


# Linux on IBM Z News

—

Stefan Raspl

Linux on IBM Z Development



# Trademarks

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

AIX*	DB2*	HiperSockets*	MQSeries*	PowerHA*	RMF	System z*	zEnterprise*	z/VM*
BladeCenter*	DFSMS	HyperSwap	NetView*	PR/SM	Smarter Planet*	System z10*	z10	z/VSE*
CICS*	EASY Tier	IMS	OMEGAMON*	PureSystems	Storwize*	Tivoli*	z10 EC	
Cognos*	FICON*	InfiniBand*	Parallel Sysplex*	Rational*	System Storage*	WebSphere*	z/OS*	
DataPower*	GDPS*	Lotus*	POWER7*	RACF*	System x*	XIV*		

\* 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](http://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.

# Agenda

- Linux on IBM Z Distributions
- GA2 News for IBM z14 and LinuxONE
- SMC for Linux on Z
- Pervasive Encryption
- Kernel News
















































# Supported Linux Distributions

 **Certified by Linux partner**

*Overview shows Linux distributions in service.*

*Extended support is available for Linux distributions that are out of service.*

	Emperor II	Rockhopper II	Emperor	Rockhopper	zEnterprise – zEC12, zBC12	zEnterprise – z196, z114	System z10, System z9
	z14 M0x	z14 ZR1	z13	z13s			
RHEL 7							
RHEL 6							
RHEL 5							
SLES 15							
SLES 12							
SLES 11							
Ubuntu 18.04							
Ubuntu 16.04							

Last update 04/12/2019

See [www.ibm.com/systems/z/os/linux/resources/testedplatforms.html](http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html) for latest updates and details.

# z14 Hardware Certification

z14 M0x, Emperor II

**z14 ZR1, Rockhopper II**

z13, Emperor

z13s, Rockhopper

zEnterprise – zEC12, zBC12

zEnterprise – z196, z144

System z10 and System z9

## z14 ZR1, Rockhopper II

### Linux partner certifications for this hardware:

Canonical hardware certifications:

- Ubuntu 18.04 LTS: [z14 ZR1](#) and [Rockhopper II](#)
- Ubuntu 16.04 LTS: [z14 ZR1](#) and [Rockhopper II](#)

Red Hat hardware certifications for z14 ZR1:

- [RHEL 7.3](#) and [RHEL 6.9](#)

SUSE hardware certifications:

- SLES 12 SP2: [z14 ZR1](#) and [Rockhopper II](#)
- SLES 11 SP4: [z14 ZR1](#) and [Rockhopper II](#)

### IBM tested distributions on this hardware:

Canonical distribution with the kernel-level and cryptography libraries:

- *Ubuntu 16.04.3 LTS*: LTS kernel 4.4.0-116; *Crypto libs*: TKE9.0, csulcca-5.2.23-12, libep11\_1.3.2-1, libep11-dev\_1.3.2-1

Red Hat distributions with the kernel-level and cryptography libraries:

- *RHEL 7.4*: 3.10.0-693.21.1.el7; *Crypto libs*: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.2-0, ep11-host-devel-1.3.2-0
- *RHEL 2.6.32-696.23.1.el6*; *Crypto libs*: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.2-0, ep11-host-devel-1.3.2-0

SUSE distributions with the kernel-level and cryptography libraries:

- *SLES 12 SP3*: kernel 4.4.120-94.17; *Crypto libs*: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.2-0, ep11-host-devel-1.3.2-0
- *SLES 11 SP4*: kernel 3.0.101-108.35; *Crypto libs*: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.2-0, ep11-host-devel-1.3.2-0; RoCE Express2 is not supported

# Linux on IBM Z Distributions: SUSE

- **SUSE Linux Enterprise Server 11**
  - 03/2009 SLES11 GA: Kernel 2.6.27, GCC 4.3.3
  - 07/2015 SLES11 SP4: Kernel 3.0, GCC 4.3.4
    - **EOS 31 Mar. 2019**; LTSS: 31 Mar. 2022
- **SUSE Linux Enterprise Server 12**
  - 10/2014 SLES12 GA: Kernel 3.12, GCC 4.8
  - 12/2018 SLES12 SP4: **Kernel 4.12**, GCC 4.8
    - EOS 31 Oct. 2024; LTSS: 31 Oct. 2027
- **SUSE Linux Enterprise Server 15**
  - 07/2018 SLES 15 GA: Kernel 4.12, GCC 6.4 / 7.1
  - EOS 31 July 2028; LTSS: 31 July 2031

# Linux on IBM Z Distributions: Red Hat

- **Red Hat Enterprise Linux 5**

- 03/2007 RHEL 5 GA: Kernel 2.6.18, GCC 4.1.0
- 09/2014 RHEL 5 Update 11
- EOS 31 Mar. 2017; ELS: 30 Nov. 2020

- **Red Hat Enterprise Linux 6**

- 11/2010 RHEL 6 GA: Kernel 2.6.32, GCC 4.4.0
- 06/2018 RHEL 6 Update 10
- EOS 30 Nov. 2020; ELS: 30 June 2024

