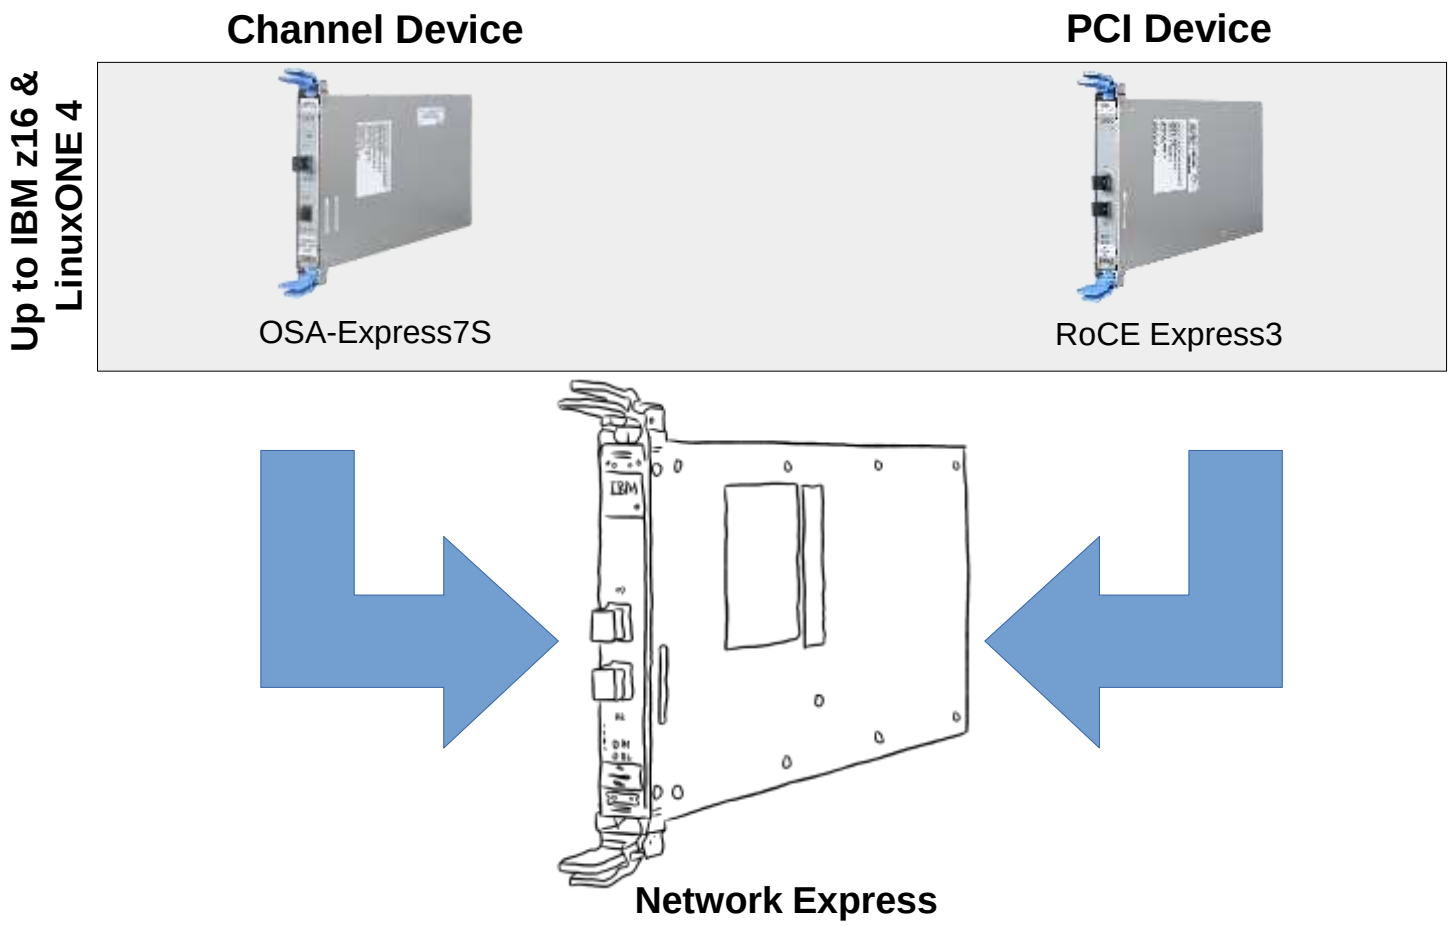

- The new **Network Express adapter** can be configured as both at the same time:
  1. OSH channel device for z/OS and z/VM VSWITCH using Enhanced QDIO (EQDIO)
  2. NETH PFID (PCI device) for Linux
- Support for 1GbE, 10 GbE and 25 GbE planned

Host Operating System Converged Network Adapter Usage

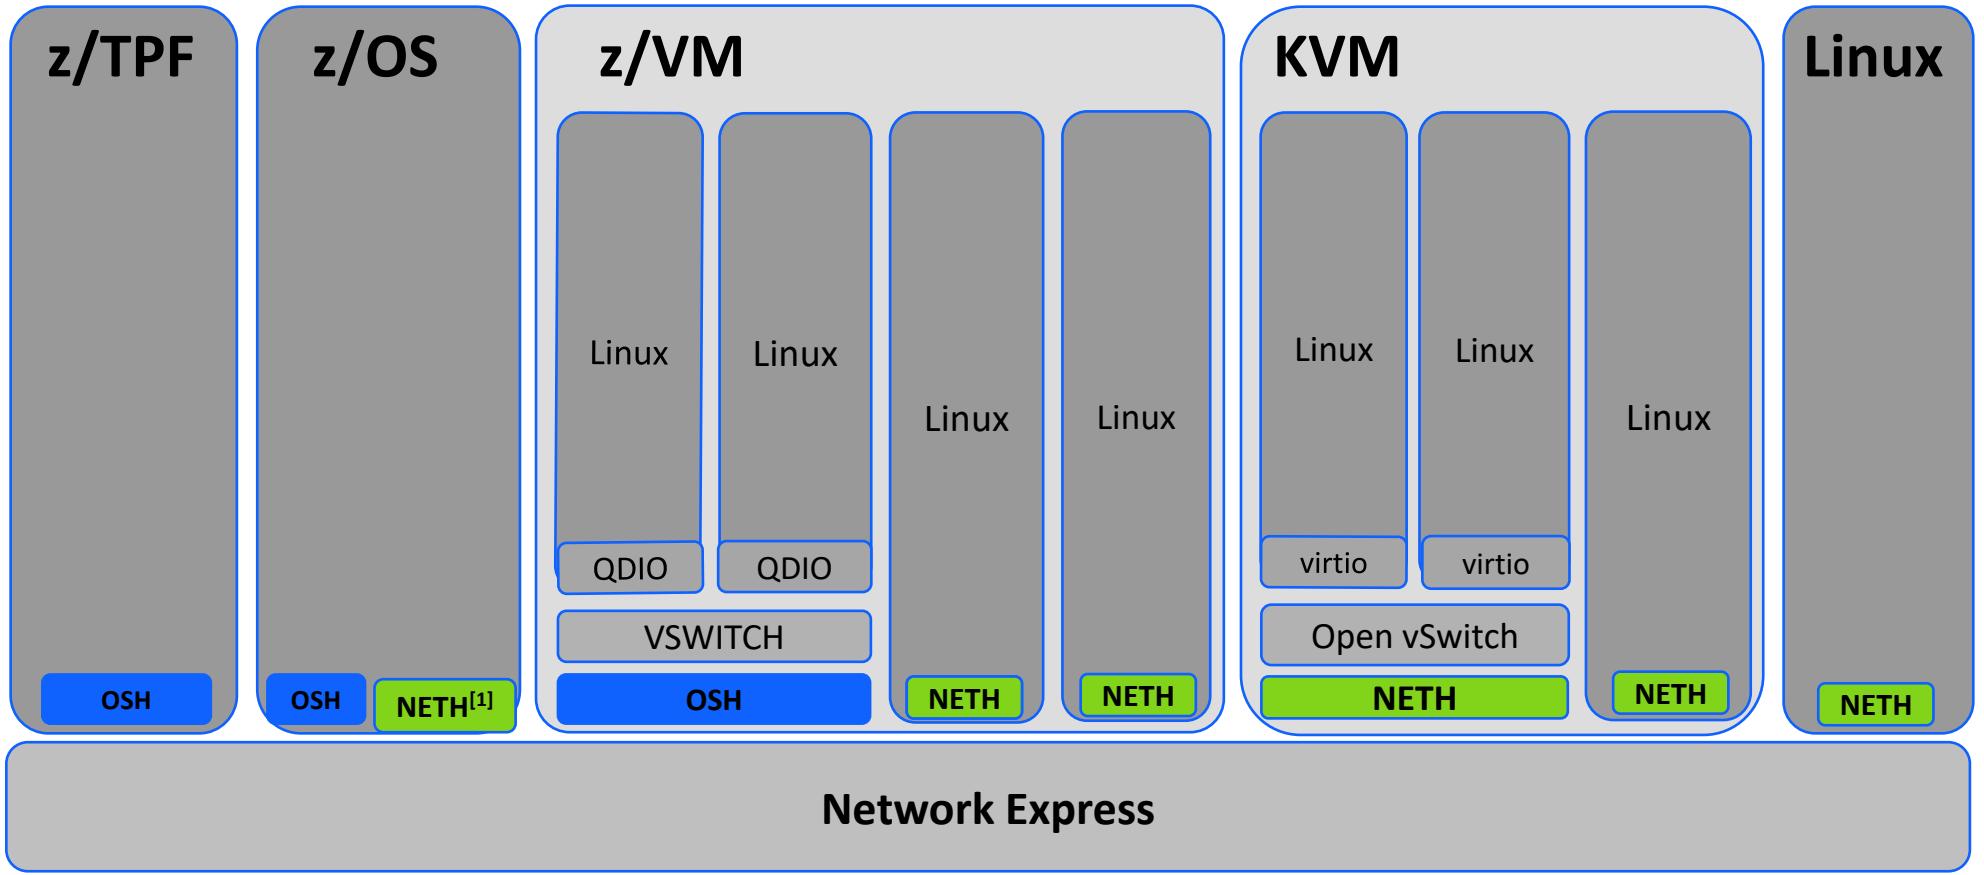

[1] SMC-R only

z/VM Adapter Usage

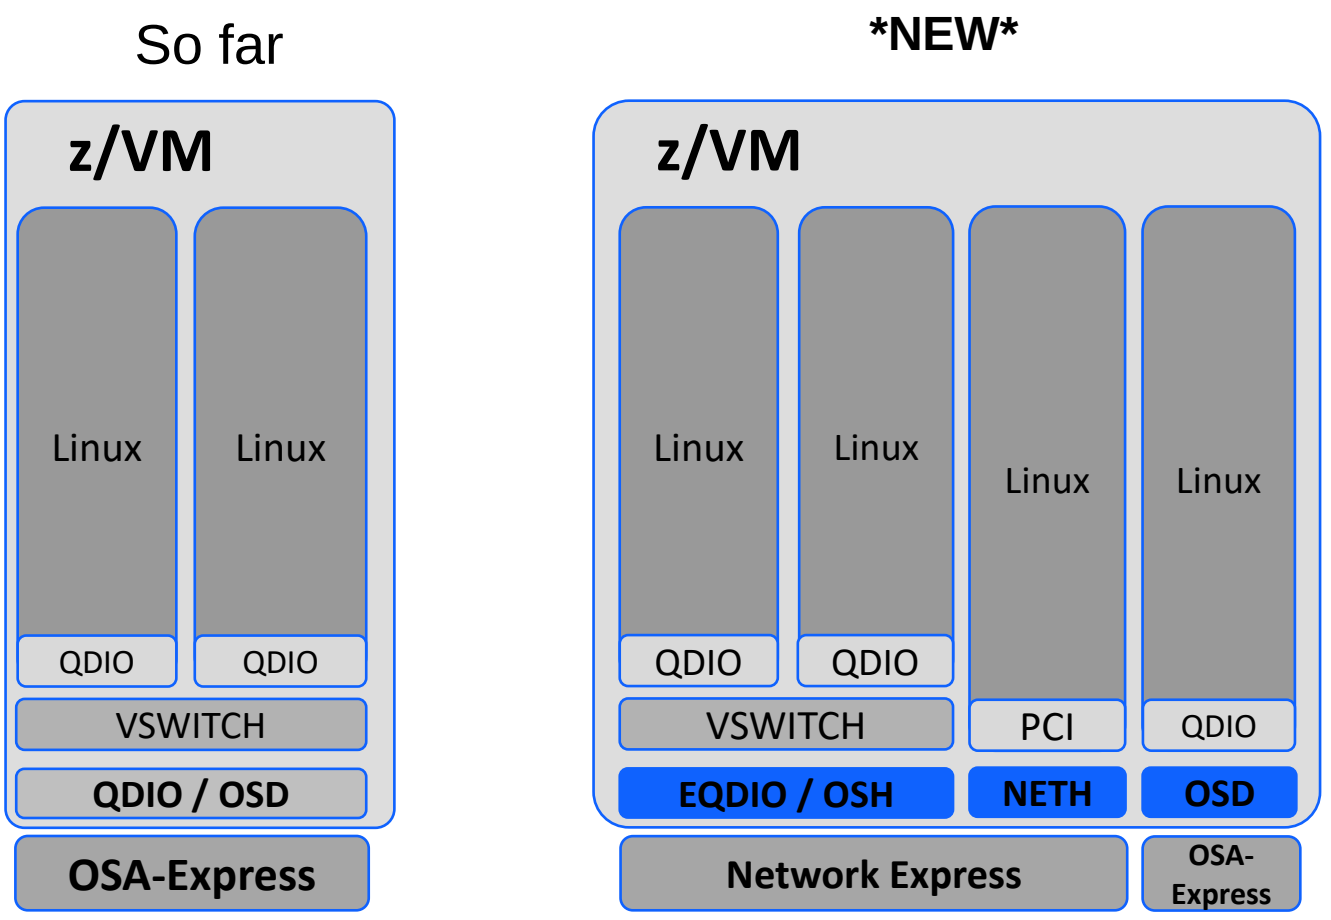

- VSWITCH-attached Linux guests can attach either
  - traditional OSA-Express cards with CHPID type OSD, or
  - A Network Express adapter using CHPID type OSH via EQDIO
- In both cases, z/VM Linux guests will continue to use the QDIO interface through the qeth device driver
- Direct-attached networking adapters to a z/VM Linux guest can use either of:
  - traditional OSA-Express cards with CHPID type OSD, or
  - A Network Express adapter using FID type NETH
- Note:
  - LGR will not work with PCI devices like NETH
  - Recommendation: Use VSWITCH instead with EQDIO/OSH

Linux / Networking

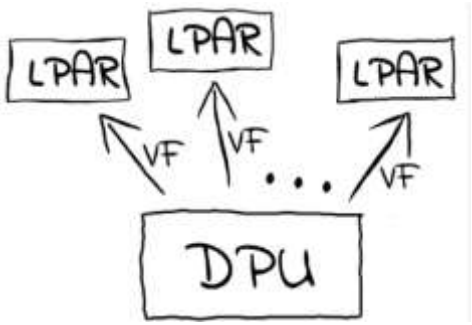

Data Processor Unit - fast I/O

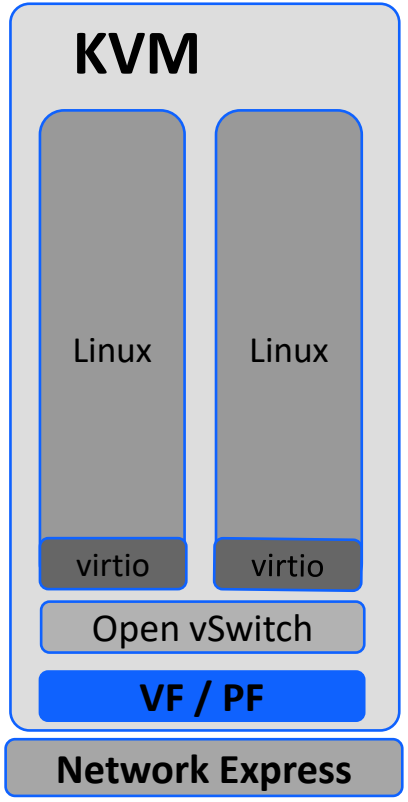

PCI-based net-working devices make for a better user experience

- Smoother transition from other platforms
- Eliminates need for IBM Z-specific tooling
- Better integration with Linux on IBM Z distributions
- Less training effort for personnel
- Native adapter virtualization
  ⇒ *Better integration with Linux ecosystem*
  ⇒ *Facilitates use of operators designed for full cloud-native SDN experience*

