

How to Virtualize IBM zSystems Hardware Cryptography with z/VM

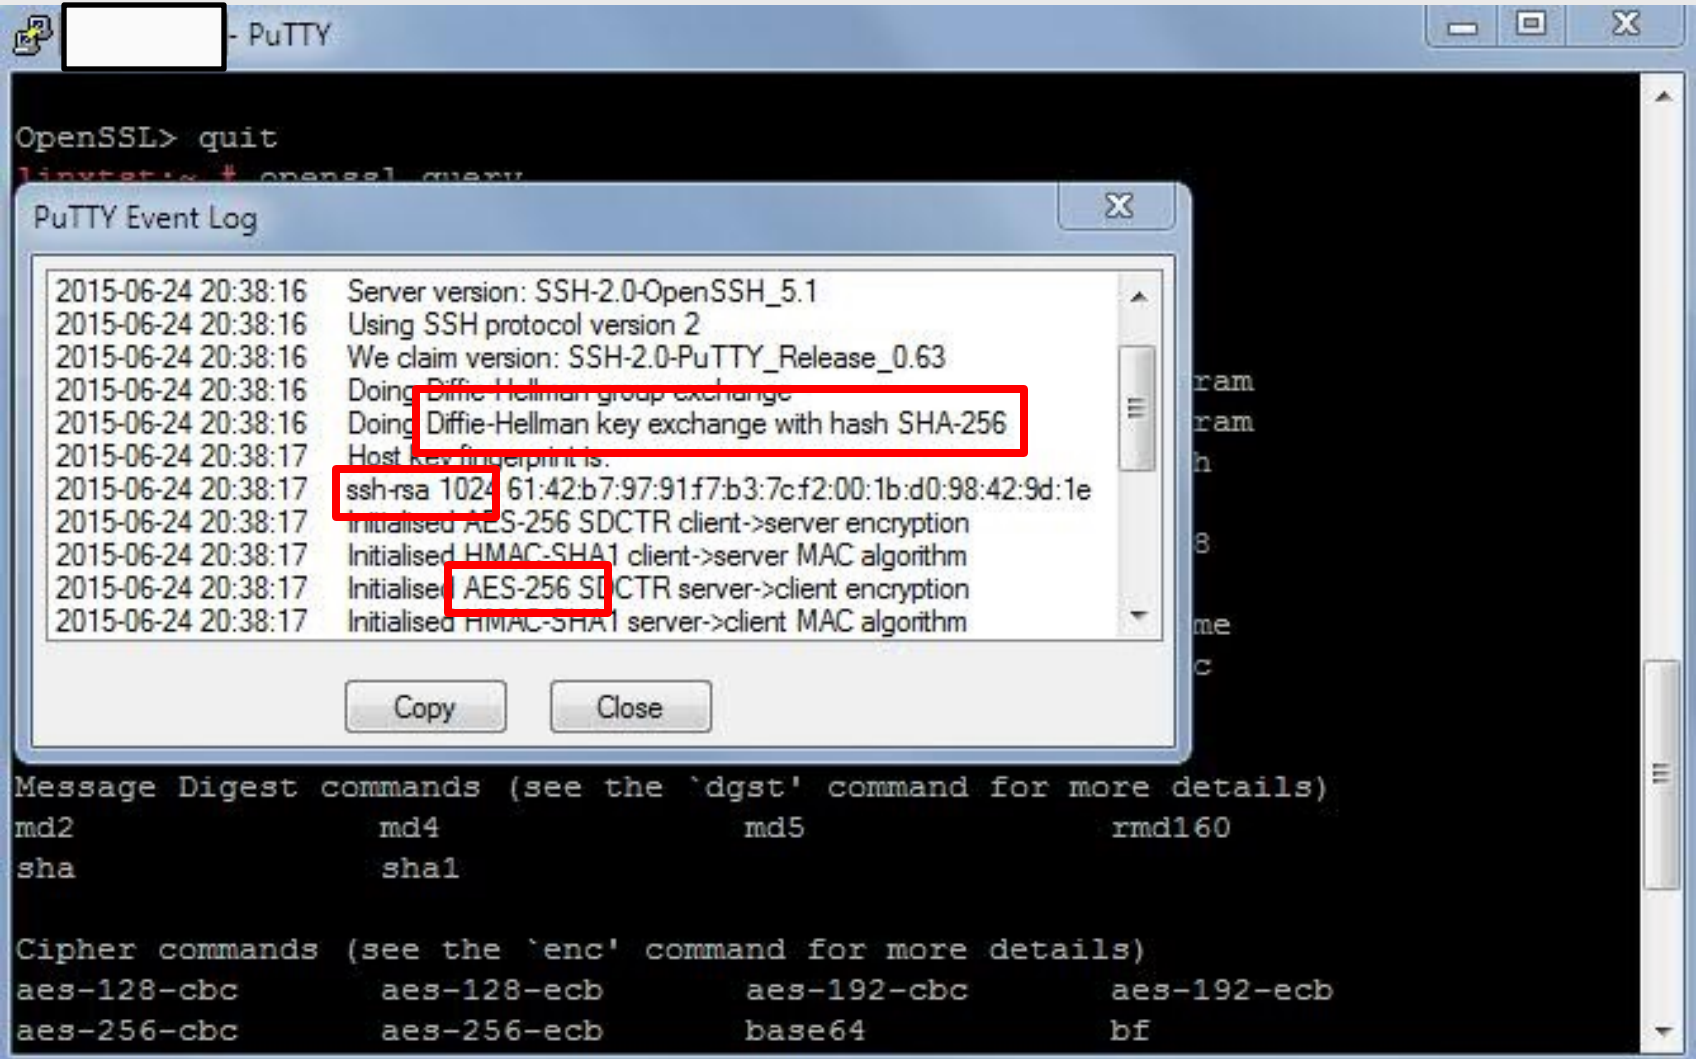
*With a discussion on workloads, keys, and
Dynamic Crypto Support for z/VM*

Brian W. Hugenbruch, CISSP
IBM zSystems Security for Virtualization & Cloud
bwhugen@us.ibm.com –  @Bwhugen

Cryptography acronyms and terms. There will be a quiz later.

AES	Advanced Encryption Standard	KGUP	If You Can Read This Raise Your Hand
ARL	Authority Revocation List	LDAP	Lightweight Directory Access Protocol
CA	Certification Authority	MAC	Message Authentication Code
CBC	Cipher Block Chaining	MDC	Message Detection Code
CCA	IBM Common Cryptographic Architecture	MD5	Message Digest 5
CCF	Cryptographic Coprocessor Facility	OAEP	Optimal Asymmetric Encryption Padding
CDSA	Common Data Security Architecture	OCSF	OS/390 Open Cryptographic Services Facility
CEX7S/6A	Crypto Express7S Accelerator Mode	OCSP	Online Certificate Status Protocol
CEX7S/6C	Crypto Express7S CCA Coprocessor Mode	PCICA	PCI Cryptographic Accelerator
CFB	Cipher Feedback	PCICC	PCI Cryptographic Coprocessor
CKDS	Cryptographic Key Data Set	PCIXCC	PCIX Cryptographic Coprocessor
CPACF	CP Assist for Cryptographic Facilities	PKA	Public Key Architecture
CRL	Certificate Revocation List	PKCS	Cryptographic Standards, e.g. PKCS #1 / EP11
CRT	Chinese Remainder Theorem	PKDS	Public Key Data Set
CVC	Card Verification Code	PKI	Public Key Infrastructure
CVV	Card Verification Value	RA	Registration Authority
DES	Data Encryption Standard	RACF	Resource Access Control Facility
DSA	Digital Signature Algorithm	RSA	Rivest-Shamir-Adleman
DSS	Digital Signature Somethhing	SET	Secure Electronic Transaction
ECB	Electronic Code Book	SHA	Secure Hash Algorithm
ECC	Elliptic Curve Cryptography	SLE	Secure Cookie Monster Encryption
FIPS	Federal Information Processing Standard	SSL	Secure Sockets Layer. See TLS
GCM	AES Galois/Counter Mode	TKE	Trusted Key Entry
ICSF	Integrated Cryptographic Service Facility	TLS	Transport Layer Security. See SSL.
IETF	Internet Engineering Task Force	VPN	Virtual Private Network
IPKI	Is Anyone Reading This Line		

Security looks complicated, and it happens quickly.



What just happened?

SSH connections and TLS connections use:

- **Asymmetric** key exchange to establish a connection
- **Symmetric** keys to encrypt bulk traffic
- **Hashing** to validate content integrity between source and target



That's a lot of math ... and it's processing power that adds up

- Happens for every secure operation (connection, application math, etc.)
- The bigger (more secure) the keys, the longer it takes
- Costs time, money

Why Use IBM zSystems Hardware Cryptography?

- Maximize Trust & reliability – proven hardware implementations
- Minimize Cost
 - Save money: offload expensive CPU workload
 - Save time: Faster crypto algorithms
- Industry-leading security
 - special built-in functions for banking and financial applications (secure key)
- Regulatory compliance starts at hardware



How should I configure my IBM zSystems Crypto (or: why you're here)

Let me counter with a few questions:

- *What workload are you running?*
- *What are your security requirements?*
 - *“Know your rules”*
 - *This changes if you're running a cloud*
 - *This changes if you're supporting multinational customers*
- *What are your mobility requirements?*
 - *Do you need z/VM Live Guest Relocation?*
 - *Do you have other plans for failover, planned outages, and HA/DR?*
 - *Who'd win in a fight – your business continuity team, or security team?*



Today's Topics

IBM zSystems Hardware Cryptography
(the parts of it, why it matters, what it means)

z/VM Virtualization of IBM zSystems Cryptography
(and how to use it)

Guest Support: Operating Systems Running on z/VM

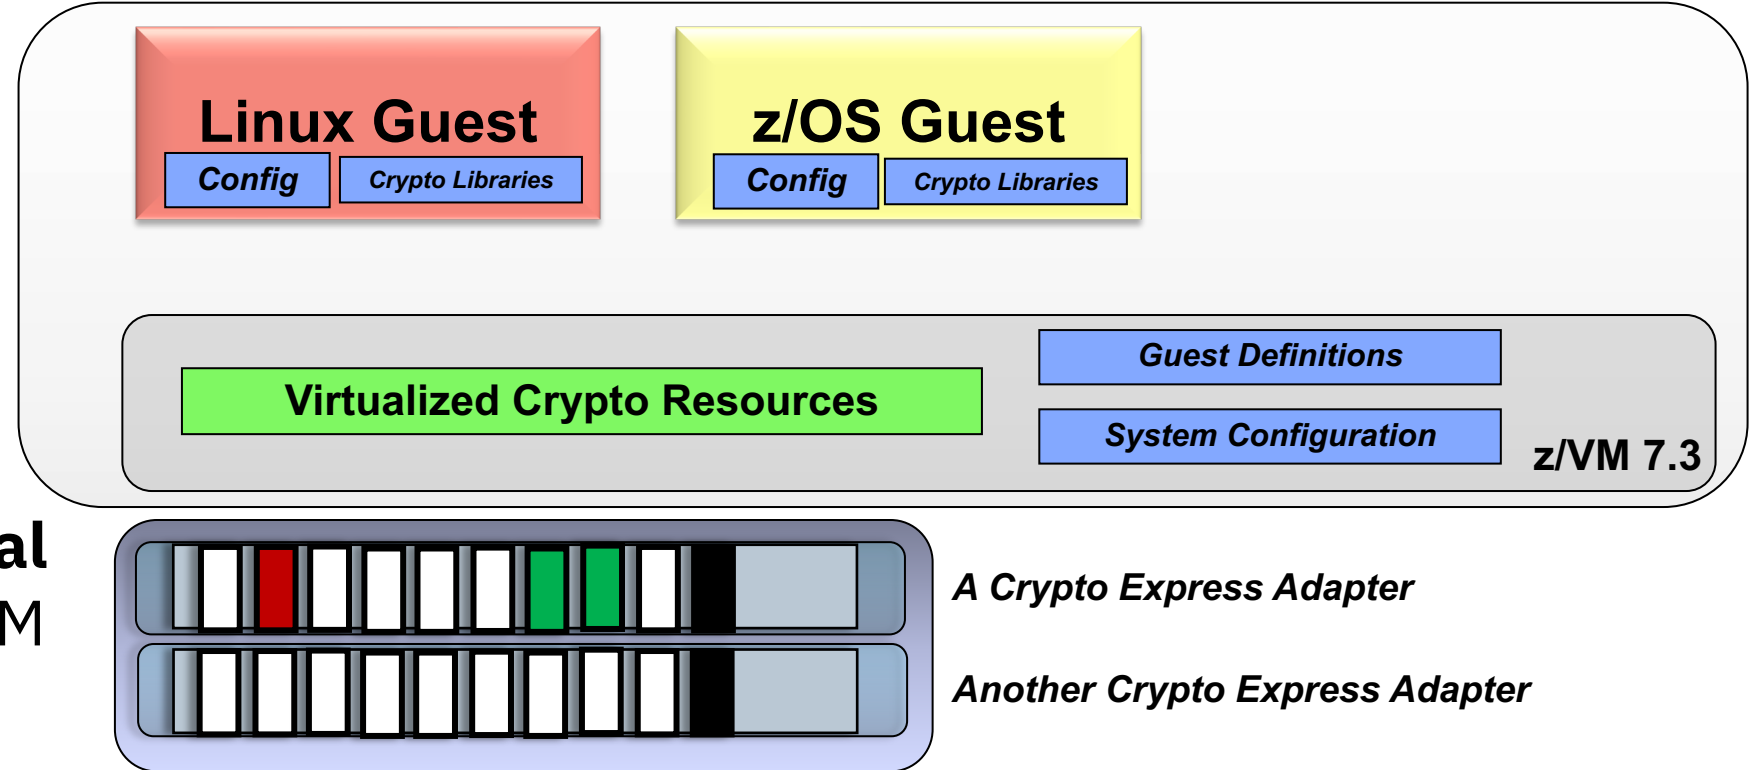
Extra: Frequently Asked Questions (if you don't ask them first)



IBM zSystems Hardware Crypto Support

How To: Configure your Crypto on IBM zSystems and LinuxONE

1. Install the **features**
2. Configure **adapters** on HMC/SE
3. Configure your **hypervisor**
4. Configure your **virtual machines** at the z/VM level
5. Configure your **guest operating system(s)**




IBM zSystems and LinuxONE Integrated Cryptographic Hardware

CP Assist for Cryptographic Functions (CPACF)

- Hardware accelerated encryption on every microprocessor core


Suited for high-speed bulk symmetric encryption



Crypto Express7S

- Next generation PCIe Hardware Security Module (HSM)
- Industry leading FIPS 140-2 Level 4 Certification

Suited for high value transactions, key protection and asymmetric acceleration



Why is it valuable:

- More performance = lower latency + less CPU overhead for encryption operations
- Highest level of protection available for encryption keys
- Industry exclusive “protected key” encryption

CP-Assisted Cryptographic Facility (CPACF)

No-charge feature on IBM zSystems hardware (Feature 3863)

On-chip cryptographic acceleration and operations

Enablement required to use the Crypto Express hardware

The screenshot shows the 'CETUS Details - CETUS' window with the following configuration details:

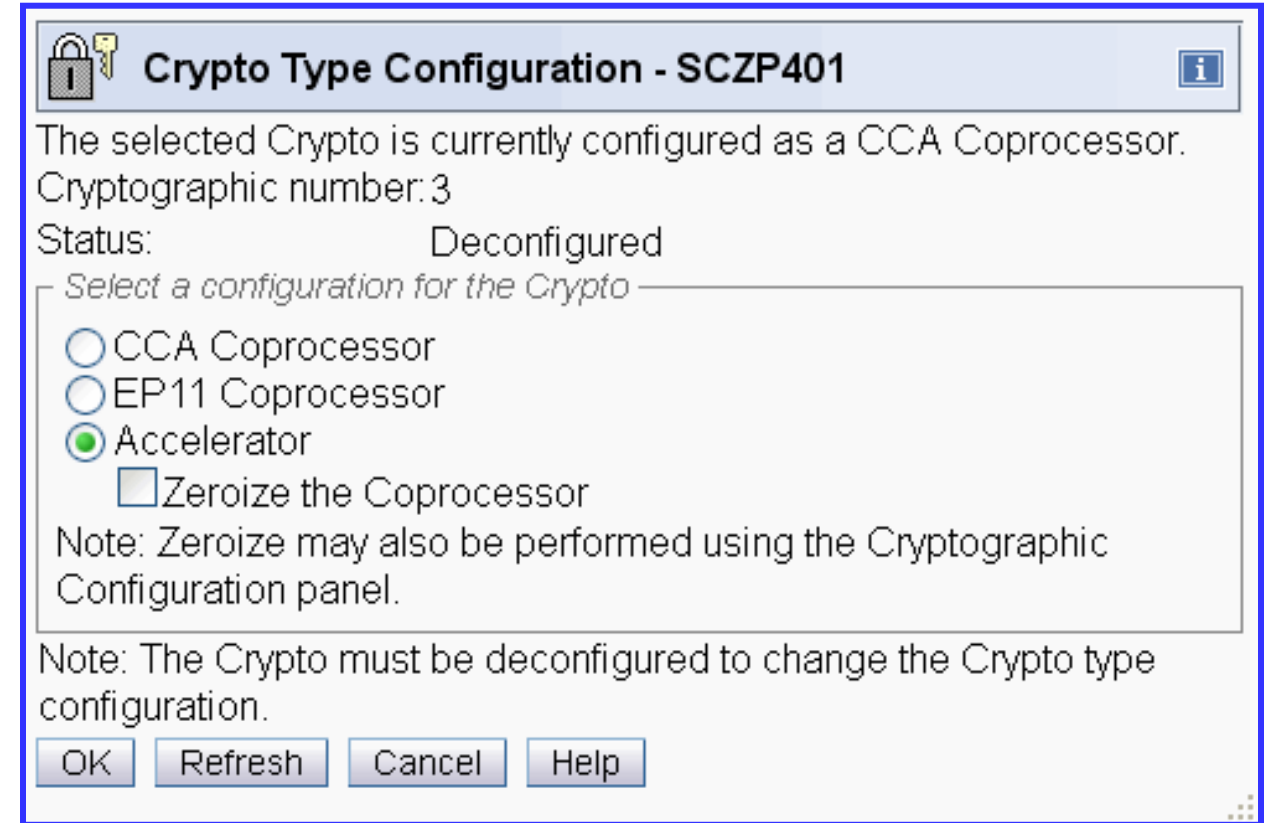
Instance Information	Product Information	Acceptable CP/PCHID Status	STP Information	Energy Management
Group:				CPC
CP status:				Operating
Channel status:				Exceptions
Crypto status:				Exceptions
Alternate SE status:				Operating
Activation profile:				DEFAULT
Last profile used:				DEFAULT
IOCDS identifier:				A1
IOCDS name:				IODF00
System mode:				Logically Partitioned
Service state:				false
Number of CPs:				41
Number of ICFs:				0
Number of IFLs:				48
Number of zIIPs:				16
Dual AC power maintenance:				Fully Redundant
CP Assist for Crypto functions:				Installed
Primary Licensed Internal Code security mode:				notification
Alternate Licensed Internal Code security mode:				notification
Lock out disruptive tasks:				<input type="radio"/> Yes <input checked="" type="radio"/> No

Buttons at the bottom: OK, Apply, Change Options..., Cancel, Help.