- **Red Hat Enterprise Linux 7**

- 06/2014 RHEL 7 GA: Kernel 3.10, GCC 4.8
- 10/2018 RHEL 7 Update 6
- EOS 30 Jun. 2024; ELS: tbd

- **Red Hat Enterprise Linux 8**

- 05/2019 RHEL 8 GA: Kernel 4.18, GCC 8.2.1
- EOS: Expected for 2029; ELS: tbd


# Linux on IBM Z Distributions: Canonical

- **Ubuntu 16.04 (Xenial Xerus)**
  - 04/2016 Ubuntu 16.04 GA: Kernel 4.4, GCC 5.3.0+ LTS-Release
  - 08/2017 Ubuntu 16.04.03 LTS
- **Ubuntu 18.04 (Bionic Beaver)**
  - 04/2018 Ubuntu 18.04 GA: Kernel 4.15, GCC 7.2.0 LTS-Release
- Lifecycle:
  - Regular releases every 6 months and supported for 9 months
  - LTS releases every 2 years and supported for 5 years
  - LTS enablement stack will provide newer kernels within LTS releases
  - <http://www.ubuntu.com/info/release-end-of-life>



# **GA2 for z14 & LinuxONE**

# IBM Adapter for NVMe



- Two cards have been tested
  - Intel PN SSDPE2KX010T801 (1TB)
  - Intel PN SSDPE2KX040T801 (4TB)
- Customer is free to use any NVMe card
- Uses the standard NVMe Linux kernel driver (`nvme.ko`)
- Support available for
  - SLES 12 SP4
  - SLES 15
  - RHEL 7.6 (\*)
  - RHEL 8
  - Ubuntu 18.04 (\*)

\* please update the kernel to the latest level  
For a whitepaper on NVMe performance see


<https://www.ibm.com/us-en/marketplace/linuxone-emperor-ii/resources>

# IBM FCP Express 32S



- High-speed network technology for SANs
- 32 Gigabits per second fibre-channel
- The IBM Z specific zfcpl Linux driver is used to drive traffic over the FCP Express 32S cards
  - The existing driver works without changes
  - To report the new link speed of 32 Gbps one small kernel patch has been added in 2017

# IBM OSA-Express7S



- 25 GbE networking adapter
- Strictly requires 25GbE optics and switches - no dynamic switching to 10GbE!
- Availability: April 9, 2019
- The IBM Z specific `qeth` Linux driver drives traffic
  - Existing driver works without changes
  - Small patch required to report correct link speed
  - Likewise `qethqoat` (part of `s390-tools`) to report correct model generation