FID
Type
NETH

- Represents a virtual function (VF), similar to RoCE express adapters with on IBM z16 and earlier
- Can be defined in parallel to OSH CHPIDs on the same port, and sets entire card into hybrid mode
- Now supports promiscuous mode: In IOCDS, set FIDPARM for the respective device to "1"
- Usage Scenarios: Connectivity for Linux guests and virtual switches

Linux / Networking

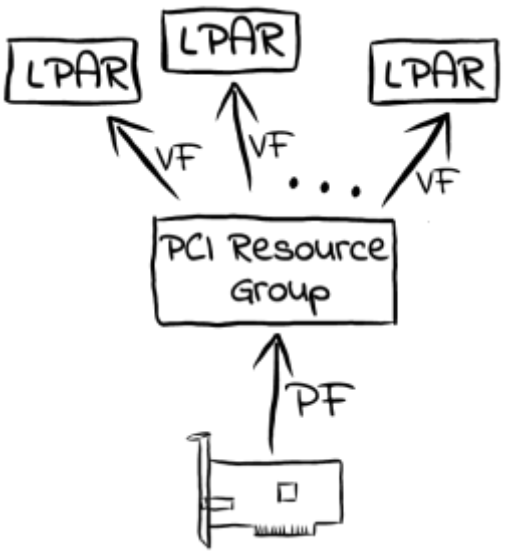

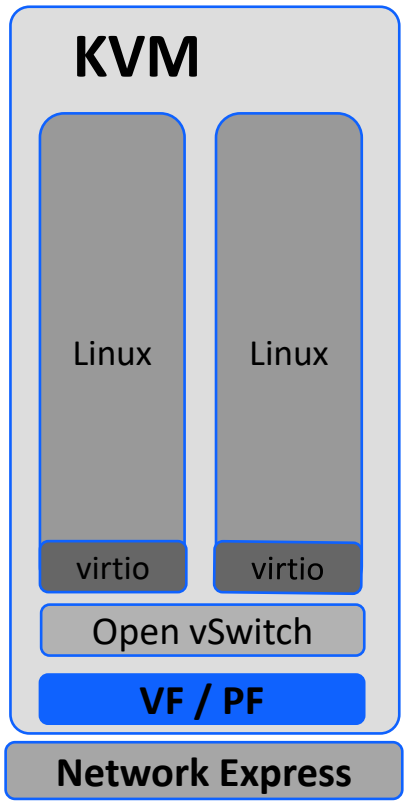

PCI-based net-working devices make for a better user experience

- Smoother transition from other platforms
- Eliminates need for IBM Z-specific tooling
- Better integration with Linux on IBM Z distributions
- Less training effort for personnel
- Native adapter virtualization
  ⇒ *Better integration with Linux ecosystem*
  ⇒ *Facilitates use of operators designed for full cloud-native SDN experience*

FID Type NETH

- Represents a virtual function (VF), similar to RoCE express adapters with on IBM z16 and earlier
- Can be defined in parallel to OSH CHPIDs on the same port, and sets entire card into hybrid mode
- Now supports promiscuous mode: In IOCDS, set FIDPARM for the respective device to "1"
- Usage Scenarios: Connectivity for Linux guests and virtual switches

Co-Location: SMC-Dv2

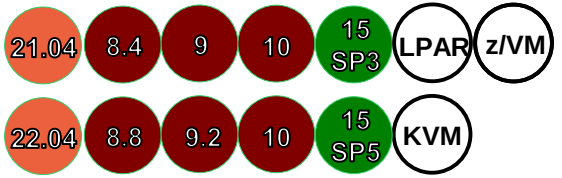

- **What it does**: Provides acceleration for TCP traffic
- **Why you should care**: v2 lifts limitations and greatly simplifies usage
- **Recap**
  - **Shared Memory Communications – Direct** provides intra-CEC acceleration for TCP traffic using *Internal Shared Memory* (ISM) devices
  - **Superior performance** (low latency, high throughput) at reduced CPU consumption
  - *However, SMC-Dv1 had limitations*:
    - Peers must be in **same IP subnet**
    - Devices need to be **paired using PNET IDs**
- **SMC-Dv2**
  - Peers can be in ***any* IP subnet**
  - No PNET IDs required
    ⇒ Simplified configuration!
  - Requires z15 or LinuxONE III
  - As with SMC-Dv1: Full **z/OS compatibility**
- **New performance paper available [here]**

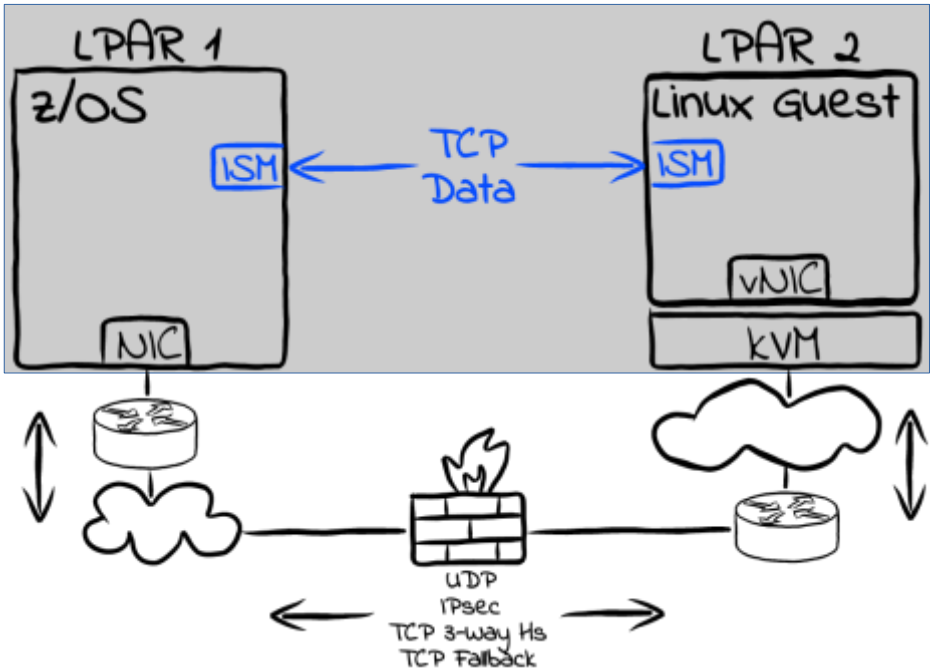

Fig.1: *Traffic flows with SMC-Dv2*

IBM Z © 2022 IBM Corporation

Eating our own dogfood: Leveraging SMC-Dv2 in IBM StorageScale

- **What we did**: Changed IBM Storage Scale to utilize SMC-Dv2 to benefit from superior performance and less CPU utilization
- **Why you should care**: Not only did this conversion yield strong results – it was also very easy to apply!

C++ code changes

- AF_INET sockets have been replaced with AF_SMC sockets for GPFS daemons: mmfsd and mmsdrserv (performance critical GPFS components)

```
socket(AF_INET, SOCK_STREAM, 0);
=>
socket(AF_SMC , SOCK_STREAM, 0);
```

- Other GPFS binaries, utilities, python/shell scripts are still using TCP

Python scripts changes

- New SMC-D Prerequisites Verification Tool: `tssmcdnodeverify`
- Enhanced `mmnetverify` tool: added SMC-D connections verification

IBM Z © 2022 IBM Corporation

Configuring IBM Storage Scale for SMC-Dv2 with Linux on Z nodes

1

Verify the Internal Shared Memory (ISM) device availability

```
# lspci | grep ISM
1014:00:00.0 Non-VGA unclassified device: IBM
Internal Shared Memory (ISM) virtual PCI device
```

2

Install `smc-tools` and `qclib` OS packages:

SLES: `zypper install smc-tools qclib`

RHEL: `dnf install smc-tools qclib`

3

Verify that SELINUX is set to "`permissive`" or "`disabled`"

SMC sockets are not included into the standard SELINUX policies and therefore SMC-D does not support SELINUX "enforcing" mode.

Co-Location: SMC-Rv2

22.04 8.7 9.1 10 15 SP4 LPAR z/VM

- **What it does**: Provides acceleration for TCP traffic
- **Why you should care**: v2 lifts same subnet limitation
- **Recap**
  - **Shared Memory Communications – Remote** (SMC-Rv1) provides acceleration for TCP traffic using *RDMA via RoCE Express adapters* devices
  - **Superior performance** (low latency, high throughput) at reduced CPU consumption
  - *However*: Peers must be in **same IP subnet**
- **SMC-Rv2**
  - Peers can be in ***any* IP subnet**
  - Can run with a single RoCE Express adapter (TCP and RDMA traffic)
  - Requires a separate pair of OSA and RoCE devices in case of redundancy (utilizing bonding and link groups respectively)
  - Requires RoCE Express2 cards
  - Full **z/OS compatibility**

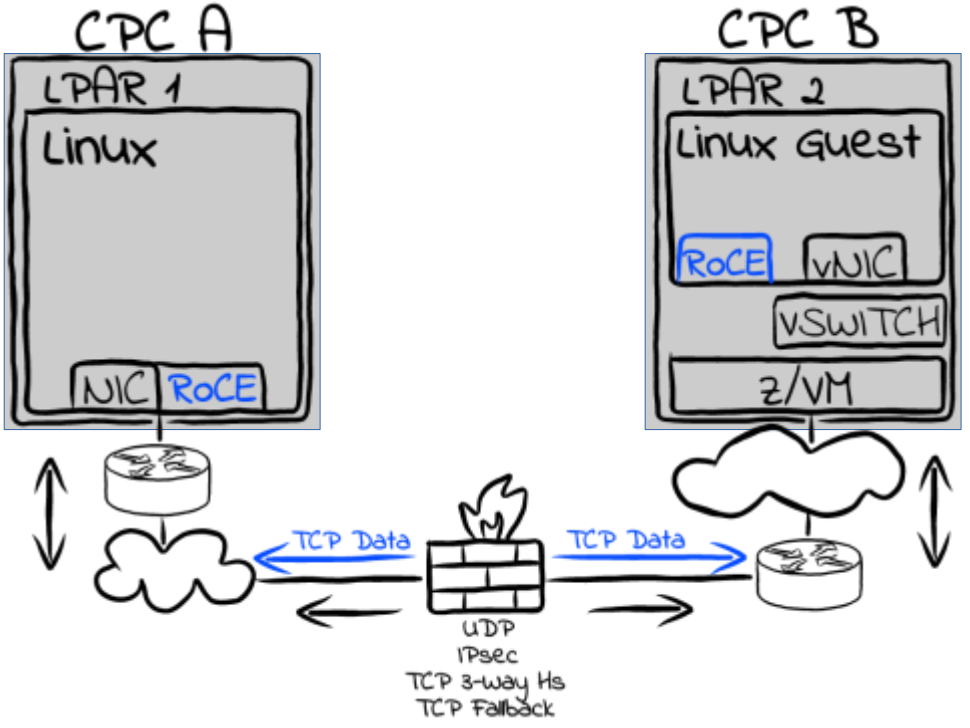

Fig.1: *Traffic flows with SMC-Rv2*

IBM Z © 2022 IBM Corporation

Co-Location: HiperSockets Converged Interface Multi-MAC Support

22.04 8.6 9 10 15 SP4 LPAR z/VM KVM

- **What it does**:
  - Forms a **single logical LAN segment** across
    - external LAN, and
    - HiperSockets
  - I.e. one IP to rule them all!
- **Why you should care**:
  - Simplifies network architecture
  - Provides HiperSockets Layer 2 connectivity to z/OS hosts
- **The best of both worlds**: Prefer HiperSockets over external LAN, taking advantage of the low latencies of HiperSockets in a simple setup
- Can register ***multiple*** **MACs now** – i.e. usable with Open vSwitch
- **Requirements/compatibility**:
  - z15 or LinuxONE III
  - HiperSockets VCHID with External Bridge option set
  - All participants accessing the same HiperSockets channel need to setup HSCI, too
  - Compatible with z/OS, providing layer 2 connectivity to z/OS over HiperSockets
  - In z/VM, use VSWITCH Bridgeport
  - In KVM uses macvtap, OVS or Linux bridge

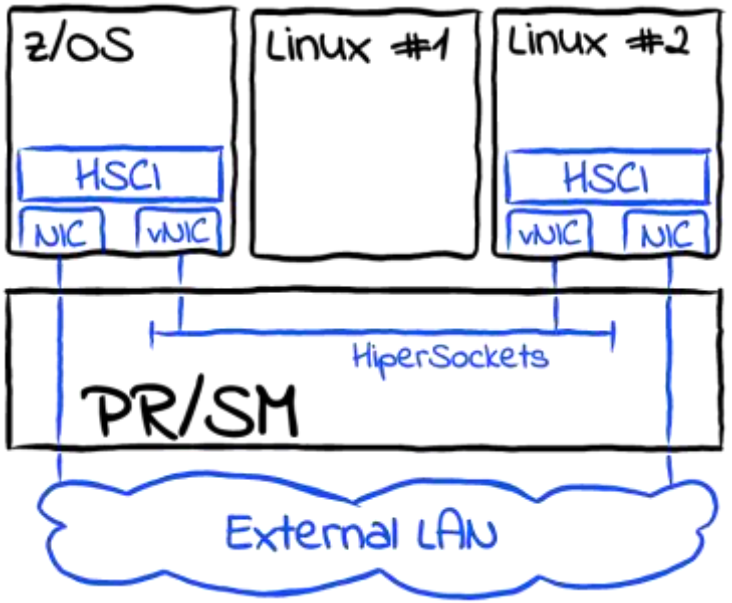


```
$ hsci add enc8410 encb040
Verifying net dev encb040 and HiperSockets dev enc8410
Adding hsci8410 with a HiperSockets dev enc8410 and an external dev encb040
Set enc8410 MAC fe:c2:f4:35:00:12 on encb040 and hsci8410
Successfully added HSCI interface hsci8410

$ ip addr add 192.1.2.3/16 dev hsci8410

$ hsci show
HSCI      PNET_ID  HiperSockets  External
-----------------------------------------
hsci8410  NET1     enc8410       encb040
```


IBM Z © 2022 IBM Corporation

Linux / Storage

Secure Boot for ECKD DASD

LPAR z/VM

- **What it does:** Linux can boot from ECKD DASD in Secure Boot mode
- **Why you should care:** Secure Boot is a prerequisite for the NIAP certification, and deployment of Linux in environments with extra high security requirements
- **What you need:**
  - IBM z16 with GA1.5 firmware
  - **For Basic boot support:**
    - s390-tools v2.25

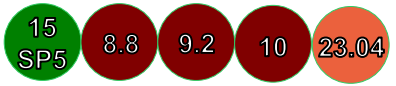

  - **For Reboot and dump support:**
    - s390-tools v2.26
    - Linux kernel v6.2

- **How to use it:** With the new support, Linux DASDs contain 2 types of boot loader:
  - CCW IPL: Standard boot
  - LD-IPL ("List-Directed IPL"): Supports Secure Boot
- **Note: Secure Boot can only be enabled/disabled on the HMC Load panel**

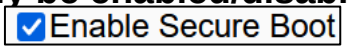

- **zipl will always install both boot loader types:**

```
$ zipl
Using config file '/etc/zipl.conf'
...
Preparing boot device for CCW- and LD-IPL: dasda (1234).
Done.
```

- **For reboot, IPL-type must be chosen manually**
  - `chreipl eckd` for DASD LD-IPL with Secure Boot support
  - `chreipl ccw` for DASD CCW-IPL with standard boot

```
$ chreipl eckd 0.1.1002
Re-IPL type: eckd
Device:      0.1.1002
bootprog:    0
br_chr:      auto
Bootparm:    ""
Loadparm:    ""
clear:       0
```

```
$ chreipl ccw 0.1.1002
Re-IPL type: ccw
Device:      0.1.1002
Loadparm:    ""
clear:       0
```

IBM Z © 2022 IBM Corporation

IBM LinuxONE support for NVMe drives

- **What it does:** Provide adapter for NVMe
  - Carrier card for industry standard U.2 NVMe drives
    - Common capacities up to 16 TB per drive
    - 1 drive per carrier, up to 16 cards per CEC
  - Available for IBM LinuxONE starting with Emperor II and Rockhopper II

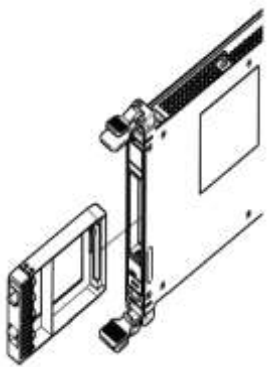

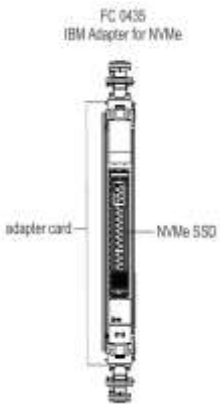

- **Why you should care**:
  - Low-cost
  - Low-latency
  - High-throughput
- **NVMe drive characteristics**
  - PCI direct-attached (no SAN)
    - No cabling, switches, etc. required
  - No classic virtualization or shared access: can use one drive only in one LPAR/VM
  - KVM can split a single NVMes into multiple partitions/LVMs for multiple guests
- **Linux on Z support for NVMe**
  - Uses standard Linux NVMe driver
  - Always apply latest service levels!