Setting Operational Mode for a Crypto Express Adapter (1/2)

Configuration for a Crypto Express feature is done on the **Hardware Management Console (HMC)**

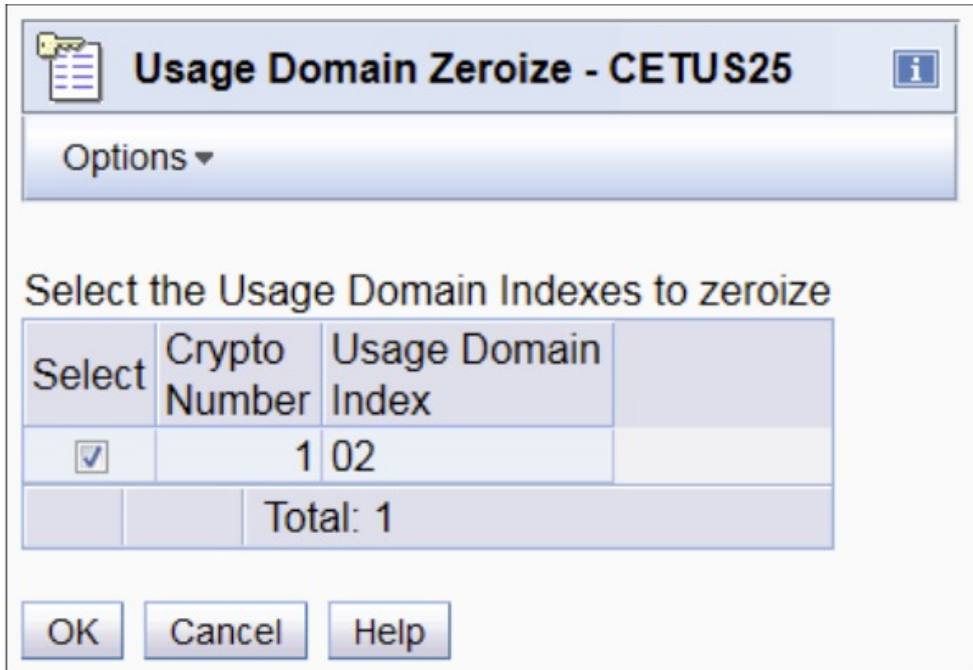
- **Step 1:** Make sure CPACF is enabled.
- **Step 2:** Select adapter, then choose the operational mode
 - Accelerator (clear key only)
 - CCA Coprocessor (more security, HSM features)
 - EP11 (open-source crypto framework, also has HSM features)



Setting Operational Mode for a Crypto Express Adapter (2/2)

Step 3: Validate option selection

- May zeroize existing keys in the process (destroy any residual secrets)



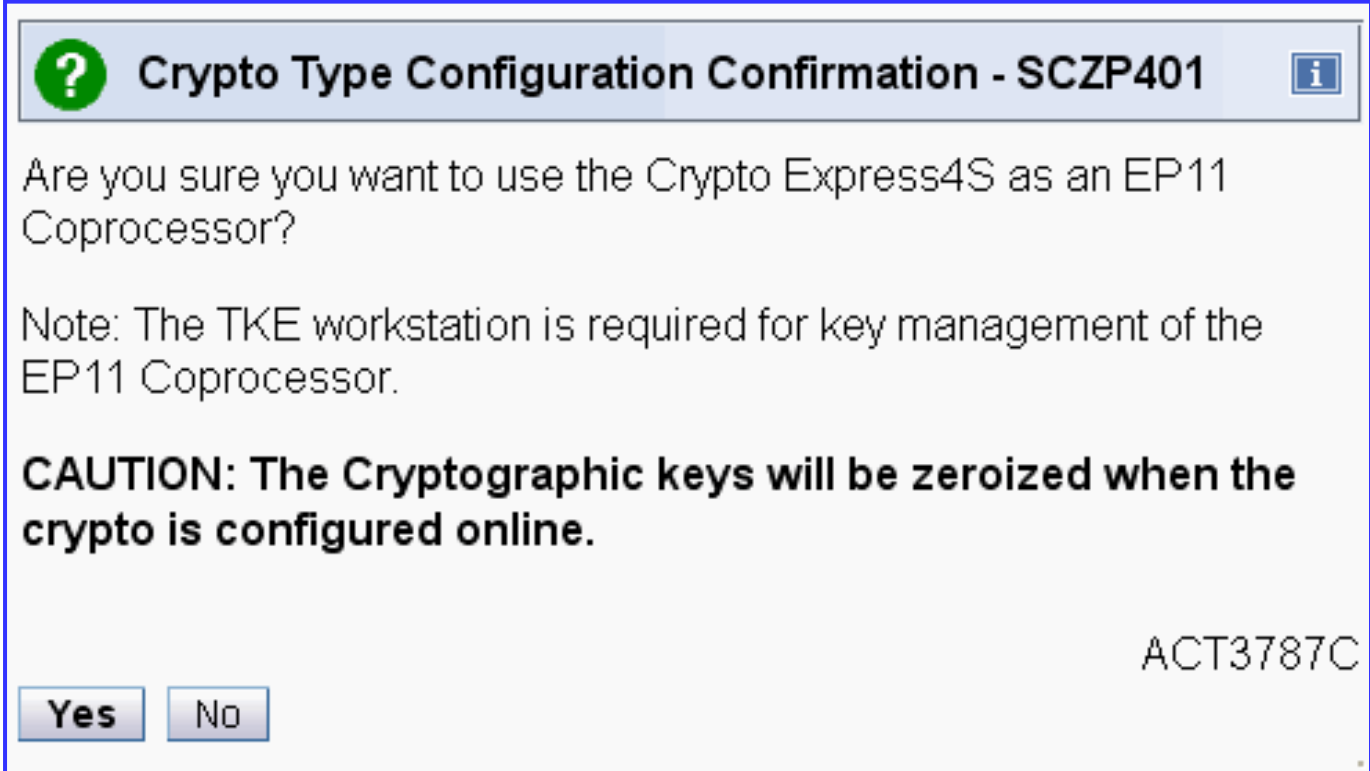
Usage Domain Zeroize - CETUS25

Options ▾

Select the Usage Domain Indexes to zeroize

Select	Crypto Number	Usage Domain Index	
<input checked="" type="checkbox"/>	1	02	
		Total: 1	

OK Cancel Help



Crypto Type Configuration Confirmation - SCZP401

Are you sure you want to use the Crypto Express4S as an EP11 Coprocessor?

Note: The TKE workstation is required for key management of the EP11 Coprocessor.

CAUTION: The Cryptographic keys will be zeroized when the crypto is configured online.

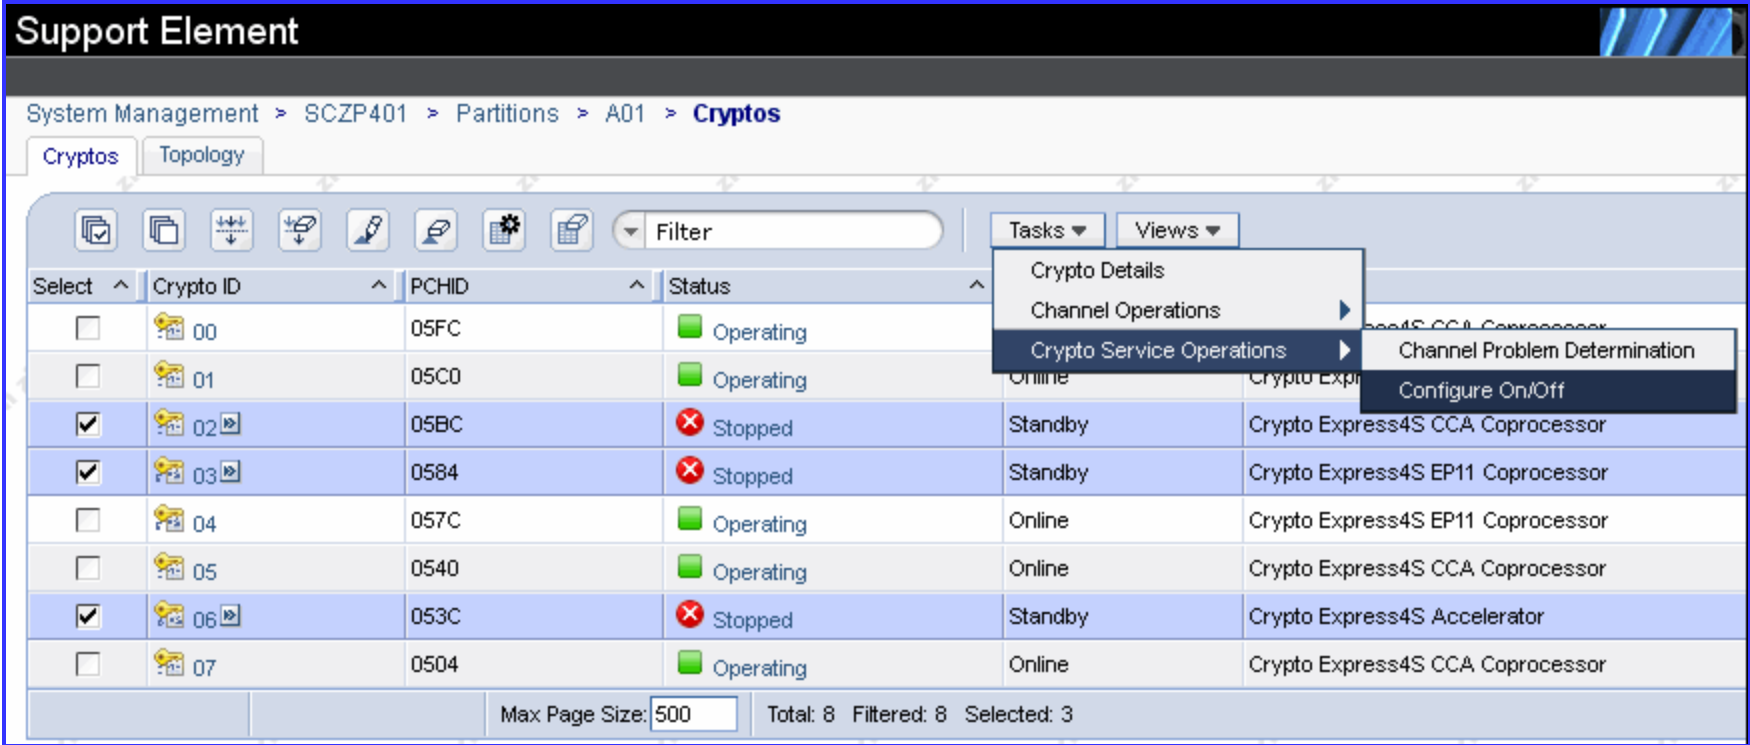
ACT3787C

Yes No

Activating a Crypto Express Adapter

Hardware activation is done from the Support Element

Select pertinent feature, "Configure On/Off"



Attaching a Crypto Express logical domain to an LPAR

LPAR assignment is done from the HMC (building an activation profile)

- **Candidate list:** domains on this adapter which are **eligible to be accessed** by this partition
- **Online List:** crypto resources automatically brought online at LPAR startup.
- **Usage Domain:** bundles domains across assigned adapters inside a common cryptographic boundary
- **Control Domain:** identifies domain index pertinent to TKE control of the LPAR. *If the Usage Domain is checked, the Control Domain must also be checked.*

z/VM will only detect those adapters and domains assigned to the LPAR

Customize Image Profiles: CETUS : CETUS24 : Crypto

CETUS
CETUS24
General
Processor
Security
Storage
Options
Load
Crypto

Assigned Domains

Select	Index	Control	Control and Usage
<input type="checkbox"/>	1		✓

Assigned Cryptos

Select	Number	Candidate	Candidate and Online
<input type="checkbox"/>	0		✓
<input type="checkbox"/>	1	✓	
<input type="checkbox"/>	3		✓
<input type="checkbox"/>	7		✓

Attention: You must install the 'IBM CP Assist for Cryptographic Functions' (CPACF) feature if a cryptographic candidate is selected from the list box. Otherwise, some functions of Integrated Cryptographic Service Facility (ICSF) may fail.

Cancel Save Copy Profile Paste Profile Help

Getting Keys into Your Crypto Express features

Trusted Key Entry (TKE) Workstation – an optional priced feature which communicates directly with the Crypto Express features over a secure TCP/IP connection.

- Functions as a separate physical device to the side of your IBM zSystems or LinuxONE
- Card reader for crypto secret storage
- Generates new secrets, stores data in Crypto Express domains
- Required if running Crypto Express features in EP11 mode!

z/OS Integrated Cryptographic Services Facility (ICSF) – a base component which allows interaction with Crypto Express features. (Requires z/OS; only for installing keys in CCA mode).

Panel and Catcher Utilities for Linux – **Panel** is a Linux package installed as part of the IBM .rpms which allows for key management function. **Catcher** is the Linux daemon for communicating with TKE.

- `/opt/IBM/CEX8C/bin/panel.exe`

IBM Enterprise Key Management Foundation (EKMF) – an IBM Lab Services offering for flexible and secure key management services.

- See also *Advanced Crypto Service Provider*
- <http://www-05.ibm.com/dk/security/cccc/products/acsp.html>

IBM HyperProtect Crypto Services – a cloud offering for key storage and retrieval.

- <https://www.ibm.com/cloud/hyper-protect-crypto>

Create an HSM as a Service

☰ IBM Cloud

← View all

Hyper Protect Crypto Services

Lite **Experimental**

Attention: This service is experimental. It might not yet be stable and might change in ways that make it incompatible with earlier versions. This service is not recommended for production environments.

IBM Cloud Hyper Protect Crypto Services is a complete set of encryption and key management services backed by IBM Z technology. These services bring the security and integrity of IBM Z to the cloud. The same state of the art cryptographic technology relied upon by banks and financial services is now offered to cloud users via IBM Cloud. The network addressable Hardware Security Module provides safe and secure PKCS#11 cryptography via industry standard open source application programming interfaces. It supports secure key operations and random number generation via IBM Z cryptographic hardware, FIPS-140-2 level 4 certified technology. This is the industry's first and only FIPS 140-2 Level 4 certified technology in the public cloud market today and is the same technology that is the backbone of the IBM Enterprise Blockchain solution.