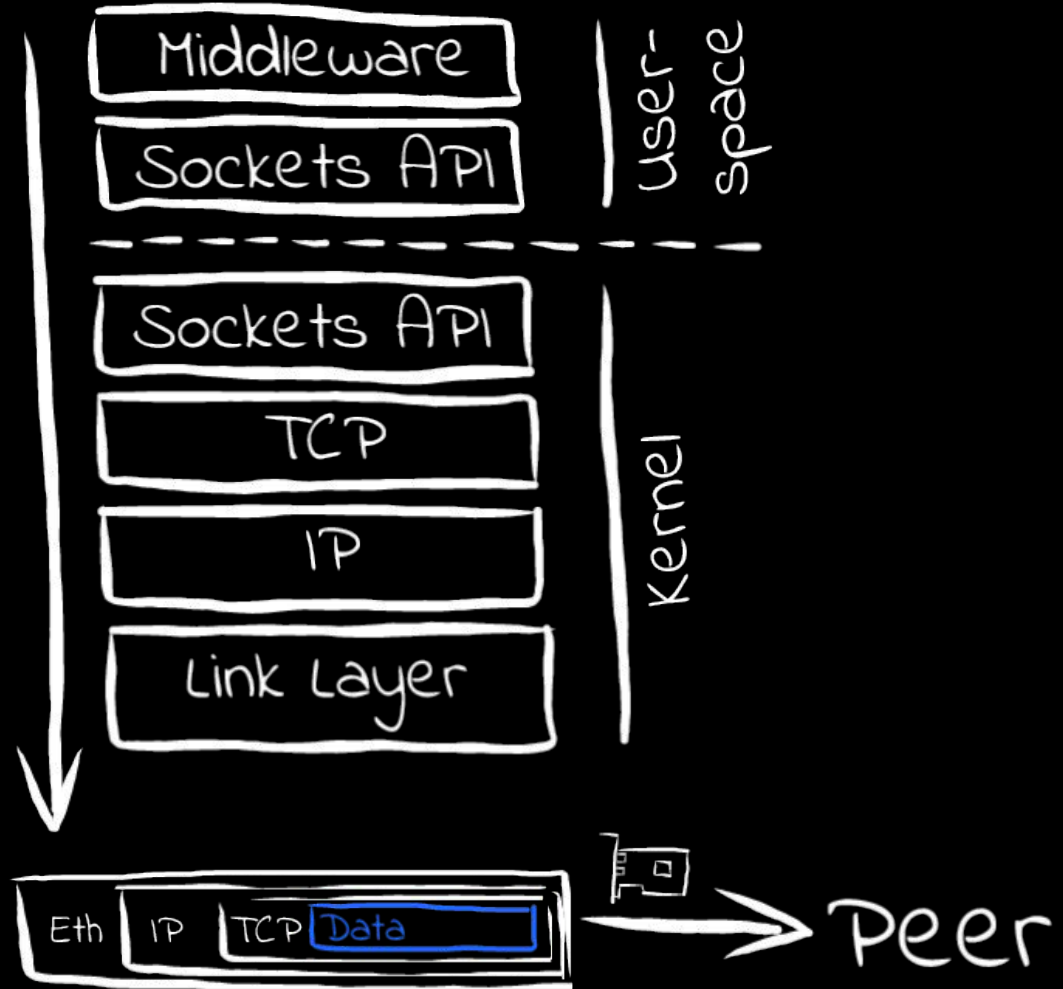
# IBM RoCE Express2




- 25 GbE networking adapter
- Strictly requires 25GbE optics and switches - no dynamic switching to 10GbE!
- Vendor-provided mlx5 Linux driver required

# SMC

# What sending data through BSD sockets looks like



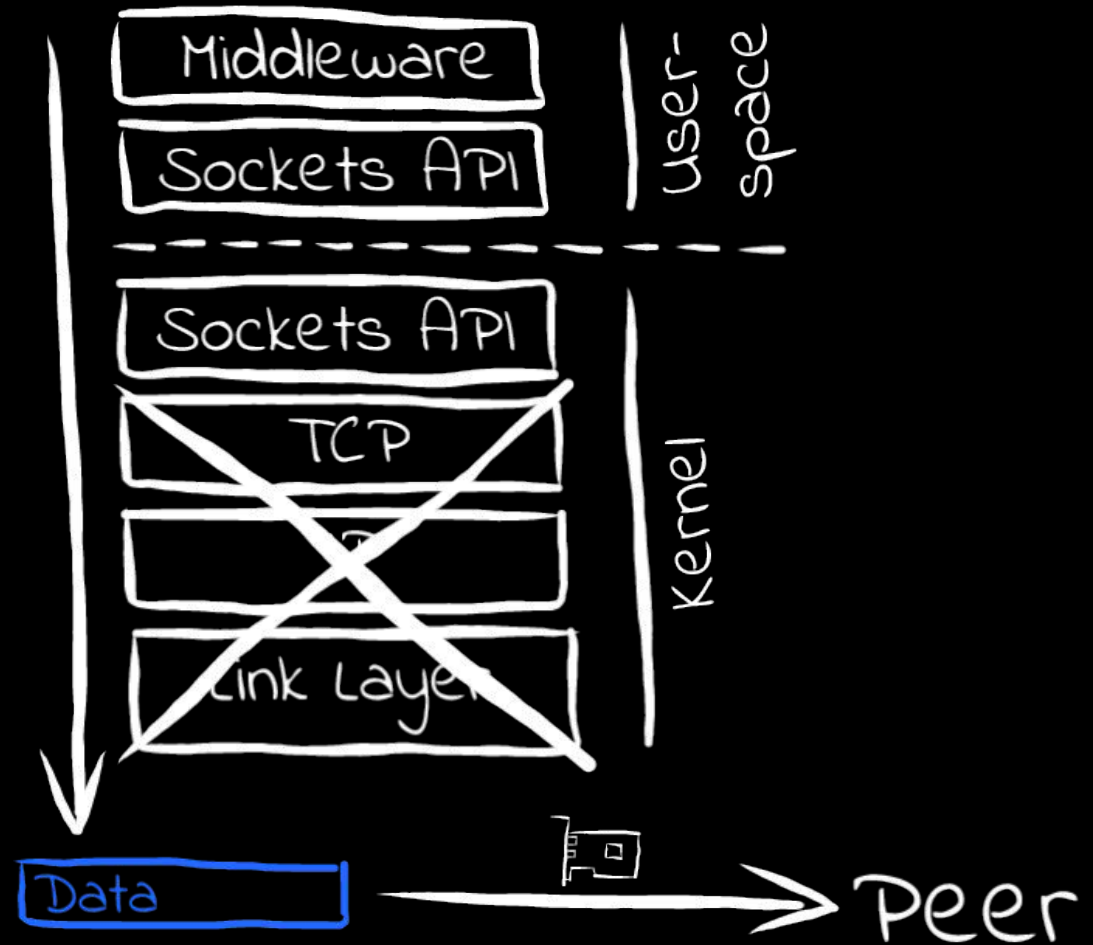
**What if we had a simple buffer to write data to and let hardware do the rest...?**







**What if we had a simple buffer to write data to and let hardware do the rest...?**

- Less latency
- Lower CPU usage




# SMC-R Overview

- Cross-CEC connectivity using **RoCE Express** cards
- IBM Z hardware requirements
  - IBM zEC12 and zBC12, or later
  - LinuxONE Emperor and Rockhopper, or later
  - Classic and DPM mode supported
- RoCE Express cards
  - RoCE Express & RoCE Express2 cards supported
  - Switches need to support and enable *Global Pause* (standard Ethernet switch flow control feature as described in IEEE 802.3x)
- PNET ID configuration
  - IOCDs or smc\_pnet (→see *smc-tools* package)
  - Use OSA or RoCE card for regular connectivity
- **Notes:**
  - Linux on Z can use a single RoCE card for regular and RDMA traffic!
  - No link failover!