Working with NVMe drives in Linux

- **Listing available devices**
  - Use `lspci` to show PCI device information
  - Use `lsblk` to list NVMe block devices:
  -
```
$ lsblk
NAME         MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
nvme0n1      259:0    0  3.6T  0 disk
├─nvme0n1p1  259:1    0  300M  0 part
├─nvme0n1p2  259:2    0   40G  0 part
├─nvme0n1p3  259:3    0  3.6T  0 part
└─nvme0n1p4  259:4    0    2G  0 part
```
- **Storing data**
  - Use NVMe for swap, root, boot* and data file systems (*=requires IPL support)
  - Use of software RAID recommended

- **Management tools**
  - `nvme-cli`: query and manage NVMe device functions
  - `zpcictl`: recovery and service actions
- **IPL support** — 20.04, 8.3, 9, 10, 15 SP3, LPAR
  - Available on LinuxONE III with latest firmware
    - LPAR only
  - KVM supports IPL via virtio-blk — KVM
  - NVMe IPL support added with
    - Linux kernel v5.8
    - s390-tools v2.14
  - Distribution installer support available
- **Load normal support** — 21.04, 8.4, 9, 10, 15 SP3, LPAR
  - Reduces boot time for large LPARs
- **Dump support** — 22.10, 8.7, 9.1, 10, 15 SP5, LPAR
  - Full tools integration, e.g. `makedumpfile`
  - Requires IBM z15 or later

IBM Z © 2022 IBM Corporation

Co-Location: zdsfs

19.04 8.2 9 10 12 SP5 15 SP2 LPAR z/VM KVM

- **What it does**: Provides read-only access to z/OS data sets on DASD
- **Why you should care**:
  - Transfer speed closer to FICON limits
  - No CPU cycles spent in z/OS during data transfer
- **Latest Improvements**
  - **Robustness**
    - Use storage mechanisms to prevent concurrent access
    - Get an updated dataset list without need to unmount/mount DASDs
  - **Coordinated read access** 20.04 8.4 9 10 15 SP3
    - Establish authorization checking and audit capabilities
    - DASD no longer required offline in z/OS during Linux access
  - **Transparent dataset conversion** 21.10 8.5 9 10 15 SP4
    - Read EBCDIC encoded datasets as ASCII without any add'l codepage conversion required

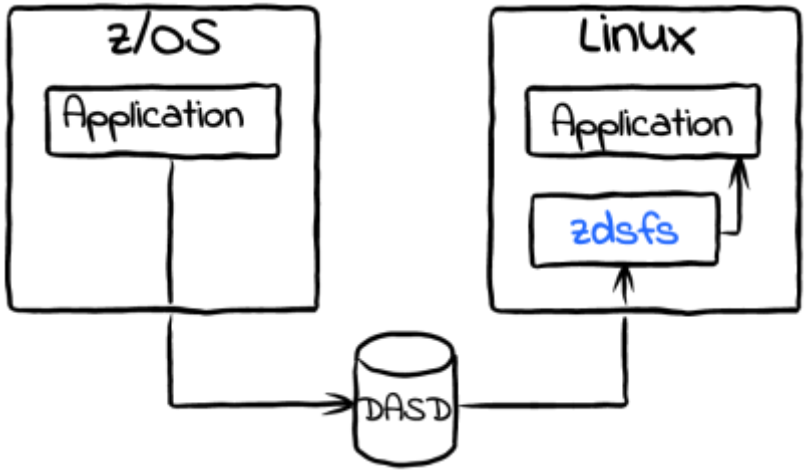

Fig.1: *Architecture overview*

```
$ zdsfs /dev/disk/by-path/ccw-0.0.edc0 /mnt/
Using z/OSMF REST services on https://192.168.1.1/zosmf/

$ cat /mnt/TEST.ZDSFSDV.LARGE.TEST2
[...]

$ fusermount -u /mnt
```

Fig.2: *Usage example*

IBM Z © 2022 IBM Corporation

Multi-Path Re-IPL

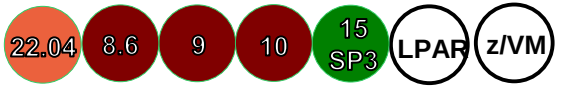

- **What it does**: Keeps re-IPL path up-to-date with working path for next re-IPL
- **Why you should care**: Protects *vastly* against re-IPL issues due to IPL path becoming unavailable.
- The `chreipl-fcp-mpath` toolset monitors udev events about paths to the re-IPL volume
- If currently configured FCP re-IPL path becomes unavailable, re-configures the FCP re-IPL settings alternative operational path to same volume
- Thus, re-IPL from an FCP-attached SCSI volume can be successful despite path failures on a running Linux instance if at least one path to the re-IPL volume remains operational
- **See man-page for `chreipl-fcp-mpath` for further details**

```
# Activate by installing respective package
$ apt install s390-tools-chreipl-fcp-mpath

$ lsreipl
Re-IPL type: fcp
WWPN:        0x5005076309005430
LUN:         0x4018401600000000
Device:      0.0.1700
bootprog:    0
br_lba:      0
Loadparm:    ""
Bootparms:   ""
# Cable Pull/Switch Port Toggle/Path Goes Away

$ journalctl --boot --identifier=chreipl-fcp-mpath
May 02 10:04:42 t3545003 chreipl-fcp-mpath[46089]:\
Changed re-IPL path to: 0.0.1740:0x5005076309045430\
:0x4018401600000000.

$ lsreipl
Re-IPL type: fcp
WWPN:        0x5005076309045430
LUN:         0x4018401600000000
Device:      0.0.1740
bootprog:    0
br_lba:      0
Loadparm:    ""
Bootparms:   ""
```

Fig.1: *Sample output*

IBM Z © 2022 IBM Corporation

Software Defined Storage Options for IBM Z and LinuxONE

IBM Storage Fusion

IBM Storage Scale

Storage for Data and AI

- **Global Data Platform**
- High performance AI/ML and Analytics
- Multiprotocol file / object
- File Store (w/ object support)

IBM Fusion Data Foundation

Storage for Hybrid Cloud

- Open-Source model
- Hybrid Cloud Storage
- S3 protocol fidelity
- Object, Block and File
- Ships with
  - **Storage Scale**
  - **IBM Fusion Data Foundation**
- Includes **CNSA plugin** to access Scale

IBM Storage Fusion: Software Stack

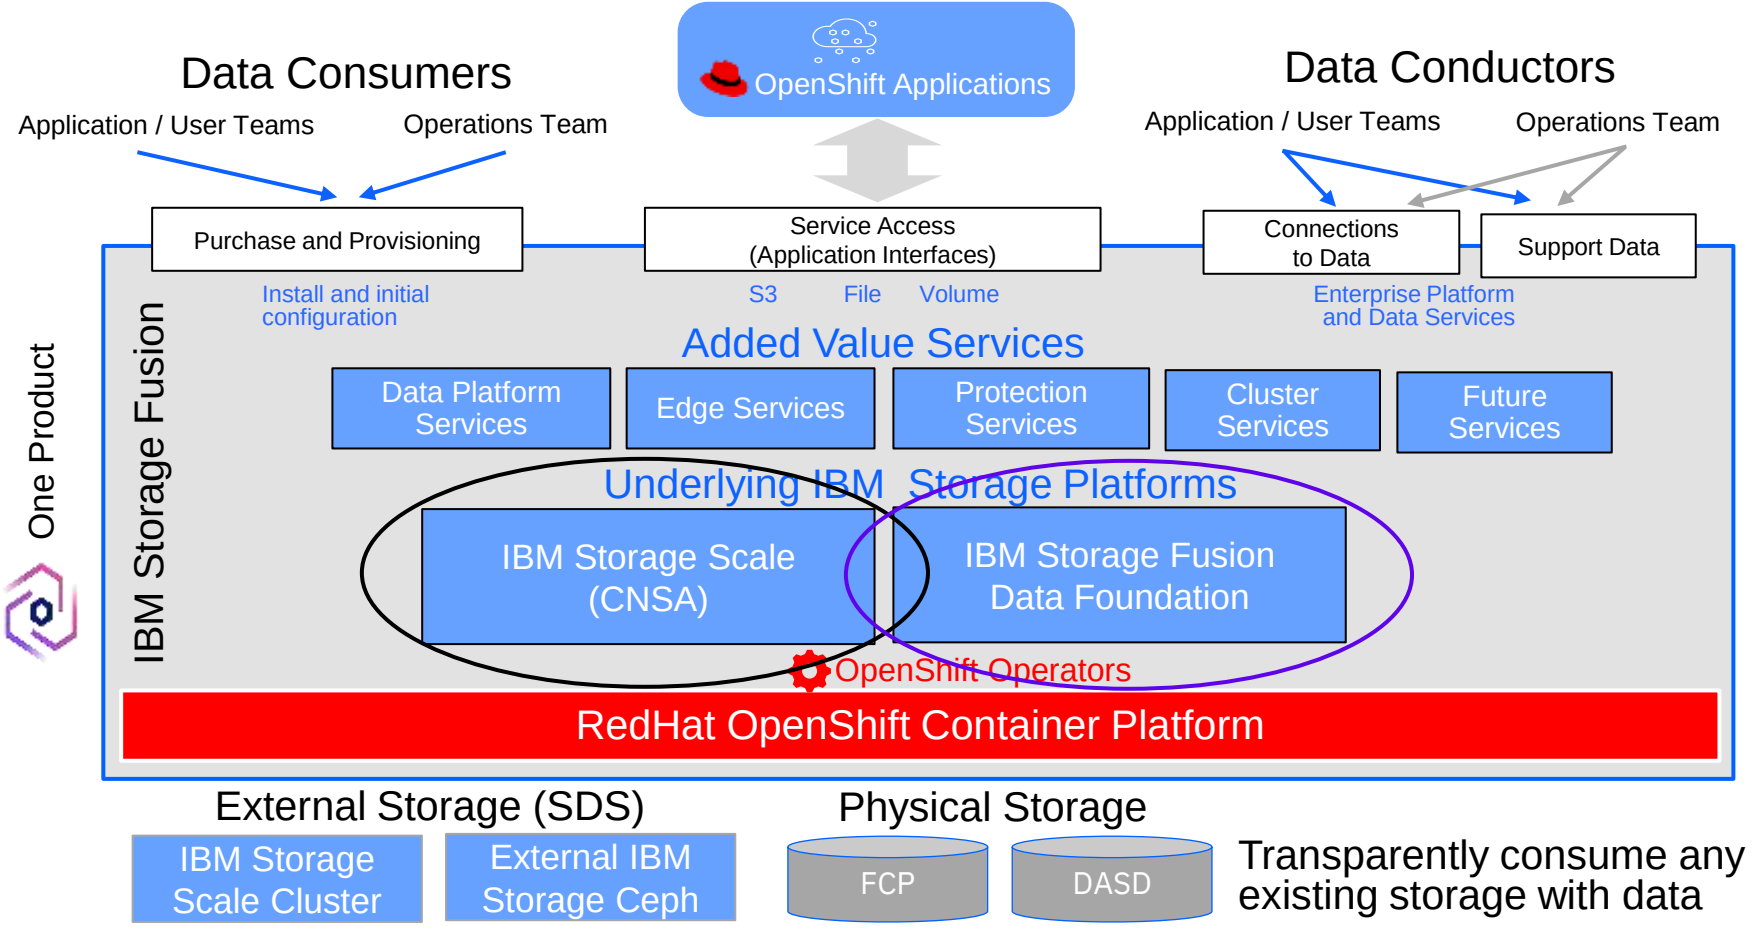

Linux / Crypto

Linux Exploitation of New CPACF Functions

- Improved AES-XTS support
  - Acceleration for smaller buffers
  - Improved security for protected key function
  - Usage scenarios: dm-crypt, DBs
- New SHA2-HMAC support
  - Acceleration for small buffers, key derivation
  - Enablement for protected key HMAC
  - Usage scenarios: dm-integrity, TLS
- Improved SHA3 & SHAKE support
  - Acceleration for small buffers
  - Usage scenarios: QS crypto deployments
- CPACF FW versioning
  - hash of CPACF FW code
  - Usage scenarios: Crypto certification labs

Linux In-Kernel Crypto

- improved AES-XTS clear key support
- improved PAES-XTS protected key support
- improved SHA2-HMAC clear key support
- new SHA2-HMAC protected key support
- improved SHA3 support

openSSL

- improved AES-XTS clear key support
- improved SHA2-HMAC clear key support
- improved SHA3 and SHAKE support

PAI counter support for new CPACF functions in IBM z17 / LinuxONE 5

`cpacfstats` tool update

- Usage scenario: Monitor in-use ciphers, e.g. to detect usage of deprecated ciphers
- Ships with s390-tools
- Needs to run `cpacfstatsd` daemon with root permissions
- User must be in `cpacfstats` group
- User can access information provided by the daemon through the `cpacfstats` tool
  - Enable/disable CPACF Crypto Activity counters
  - Display counters after use (with optional filter to skip zero counters) to show usage of the CPACF functions
  - Reset counters to 0
  - Convert output to JSON format

```
$ openssl speed -hmac SHA-256
[snip]
$ cpacfstats -n
 des counter: 0
 aes counter: 0
 sha counter: 15445553
 rng counter: 0
 ecc counter: 0
 pai_user   : enabled
  (162) KMAC HMAC SHA 256       : 1321590
```

New `cpacfinfo` tool

Usage scenario: Inspect crypto HW capabilities

Ships with s390-tools

Display Query Authentication Information (introduced with MSA 13)

Display available CPACF instructions and functions

Display detailed MSA information

Filter output by CPACF instruction

Filter output to only show enabled/disabled functions

Output in raw JSON format available

Uses new kernel interfaces

- `/sys/devices/system/cpu/cpacf/<instruction>_query_raw`
- `/sys/devices/system/cpu/cpacf/<instruction>_query_auth_info__raw`

```
$ cpacfinfo -fi km
Cipher Message (KM)
  Format: 0; Hash length: 32; Hash version: 24
  Hash:
    47 bb 52 ba a7 4e b0 a0 13 e5 f6 7e 04 f5 68
d5
    5c d1 d0 b3 ae ad 20 da 32 7a 17 0a 04 64 0c
af
        (  0) [    AVAILABLE] KM-Query
        (  1) [    AVAILABLE] KM-DEA
        (  2) [    AVAILABLE] KM-TDEA-128
        (  3) [    AVAILABLE] KM-TDEA-192
        …
```

CEX8S exploitation (including Quantum Safe crypto)