[View Docs](#) [Terms](#)

AUTHOR IBM
 PUBLISHED 05/30/2018
 TYPE Service

Service name:

MyHSM

Choose a region/location to deploy in: **Select a resource group:**

US South Default

Pricing Plans

Monthly prices shown are for country or region: [Germany](#)

PLAN	FEATURES	PRICING
✓	Hyper Protect Crypto Services - Lite Plan 10 Crypto Slots <hr/> This Lite Plan enables you to use Hyper Protect Crypto Services for free as an Experimental offering Lite plan services are deleted after 30 days of inactivity.	Free

FEEDBACK

Hyper Protect Crypto Services

IBM Cloud Hyper Protect Crypto Services provides cryptographic functions from a high

Lite **Experimental**

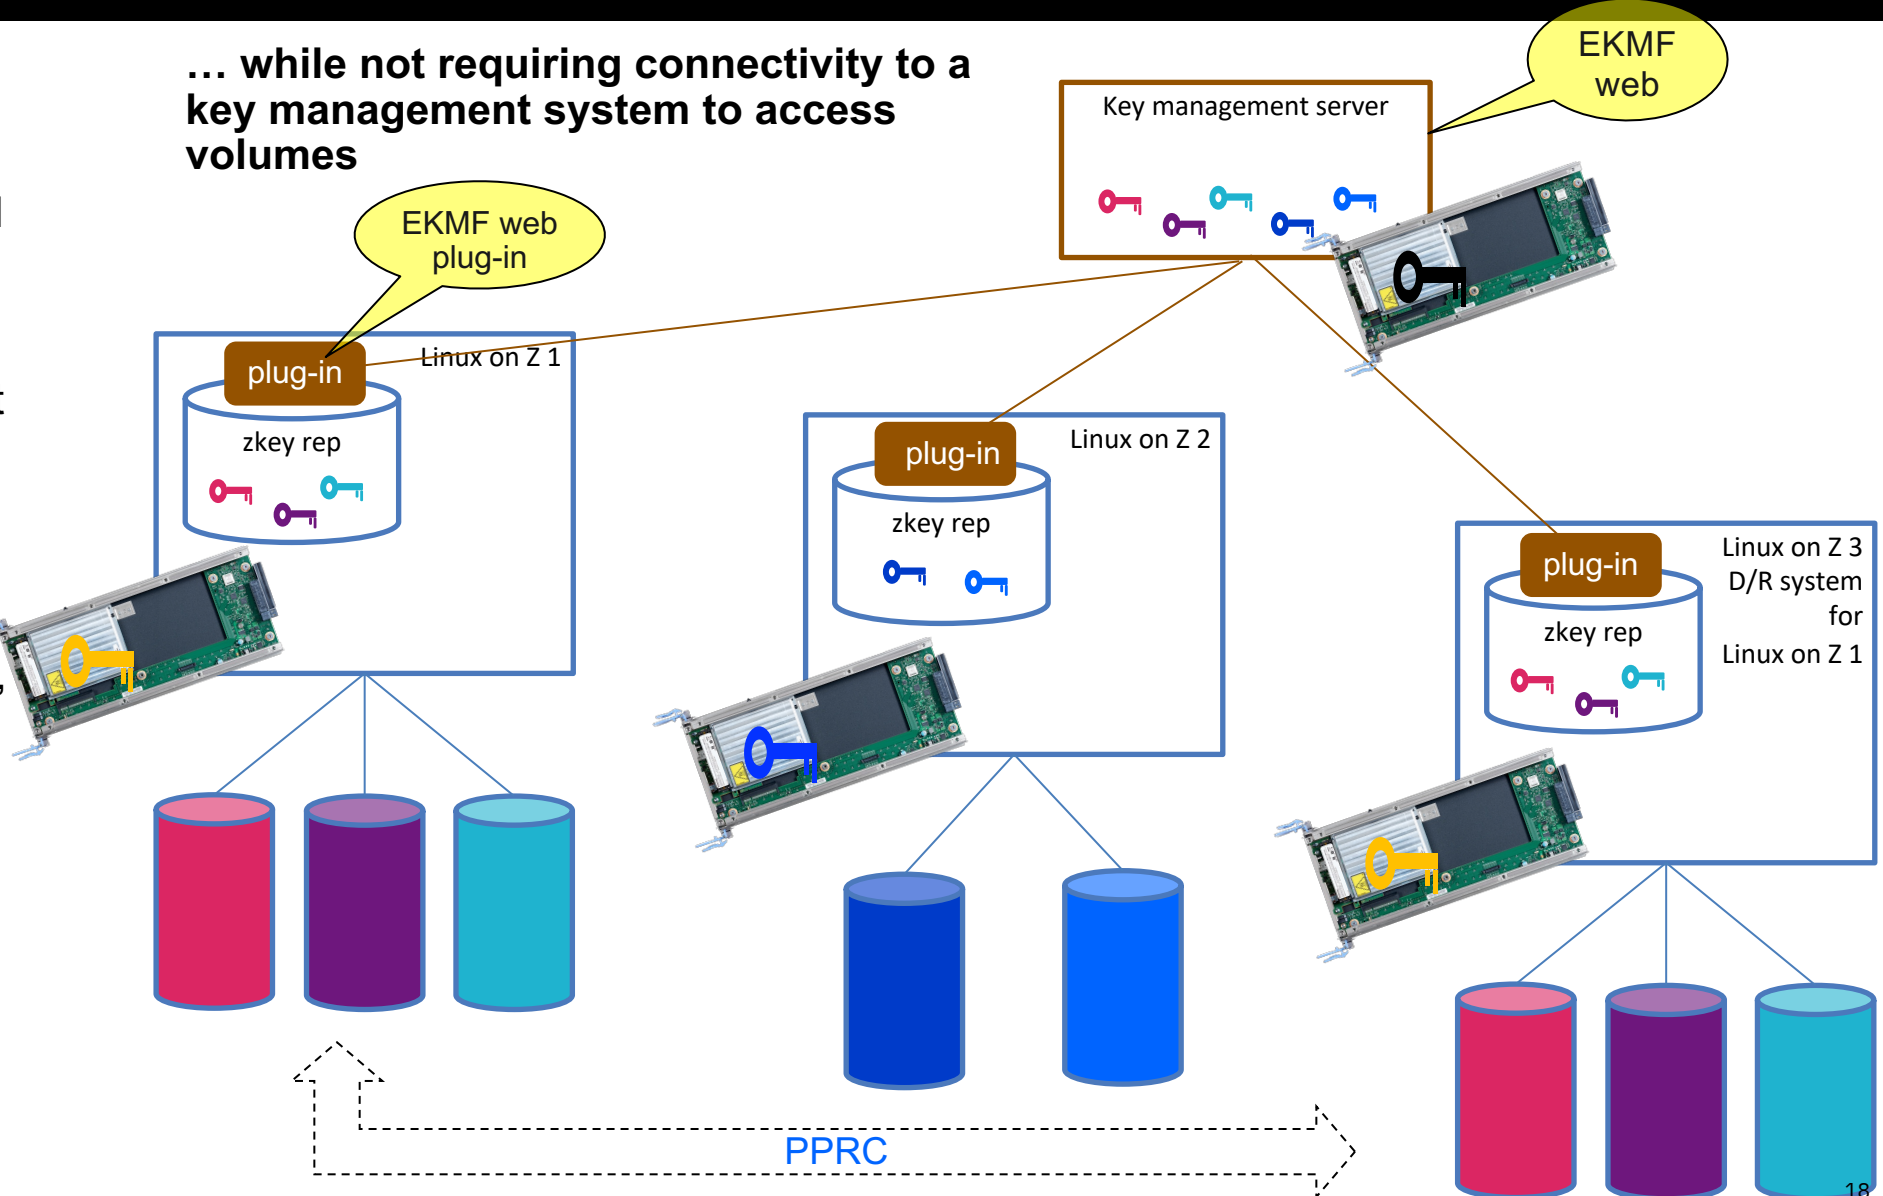
zkey now connects to EKMF Web

upstream: new KMIP plug-in for KMIP server (GKLM, GDE)

Goal:
 Allow for central enterprise key management system managing keys in all zkey repositories

- Generate and manage all keys on key management server
- Access control by key server:
- restrict key access to set of authorized clients (e.g., primary and backup systems)
- Simplify MK role procedures: just reimport from key management server

... while not requiring connectivity to a key management system to access volumes



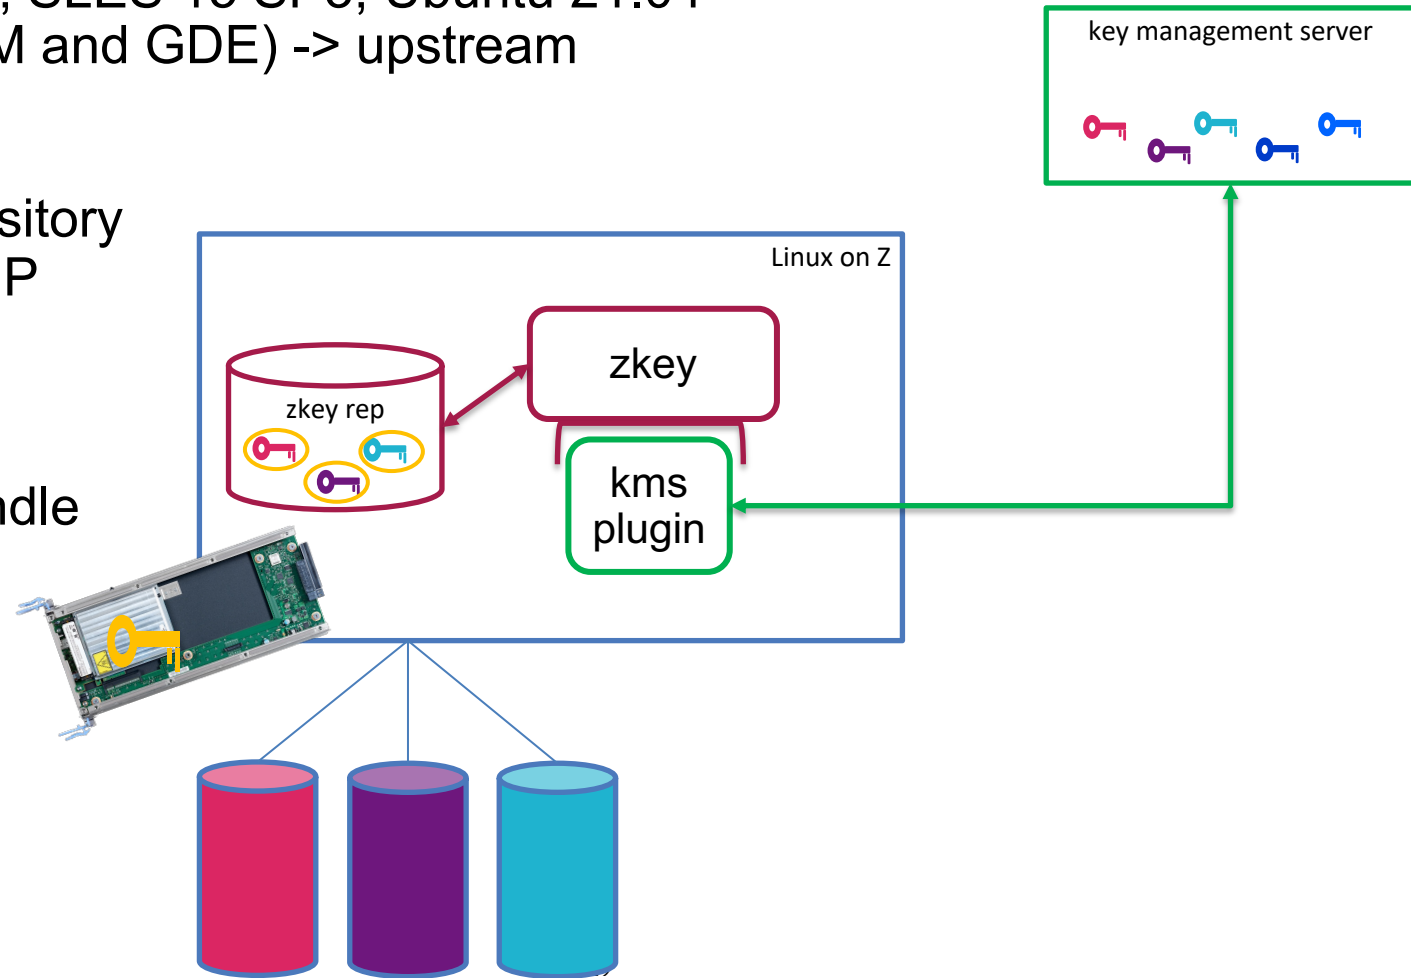
zkey key management server (kms) plug-in

kms plugins available for zkey

- EKMFweb plugin -> RHEL 8.4, SLES 15 SP3, Ubuntu 21.04
- KMIP plugin (tested with GKLM and GDE) -> upstream


using kms plugins

1. bind kms plug-in to zkey repository
 - choose EKMFweb or KMIP
2. configure bound kms plug-in
 - local APQNs
 - connection to server
 - server certificate / CA bundle
 - identity key
3. use zkey kms commands
 - generate
 - kms list
 - kms import
 - kms refresh
 - kms reencipher
 - remove




IBM zSystems Operational Keys: Clear, Protected, or Secure

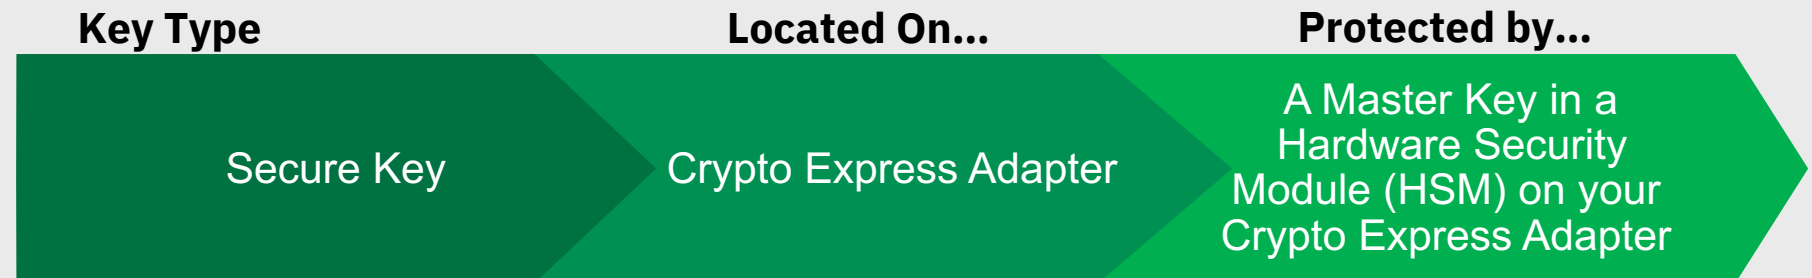
- Clear Keys are not encrypted. **Crypto operations may be performed in CPACF or on a Crypto Express adapter**



- Protected keys are encrypted under a CPACF wrapping key. **Crypto operations are performed only using CPACF**



- Secure keys have key values that are encrypted by a Master Key on a tamper-responding **Crypto Express adapter**.