# SMC-D Overview

- Intra-CEC connectivity using **Internal Shared Memory (ISM)** devices
- IBM Z hardware requirements
  - IBM z13 (requires driver level 27 (GA2)) and z13s, or later
  - LinuxONE Emperor and LinuxONE Rockhopper, or later
  - Classic mode only (i.e. DPM not supported)
- ISM devices
  - Virtual PCI network adapter of new VCHID type ISM
    - No PCI bus usage
    - No extra hardware required
  - Provides access to memory shared between LPARs
  - 32 ISM VCHIDs per CPC, 255 VFs per VCHID (8K VFs per CPC total)  
I.e. the maximum no. of virtual servers that can communicate over the same ISM VCHID is 255
  - Each ISM VCHID represents a unique (isolated) internal network, each having a unique Physical Network ID
- PNET ID configuration
  - IOCDs only
  - Use HiperSockets devices, OSA, or RoCE cards for regular connectivity



# SMC-R vs. SMC-D Comparison

Feature	SMC-R	SMC-D
Intra-CEC	yes	yes
Cross-CEC	yes	no
RDMA Device	RoCE	ISM
Interface Type	PCI	PCI
Bus used	PCI	-
PNET ID Definition	IOCDS, or smc_pnet	IOCDS
Failover	tbd	n/a
Upstream Status	Initial code upstream in Linux kernel 4.18	Initial code upstream in Linux kernel 4.19
Linux Distributions	SLES 12 SP4 / Ubuntu 18.10 / RHEL 8	

# Pervasive Encryption

# Technical Foundation

- IBM z14 and LinuxONE – Designed for Pervasive Encryption
  - CPACF – Dramatic advance in bulk *symmetric* encryption performance
  - Crypto Express6S – Doubling of *asymmetric* encryption performance for TLS handshakes
- Linux on IBM Z – Ecosystem combined with IBM z14 and LinuxONE capabilities
  - Network security – Enterprise scale encryption and handshakes using z14 CPACF and SIMD
  - `dm-crypt` – Transparent volume encryption using *industry unique CPACF protected-keys*
  - Secure Service Container – Automatic protection of data and code for virtual appliances
- z/VM – New: Encrypted paging support
  - z/VM 6.4 APAR VM65993

# CPACF Performance

- CPACF performance depends on the size of buffers being en-/decrypted
- The larger the buffer the better the performance
  - Best performance with  $\geq 4\text{kB}$  buffers
- `dm-crypt` always used 512 byte buffers
- Starting with Linux kernel 4.12 and `cryptsetup 2.0`:
  - `dm-crypt` can use 4kB buffers

Relative AES GCM/XTS  
in memory performance

	512B	4kB
z13	1	1.5
z14	4	13

# Data in Flight

- openSSL and libcrypto
  - De-facto standard TLS and crypto libraries used by many projects, no IBM Z specific configuration required
  - Exploitation of IBM Z CPACF and SIMD code by libcrypto (w/o ibmca engine)
  - Focus on TLS 1.2 and 1.3 ciphers
  - Support for z14 AES-GCM accepted for openSSL version 1.1.1-alpha2
- IPsec
  - De-facto standard TLS and crypto libraries used by many projects
  - Transparently uses CPACF through the in-kernel crypto API
  - Kernel 4.15 and later use new CPACF instruction for AES-GCM
- GSKit
  - IBM C library for TLS and crypto, e.g., used by IBM HTTP Server (IHS) and many IBM Software products
  - Uses IBM Z CPACF
  - Release 8.0.50.82 and later use the new z14 CPACF instructions
- Java 8 / JCE
  - De-facto standard TLS and crypto libraries used by many projects
  - Java 8 service refresh 5 and later use z14 CPACF instructions