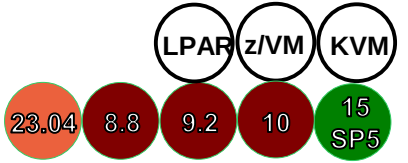

- Download page for latest **CCA host package** and **EP11 support program** is here
- Provide libraries to communicate with CEX cards, tools and a TKE proxy
- Free access after license acceptance

CCA CEX8S toleration

- Requires CCA host package 7.3.45

CCA CEX8S exploitation

- Requires CCA host package 8.0
- Provides support for full CEX8S CCA functions includes quantum safe algorithms:
  - Round 2 Kyber (key exchange)
  - Round 3 Dilithium (signature)
- Provides support for TLS connection between the TKE and catcher daemon
- See further documentation here

EP11 CEX8S toleration

- Requires EP11 support program 3.1.0

EP11 CEX8S exploitation

- Requires EP11 support program 4.0
- Includes quantum safe algorithms:
  - Round 2 Kyber (key exchange)
  - Round 3 Dilithium (signature)
  - Provides support for TLS connection between TKE and ep11TKEd daemon
- Requires openCryptoki 3.20 (upstream)
- EP11 token supports IBM specific mechanisms for Kyber and Dilithium

IBM Z © 2022 IBM Corporation

OpenSSL pkcs11 Signing Provider

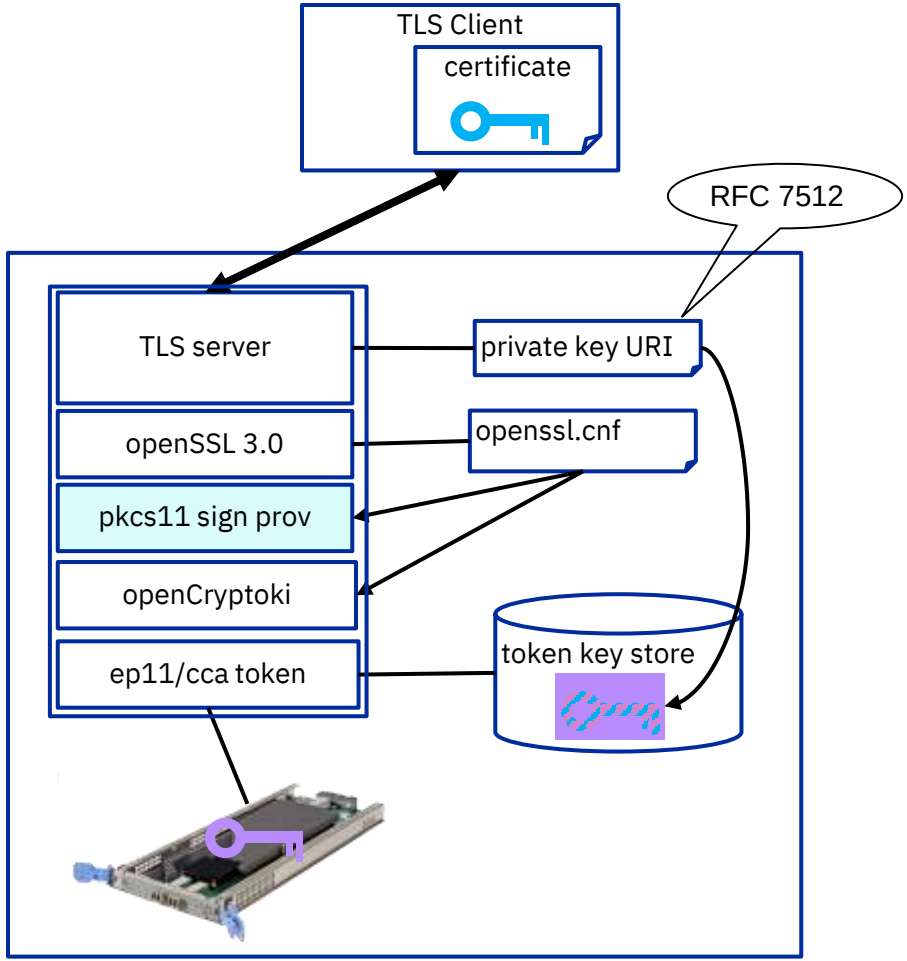

15 SP6 | 10 | 23.10 | LPAR | z/VM | KVM

Problem

– a hacker who steals the private signing key of a TLS server can impersonate the TLS server

Solution

– protect the private signing key with an HSM

Note

– all other keys of a TLS connection are ephemeral and therefore less critical

Release 1.0 of openssl-pkcs11-sign-provider released on https://github.com/opencryptoki/openssl-pkcs11-sign-provider

p11sak from openCryptoki 3.21 supports key URIs

supports

- ECDSA
- RSA sign
- RSA decrypt

restriction

- process must not fork

IBM Z © 2022 IBM Corporation

Master Key Change Protocol for openCryptoki – so far

Disruptive Procedure!

CCA

Re-encipher token key repository of openCryptoki CCA token

1. *! Stop all processes using openCryptoki CCA token*
2. *Perform MK change on HSM*
3. *Use the `pkcscca` tool to re-encipher token keys*
4. *Restart processes using openCryptoki CCA token*

Note: CCA has 4 different MKs for different key types: (3)DES, AES, RSA, ECC

- ***Each MK can be changed independently***

Disruptive Procedure!

EP11

Re-encipher token key repository of openCryptoki ep11 token

1. *! Stop all processes using openCryptoki EP11 token*
2. *Commit new MK on HSM*
3. *Use the `pkcsep11_migrate` tool to re-encipher token keys*
4. *Activate new MK on HSM*
5. *Restart processes using openCryptoki EP11 token*

IBM Z © 2022 IBM Corporation

Linux / AI

Improved On-processor AI Acceleration

- Integrated as a CISC instruction
- Remote AI accelerator
- Support for LLM compute primitives
- Int8, FP16 datatypes

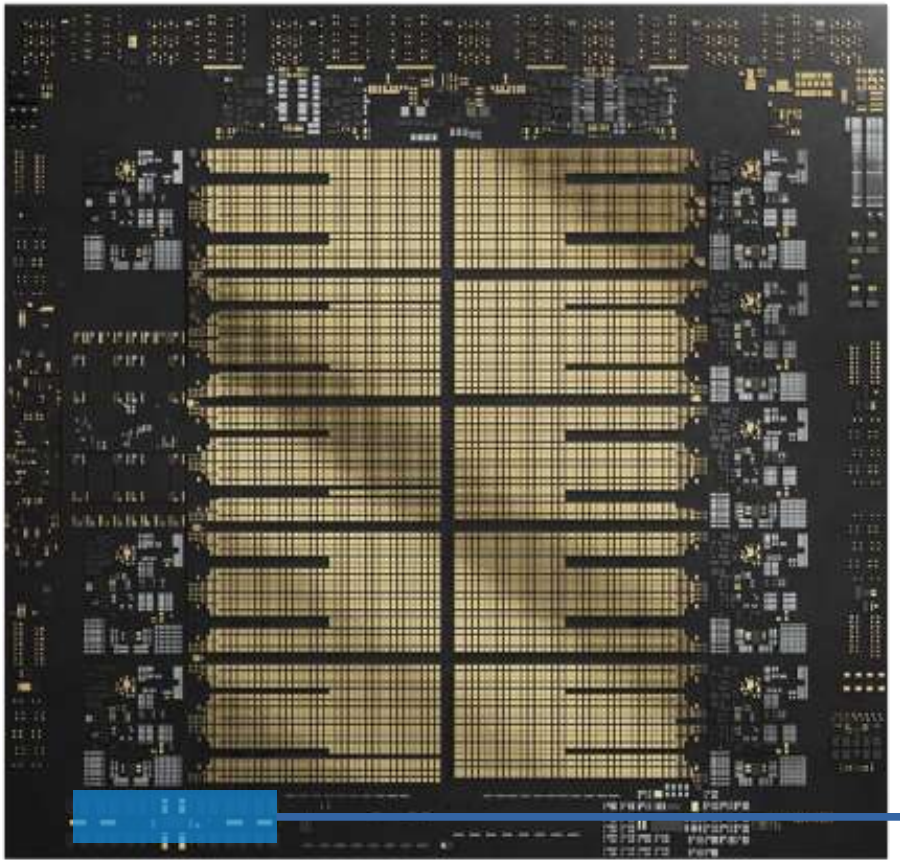

Telum II: Improved On-Chip AI Acceleration

- Integrated as a CISC instruction
- Support for LLM compute primitives
- Int8, FP16 datatypes
- Remote AI accelerator

- On z16, all cores on a chip (up to 6 user cores) would share the local on-chip accelerator; what can lead to serialization of AI work running on the same chip
- There are 8 on-chip accelerators available within a drawer
- z17 and LinuxONE 5: these accelerators are shared across all cores in that drawer, allowing transparent load balancing of AI work on all 8 accelerators
- AI work is spread across accelerators transparently always using the closest to the AI workload accelerator and giving on average equal runtime on the accelerator to all AI workloads
- Note: There is some penalty for using a remote accelerator

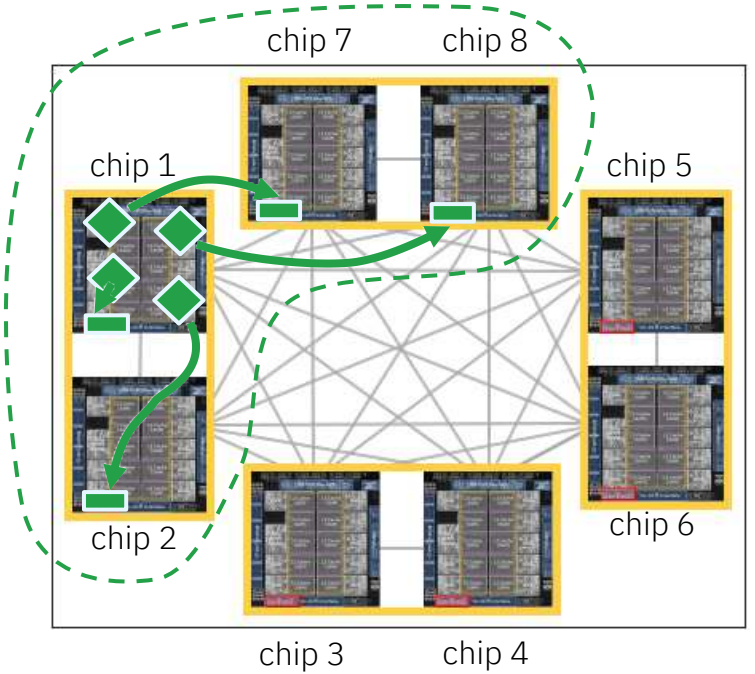

Unlock potential with accelerating computing on IBM® LinuxONE Emperor 5

In-transaction AI with encoder LLMs and multiple AI model techniques

2nd Gen on-chip AI accelerator in Telum II

- Support for LLM compute primitives
- Improved quantization and matrix operations
- Improved AI processing over IBM® LinuxONE 4