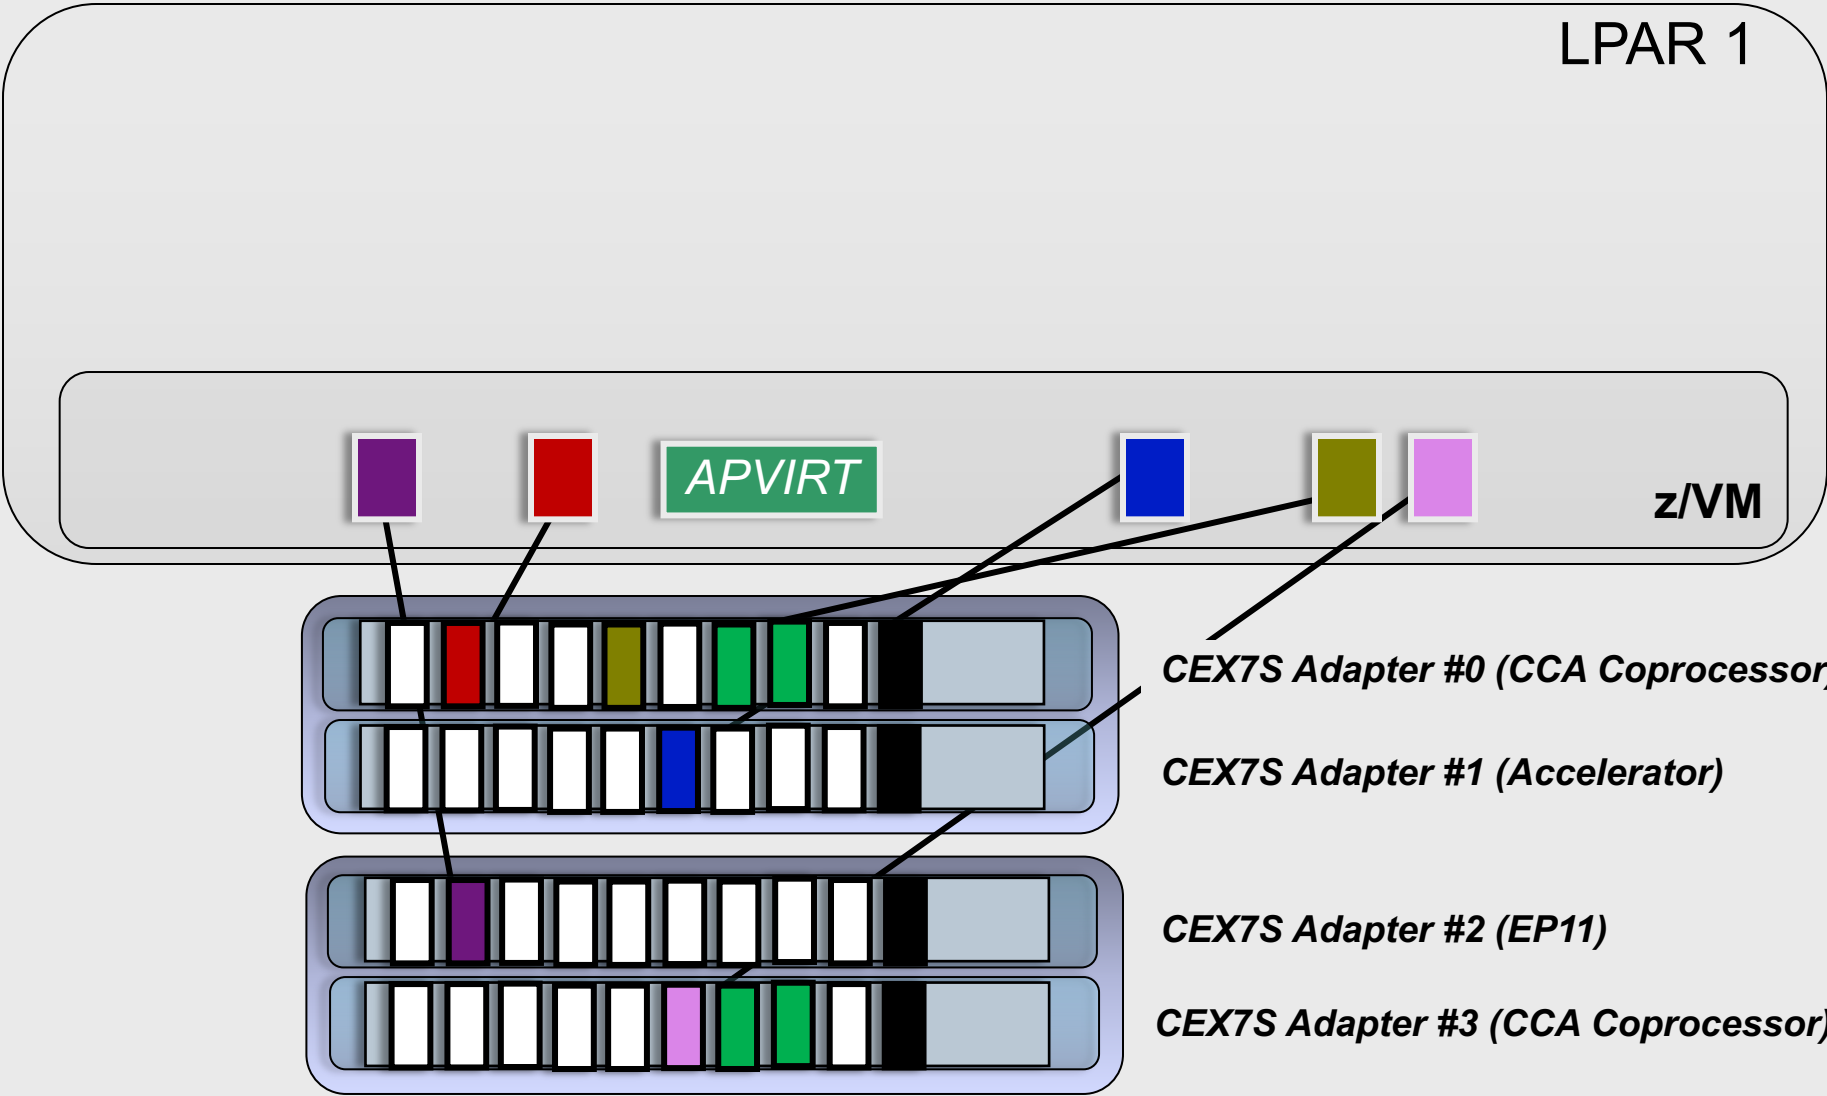


z/VM Support for Hardware Crypto

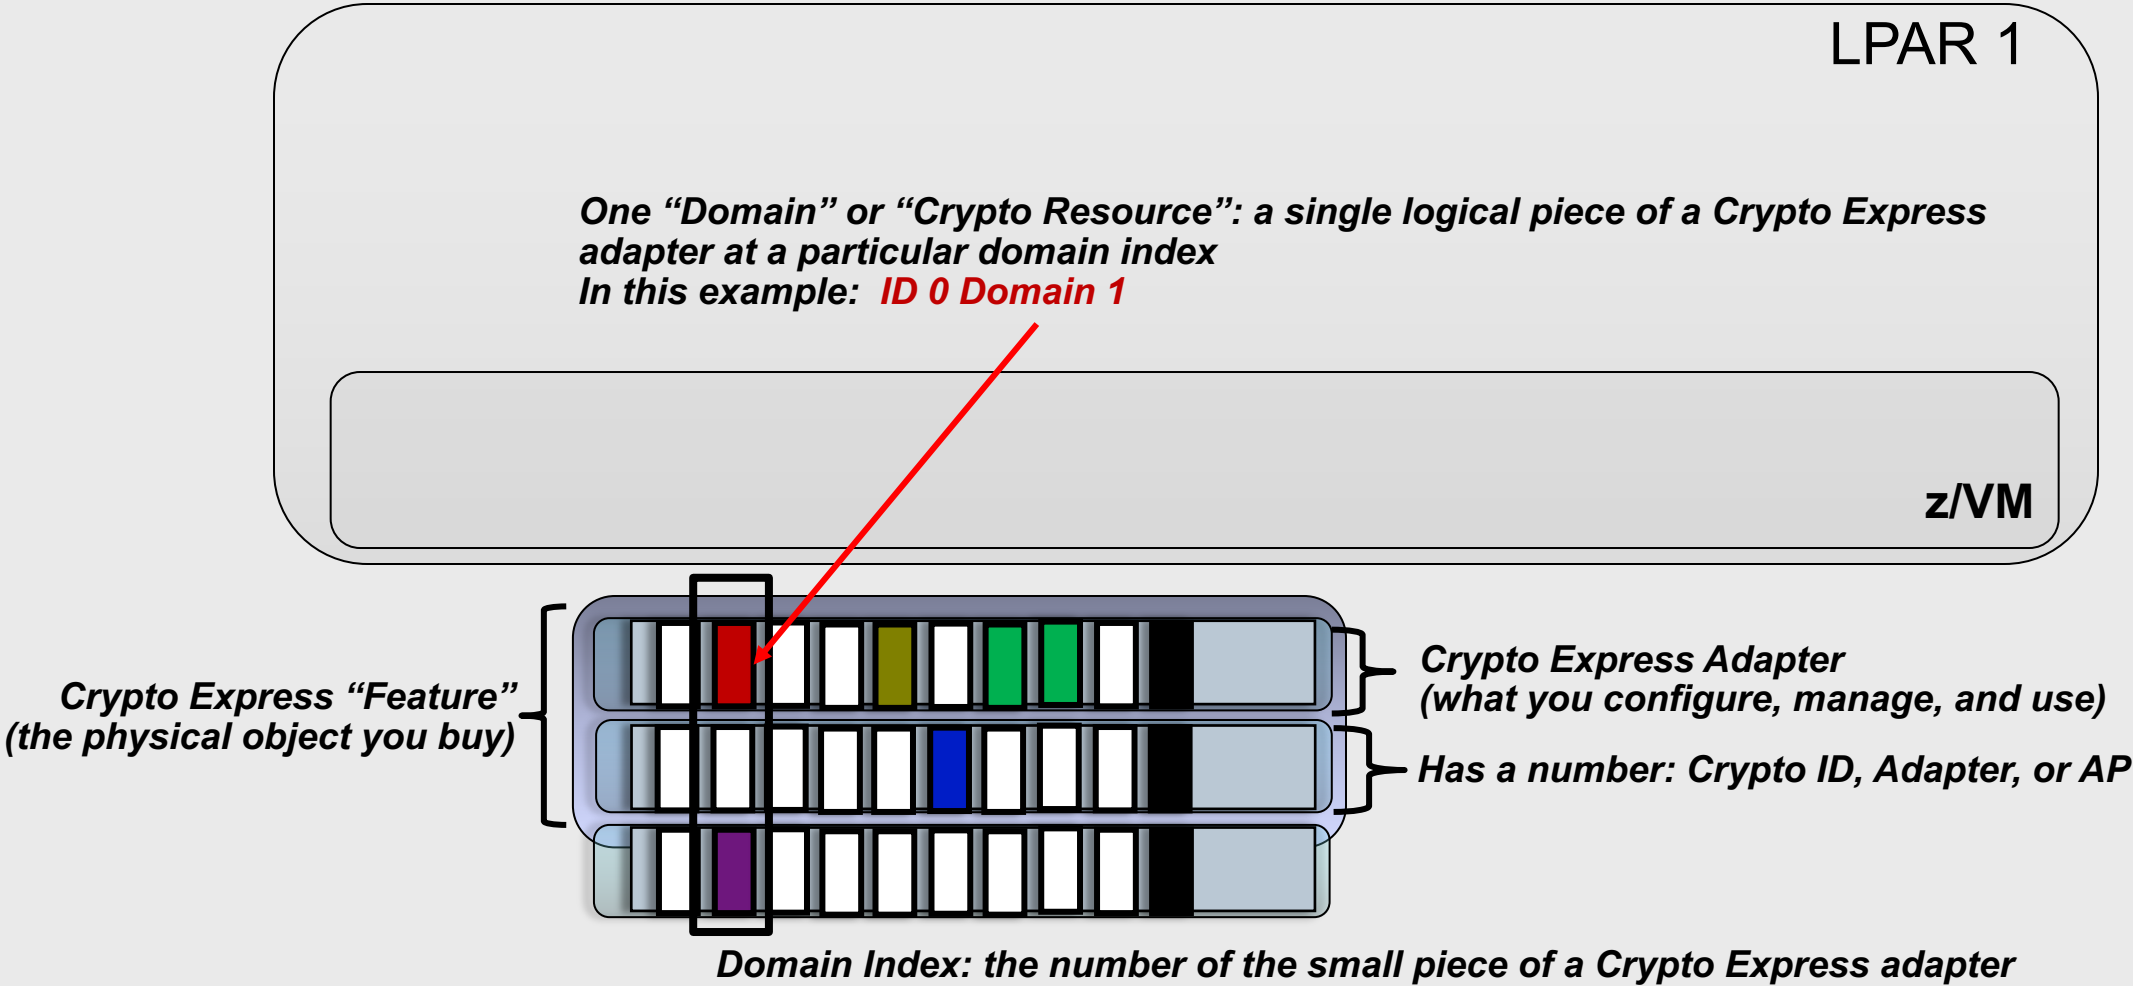
z/VM Virtualization of Hardware Cryptography (z/VM's view)

Once domains are assigned to a z/VM LPAR for use, they appear to the hypervisor and can be used by virtual machines.

z/VM sees crypto resources as virtual devices represented by a **Crypto ID** and a **domain index**.

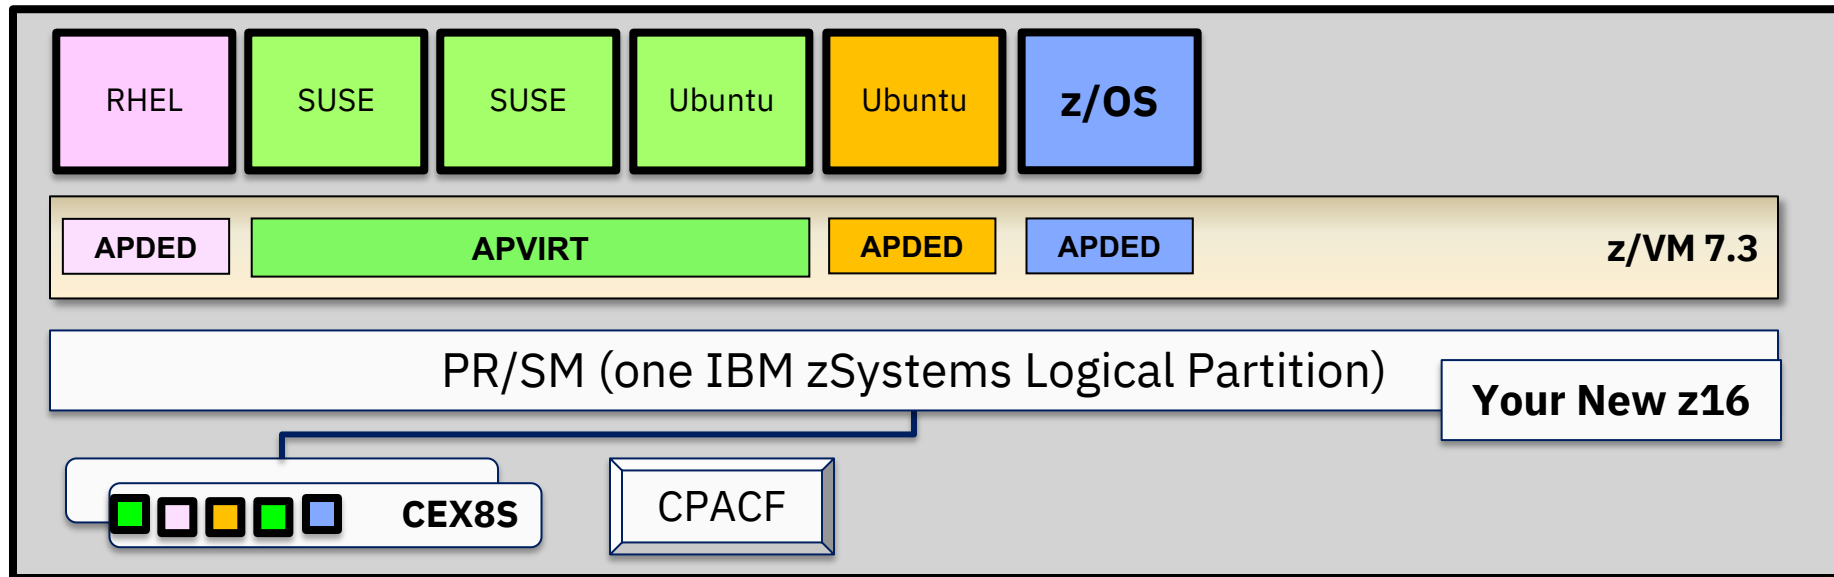


Your Crypto Lexicon (what the terms mean)



Intro to z/VM and Cryptographic Virtualization

Crypto Express adapters attached to your z/VM partition are **virtualized for the benefit of your guests**:



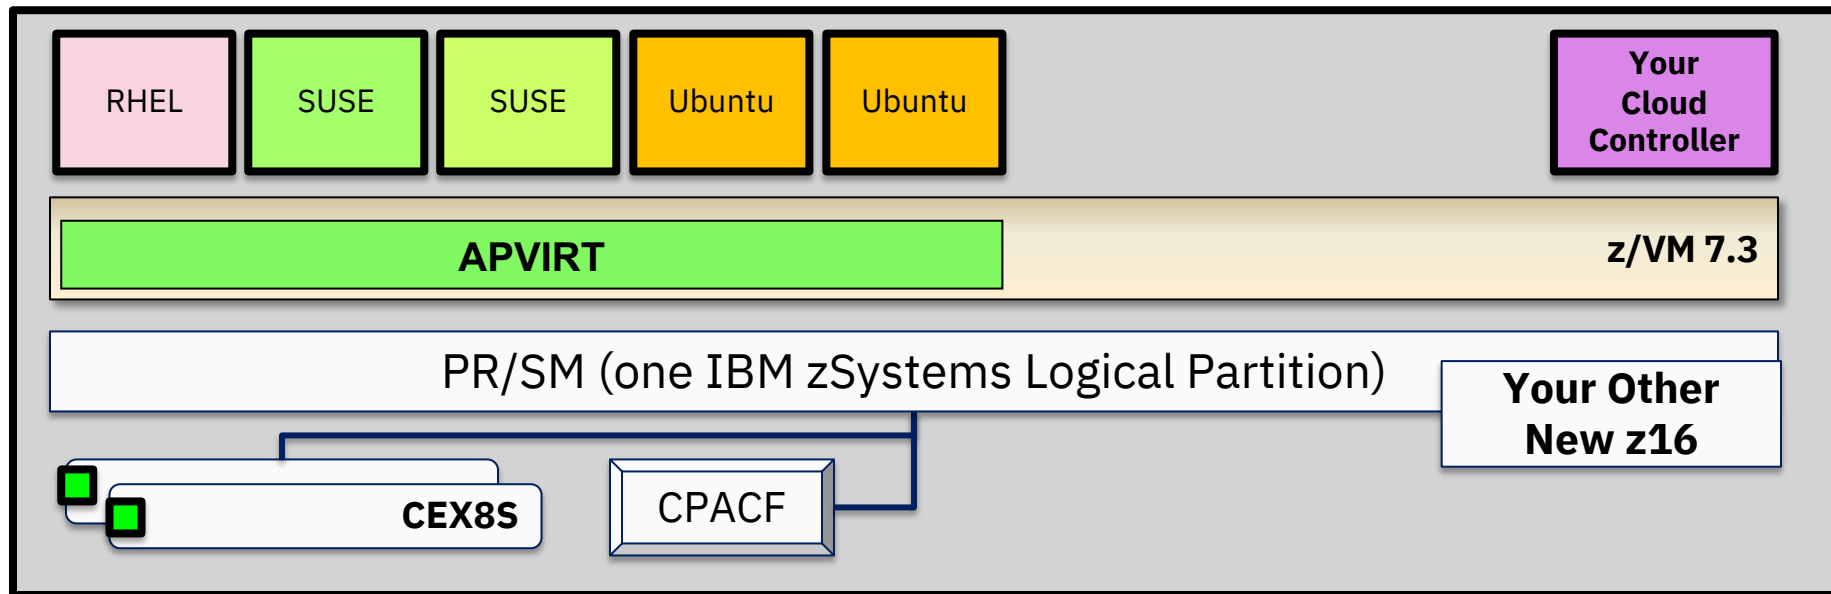
Dedicated (“APDED”)

Connects a particular Crypto Express domain (or multiple crypto resources) directly to a virtual machine – no hypervisor interference
All card functions are available to the guest

Shared (“APVIRT”)

Virtual machine can access a collection of domains controlled by the hypervisor layer
Restricted to **clear-key operations only** – sharing crypto material might break security policy.