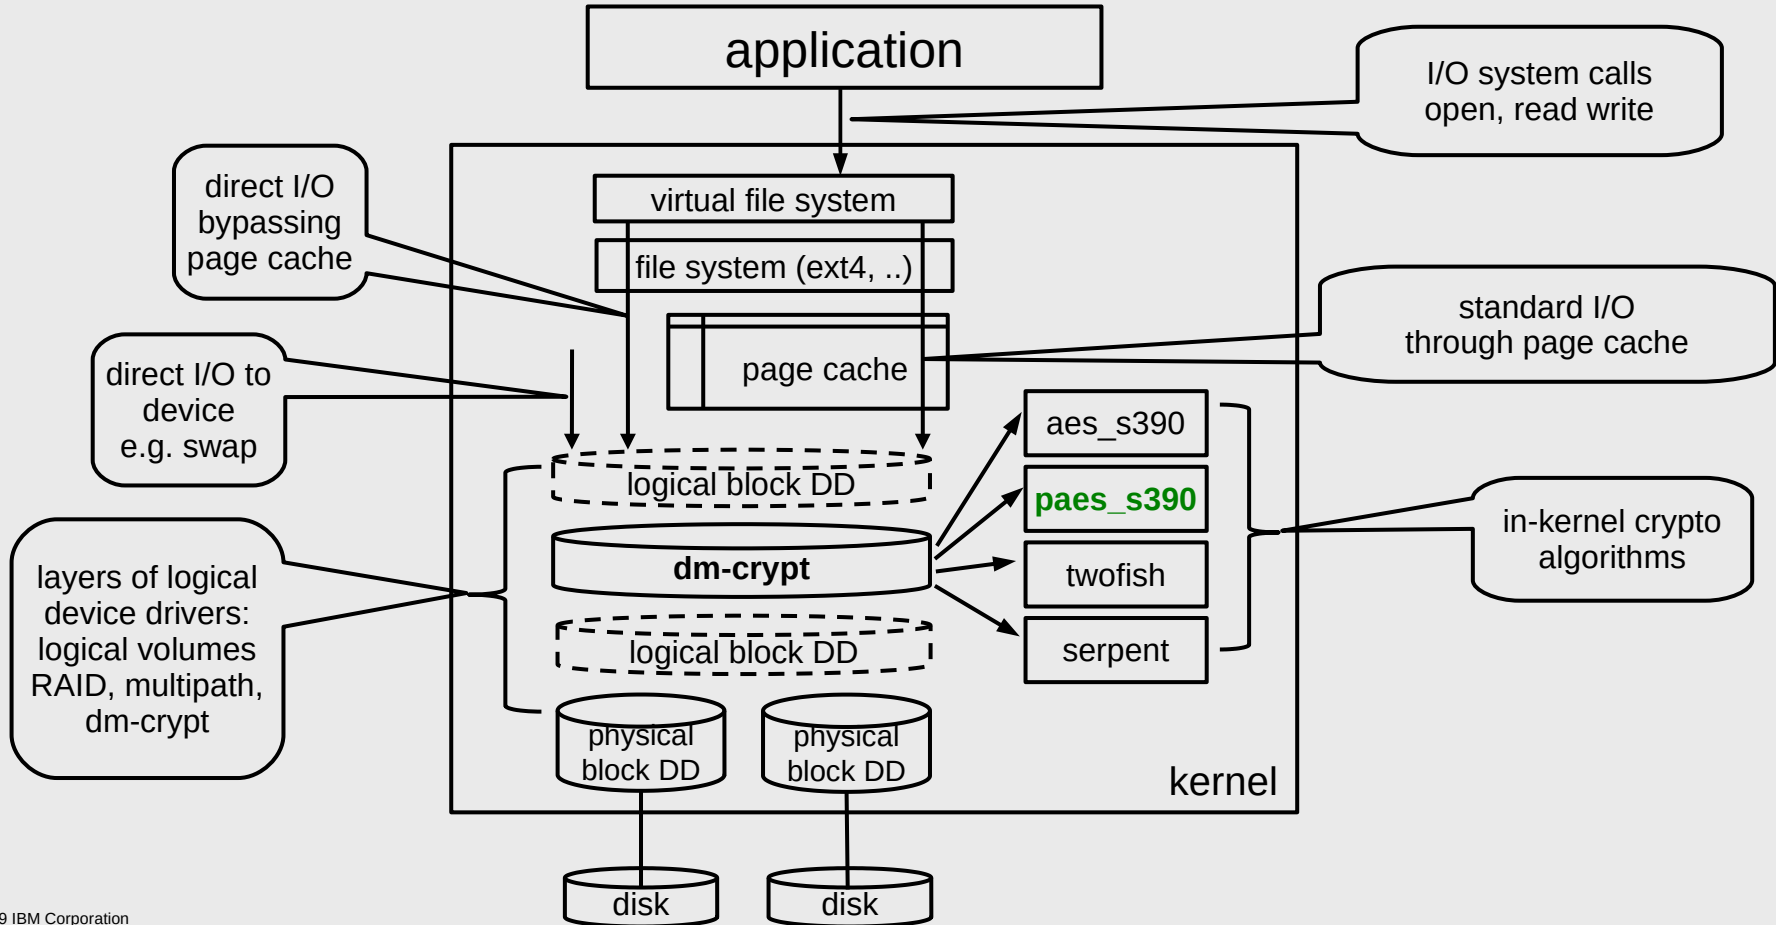
# Data at Rest

- dm-crypt: block device / full volume encryption
  - Uses in-kernel crypto API
  - Granularity: disk partition / logical volume
  - [New protected key option](#)
- ext4fs with encryption option: file system encryption
  - Uses in-kernel crypto API
  - Granularity: file, directory, symbolic link
- NFS v4 with encryption option: encryption of file transport
  - Uses in-kernel crypto API
- SMB v3.1: encryption of file transport
  - Uses in-kernel crypto API
- Spectrum Scale (GPFS) with encryption option: file encryption
  - Uses GSKIT or CLiC in user space
  - Granularity: file
- DB2 native encryption: data base encryption
  - Uses GSKIT in user space