AI workload balancing during peak usage

Intelligent routing

- Remote AI processing
- Up to 8x AI processing available

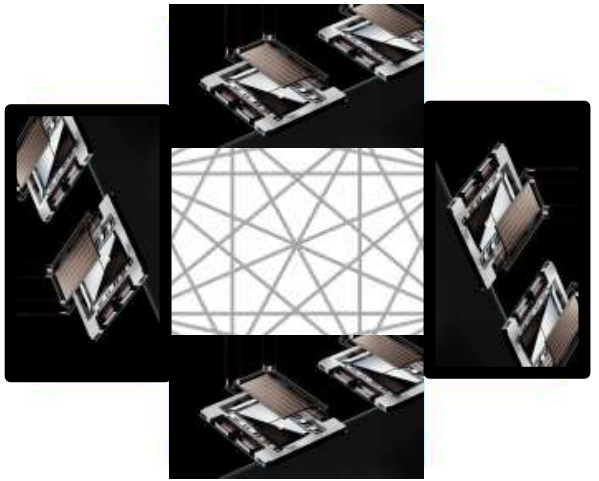

Optimize generative AI and LLM use cases

IBM Spyre Accelerator cards, available via PCIe card in 4Q2025

- 32 Gen AI-ready cores per adapter card
- Up to 48 adapter cards per syst

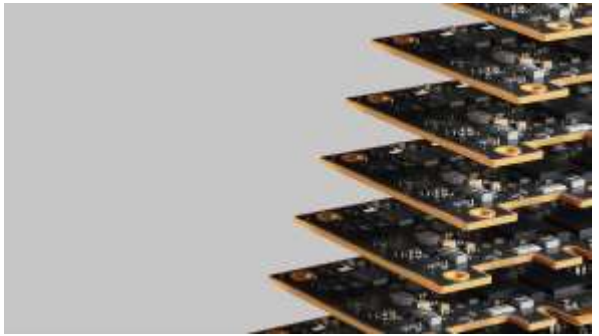

© 2025 IBM Corporation

AI Ecosystem – Hardware Exploitation

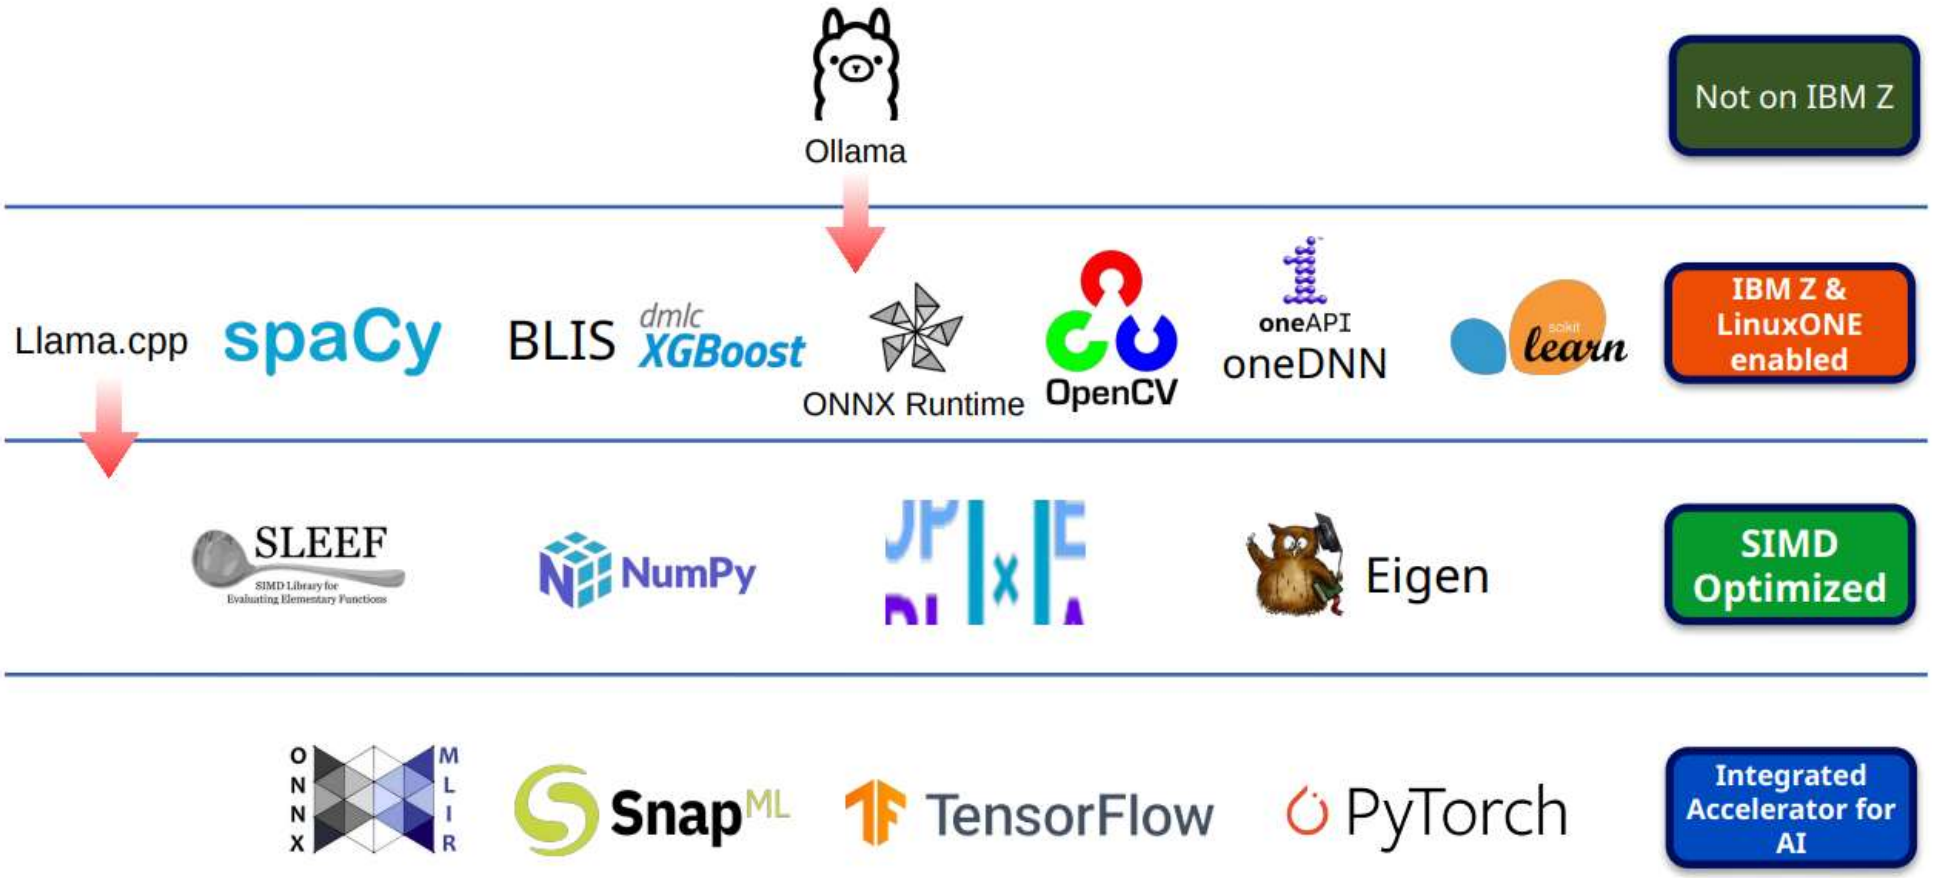

Linux / Misc

HiperDispatch Support aka Vertical CPU Polarization

- **What it does:** Prioritize process scheduling to CPUs with more consistent processing guarantees to avoid steal time
- **Why you should care:** Can yield substantial performance improvements for CPU-intensive workloads on highly utilized CECs
- Platform differentiates between *vertical high*, *medium* and *low* IFLs, with varying capacity grants
- Basically no steal time on *vertical high* IFLs
- Modifies the scheduler to prefer *vertical highs* and *mediums* for CPU-intensive workloads
- Workloads running large numbers of small tasks might perform better with horizontal CPU polarization
- **How to use:**
  - Enabled by default
  - Use sysctl `s390.hiperdispatch` to enable or disable:

```
sysctl -w s390.hiperdispatch=[0|1]
```


Tunables:

- `/sys/devices/system/cpu/hd_steal_threshold`
- `/sys/devices/system/cpu/hd_delay_factor`
  Steal time evaluation period. Reducing this value improves responsiveness to changes in workload behavior. Increasing it delays reaction to sudden changes in steal time.

IBM Z © 2022 IBM Corporation

Kernel Live Patching

22.04 15 SP1 LPAR z/VM KVM

- **What it does**: Applies patches to a running Linux kernel by dynamically modifying the kernel code in memory.
- **Why you should care**: Further reduces downtime: Apply critical security fixes to a live system without the need of a reboot, reducing service disruptions
- Hooks into the kernel function call sites and redirects them to the patched versions of functions.
- No need to wait till the next service window or to implement unplanned downtimes: Apply the latest security vulnerabilities fixes without need for reboot
- Less impact to critical systems
- Variants to live-patch userspace code exist – no disruptions due to restarts of your most critical applications!
- Check your Linux distribution's manual for further details on how to use, e.g. SLES 15, and Ubuntu 22.04

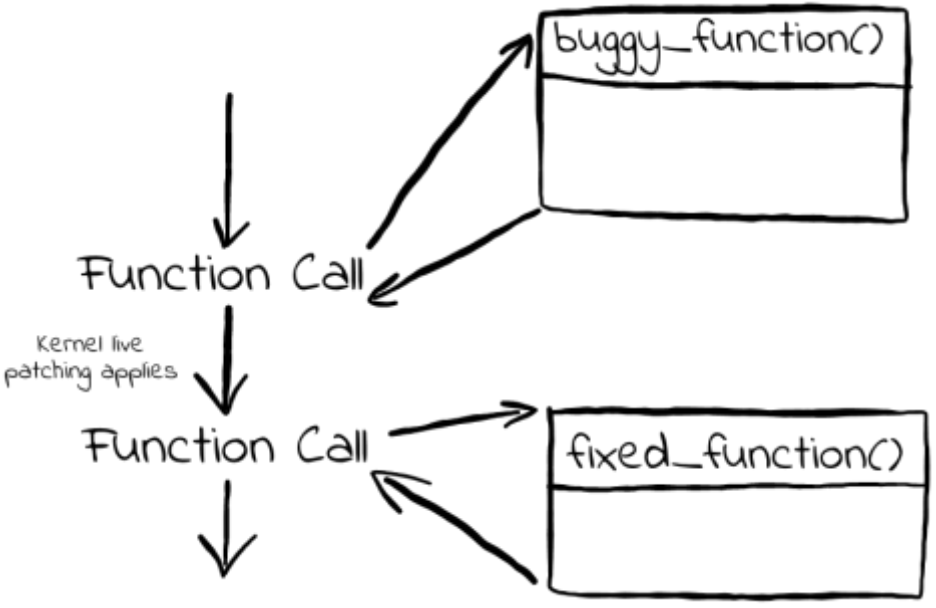

Fig. 1: *Kernel Live Patching schematic illustration*

IBM Z © 2022 IBM Corporation

Removal of 32-bit Support

LPAR z/VM KVM

- **What it does**: Remove the ability to run ***any*** 32-bit application in forthcoming Linux distributions
- **Why you should care**: If you are still deploying 32-bit applications in ***any*** shape or form, it is about time to plan for a migration
- If kernel still supports 32-bit, userspace can run in 32-bit, provided all necessary libraries are available in 32-bit, too!
- Ways to provide 32-bit libraries:
  - Use in distro (see `compat` packages)
  - Link statically
  - Package in containers

Distro	Properties	RHEL	SLES	Ubuntu
32-bit distro	Full 32-bit distro available	≤ RHEL 4	≤ SLES 9	-
64-bit distro	32-bit userspace	RHEL 5 RHEL 6 RHEL 7	-	-
	32-bit compat packages only	-	SLES 10 SLES 11 SLES 12	≤ Ubuntu 22.04
	32-bit compat kernel support, statically linked 32-bit & containers still work	RHEL 8 RHEL 9	SLES 15	-
	No 32-bit support at all	**RHEL 10**	**SLES next**	**Ubuntu 24.04**

Fig. 1: *Distro support overview*

IBM Z © 2022 IBM Corporation

Linux / Power Consumption

25.04 LPAR

- **Usage scenario**: Determine carbon footprint
- Data available in binary format through `IOCTL DIAG324_GET_PIBBUF` via `/dev/diag`
- Use convenience tool **`zpwr`** (comes as part of s390-tools) to display data in human-readable format
- Used data as-is, or perform further processing, e.g. to derive the per-guest portion in a KVM environment.
- The power information will consist of:
  - CPC total power
  - CPC unassigned resources power
  - CPC infrastructure power
  - LPAR processor power
  - LPAR I/O power
  - LPAR memory power
- See man page of **`powerinfo`** for further details

```
Usage: zpwr [OPTIONS]

Power readings of s390 computing environment

OPTIONS
--format FORMAT    List data in specified FORMAT (json
                   json-seq pairs csv)
-d, --delay NUMBER Power readings after delay
                   (seconds)
-c, --count NUMBER Number of power readings
-s, --stream Power readings in stream mode
-h, --help         Print this help, then exit
-v, --version      Print version information, then exit
```

Installation Assistant for Linux on IBM Z & LinuxONE

- **What it does**: Creating configuration files for starting Linux on IBM Z and LinuxONE installations
- **Why you should care**: Writing parameter files can be a challenge, with bugs triggering cycles with lengthy turnaround times
- Generates installer parameter files for the latest RHEL and SLES Linux distributions
- Supports OSA and PCI networking devices, IPv4/v6, and VLAN installations
- Provides easy to follow step by step instructions and context help
- Access at https://ibm.github.io/liz/

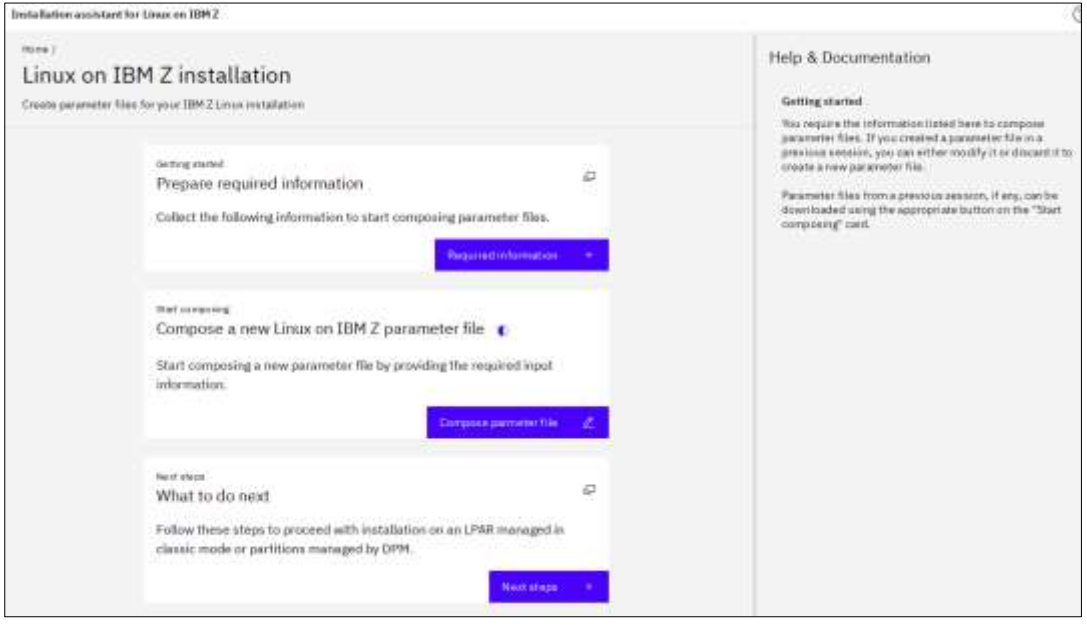

Fig. 1: *Landing page of the installation assistant*

IBM Z © 2022 IBM Corporation

Linux on IBM Z – Content Design
Distribution information

- Distribution-specific versions of Device Driver book discontinued
- However: Upstream-version maps features to Linux distribution releases
- On the Distributions page, click Distribution information

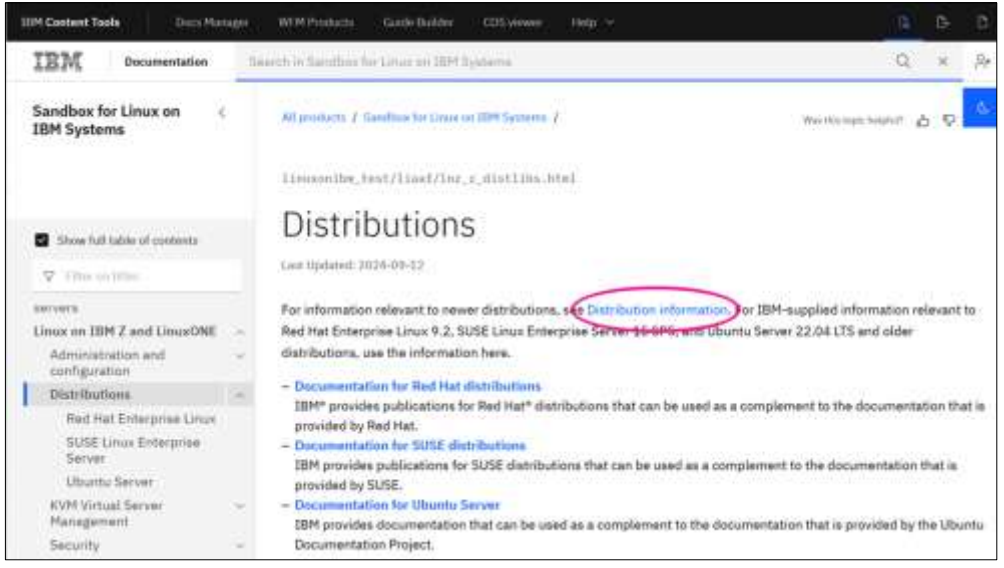

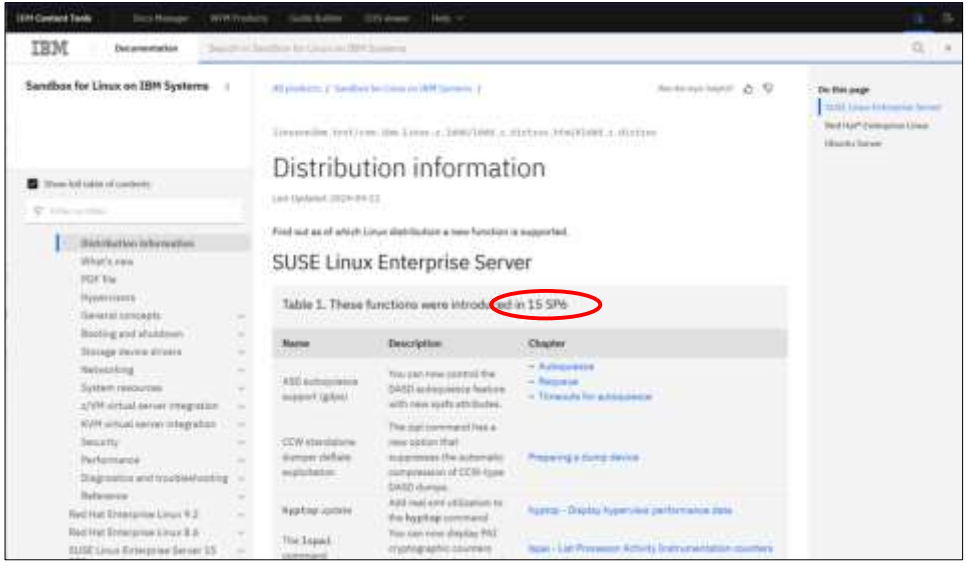

Compilers

- **Usage Scenarios**: Optimize homegrown applications to get the most out of the new HW generation
- Development focus on ***GCC*** and ***LLVM***
- **IBM z17 & LinuxONE 5 support** available in GCC 15.1 and LLVM 20.1 or later as follows:
  - **`-march=arch15`** enables GCC and LLVM z17 & LinuxONE 5 intrinsics
    - Vector-Enhancement Facility 3
    - Miscellaneous-Instruction-Extensions Facility 4
  - **`-mtune=arch15`** for z17 & LinuxONE 5 specific instruction scheduling
    - No new instructions used (does not require a zNext to run)

IBM Z © 2025 IBM Corporation

Libraries & Debuggers

- **GNU C Library Support**
  - Hardware capability flags indicate z16 facilities
  - Included in Glibc 2.34
  - Ubuntu 22.04 or later
- **GNU Debugger GDB**
  - z15 instruction set support for record and replay (reverse debugging)
    - Upstream in GDB 9.1
    - Available via RHEL 8.3, SLES 15 SP2, and Ubuntu 20.04
  - z16 support in progress
- **Valgrind - Memory Debugger**
  - Full instruction set translation to and from intermediate language required
  - z15 support available as of Valgrind v3.17, available with RHEL 8.5, Ubuntu 21.10
  - z16 support in progress
- **Perf – Performance Profiling**
  - Includes support for IBM Z hardware sampling and counter facilities
  - z15 support available with RHEL8.0, RHEL7, SLES15, SLES12 & Ubuntu19
  - z16 support in progress (upstream 5.17)

Activation profile settings for HW sampling and counter support

IBM Z © 2022 IBM Corporation

Red Hat Ansible Automation Platform

Deployment options on IBM Z / IBM LinuxONE

Red Hat Ansible Automation Platform runs on

- Linux
- Red Hat OpenShift in virtual machines
- IBM zCX Foundation for Red Hat OpenShift
- IBM z/OS Container Extensions

on IBM Z / IBM LinuxONE.