Sample of Virtualization: LinuxONE Developer Cloud



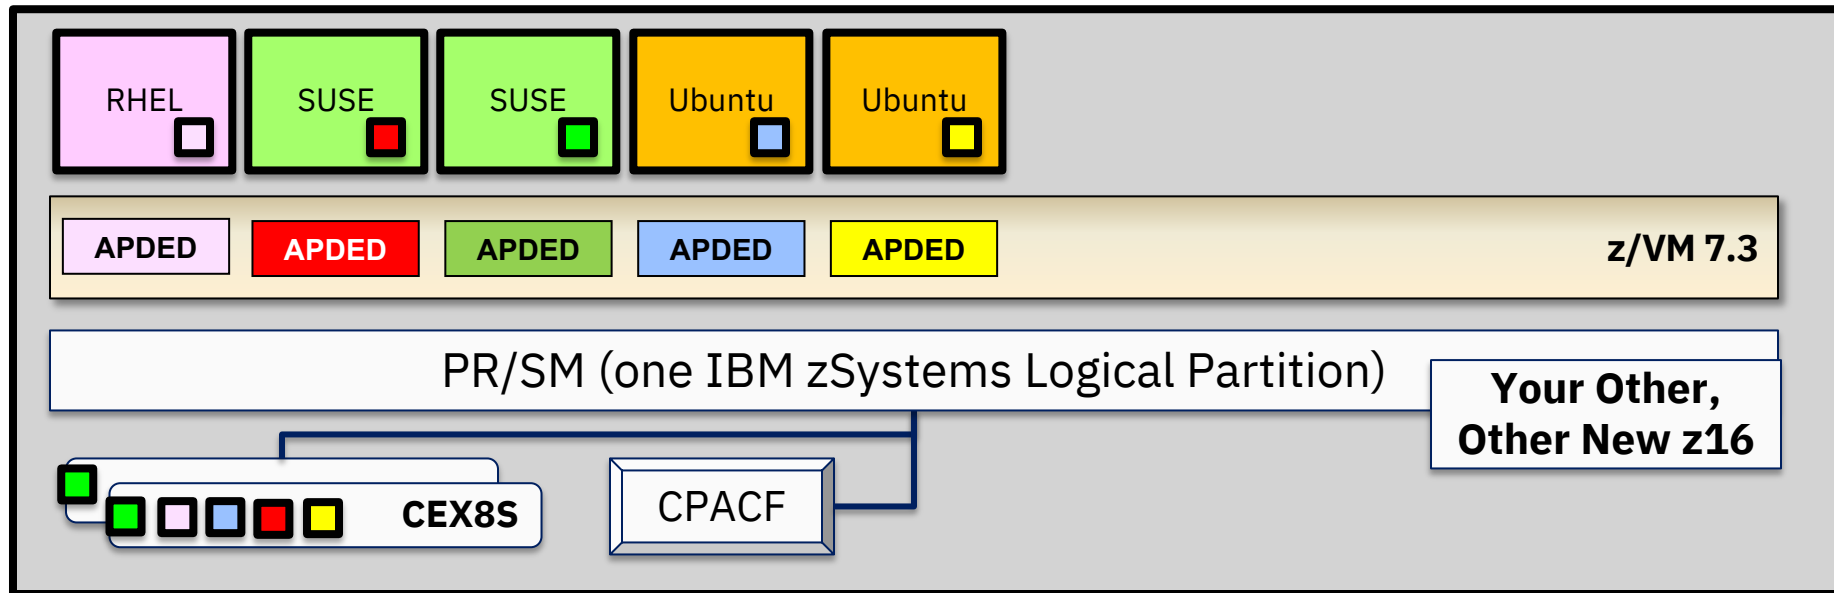
Crypto operations: SSH (RSA, SHA-2, AES), and *whatever data handled inside the guests*

Environmental Requirements: Guests must be relocatable (it's a cloud)

Recommended Hardware: CPACF and a Crypto Express Accelerator in shared configuration ("APVIRT")

- Assign 1 domain from 2-3 different adapters (for hardware failover and better performance)

Sample of Virtualization: Linux on IBM zSystems Blockchain (*not* HSBN)



Crypto operations: A lot. It's a Blockchain

Environmental Requirements: Protection of key material. (It's a Blockchain.)

Recommended Hardware: CPACF and Crypto Express adapters in EP11 (PKCS #11) mode

- One domain per guest participating in the Hyperledger fabric

z/VM Virtualization of Hardware Cryptography

The **CRYPTO User Directory statement** grants a z/VM userid access to crypto resources associated with the Crypto Express adapters:

```
                v-----+                v-----+
CRYPTO +- DOMAIN ---+domains +- APDEDicated +- ids ---+--><
      |
      +- APVIRTual-----+-----^
```

Notes:

- Guests should not try to dedicate the same domains (*first to IPL wins, all others complain*)
- Guests with a dedicated crypto resource may not be relocated
- Guests may not have both dedicated and shared crypto resources
- Shared crypto resources are limited to clear-key (Accelerator) mode only

Assigning Domains to APVIRT

The CRYPTO APVIRT statement in your System Configuration file allows you to request particular crypto resources (by Crypto ID and domain index) to be assigned to hypervisor's list of **shared crypto resources**:

```
CRYPTO APVIRT AP 1 DOMAIN 0 1  
CRYPTO APVIRT AP 0 DOMAIN 22
```

Usage Notes:

- z/VM will designate the first available domain in this list as the designated type (hardware + mode)
- Any other available domains in SYSTEM CONFIG also of that type are designated for shared usage
- Domains that do not meet criteria are ignored.
- If no domains meet criteria, no APVIRT usage will be allowed
- EP11 domains (and adapters) may not be used for shared use or assigned to APVIRT

If this statement is not present in the System Configuration file, z/VM will select two available domains, with a preference for Accelerator mode domains on the latest hardware.

Assigning Domains to APVIRT

Given the following System Configuration:

```
CRYPTO APVIRT 1 2 DOMAIN 7 8  
CRYPTO APVIRT 4 DOMAIN 9
```

... z/VM will check domains in the following order:

```
AP 1 DOMAIN 7          /* CEX6A */  
AP 1 DOMAIN 8          /* CEX6A */  
AP 2 DOMAIN 7          /* CEX5A */  
AP 2 DOMAIN 8          /* CEX5A */  
AP 4 DOMAIN 9          /* CEX6C */
```

If **AP 1 DOMAIN 7** is available at system initialization, it will be APVIRT.

- APVIRT must then use type CEX6A
- Only AP 1 DOMAIN 8, with a matching type and mode, is set as APVIRT
- If a guest lists AP 1 DOMAIN 7 as **APDED**, the guest will be denied access

Example: Static Assignment of Domains for z/VM Guests

System Configuration: CRYPTO APVIRT AP 1-2 DOMAIN 15-16

Guest A: CRYPTO DOMAIN 13-18 APDED 0-3

/* Conflicts on AP 1-2; no domains granted on AP 1 or 2. */


Guest B: CRYPTO DOMAIN 11-14 APDED 0

/* Conflict at Domain 14. No Domains granted on this AP. */

Guest C: CRYPTO DOMAIN 2 APDED 0-3

/* No conflicts. */

Reverse the logon order of **Guest A** and **Guest B** ...



z/VM Virtualization of Hardware Cryptography

```
QUERY CRYPTO DOMAINS USERS
```

<u>Crypto ID</u>	<u>device</u>	<u>Domain Index</u>	<u>device status</u>	<u>config state</u>	<u>planned crypto resource usage</u>
AP 001	CEX5A	Domain 010	operational	online	free, dedication planned
AP 001	CEX5A	Domain 011	operational	online	free, dedication planned
AP 001	CEX5A	Domain 084	operational	online	free
AP 002	CEX5C	Domain 010	resetting	online	free, dedication planned
AP 002	CEX5C	Domain 011	resetting	online	free, dedication planned
AP 002	CEX5C	Domain 084	resetting	online	shared
AP 003	CEX5C	Domain 010	operational	online	attached to BWHUGEN
AP 003	CEX5C	Domain 011	operational	online	attached to BWHUGEN
AP 003	CEX5C	Domain 084	operational	online	shared

Notes:

- **Device Status** can be operational, resetting, checkstop, deconfigured, busy, revoked, unsupported
- **Configuration State** can be online or offline. **These are logical states (how the card looks to z/VM)**
- **Device assignment** can be **free, dedication planned; attached to *userid*; free; shared**

z/VM Virtualization of Hardware Cryptography

QUERY VIRTUAL CRYPTO

(Class G) will display virtual crypto resources **for your guest.**

Keyword "virtual" required for Guests with A, B, C, or E privileges.

```
          ,--Virtual---,  
>>-Query---+-----+---CRYPTo-----<<
```

```
QUERY VIRTUAL CRYPTO
```

```
AP 03 CEX6A Domain 06 dedicated
```

```
Ready;
```

Assigning AP Domains to z/VM Guests

The Big Question: Which type of adapter do I need, and what domains do I want to assign to my guest?

It depends:

- Do you need secure key operations? (APDED)
- Does your security policy require physical isolation? (APDED)
- Do your guests need to exploit EP11 mode? (APDED **only**)
- Do you need to relocate your guest? (APVIRT*)
- Can you share your domains without impact to security or performance? (APVIRT)
- Are you running out of domains attached to the LPAR?
- Are your guests similar, cloned, or tied to HA solutions?
- Does your guest operating system have particular restrictions?

Different guests will have different needs, based upon their drivers and configuration requirements.

And, until recently, this meant a lot of planning, because changing your config at the LPAR level, or changing shared crypto resource assignments, meant a re-IPL of your z/VM system...

**Note: some restrictions apply. Consult the [CP Planning and Administration Guide](#) or [Getting Started With Linux](#) manuals.*

Dynamic Crypto Support for z/VM

http://www.vm.ibm.com/newfunction/#dynamic_crypto

z/VM 7.1
PTF for APAR VM66266

Dynamic Crypto support enables **changes to the z/VM crypto environment** without requiring an IPL of z/VM or its **guests** (e.g. Linux on Z).

This allows:

- Less disruptive addition or removal of Crypto Express hardware to/from a z/VM system and its guests
- Less disruptive maintenance and repair of Crypto Express hardware attached and in-use by a z/VM system
- Reassignment and allocation of crypto resources without requiring a system IPL or user logoff/logon
- Greater flexibility to change crypto resources between shared and dedicated use.

Additionally, there are RAS benefits for shared-use crypto resources:

- Better detection of Crypto Express adapter errors with "silent" retrying of shared pool requests to alternative resources
- Ability to recover failed Crypto Express adapters
- Improved internal diagnostics for IBM service
- Improved logoff and live guest relocation latency for users of shared crypto.

z/VM Dynamic Crypto – Commands

VARY ONLINE CRYPTO (B)

- Bring a Crypto Express adapter online

VARY OFFLINE CRYPTO (B)

- Take a Crypto Express adapter offline (device associations remain in place)

ATTACH CRYPTO (B)

- Add crypto resource(s) to your z/VM guest (or APVIRT)

DETACH CRYPTO (B or G)

- Remove dedicated crypto resources from a guest
- Remove crypto resources from the shared crypto pool
- Remove guest access to the shared crypto pool

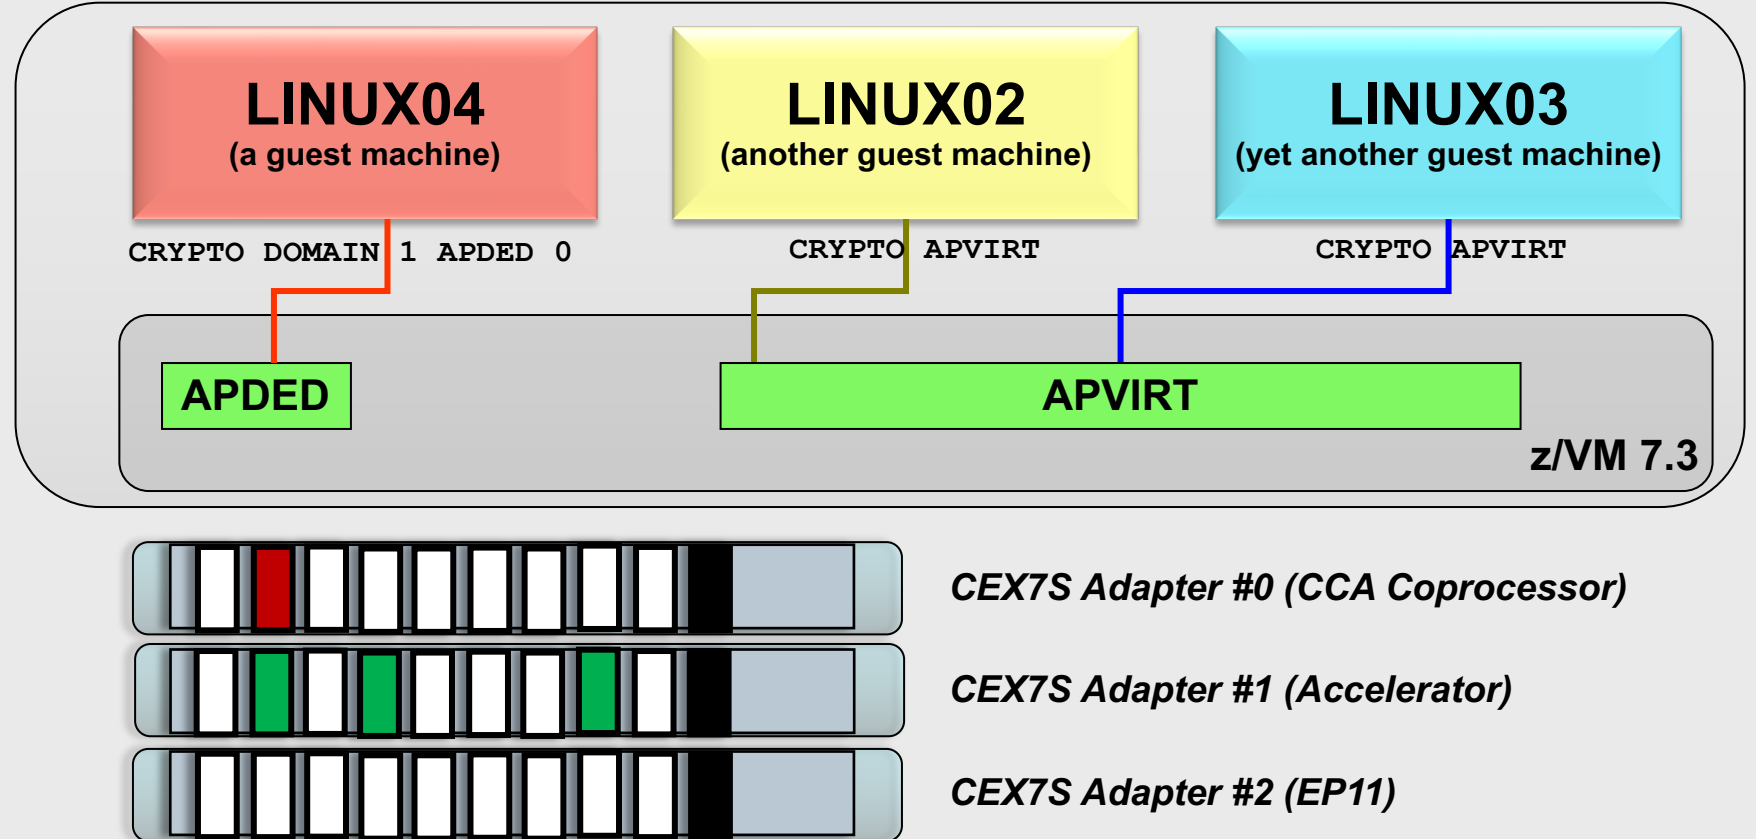
– **DEFINE CRYPTO** APVirtual (G)