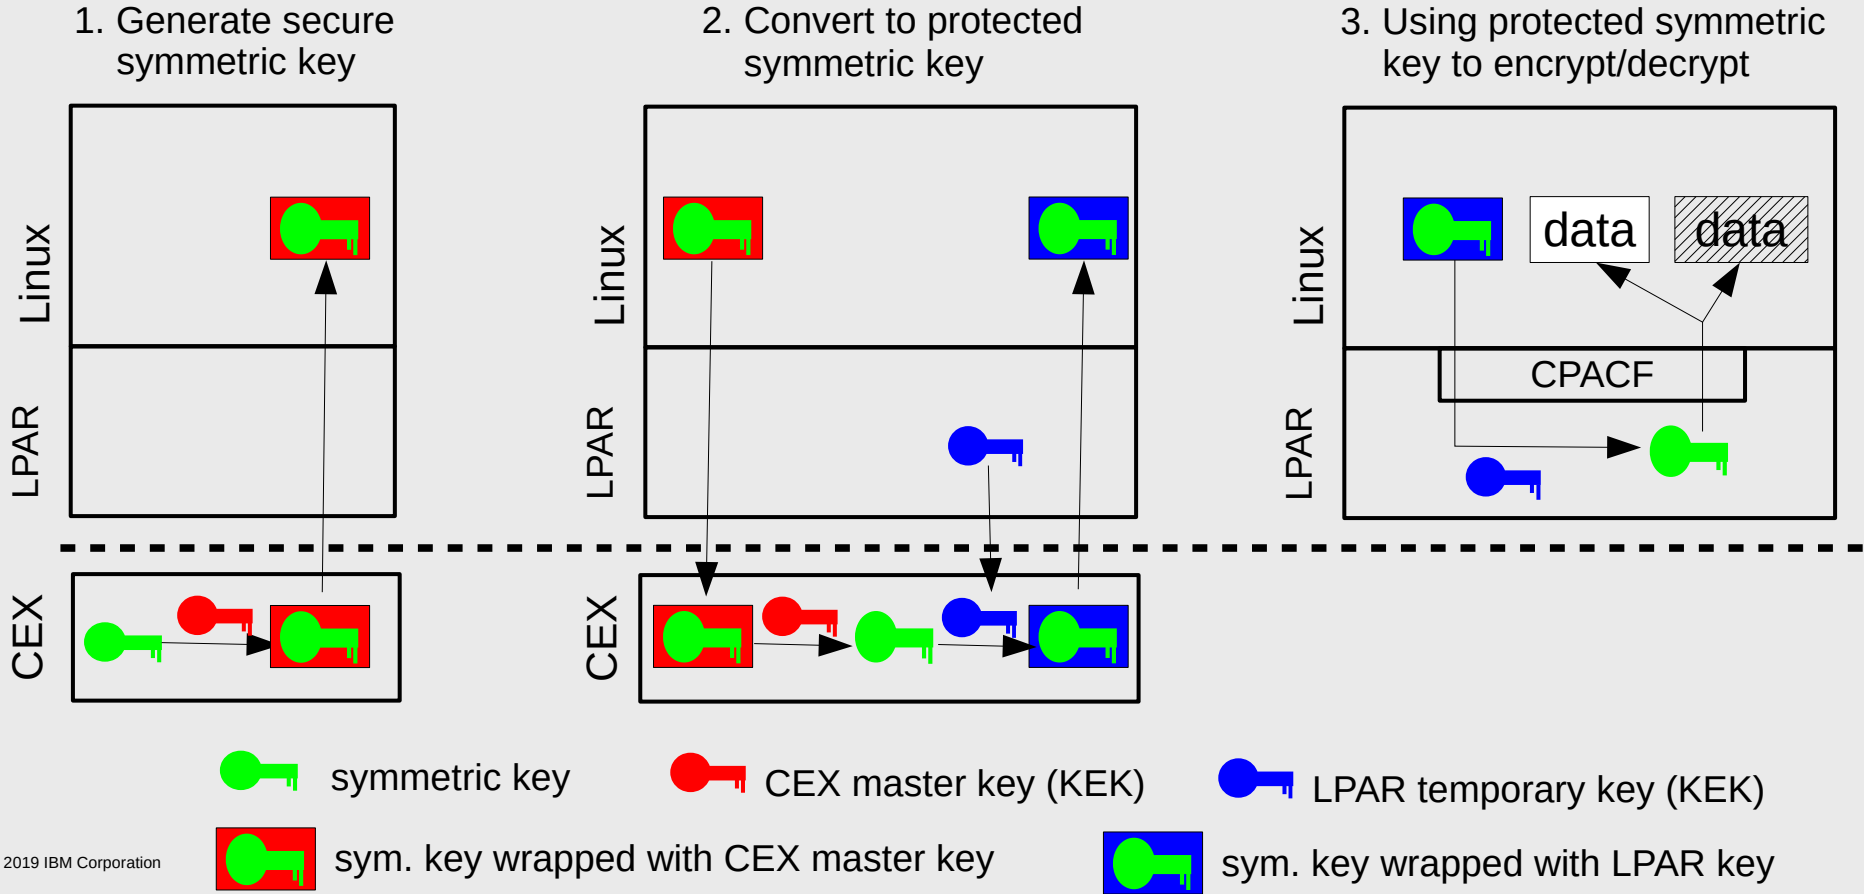
Kernel crypto automatically uses CPACF for AES if the module `aes_s390` is loaded