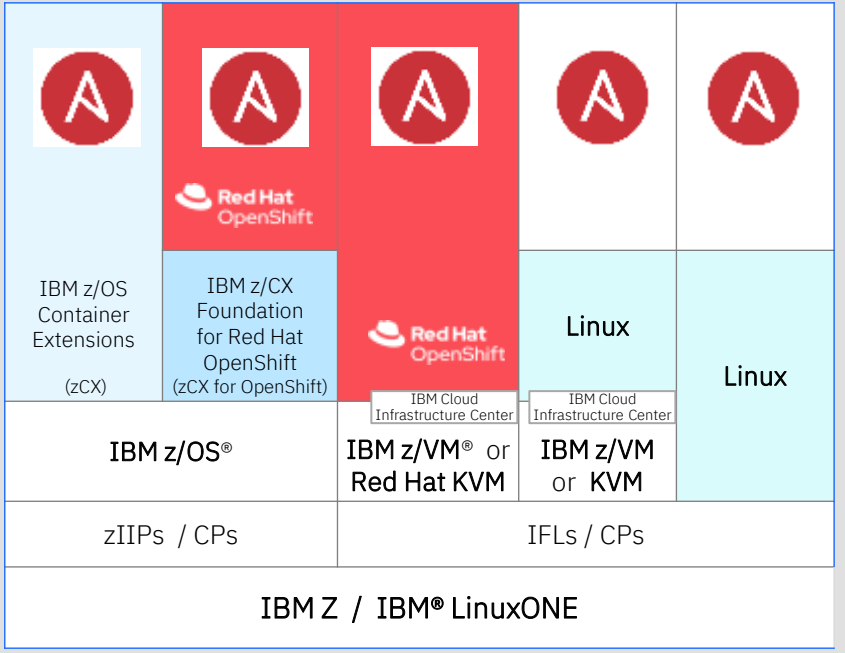

© 2025 IBM Corporation

Red Hat Ansible usage with IBM Z and IBM LinuxONE

Ansible can be helpful to work with/automate the following IBM Z / IBM LinuxONE environments as well:

- Linux on IBM Z / IBM LinuxONE
- IBM z/VM
- KVM on IBM Z / IBM LinuxONE
- Red Hat OpenShift
- IBM Cloud Infrastructure Center

Developers, administrators, and operators can benefit from **pre-existing certified content** to build from, for both building and testing.

Ansible Content Collections

Ansible content can be created and managed internally for your organizations to use. However, curated content is also available from Red Hat through Ansible Content Collections. These collections provide developers with the option of building on curated automation content, which includes more than 100 certified collections and more than 40,000 modules.

- Ansible Content Collections
- Getting Started With Ansible Content Collections

Ansible automation hub

This hosted service is the place for users to find and use supported Ansible Content Collections, which contains modules, roles, and plug-ins, along with the documentation needed to get started.

- Ansible automation hub

© 2025 IBM Corporation

Ansible Galaxy – Collections

https://galaxy.ansible.com/ui/collections/?page_size=10&view_type=list&sort=name&page=1&keywords=HMC

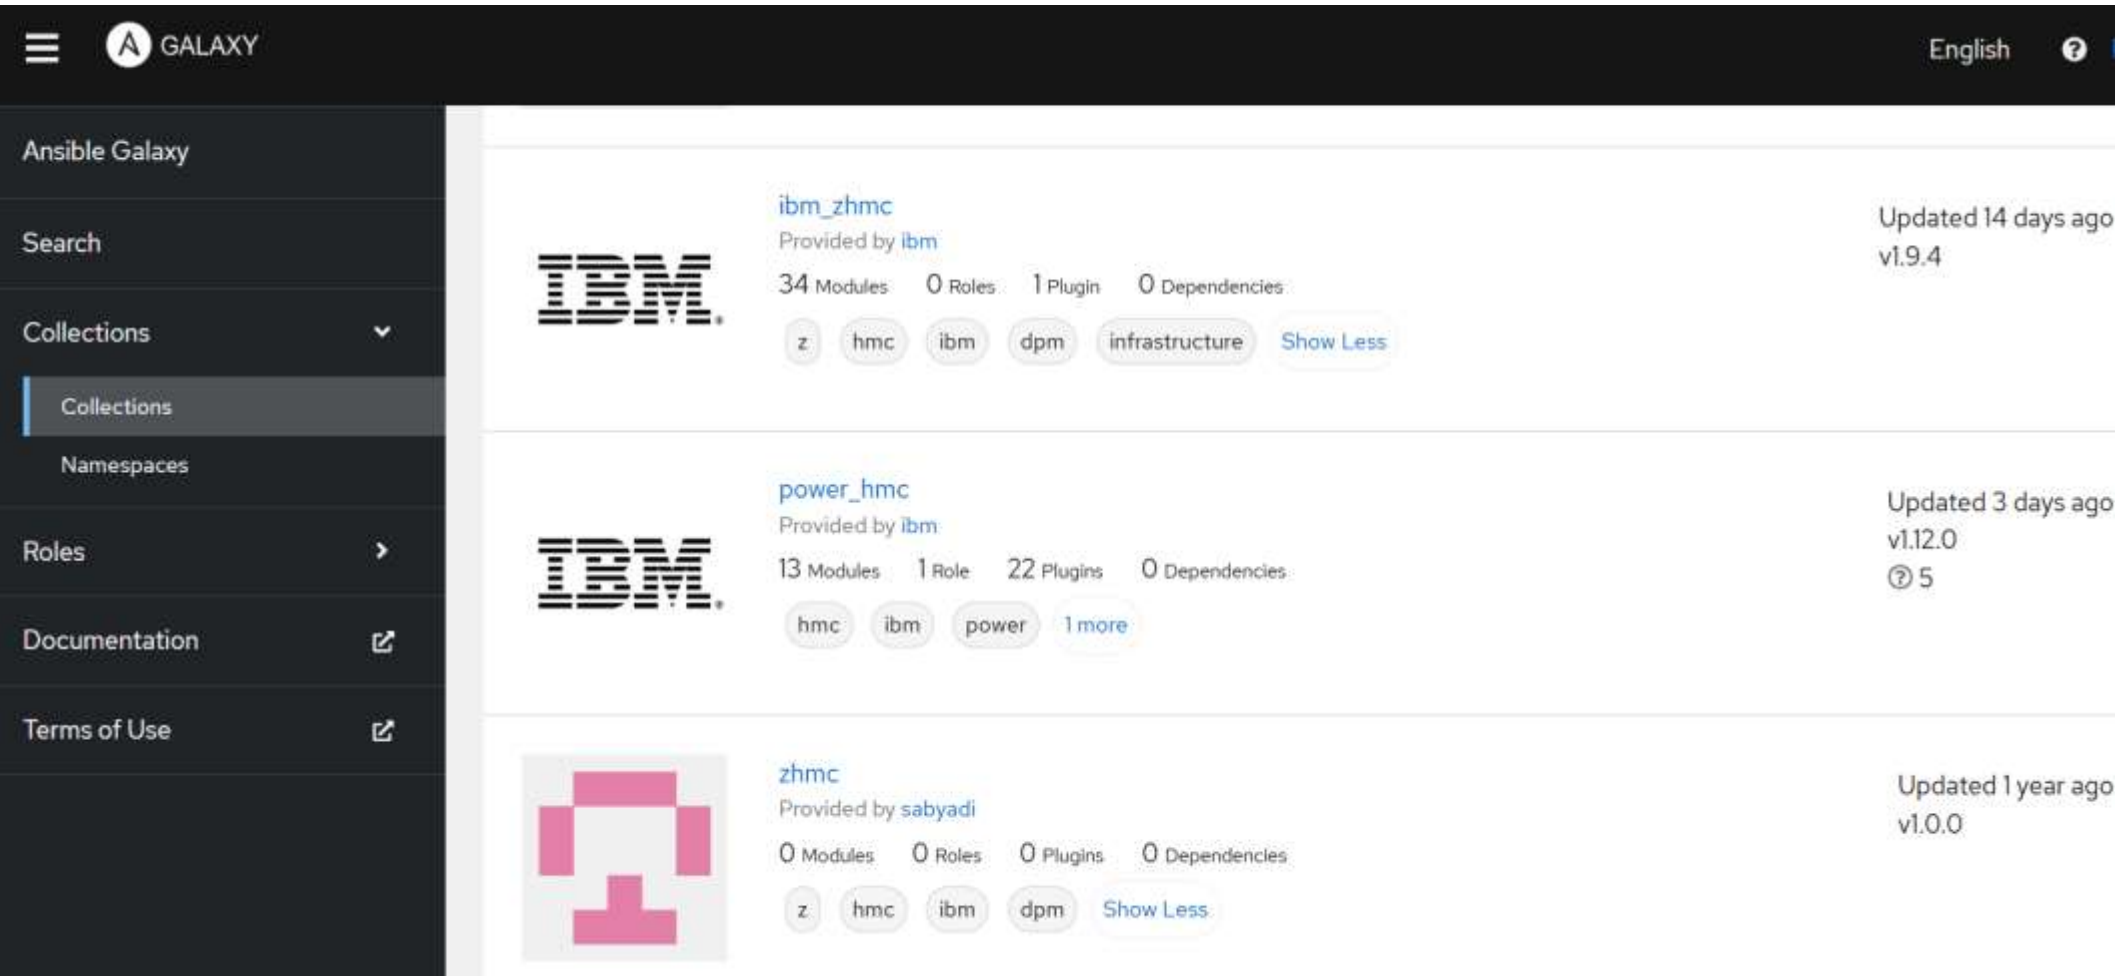

Red Hat Ansible Certified Content for

Hardware Management Console

The available **IBM Z HMC collection** provides Ansible modules and sample playbooks for automating management tasks on the Hardware Management Console (HMC) of IBM Z and IBM LinuxONE machines, such as creating, updating or deleting partitions and other resources.

LPAR automation example – in the RH OpenShift installation using Ansible: https://ibm.github.io/Ansible-OpenShift-Provisioning/set-variables-group-vars/

available at
catalog.redhat.com/software/collection/ibm/ibm_zhmc

Modules (26)

- zhmc_adapter
- zhmc_adapter_list
- zhmc_console
- zhmc_cpc
- zhmc_cpc_list
- zhmc_crypto_attachment
- zhmc_hba
- zhmc_ldap_server_definition
- zhmc_ldap_server_definition_list
- zhmc_lpar
- zhmc_lpar_list
- zhmc_nic
- zhmc_nic_list
- zhmc_partition
- zhmc_partition_list
- zhmc_password_rule
- zhmc_password_rule_list
- zhmc_session
- zhmc_storage_group
- zhmc_storage_group_attachment
- zhmc_storage_volume
- zhmc_user
- zhmc_user_list
- zhmc_user_role
- zhmc_user_role_list
- zhmc_virtual_function

© 2025 IBM Corporation

Red Hat Ansible Certified Content for IBM z/OS environment

IBM CICS® TS Operator collection provides automation for provisioning CICS TS on one or more z/OS endpoints and managing its lifecycle in a hybrid cloud environment.

IBM z/OS IMS collection supports tasks such as generating IMS Database Descriptors (DBD), Program Specification Blocks (PSB), Application Control Blocks (ACB), and running IMS type-1 & type-2 commands.

IBM Operator Collection SDK provides the automation to deploy an operator in your namespace that contains your latest Ansible collection modifications, quickly redeploy your local modifications in seconds, and delete the operator once development is complete.

IBM Z Open Automation Utilities Operator collection provides automation for installing the ZOAU language on one or more z/OS endpoints and managing its lifecycle in a hybrid cloud environment. It uses the z/OS Package Manager to install the software on to z/OS and manage its lifecycle.

IBM z/OS core collection supports automation tasks submitting / querying jobs, creating / fetching / copying data sets, executing operator / TSO commands, ping, querying operator actions, backing up and restoring data sets / volumes, APF authorizing libraries, mounting file systems, running z/OS programs without JCL, initializing volumes, archiving / unarchiving / templating with Jinja, etc.

IBM Z System Automation collection supports operational tasks using the IBM Z System Automation Operations API such as creating and deleting dynamic resources from a template defined in the current active policy of an IBM Z System Automation environment. It interacts with IBM Z System Automation using the SA Operations API provided by the SA Operations REST Server.

IBM z/OS Package Manager collection provides automation for installing z/OS Package Manager and the z/OS products on one or more z/OS endpoints and managing their lifecycle in a hybrid cloud environment. IBM z/OS Package Manager is a utility that can install any z/OS software that is packaged as an OCI artifact on z/OS.

IBM z/OSMF collection supports automation tasks such as operating z/OS workflows, provisioning and managing z/OS middleware / software, via z/OSMF RESTful services.

© 2023 IBM Corporation

Apache CloudStack for Managing Linux on IBM Z and LinuxONE

The leading open-source Infrastructure-as-a-Service (IaaS) platform, available on Linux on IBM Z and LinuxONE

Combines efficiency and scalability of Linux on IBM Z and LinuxONE with flexibility of CloudStack, empowering users to effectively deploy, automate and manage their entire infrastructure for deploying cloud environments on cutting-edge hardware

See here for further information

User Experience:

- Enhanced Efficiency: The energy-efficient s390x architecture ensures reduced operational costs without compromising performance.
- Robust Compatibility: Apache CloudStack's KVM support and advanced networking capabilities integrate seamlessly with Linux on IBM Z and LinuxONE.
- Scalability: Designed to handle complex cloud environments, Apache CloudStack enables efficient resource management and multi-tenant orchestration.
- Add it in! : Enhance existing Apache CloudStack environment with the capabilities of Linux on IBM Z and LinuxONE by adding it as a KVM host.

Streamlined build and deployment of Apache CloudStack on Linux on IBM Z and LinuxONE:

1) Follow our build instructions guide to build Apache CloudStack on Linux on IBM Z and LinuxONE servers.
2) Refer to the quick installation guide to install and set up Apache Cloudstack with the s390x packages you build in step 1.
3) Use the CloudStack management UI to add s390x-based KVM hosts
4) Configure NFS or other secondary storage options.
5) Manage virtual machines on s390x KVM hosts.