- assign or reassign shared crypto resource access to a z/VM guest

– **QUERY CRYPTO DOMAINS** (per previous slide)

How To: Make a new adapter available to z/VM

VARY ON CRYPTO 2

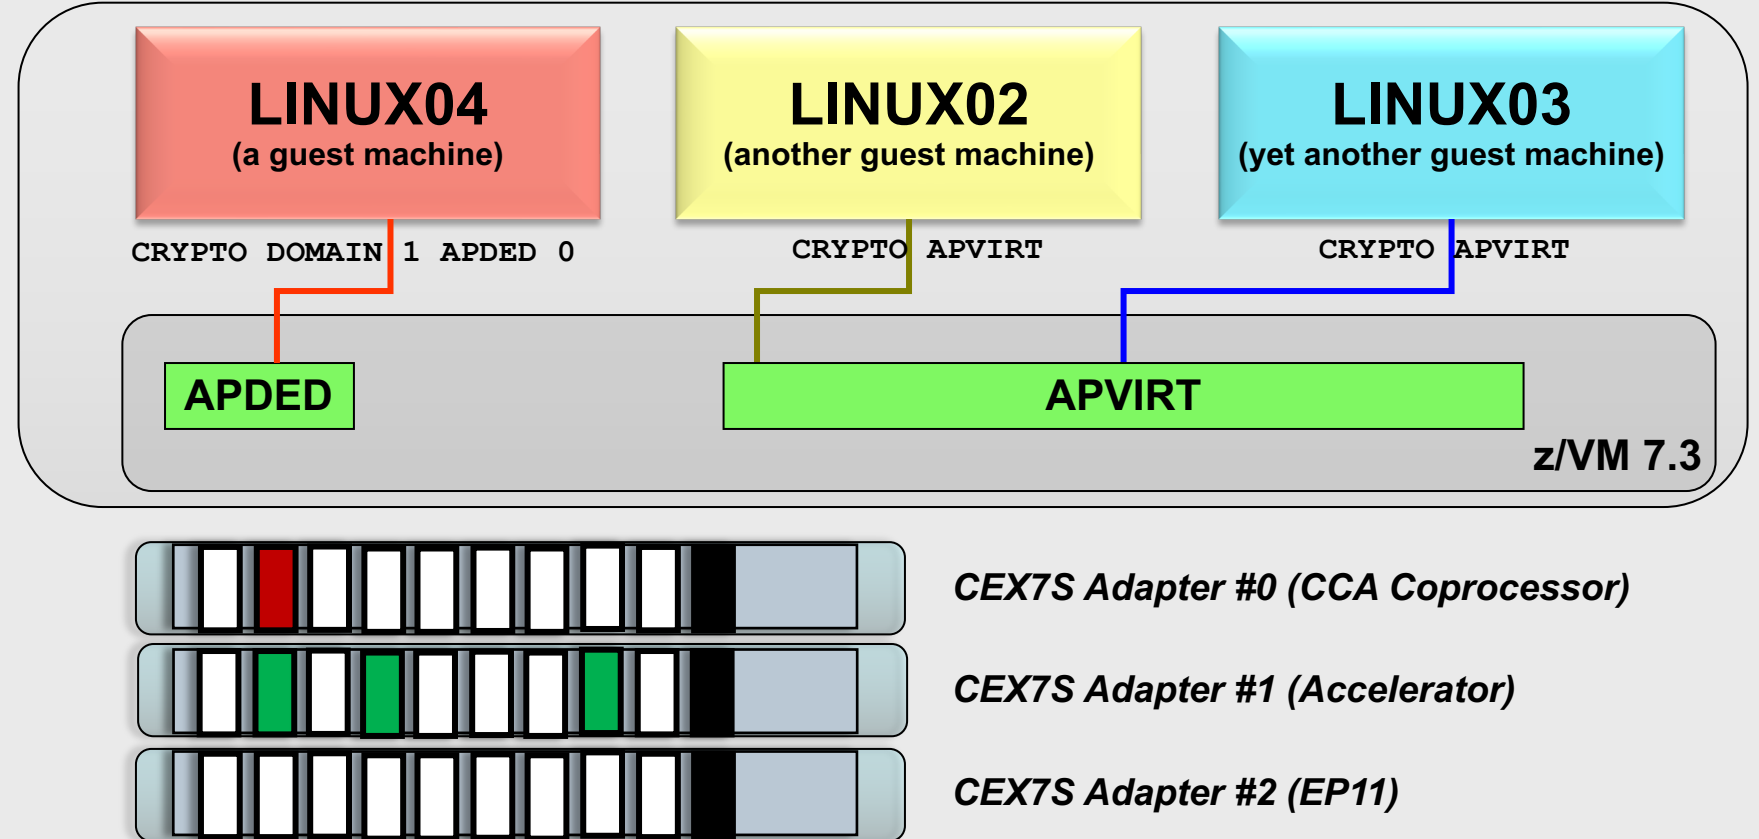


How To: Assign a crypto resource to a user

ATTACH CRYPTO AP 2 to LINUX04

Warning: does not change your z/VM User Directory... so static configuration does not update automatically.

Don't forget to update your defaults!


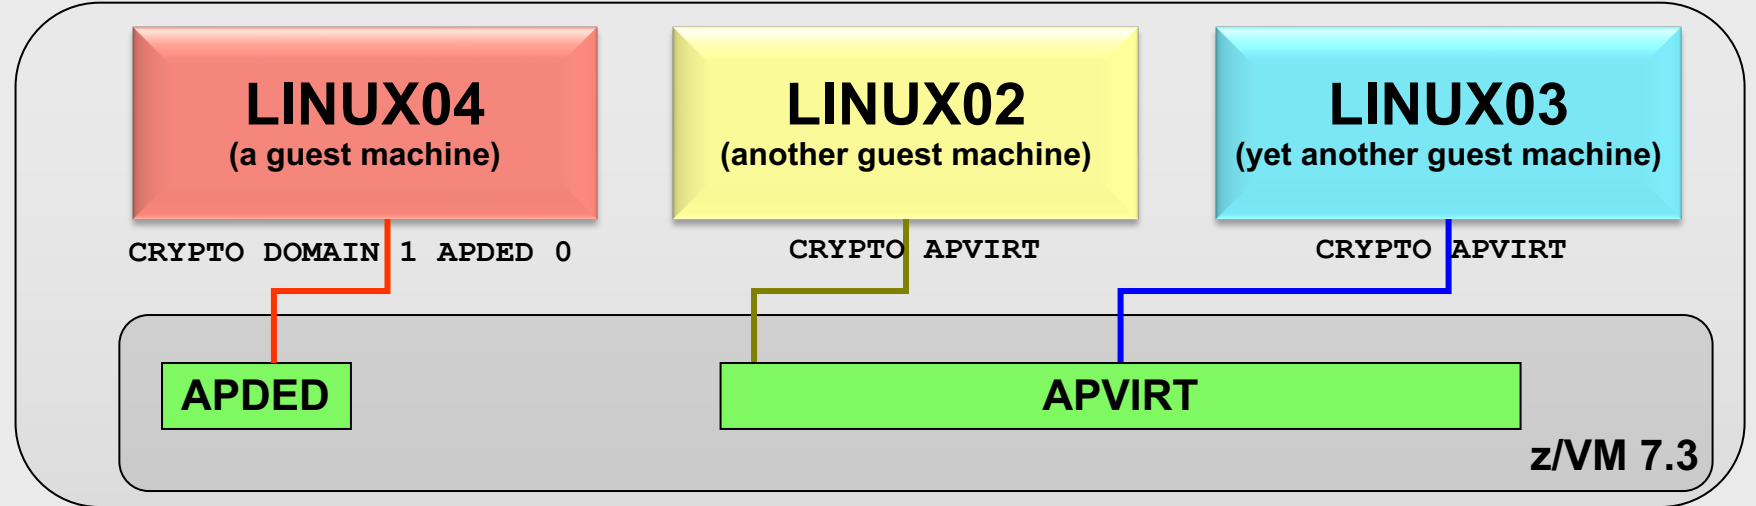


How To: Assign new crypto resources for sharing

z/VM 7.1
PTF for APAR VM66266

VARY ON CRYPTO 3

ATTACH CRYPTO AP 0 3 DOMAIN 6 7 to SYSTEM



CEX7S Adapter #0 (CCA Coprocessor)

CEX7S Adapter #1 (Accelerator)

CEX7S Adapter #2 (EP11)

CEX7S Adapter #3 (CCA Coprocessor)


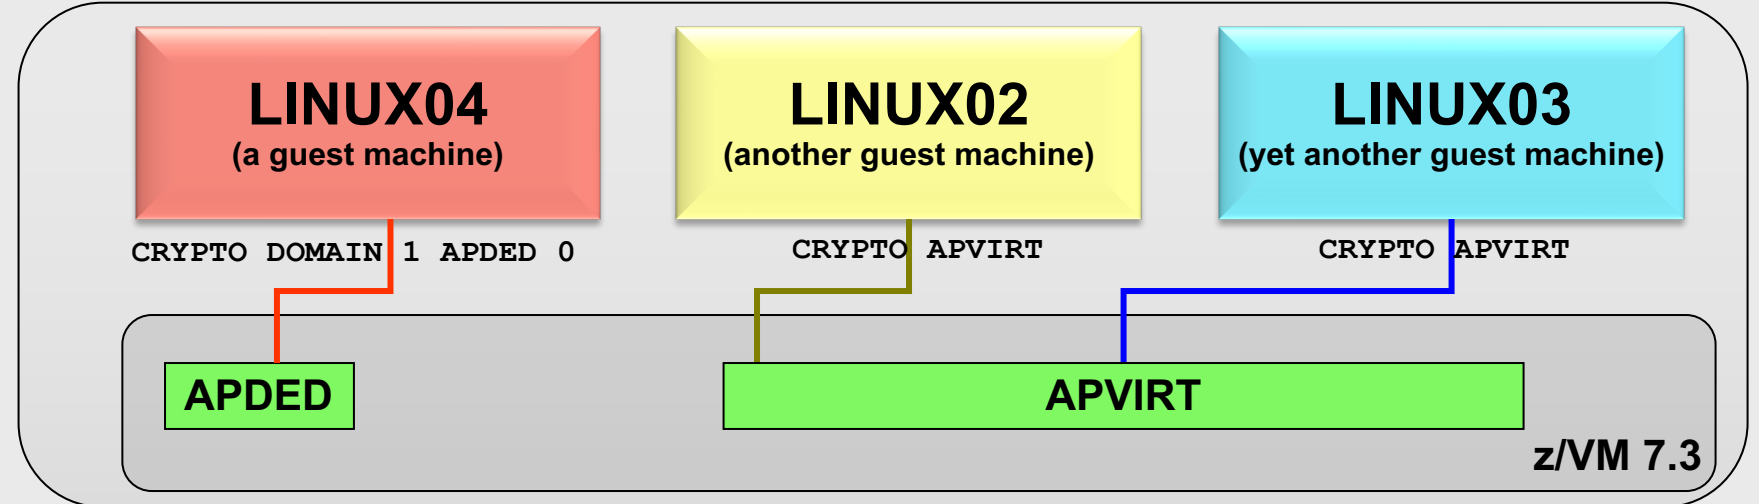
How To: remove crypto resources from shared pool

z/VM 7.1
PTF for APAR VM66266

DETACH CRYPTO AP 1 DOMAIN 1 3 7 from SYSTEM FORCE

Change does not remove APVIRT access from the guests.

Note: this is an extreme example, you may not want to remove these all at once.




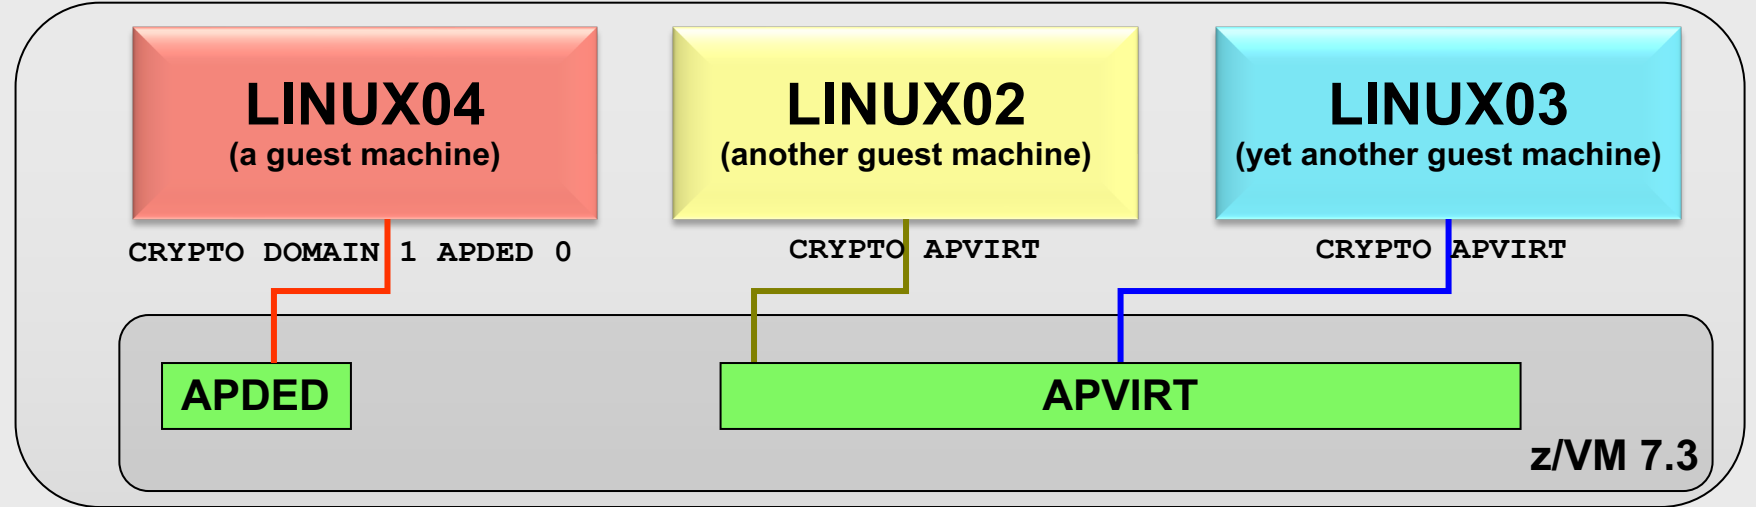
How To: Take an adapter offline

VARY OFF CRYPTO 1

Adapter will be listed as offline, and will not be available for use.

VARY ON the adapter to bring it back to active configuration...

...no IPL required.



CEX7S Adapter #0 (CCA Coprocessor)

CEX7S Adapter #1 (Accelerator)

CEX7S Adapter #2 (EP11)

CEX7S Adapter #3 (CCA Coprocessor)

z/VM Dynamic Crypto – Usage Notes

Attachments persist even when a device is taken offline

Resource assignment (dedicated/shared) does not change when an adapter is varied on/off

FORCE option:

- Not required when DETACHing from a user
- Required when VARYing OFF an adapter with crypto resources in use
- Required when detaching last resource from shared pool
- Exercise caution when using

The Importance of Cryptographic Hygiene

Dynamic Crypto gives you a lot of power to modify the environment

- This is a good thing and a bad thing
- **“With great power comes great responsibility.”**

z/VM does not zeroize domains before reassigning to a guest (or to APVIRT)

- We don't want to make that assumption (traditionally, this is HMC territory)
- **This might lead to “residual crypto” (Ewww)**

Basic guidelines:

- Zeroize (at HMC) when changing adapter modes or changing security zones
- Changes between unused and APVIRT: safe (no key material involved)
- Changes involving clear-key APDED: consider zeroizing
- Changes involving secure-key APDED: definitely zeroize

See z/VM CP Planning & Customization, Chapter 5. Crypto Planning and Management

<https://www.vm.ibm.com/library/720pdfs/72627110.pdf>

z/VM Dynamic Crypto – Summary

z/VM 7.1
PTF for APAR VM66266

Now available via PTF for APAR VM66266 for z/VM 7.1, or in later releases

- *Prereq VM66206 for z/VM 6.4 and z/VM 7.1 (installed on all SSI members before dynamic crypto is applied.)*

Dynamic Crypto support enables changes to the z/VM crypto environment without requiring an IPL of z/VM or its guests (e.g. Linux on Z).