GSKit and latest versions of CLiC use CPACF for AES

# Linux File System Stack with dm-crypt



# Protected keys with CPACF



## pkey kernel module (kernel 4.11)

- Activate with: `# modprobe pkey`
- Provides IOCTLs via `/dev/pkey`
  - + `PKEY_GENSEC`
    - Generate a random CCA secure key
  - + `PKEY_CLR2SEC`
    - Generate a CCA secure key from a clear key
  - + `PKEY_SEC2PROT`
    - Generate a protected key from a CCA secure key
  - + `PKEY_FINDCARD`
    - Find an adapter and domain associated with a given secure key
  - + `PKEY_SECK2PROTK`
    - First `PKEY_FINDCARD`, then `PKEY_SEC2PROT`

## zkey user space tool (s390-tools 1.39)

- Requires `pkey` kernel module
- **Generate** option
  - + Generates file with AES secure key (AESDATA)
  - + Random key or from clear key
  - + Single key for CBC (64 bytes) or two keys for XTS (128 bytes)
- **Validate** option
  - + Checks if input file contains a valid AES secure key
  - + If yes, displays key attributes
- **Re-encipher** option, requires the CCA package
  - + Support master key change on Crypto Express adapter
  - + Transforms a valid secure key wrapped by an old master key into a secure key wrapper with a new master key

# Function Overview

Function	Component and Upstream versions	Distribution releases
PAES support for dm-crypt: plain format	kernel 4.11: pkey, paes_s390 s390tools 1.39: zkey	RHEL 7.5, RHEL 8 SLES 12 SP4, SLES15 Ubuntu 18.04
Master key rolling support for PAES keys for dm-crypt with plain format	kernel 4.11: pkey, paes_s390 s390tools 1.39: zkey	RHEL 7.5, RHEL 8 SLES 12 SP4, SLES15 Ubuntu 18.04
fast dm-crypt using 4kB sectors	kernel 4.12 cryptsetup 2.0.0	RHEL 7.6, RHEL 8 SLES 15 Ubuntu 18.04
4kB sector support for plain volumes in /etc/crypttab	systemd patches	RHEL 8 Ubuntu 18.10
PAES key repository for dm-crypt with plain format	s390tools 2.4: zkey	RHEL 7.6, RHEL 8 SLES 12 SP4 Ubuntu 18.10
PAES support for dm-crypt: LUKS2 format	kernel 4.11: pkey, paes_s390 cryptsetup 2.0.3	RHEL 7.6, RHEL 8 Ubuntu 18.10
PAES key repository for dm-crypt with LUKS2 format	s390tools 2.6: zkey	RHEL 8 Ubuntu 18.10
Master key rolling support for PAES keys for dm-crypt with LUKS2 format	s390tools 2.6: zkey-cryptsetup	RHEL 8 Ubuntu 18.10

# Kernel News

# Linux Kernel – Base IBM Z support


- Add LIC identifier to sysinfo (kernel 4.16)
  - Display licensed internal code identifier to the output of `/proc/sysinfo`
  - Can be a valuable debugging aid
- IBM z14 Model ZR1 base support (kernel 4.17)
  - Translate machine type 0x3907 to the ELF platform name “z14”
- I/O device pre-configuration (kernel 4.17)
  - An LPAR managed via DPM mode can retrieve a list of I/O devices via SCLP Store Data
  - FCP devices, OSA-Express, and HiperSockets devices can be enabled automatically
  - This helps with the device selection at installation time



# Linux Kernel – Base IBM Z support

- Add `kexec_file_load` system call (kernel 4.17)

- The existing `kexec` interface uses a binary image to boot into a new kernel
- The `kexec_file_load` interface uses a kernel image in ELF format
- The ELF format allows you to add additional information, e.g., a signature to verify the integrity of the kernel



- Rework initial boot sequence (kernel 4.19)

- Move the memory detection to the initial boot phase, some features need the exact memory layout very early
  - Kernel Address Sanitizer (KASAN)
  - Kernel Address Layout Randomization (KASLR)



- Add support for restartable sequences (kernel 4.19)

- Restartable sequences are a lightweight interface to run user code atomically relative to scheduler preemption and signal delivery
- Typically used to implement fast per-cpu operations in user space





# Linux Kernel – Base IBM Z support

- Kernel Address Sanitizer (kernel 4.20)
  - Compile time option for a debug kernel to detect dynamic memory errors
    - Relies on the gcc option `-fsanitize=kernel-address`
- Virtually mapped kernel stacks (kernel 4.20)
  - Allocate the kernel stack for tasks from the vmalloc space
  - Use a guard page to trap kernel stack overflows instead of instructions
  - Slightly faster kernel code and a ~200K size reduction of the kernel



CONFIG\_VMAP\_STACK=y





```
0000000000000000 <fn_a>:
000:  c0 04 00 00 00 00      brcl   0,0 <fn_a>
006:  eb bf f0 70 00 24      stmg   %r11,%r15,112(%r15)
00c:  b9 04 00 ef           lgr    %r14,%r15
010:  e3 f0 ff d0 ff 71      lay    %r15,-48(%r15)
016:  e3 e0 f0 98 00 24      stg    %r14,152(%r15)
...
```