IBM Z © 2025 IBM Corporation

IBM Event Automation

A complete and composable set of capabilities for building and managing event-driven architectures

Event Streams (powered by Apache Kafka)

Distribute a continuous flow of events from where they are created to where they are needed, made easy to deploy and manage in an enterprise environment

Event Endpoint Management

Manage events, making them discoverable across the enterprise, with self-service access alongside controls to govern how events are being used

Event Processing (powered by Apache Flink)

Process events in real-time and enrich with AI for deeper insights, using a low code user interface designed to simplify the build and iterative testing of solutions

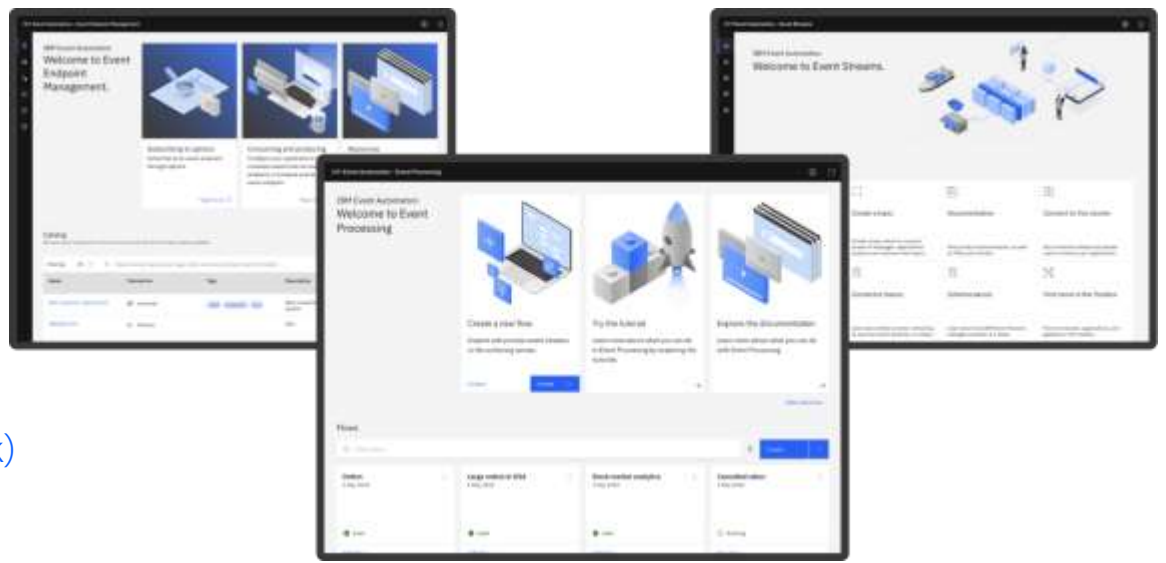

© 2025 IBM Corporation

KVM for
Linux on IBM Z & LinuxONE

IBM Z © 2022 IBM Corporation

KVM Availability

KVM is available and supported in
- SLES12 SP2 and later
- RHEL 8 starting with RHEL 8.4 via Advanced Virtualization repository
- Ubuntu 16.04 and later

Community distributions with KVM support:
- Debian
- Fedora
- OpenSUSE

Documentation: KVM Virtual Server Management available here
Nov 2022 update now also covers
- Persistently configure VFIO mediated devices for both DASD and cryptographic resources
- Share virtual servers' file system with a virtual server hosted on the KVM host file system
- Dump automation improvements on the KVM host

Package versions

- Red Hat

	kernel	QEMU	Libvirt
RHEL 8.10	4.18	6.2	8.0
RHEL 9.6	5.14	9.1	10.10
RHEL 10.0	6.12	9.1	10.10

- SUSE

	kernel	QEMU	Libvirt
SLES12 SP5	4.12	3.1	5.1
SLES15 SP6	6.4	8.2	10.0

- Ubuntu

	kernel	QEMU	Libvirt
16.04 ESM-only	4.4	2.5	1.3.1
18.04 LTS	4.15	2.11	4.0
20.04 LTS	5.4	4.2	6.0
22.04 LTS	5.15	6.2	8.0
24.04 LTS	6.8	8.2	10.0

IBM Z © 2025 IBM Corporation

Recent Improvements

- **Full-featured pass-through of PCI devices to KVM guests**
  Use PCI devices like RoCE Express adapters, or NVMe devices (LinuxONE only) in KVM setups
- **Crypto Passthrough KVM - Performance & Dynamic Management**
  Exploit IBM Z virtualization features to increase performance and allow for dynamic configuration of crypto adapters to KVM guests.
- **Bus id for subchannels**: Allows you to identify passthrough CCW devices by their device bus id in the host without going through hoops.
- **Driverctl now lists persisted overrides**: Makes it easier to identify and manage passthrough devices.
- **Persistent configuration for vfio-ap**: The s390-tools command zdev can now be used to persist Crypto passthrough configurations.
- **Dynamic configuration updates for vfio-ap**: Allows you to hot plug and unplug Crypto domains of a Crypto passthrough configuration for running KVM guests.
- **Support of long kernel command lines of up to 64 KB length**: For example, allows you to specify plenty of I/O devices.

IBM Z © 2025 IBM Corporation

KVM Hardware Support: Selecting the right CPU Model

```
<os>
    <type arch='s390x' machine='s390-ccw-virtio'>hvm</type>
    [...]
</os>
<cpu mode='host-model'/>
```

- E.g. z16 support provided by new model `gen17a`, enabling all z16 features per default
- Choose among the following CPU models:
- z17 & LinuxONE 5 CPU model readily available in RHEL 9.6, RHEL 10, and Ubuntu 25.04, and later.

Mode	Feature Set	Migration Safe	Syntax

Pre-defined	Static	✔ ✔	`<cpu mode='custom'>`
			`  <model fallback='allow'>gen15a</model>`
			`</cpu>`

Host model (***recommended and default on new distros***)	Maximum (based on current host)	✖	`<cpu mode='host-model'/>`
Host passthrough	Maximum		`<cpu mode='host-passthrough'/>`

IBM Z © 2025 IBM Corporation

Co-Location: SMC-Dv2

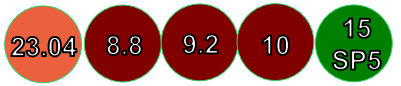

- **What it does**: Provides acceleration for TCP traffic
- **Why you should care**: v2 lifts limitations and greatly simplifies usage
- **Recap**
  - **Shared Memory Communications – Direct** (SMC-Dv1) provides intra-CEC acceleration for TCP traffic using *Internal Shared Memory* (ISM) devices
  - **Superior performance** (low latency, high throughput) at reduced CPU consumption
  - *However*:
    - Peers must be in **same IP subnet**
    - Devices need to be **paired using PNET IDs**
- **SMC-Dv2**
  - Peers can be in ***any* IP subnet**
  - No PNET IDs required
    ⇒ Simplified configuration!
  - Requires z15 or LinuxONE III
  - Full **z/OS compatibility**

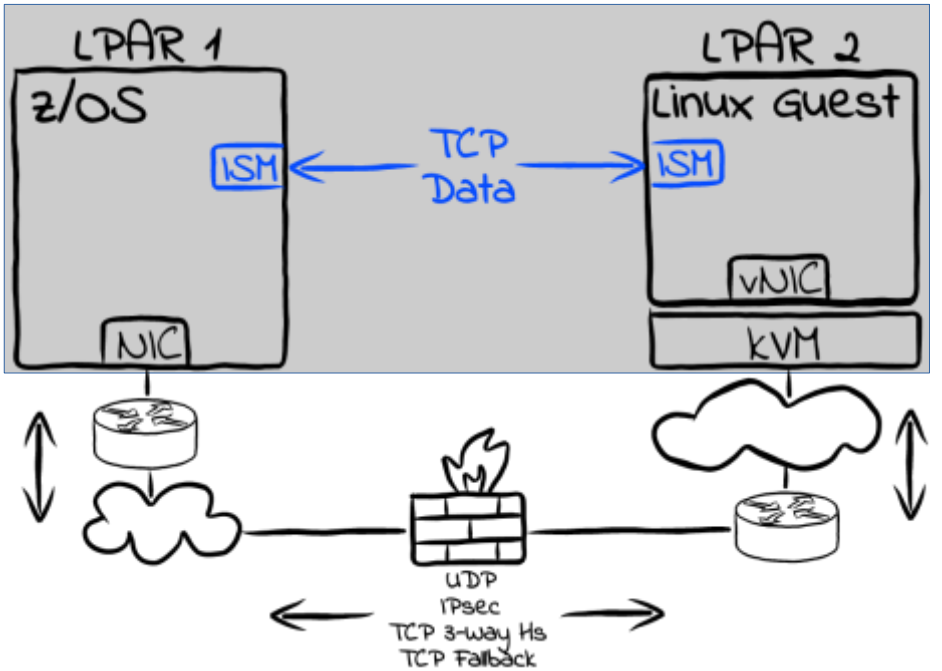

Fig.1: *Traffic flows with SMC-Dv2*

IBM Z © 2025 IBM Corporation

Tools for Basic VM Management

- For the power user: `virsh` + libvirt XML
  - the swiss army knife for KVM
  - Ideal for scripting
  - Requires some knowledge

```
virsh # list --all
 Id    Name                           State
----------------------------------------------------
 -     susanne                        shut off
 -     raul                           shut off
 -     udo                            shut off
```

- The Linux UI: virt-manager
  - Easy to use graphical user interface
  - Can manage guests on multiple KVM hosts

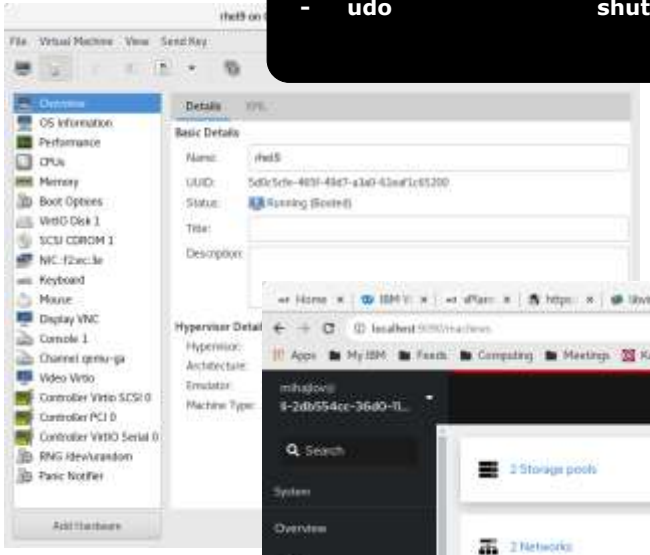

- The web client: cockpit
  - Allows management of VMs using a web browser

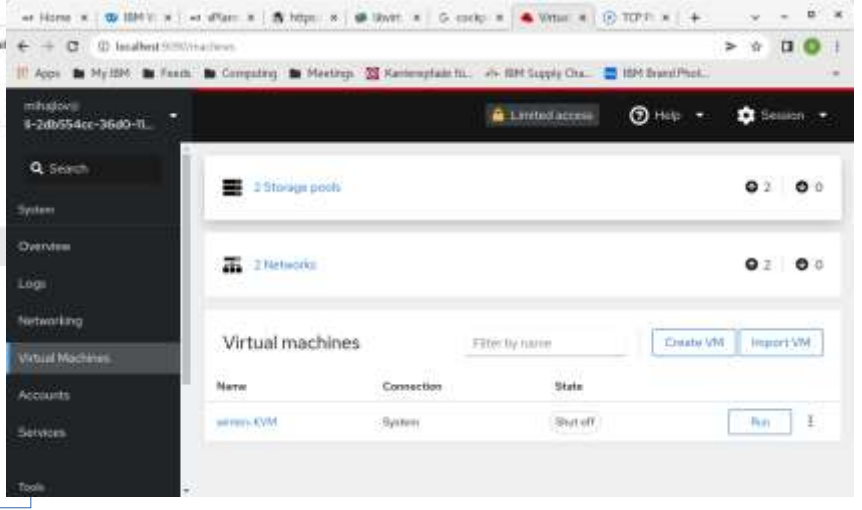

IBM Z © 2025 IBM Corporation

Need something else for Linux & Virtualization?

Linux and KVM

A) Use the Request for Enhancements (RFE) database:
- enter in your IBM ID
- select Brand "*Servers and System Software*"
- select Product "*Linux on System z*" (includes KVM)

B) Reach to us at conferences, e.g.
SHARE GSE Conferences Tech. Univ VM Workshop

How the Linux Distro Partners handle requirements

Red Hat defined RFE process for customers

SUSE requirements can be submitted to their sales reps as well as using the "feedback" button at the bottom of the SUSE Linux Enterprise Server for IBM Z and LinuxONE web site

Canonical is handling requirements for Ubuntu through Launchpad: Open a bug, put requirement in title and tag with s390x

More information about Linux & KVM

- Official web site
  https://www.ibm.com/it-infrastructure/z/os/linux
- Linux & KVM (see Backup)
  Key Documentation Links
- Secure Execution & Compression (see Backup)
  Videos & books
- Enterprise Key Management for Linux (see Backup)
- Videos & books

User forums

- Mailing lists at Maris college
- Linux on s390x forum at Open Mainframe Project

IBM Z © 2025 IBM Corporation

Staying Up-To-Date

Blogs

- Very latest news from the development team
  - KVM on Z: http://kvmonz.blogspot.com/
  - Linux on Z & containers: http://linux-on-z.blogspot.com/
- Focus primarily on upstream submissions, which will end up in Linux distributions later
- Also features in-depth articles on specific topics
- Provided by Linux & KVM on Z development teams

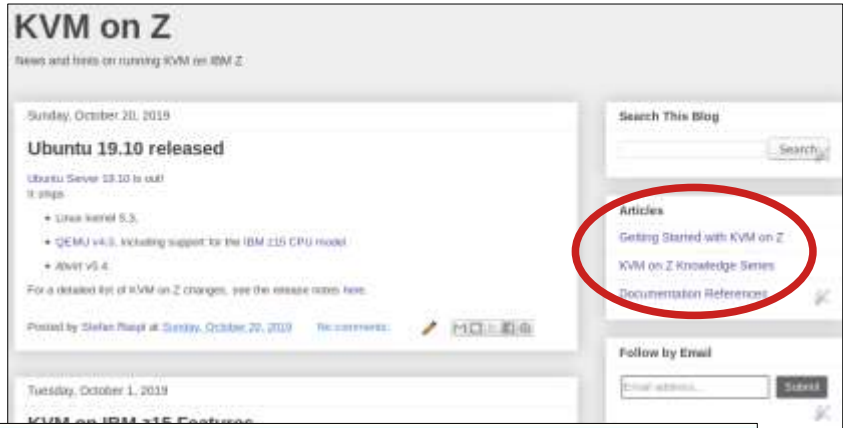

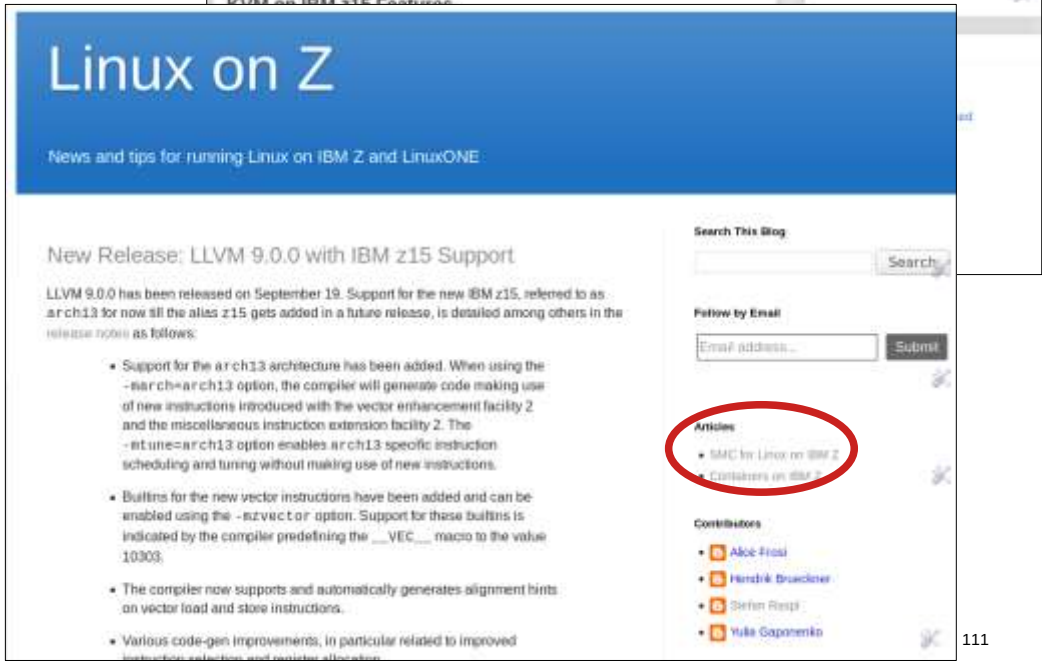

References

Documentation

- Linux on Z and LinuxONE on IBM Documentation
  https://www.ibm.com/docs/en/linux-on-systems?topic=linux-z-linuxone
- Videos explainers
  https://www.ibm.com/docs/en/linux-on-systems?topic=linuxone-video-explainers
- Solution assurance
  https://www.ibm.com/docs/en/linux-on-systems?topic=linuxone-solution-assurance
- z/VM Education Roadmap
  https://www.vm.ibm.com/education/

Webcasts

- In-depth sessions right from the Linux on Z development team
- Recordings available
  https://ibm.biz/Linux-on-IBMzSystems-LinuxONE-Webcasts