Sponsor users were engaged heavily in the process

- Design playbacks and to-be scenarios
- Usability iterations
- Demos and hands-on-code early testing

Mixed-APVIRT LGR allows **flexible crypto configurations** so guests using APVIRT **can relocate with fewer hardware restrictions.**

Removes restrictions on guest relocation in a z/VM Single System Image:

- *Then:* needed common type and mode (e.g., CEX7A) on source and target system
 - including firmware levels
- *Now:* guests in a relocation domain see lowest type of a common mode
 - E.g., a combination of CEX7A and CEX5A is seen as a CEX5A by all guests in that domain
 - Guests without a need to relocate, or in specialized domains, can see higher levels
 - Still requires common adapter “mode” (Accelerator or Coprocessor; EP11 cannot be relocated)
- **New Function APAR for z/VM V7.2 only**

Host Exploitation of Crypto Interruptions

With the PTF for APAR VM66534, z/VM V7.2 supports host crypto-interruption exploitation for APVIRT cryptographic resources in the shared pool. The host is not required to poll cryptographic resources for replies that are ready to be delivered to the guest.

- Some performance benefit may be derived from enabling this capability
- Enabled by setting APVIRT POLLING to OFF
 - Not enabled by default via z/VM V7.2 PTF (default state is “polling is on”)

Commands impacted:

- **SET CRYPTO APVIRT POLLING** – change setting for entire APVIRT pool
- **QUERY CRYPTO POLLING** – query POLLING state [ON/OFF]

```
QUERY CRYPTO POLLING
```

```
Shared-crypto polling is OFF  
Ready;
```

z/VM Support for IBM z16

z/VM 7.1, 7.2
PTF for APAR VM66532

With the PTF for APAR VM66532, z/VM® 7.1 and 7.2 provide support to enable guests to exploit function on IBM z16®. The following support is included:

- Breaking-event-address register (BEAR) enhancement facility, which facilitates the debug of wild branches.
- Reset DAT protection facility, which provides a more efficient way to disable DAT protection, such as during copy-on-write or page-change tracking operations.
- RoCE Express3 adapter, which allows guests to exploit Routable RoCE, Zero Touch RoCE, and SMC-R V2 support.
- The Crypto Express8S (CEX8S) adapter, supported as a dedicated or shared resource. Dedicated guests are able to take advantage of all functions available with the CEX8S adapters, including assorted new enhancements and use of Quantum-Safe APIs.

All crypto adapters that are configured in EP11 mode are reported with the 'P' suffix instead of the 'S' suffix (e.g., CEX8P).

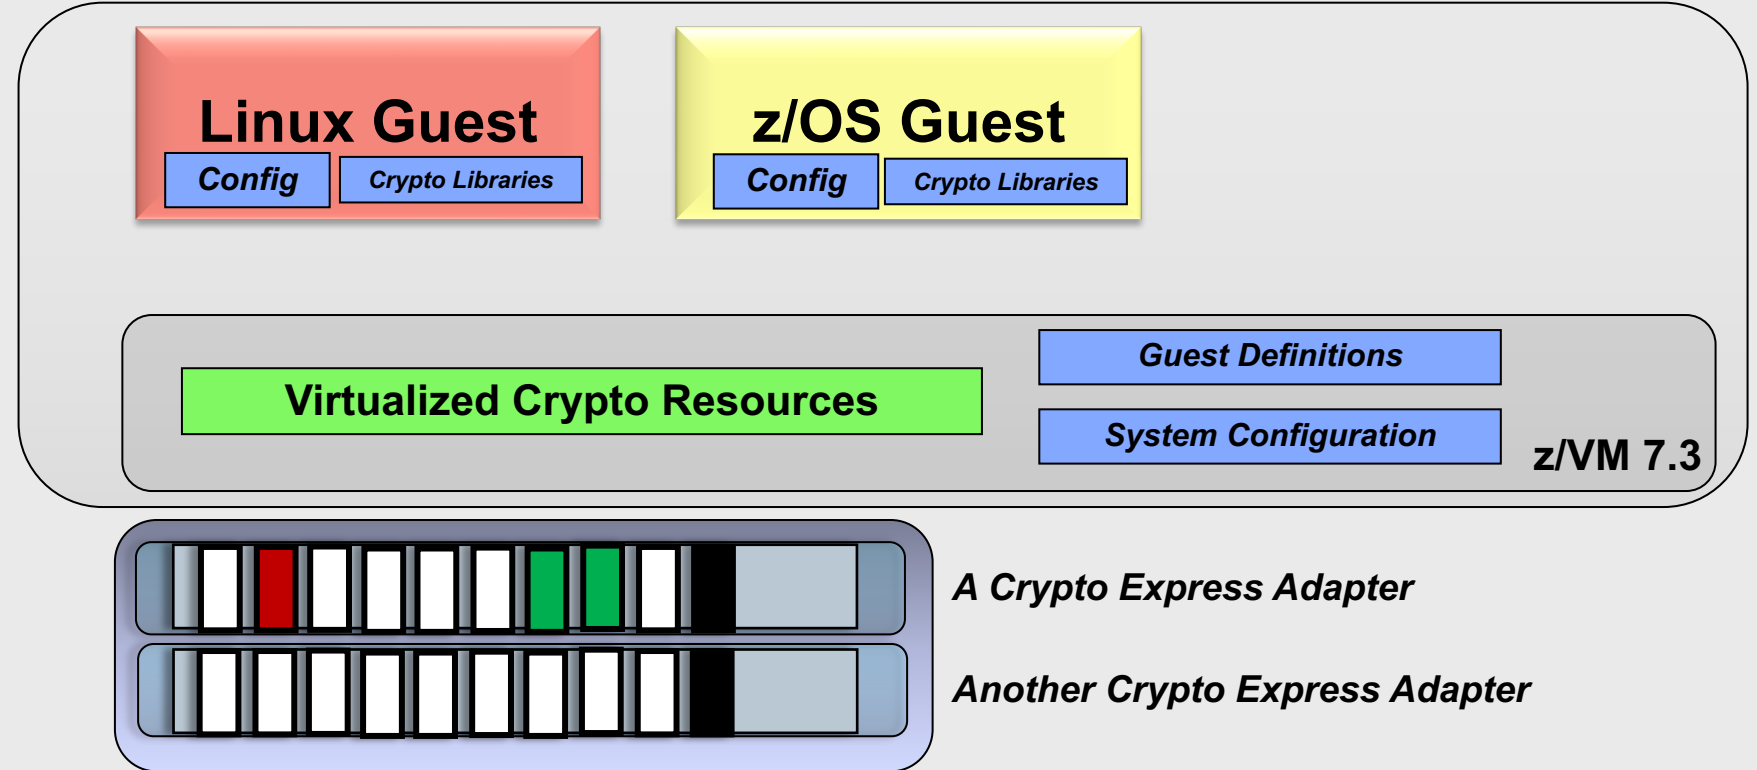
Guest Use of Hardware Crypto

How To: Configure your Crypto on IBM zSystems and LinuxONE

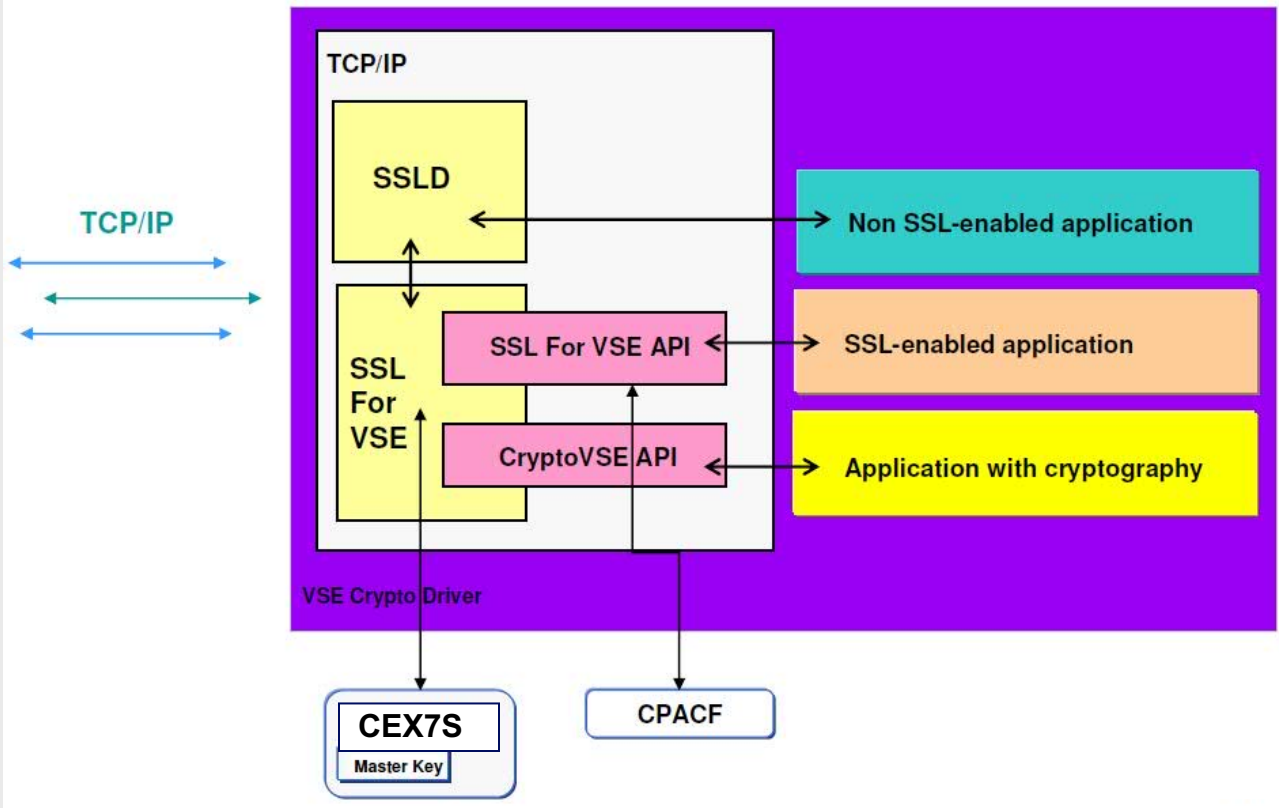
Crypto libraries will vary from OS to OS

Some may require specific configuration to make use of certain features

Consult pertinent local documentation



z/VSE Cryptographic Infrastructure



z/VSE automatically detects any Crypto Express features dedicated to (or shared with) the virtual machine in which it's running

CMS Guests Running on z/VM

CMS guests can utilize CPACF if enabled

- Need to issue appropriate machine instructions
- Some functions (Pipelines, TLS/SSL Server) use these automatically

The CMS environment does **not** have Crypto Express libraries

- Different instructions / communication paths than CPACF
- Nothing available yet for general system programmer use
- **Exception:** TLS/SSL Server for data-in-flight encryption to/from/within the hypervisor (APVIRT only)

Crypto APVIRT for the z/VM TLS/SSL Server

```
PROFILE TCPSSLU
  CRYPTO APVIRTUAL
  IPL ZCMS PARM FILEPOOL VMSYS
  IUCV ALLOW
  LOGONBY BWHUGEN
  NAMESAVE TCPIP
  OPTION ACCT MAXCONN 1024 QUICKDSP
  POSIXINFO UID 7 GNAME security
  SHARE RELATIVE 3000
  CONSOLE 0009 3215 T
  [...]
```


Add **CRYPTO APVIRT** to your SSL server's PROFILE entry

- **TCPSSLU** - the default PROFILE entry for the TLS/SSL Server
- APDED not allowed for a POOL of userids

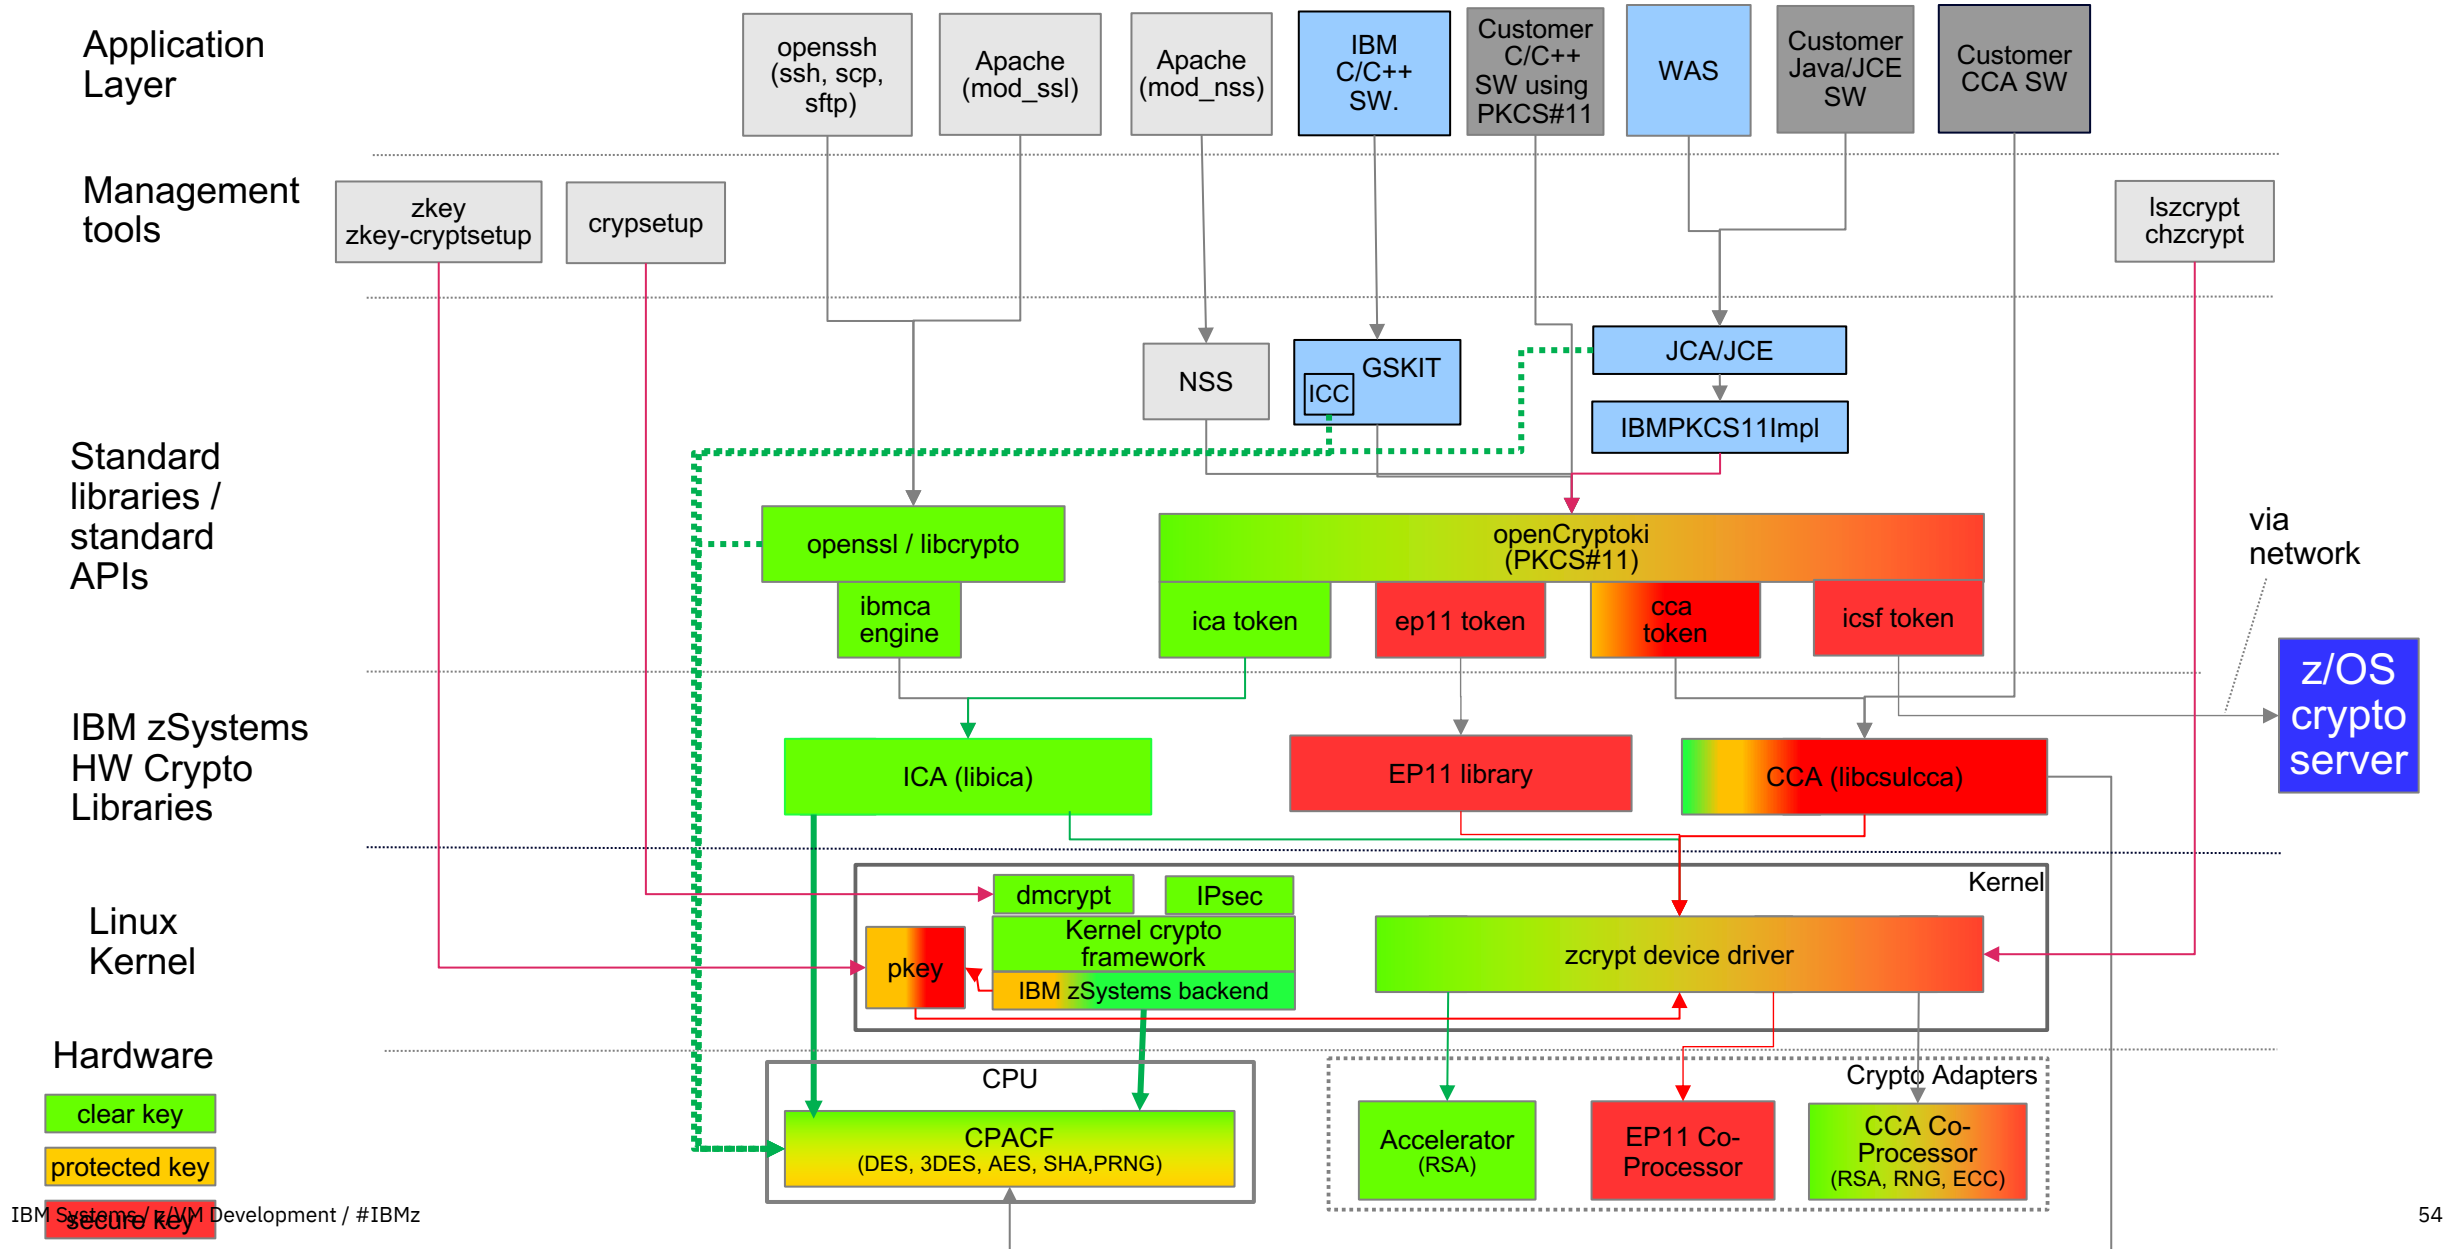
Insert directly into VM definition for:

- **LDAPSRV** - uses its own System SSL calls
- **GSKADMIN** - for certificate creation / management

z/OS Cryptographic Infrastructure



Linux on IBM zSystems Cryptographic Infrastructure



Linux Kernel and Cryptography

The Linux kernel provides a set of cryptographic functions

- Generic, platform-independent implementations of cryptographic algorithms
- Support for platform-optimized algorithms that are automatically used if available

The Linux on IBM zSystems kernel includes support for

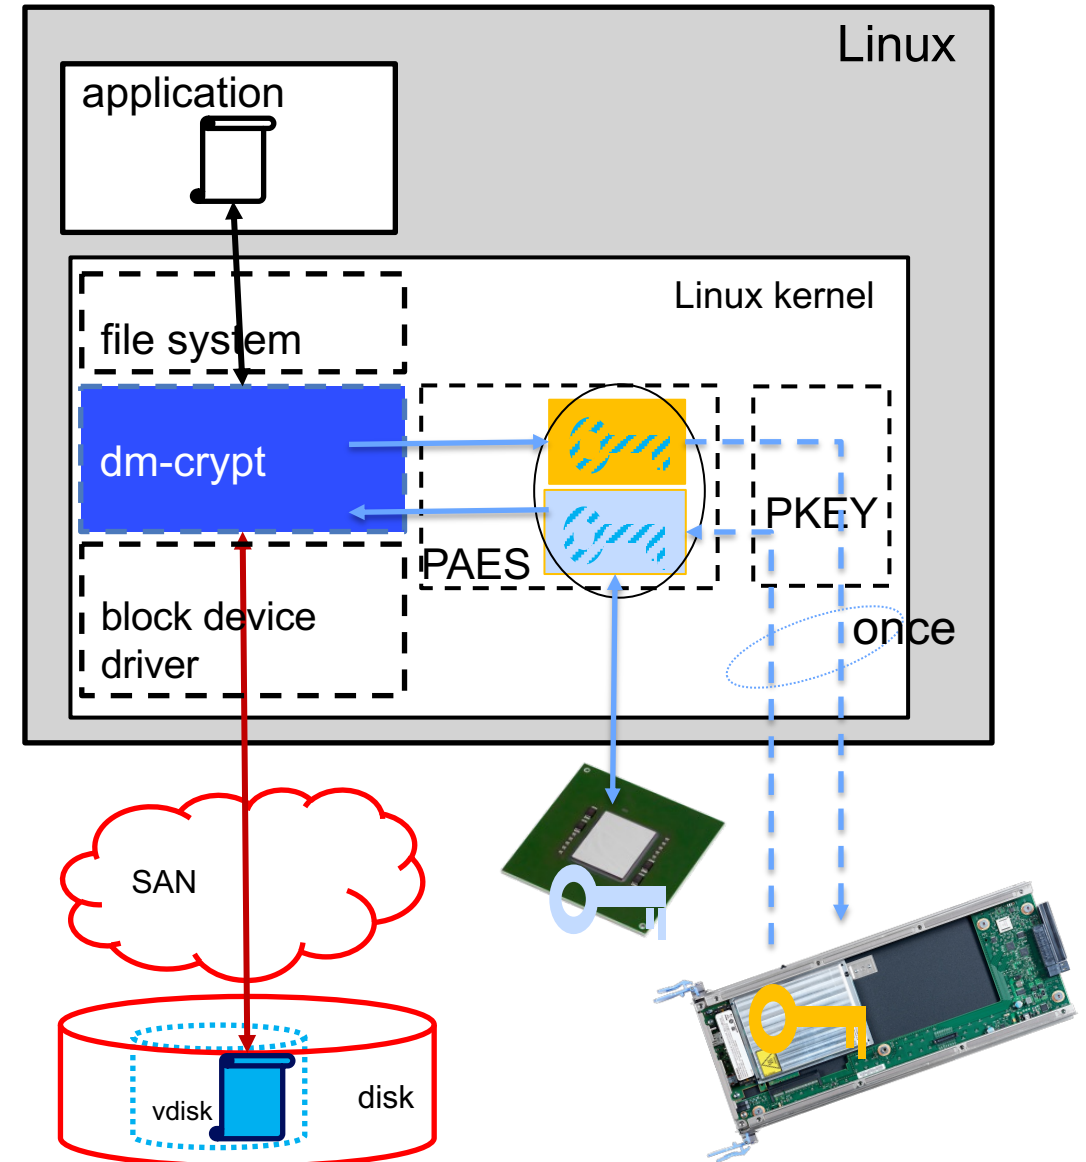
- Exploiting CPACF to optimize and accelerate symmetric cryptographic functions
- Managing Crypto Express cards with the ***zcrypt*** device driver

Which applications can benefit from accelerated in-kernel cryptographic functions?

- IPsec and ssh (from the beginning of the presentation, remember?)
- Linux device-mappers – for example, **dm-crypt** or **eCryptFS**

File System Encryption with dm-crypt for Linux on IBM zSystems

- dm-crypt
 - a mechanism for end-to-end data encryption
 - data only appears in the clear in application
- Linux kernel component that transparently
 - for all applications
 - for a whole block device (partition or LV)
 - encrypts all data written to disk
 - decrypts all data read from disk
- E2E data encryption
 - The complete I/O path outside the kernel is encrypted:
 - HV, adapters, links, switches, disks
- How it works:
 - uses in-kernel crypto
 - can use IBM zSystems CPACF Protected Key Crypto:
 - XTS-PAES as from the `paes_s390` module
 - encrypted volumes must be opened before usage
 - opening provides encryption key to kernel
 - establishes virtual volume in `/dev/mapper`



Validating Linux and z/VM Configuration

```
cert1xb:~ # cat /proc/driver/z90crypt
zcrypt version: 2.1.1
Cryptographic domain: 6
Total device count: 1
PCICA count: 0
PCICC count: 0
PCIXCC MCL2 count: 0
PCIXCC MCL3 count: 0
CEX4C count: 0
CEX4A count: 1
requestq count: 0
pendingq count: 0
Total open handles: 0
```

Validating Linux and z/VM Configuration

```
Last login: Thu Mar 28 10:18:05 2013 from nn.nn.nn.nnn
certlxb:~ # cat /proc/crypto
name          : stdrng
driver        : krng
module        : kernel
priority      : 200
refcnt        : 1
selftest      : passed
type          : rng
seedsize      : 0

name          : sha1
driver        : sha1-generic
module        : kernel
priority      : 0
refcnt        : 1
selftest      : passed
type          : shash
blocksize     : 64
digestsize    : 20
```

Validating Linux and z/VM Configuration

```
certlxb:~ # icainfo
```

```
The following CP Assist for Cryptographic Function  
(CPACF) operations are supported by libica on this  
system:
```

```
SHA-1:      yes  
SHA-256:    yes  
SHA-512:    yes  
DES:        yes  
TDES-128:   yes  
TDES-192:   yes  
AES-128:    yes  
AES-192:    yes  
AES-256:    yes  
PRNG:       yes
```

Validating Linux and z/VM Configuration

icastats – data from the libica crypto library

- SLES 12 and RHEL 7.1 onward

cpacfstats – data about CPACF on-chip usage

- On s390tools

- Works for Linux running in an LPAR directly

- CPUMF data (authorization required)

lszcrypt – statistics on Crypto Express requests

Validating Linux and z/VM Configuration

```
certlxb:~ # sudo vmcp QUERY VIRTUAL CRYPTO  
AP 001 CEX8A Domain 001 shared online
```

Remember that **QUERY VIRTUAL CRYPTO** is a Class G command

This indicates the virtual AP number and virtual Domain number provided to the guest and the type of crypto feature being shared.

Documentation & References

Linux on Z and LinuxONE Device Drivers Features and Commands:

<https://www.ibm.com/docs/en/linux-on-systems?topic=overview-device-drivers-features-commands>

in-kernel crypto and zcrypt DD: “Security” part,

crypto configuration for KVM guest: „Setting up a KVM host for VFIO passthrough“ chapter

- see also: <https://www.ibm.com/docs/en/linux-on-systems?topic=kvm-configuring-crypto-express-adapters-guests>

The screenshot shows the IBM Documentation website for 'Cryptographic hardware support'. The left sidebar contains a table of contents with the following items:

- Show full table of contents
- Security
 - Security concepts
 - Introducing IBM Secure Execution for Linux
 - How to set an AES master key
 - Pervasive encryption
 - Cryptographic hardware support**
 - Cryptographic device drivers
 - openCryptoki: Open Source PKCS #11
 - libica Programmer's Reference
 - Secure Key Solution with the Common Cryptographic Architecture Application Programmer's Guide
 - Exploiting Enterprise PKCS #11 using openCryptoki
 - Related materials about security

The main content area features the following links and descriptions:

- Cryptographic device drivers**: The security part of *Device Drivers, Features, and Commands* describes the generic cryptographic device driver and the pseudo-random number device driver.
- openCryptoki - An Open Source Implementation of PKCS #11**: Find information on how to configure openCryptoki and how to write PKCS #11 programs that exploit the provided cryptographic services.
- libica Programmer's Reference**: The Library for IBM® Cryptographic Architecture (libica) provides cryptographic functions for cryptographic applications on IBM Z® or IBM® LinuxONE, both with and without cryptographic hardware.
- Secure Key Solution with the Common Cryptographic Architecture Application Programmer's Guide**: The Common Cryptographic Architecture (CCA) provides a variety of services for cryptography and data-security.
- Exploiting Enterprise PKCS #11 using openCryptoki**: Linux on IBM Z or IBM LinuxONE applications that use a PKCS #11 API can take advantage of the Enterprise PKCS #11 (EP11) coprocessor mode of IBM® cryptographic adapters.

Parent topic: [Security](#)

setting up an CCAES master key

setting up dm-crypt with PAES

new: openCryptoki book

ica, cca and ep11 books

Summary

Summary

IBM zSystems **hardware** accelerates the hard math of cryptographic operations

- Saves **time**, saves CPU processing **power**, saves MIPS **cost**
- Secure Key operations are FIPS 140-2 Level 4 certified

z/VM **virtualizes** IBM zSystems hardware cryptography

- Architectural fidelity in all things zSystems
- A "shared" flavor as well as dedicated use of crypto resources

Guests understand they can utilize IBM zSystems cryptography

- May require configuration of the guest to exploit
- Different guests provide different options

Don't let cryptography (or its acronyms) scare you away

- Security is meant to enhance business, not impede it
- Cryptography protects your data, whether at rest or in flight



Resources



Redbook: Getting Started with Linux on IBM Z Encryption for Data At-Rest Redbook ***new***
<http://www.redbooks.ibm.com/abstracts/sg248436.html?Open>

Redbook: Security and Linux on IBM Z
<http://www.redbooks.ibm.com/abstracts/redp5464.html?Open>

New: IBM Z pervasive encryption landing page
<https://izswebpage.mybluemix.net/>

IBM Z pervasive encryption solution guide (Knowledge Center)
https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.3.0/com.ibm.zos.v2r3.izs/izs.htm

IBM Z pervasive encryption FAQ:
<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03116USEN>

IBM Crypto Education page: <https://ibm.biz/BdiAah>

zPET Test Reports:
<https://www.ibm.com/developerworks/community/groups/service/html/communitystart?communityUuid=43ea8e78-acbe-49f5-9290-379e4f4569cb>

MOP demo white paper:
<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102734>


Youtube Videos:

- Data Set Encryption: <https://www.youtube.com/watch?v=zdSXRUSmkb4>
- CF Encryption: <https://www.youtube.com/watch?v=lTmsFWuJwJU>
- zERT: https://www.youtube.com/watch?v=1CgEcCTX_o8
- MOP MPL Bank: <https://www.youtube.com/watch?v=EP488nLdGts>



THANK YOU

Contact Information:

Brian W. Hugenbruch
IBM zSystems Security for Virtualization
bwhugen at us dot ibm dot com
 [@Bwhugen](https://twitter.com/Bwhugen)




Notices and disclaimers

- © 2022 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights – use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **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.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines 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 the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.
- .



Frequently Asked Questions

Question: How do I determine how many instructions are being offloaded to CPACF or the Crypto Express features?



Answer: Depends upon your authority over the system.

If you're operating at the hypervisor administrator level, you can use CP Monitor Records to determine the number of instructions executed. Use your application of choice to examine them.

- **MRPRCAPC** – Crypto Performance Counters (Domain 5, Record 9)
- **MRPRCAPM** – Crypto Performance Measurement Data (Domain 5, Record 10)

Linux commands such as *lszcrypt* can be used to determine basic per-guest utilization, numbers of requests processed, etc.

Frequently Asked Questions

Can Linux on Z crypto tie into my z/OS crypto?

Answer: Yes it can! ICSF Token Support has enabled a Linux client to tie into the z/OS Crypto-as-a-Service mechanisms (such as EKMF or ACSP)

Available with z/OS 2.1 and RHEL 7.0

Crypto requests are forwarded to ICSF on z/OS

- Using LDAP protocol
- Simple and SASL authentication

Key objects are stored under z/OS

Requires LDAP client set-up on Linux

`pkcsicsf` utility for configuration

token directory `/var/lib/opencryptoki/icsf`

token configuration file to be referred to in `opencryptoki.conf`