VS

CONFIG\_VMAP\_STACK=n, CONFIG\_CHECK\_STACK=y

```
0000000000000000 <fn_a>:
000:  c0 04 00 00 00 00      brcl   0,0 <fn_a>
006:  eb bf f0 70 00 24      stmg   %r11,%r15,112(%r15)
00c:  a7 f1 3f c0           tml1   %r15,16320
010:  b9 04 00 ef           lgr    %r14,%r15
014:  a7 84 00 01           je     946 <fn_a+0x16>
018:  e3 f0 ff d0 ff 71      lay    %r15,-48(%r15)
01e:  e3 e0 f0 98 00 24      stg    %r14,152(%r15)
...
```

# Linux Kernel – Crypto Support










- Add support for up to 256 crypto adapters (kernel 4.17) 
- AP bus support for alternate drivers (kernel 4.19)
  - Will be needed for the virtualization of crypto cards 
- zcrypt multiple device nodes support (kernel 4.20)
  - Add new zcrypt device nodes with restrictions in regard to crypto cards, domains, and available ioctls 
  - The restricted zcrypt access is needed for container solutions like docker 
- Protected key API for random keys (kernel 4.20)
  - Introduce new ioctls and sysfs attributes to generate random protected and random secure keys
  - Useful to create encrypted paging devices with protected keys
  - The patch to the sysfs attribute can be added to `/etc/crypttab` directly

# Linux Kernel – Block Device Support

- Configurable IFCC handling for DASD (kernel 4.17)
  - Add a control how repeated IFCC/CCC errors are handled
  - The `path_autodisable` sysfs attribute can be set per DASD device
- Add zFCP port speed capabilities (kernel 4.18)
  - Decode and display the channel link speed of the FCP Express 32S card
  - The `lszfcpx -a` command can be used to show the port speed
- Support DIF (data integrity support) independently of DIX for ZFCP (kernel 5.0)
  - Configure DIF-only mode with kernel parameters: `zfcpx.dif=1 zfcpx.dix=0`



# Linux Kernel – Network Device Support

- Add diag26c support for VNIC info (kernel 4.16) z/VM   
  - When running on z/VM use the diagnose X'26C' to find network device related information
  - This is used to find out if the network connection uses layer-2 or layer-3
- Add qeth IPv6 RX/TX checksum offload support (kernel 4.18)  
  - Enable RX/TX checksum offload for inbound and outbound IPv6 network traffic
  - IPv6 checksum offload is available with OSA Express 4S cards
- Report 25Gbit link speed for qeth devices (kernel 4.20)  
  - Add the link speed of the OSA Express 7 cards to sysfs and ethtool
- Improve qeth TSO support (kernel 4.20)  
  - Add TCP segmentation offload for IPv6 and L3 devices, as well as, TSO for L2 devices

# Linux Kernel – Performance Analysis

- Provide detailed CPU information for perf record (kernel 4.17)
  - The CPU field with `perf report` has been a fixed string "IBM/S390"
  - Add CPU manufacturer, type, model-capacity, model, version, and CPU-MF authorization to the output
- Provide detailed PMU event descriptions (kernel 4.17)
  - Add JSON files for IBM z10EC, z10BC, z196, zEC12, zBC12, z13, z13s, and z14
  - The description can be retrieved with `perf list`
- Rewrite trace points of the common I/O layer for perf (kernel 4.19)
  - Make the trace point data of the common I/O layer consumable by `perf stat`
- Add the CPU-MF transaction counters to perf (kernel 4.19)
  - Display the number of transactions ends/aborts with `perf stat -T`



# Linux Kernel – Performance Analysis

- Export CPU-MF sampling data as perf auxiliary trace data (kernel 4.19)
  - Collect sampling data with `perf record -e rbd000`
  - Display sampling data as raw report `perf report -D` or GUI report `perf report`
- Add auxiliary trace data files to perf (kernel 4.20)
  - Used with CPU-MF sampling to write individual files for each CPU
  - The `perf config` command is used to specify a target directory with `auxtrace.dumpdir`
  - For each CPU a `aux.smp.##` file will be generated




# Further Reading

## Documentation

- Video about Pervasive Encryption of Data Volumes - <https://youtu.be/jDK3ZwEdX4I>
- How-to about Pervasive Encryption of Data Volumes - [http://public.dhe.ibm.com/software/dw/linux390/docu/l5n1dc00\\_a\\_quick\\_start.pdf](http://public.dhe.ibm.com/software/dw/linux390/docu/l5n1dc00_a_quick_start.pdf)
- Linux on Z on DeveloperWorks - [https://www.ibm.com/developerworks/linux/linux390/documentation\\_dev.html](https://www.ibm.com/developerworks/linux/linux390/documentation_dev.html)

## Blogs



- Linux on Z - <http://linux-on-z.blogspot.com/>
- Containers on IBM Z - <https://containersonibmz.com/>







# Tag Legend


- Supported distributions


 for SUSE SLES <X> Service Pack <Y>, e.g.  for SLES12 SP3

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

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

- Supported environments

 usable for systems running in LPAR mode

 usable for guests running on z/VM