Linux on IBM Z: IBM Virtual Development and Test Environment for z/OS

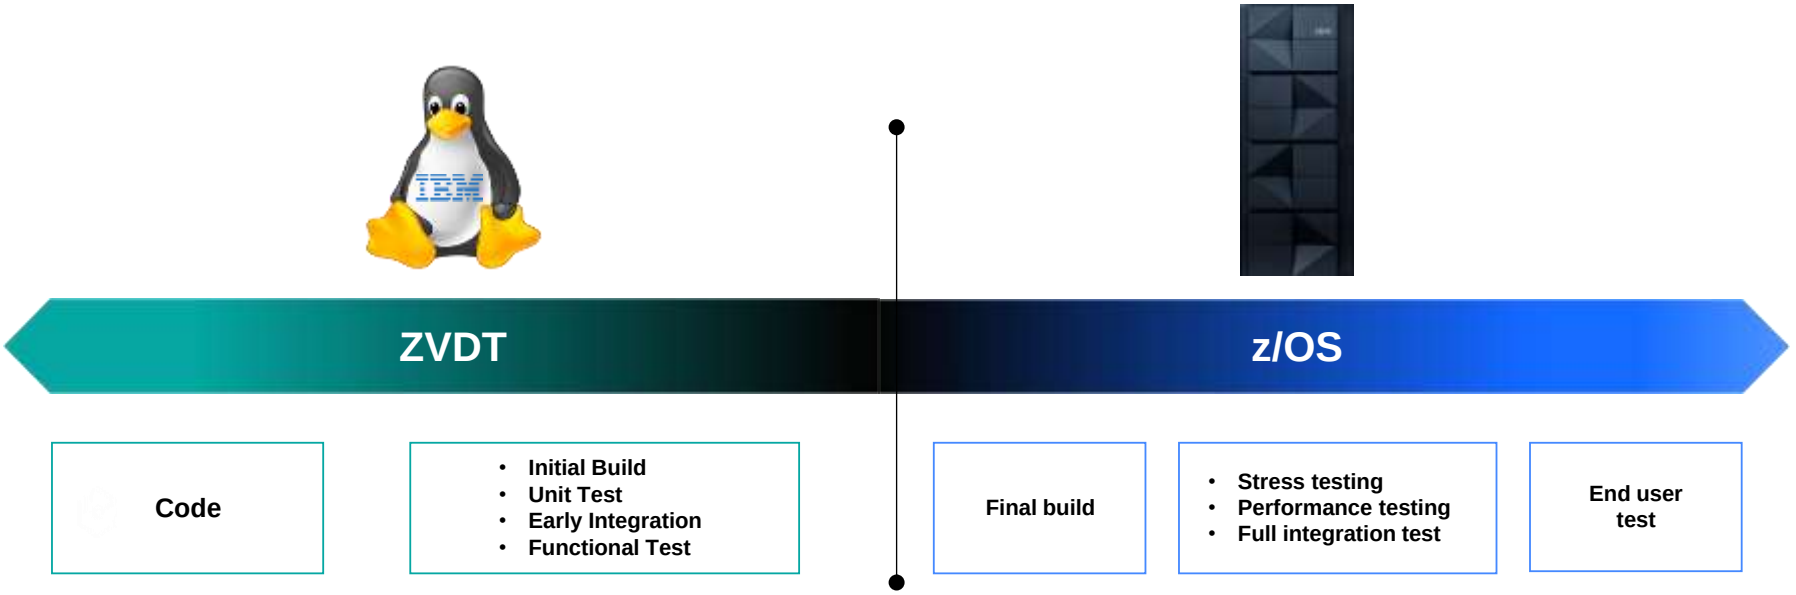

33%

of developers see a lack of skills or resource allocation inhibiting their productivity[1]

36%

of developers view a lack of collaboration between development and IT operations as inhibiting their productivity[3]

3rd

year in a row, a majority of Gitlab survey takers resoundingly pointed to testing as the area most likely to cause delays[2]

IBM Virtual Dev and Test for z/OS (ZVDT) Unique Value

Enterprise requires development and test infrastructure to be on-premises

Need of a higher level of performance compared with ZD&T solution

Capacity-based perpetual and term licensing

Mainframe team prefers to keep direct control and management of the solution infrastructure on Z hardware

Colocation with Dev LPAR simplifies the creation and maintenance of z/OS images

Product Resources

ZVDT product page

Tech Docs

Communities

IBM Z Automated Testing Community

Accelerates Application Development and Test

- Work in parallel vs serially
- Each developers can have their own z/OS Dev/Test environment

Better product quality and quicker time to market

- Shift Left testing: Better test coverage earlier in the cycle
- Enable developers to use testing resources on their own schedules

Adapt to changing business requirements

- z/OS on Linux on Z that can be virtualized and offers to developers a sandbox for prototyping and Version to Version migration
- Internal employee education and training

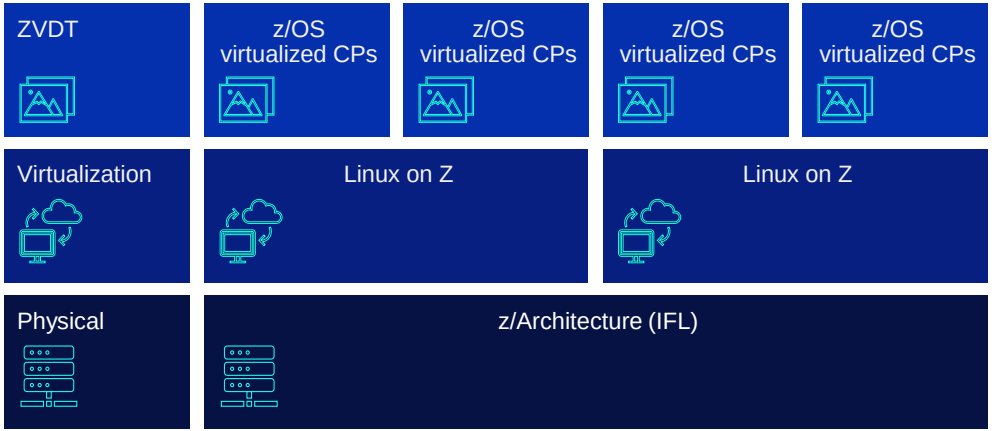

Thank You !

Wilhelm Mild

IBM Executive IT Architect

The Open Group Open Certified Distinguished Architect — IT Architecture Chief/Lead IT Architect

IBM Deutschland Research & Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany

Office: +49 (0)7031-16-3796
wilhelm.mild@de.ibm.com

Additional Redbooks:

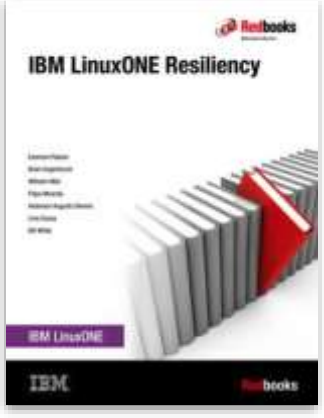

2024 IBM Corporation

Trademarks: See https://www.ibm.com/legal/copytrade for a list of trademarks

Tag Legend

- Supported distributions
  - x.y for SUSE SLES <X> Service Pack <Y>, e.g. 15 SP6 for SLES15 SP6
  - x.y for RHEL <x> Update <y>, e.g. 9.4 for RHEL9.4
  - x.y for Ubuntu x.y, e.g. 24.04 for Ubuntu 16.04 LTS
- Suppored environments
  - LPAR usable for Linux virtual servers running in LPAR
  - z/VM usable for guests running on z/VM
  - KVM usable for guests running on KVM

IBM Z © 2025 IBM Corporation

Disclaimer for IBM Spyre™ Accelerator

The IBM Spyre™ AI Accelerator will not be available with IBM general availability. The IBM Spyre AI Accelerator is currently expected to be available in 4Q of 2025. Any capabilities discussed in this presentation with respect to the IBM Spyre AI Accelerator will not be enabled by until these accelerator cards are installed in the system.

© 2025 IBM Corporation

Citations/Claims/Disclaimers

1.DISCLAIMER: IBM internal data based on measurements and projections was used in calculating the expected value. Necessary components include IBM LinuxONE Emperor 5; IBM z/VM V7.3 systems or above collected in a Single System Image, each running RHOCP 4.14 or above; IBM Operations Manager; GDPS 4.6 or above for management of data recovery and virtual machine recovery across metro distance systems and storage, including Metro Multi-site workload and GDPS Global; and IBM DS8000 series storage with IBM HyperSwap. A MongoDB v4.4 workload was used. Necessary resiliency technology must be enabled, including z/VM Single System Image clustering, GDPS xDR Proxy for z/VM, and Red Hat OpenShift Data Foundation (ODF) 4.14 or above for management of local storage devices. Application-induced outages are not included in the above measurements. Other configurations (hardware or software) may provide different availability characteristics.

2. The Cost of a Data Breach Report 2024, by Ponemon Institute and analyzed by IBM https://www.ibm.com/downloads/documents/us-en/107a02e94948f4ec

3. X-Force Threat Intelligence Index 2024, https://www.ibm.com/downloads/documents/us-en/10c31775c0d40a37

4. Dimple Ahluwalia, Gerry Parham, Kevin Skapinetz, (August 2024), Securing generative AI: What matters now. IBM Institute for Business Value: Full Data and insights deck. https://w3.ibm.com/services/lighthouse/documents/218436?ref=bmwiz

5. DISCLAIMER: Performance result is extrapolated from IBM® internal tests running on IBM Systems Hardware of machine type 9175. The Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) was deployed on Red Hat® OpenShift® Container Platform (RHOCP) 4.17. The 3 RHOCP Compute nodes ran 3 Acme Air instances in parallel, each driven remotely from Apache JMeter™ 5.2.1 with 128 parallel users. IBM Systems Hardware configuration: 8 LPARs in total with 21 dedicated and 4 shared cores (SMT). 3 LPARs running RHOCP Compute nodes each with 7 dedicated cores (SMT), 64 GB memory and DASD storage. 5 LPARs each with 4 shared cores (SMT), 16 GB memory and DASD storage, providing 3 RHOCP Management nodes and 2 RHOCP Infrastructure nodes. The Network adapters were dedicated for NETH on Linux. Results may vary.

6. DISCLAIMER: Performance result is extrapolated based on IBM® internal tests running on IBM Systems Hardware of machine type 9175. The flexible I/O tester benchmark (fio 3.35) ran with 128 parallel threads using read-only and read-write (70:30) operations on 8x 40 GB files each on a separate FCP attached LUN. IBM Systems Hardware configuration: 1 LPAR running Red Hat® Enterprise Linux® 9.4 (upstream Kernel level 6.13 with aes-256-xts exploitation patch - Commit ID 80625b670312e74512d65b19e9470184386ab265) with 13 dedicated cores (SMT), 64 GB memory, 1 FCP Express32G adapter with 4 ports connected. 1 IBM FlashSystem® 9500 (FS9500) with 8 ports connected, providing 8 LUNs of 320 GB total size, equally distributed across the two node canisters. Each luks2 encrypted LUN is connected through 8 paths and formatted with an ext4 file system. Cores utilization targeted below 50%. Results may vary.

7. Michael Osborne, Katia Moskvitch, Jennifer Janechek (August 2024) , NIST’s post-quantum cryptography standards are here. https://research.ibm.com/blog/nist-pqc-standards

© 2025 IBM Corporation

© 2025 International Business Machines Corporation

IBM, the IBM logo, and IBM Spyre are trademarks of IBM Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

THIS DOCUMENT IS DISTRIBUTED “AS IS” WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT, SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

Client examples are presented as illustrations of how those clients have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Not all offerings are available in every country in which IBM operates.

Any statements regarding IBM’s future

11.CLAIM: IBM internal tests simulating a complete IT solution running containerized WebSphere Liberty and EDB Postgres workloads, show that a single IBM LinuxONE Emperor 5 Max 136 can do the work of up to 2,944 cores of the compared x86 solution.

DISCLAIMER: IBM® internal performance tests for the core consolidation study targeted a comparison of the following servers. IBM Machine Type 9175 MAX 136 system consisting of three CPC drawers containing 136 configurable processor units and six I/O drawers to support both network and external storage. The x86 solution used a commercially available enterprise server with two 5th generation Intel® Xeon® Platinum 8592+ processors, 64 cores per CPU. Both solutions had access to the same storage. The workloads consisted of a containerized online transaction processing (OLTP) WebSphere Liberty v25 application running on Red Hat OpenShift Container Platform (OCP) v4.17, and an EDB Postgres for Kubernetes v1.25 on the same OCP cluster simulating core online banking functions. Both solutions used Red Hat Enterprise Linux v9.5 and KVM. Results may vary.

The test results were extrapolated to a typical, complete customer IT solution that includes isolated from each other production and non-production IT environments. TCO included software, hardware, energy, network, data center space, and labor costs. On the IBM z17 side the complete solution requires one IBM z17 Type 9175 MAX 136, and on x86 side, the complete IT solution requires 23 compared servers.

12. DISCLAIMER: IBM internal performance tests for the core consolidation study compared an IBM Machine Type 9175 Max136 with 136 configurable processor units with an x86 solution that used a commercially available enterprise server with two 5th gen Intel Xeon Platinum 8592+ processors and 64 cores per CPU. Workloads consisted of a containerized OLTP WebSphere Liberty v25 application running on Red Hat OCP v4.17 and an EDB Postgres for Kubernetes v1.25 on the same OCP cluster. Both solutions used Red Hat Enterprise Linux v9.5 and KVM. Test results were extrapolated to a typical, complete customer IT solution that included production and non-production IT environments isolated from each other. The IBM Machine Type 9175 solution required one Max136 and the x86 solution required 23 compared servers. Results may vary.

13.DISCLAIMER: IBM® internal performance tests for the core consolidation study targeted a comparison of the following servers. IBM Machine Type 9175 MAX 136 system consisting of three CPC drawers containing 136 configurable processor units and six I/O drawers to support both network and external storage. The x86 solution used a commercially available enterprise server with two 5th generation Intel® Xeon® Platinum 8592+ processors, 64 cores per CPU. Both solutions had access to the same storage. The workloads consisted of a containerized online transaction processing (OLTP) WebSphere Liberty v25 application running on Red Hat OpenShift Container Platform (OCP) v4.17, and an EDB Postgres for Kubernetes v1.25 on the same OCP cluster simulating core online banking functions. Both solutions used Red Hat Enterprise Linux v9.5 and KVM. Results may vary.

The test results were extrapolated to a typical, complete customer IT solution that includes isolated from each other production and non-production IT environments. TCO included software, hardware, energy, network, data center space, and labor costs. On the IBM z17 side the complete solution requires one IBM z17 Type 9175 MAX 136, and on x86 side, the complete IT solution requires 23 compared servers.

© 2025 IBM Corporation

14. Yunke Wang, Yanxi Li, Chang Xu AI Scaling: From Up to Down and Out, Feb 2025; https://arxiv.org/html/2502.01677v1

© 2025 International Business Machines Corporation

IBM, the IBM logo, and IBM Spyre are trademarks of IBM Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

THIS DOCUMENT IS DISTRIBUTED “AS IS” WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT, SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

Client examples are presented as illustrations of how those clients have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Not all offerings are available in every country in which IBM operates.

Any statements regarding IBM’s future

17.CLAIM: The IBM Integrated Accelerator for AI on IBM LinuxONE Emperor 5 at full utilization is designed to process up to 24 trillion operations per second (TOPS) shared across all cores on the chip.

DISCLAIMER: Result is the maximum theoretical number of trillion operations per second (TOPS) in 8bit precision that can be executed by a single IBM Integrated Accelerator for AI. Cores are running at 5.5GHz and have one IBM Integrated Accelerator for AI per chip. The IBM Integrated Accelerator for AI consists of 2 corelets, each with an array of 64 tensor cores capable of executing 4 integer-multiply-add operations (IMA) 8-way SIMD with no sparsity.

© 2025 IBM Corporation

© 2025 International Business Machines Corporation

IBM, the IBM logo, and IBM Spyre are trademarks of IBM Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

THIS DOCUMENT IS DISTRIBUTED “AS IS” WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT, SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

Client examples are presented as illustrations of how those clients have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Not all offerings are available in every country in which IBM operates.

Any statements regarding IBM’s future