


# Keys to the Virtual Kingdom

*Or, "An Illustrated Primer on Using z/VM to Virtualize your z Systems Cryptographic Features for the Benefit of Your Guest Operating Systems"*

Brian W. Hugenbruch, CISSP  
IBM z Systems Virtualization Security  
[bwhugen@us.ibm.com](mailto:bwhugen@us.ibm.com)  @Bwhugen



P4

# Trademarks

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

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

\* 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.

## Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

- Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at [www.ibm.com/systems/support/machine\\_warranties/machine\\_code/aut.html](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 SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

# Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

# The New z13 – “Ultimate Security”



## Did you know?

- z Systems is the only commercial operating system that has achieved **EAL 5+ certification**. This **certification** means that although different workloads are running on the same hardware, they are **protected** when running in separate partitions; one logical partition (**LPAR**) cannot reach across boundaries into the next LPAR and **compromise** its security. The LPARs are allocated their own resources and are secure and separate environments.
- Integrated **cryptographic** features provide leading cryptographic performance and functions. Reliability, availability, and serviceability (**RAS**) support for the **Crypto Express5S** is unmatched in the industry, and the cryptographic solution for the **Crypto Express4S** received the highest standardized security certification (**FIPS 140-2 Level 4**). IBM is in the process of gaining FIPS 140-2 Level 4 certification for the **Crypto Express5S** feature. With FIPS 140-2 Level 4 certified cryptographic hardware, IBM provides the most secure **tamper-sensing** and **tamper-resistant security module** that is available in the market.

From "Ultimate Security with the IBM z13"  
IBM Redbooks Solution Guide

## ... and here's your cryptography "Bingo" card.

AES	Advanced Encryption Standard	MAC	Message Authentication Code
ARL	Authority Revocation List	MDC	Message Detection Code
CA	Certification Authority	MD5	Message Digest 5
CBC	Cipher Block Chaining	OAEP	Optimal Asymmetric Encryption Padding
CCA	IBM Common Cryptographic Architecture	OCSF	OS/390 Open Cryptographic Services Facility
CCF	Cryptographic Coprocessor Facility	OCSP	Online Certificate Status Protocol
CDSA	Common Data Security Architecture	PCICA	PCI Cryptographic Accelerator
CEX2/3A	Crypto Express 2/3 Accelerator Mode	PCICC	PCI Cryptographic Coprocessor
CEX2/3C	Crypto Express 2/3 Coprocessor Mode	PCIXCC	PCIX Cryptographic Coprocessor
CFB	Cipher Feedback	PKA	Public Key Architecture
CKDS	Cryptographic Key Data Set	PKCS	Cryptographic Standards
CRL	Certificate Revocation List	PKDS	Public Key Data Set
CRT	Chinese Remainder Theorem	PKI	Public Key Infrastructure
CVC	Card Verification Code	RA	Registration Authority
CVV	Card Verification Value	RACF	Resource Access Control Facility
DES	Data Encryption Standard	RSA	Rivest-Shamir-Adleman
DSA	Digital Signature Algorithm	SET	Secure Electronic Transaction
DSS	Digital Signature Something	SHA	Secure Hash Algorithm
ECB	Electronic Code Book	SLE	Secure Cookie Monster Encryption
FIPS	Federal Information Processing Standard	SSL	Secure Sockets Layer
GSS	Generalized Security Services	TKE	Trusted Key Entry
ICSF	Integrated Cryptographic Service Facility	TLS	Transport Layer Security
IETF	Internet Engineering Task Force	VPN	Virtual Private Network
IPKI	Is Anyone Reading This Line		
KGUP	Key Generation Utility Program		
LDAP	Lightweight Directory Access Protocol		

# Agenda

- A Very Quick Intro to Cryptography (and why it matters)
- IBM z Systems Hardware Cryptography (and why it matters)
- z/VM Virtualization of z Systems Cryptography (and how to use it)
- Guest Support: Operating Systems Running on z/VM
- Frequently Asked Questions (if you don't ask them first)






## Intro to Crypto (The Really Short Version)



Rapelcgvba rkvfgr orpnhfr fbzrgvzrf  
jr yvxr gb xrrc frpergf.




Encryption exists because sometimes  
we like to keep secrets.

Cryptography is a mathematical function whereupon plaintext (“information in the clear”) is transmuted into a secret (“encrypted”) and can only be decrypted by someone who shares a common secret.

## Symmetric keys

(Examples: DES, Triple-DES, AES)

- A secret held in common by two parties
- Used to encrypt or decrypt a message in flight.
- Without the shared secret, a third party could not reasonably decrypt the message




- **The problem:** how does the secret key go from person A to person B?

## Asymmetric keys

(Examples: Diffie-Hellman, RSA, DSA, Elliptic Curve)

- Corresponding secrets used to encrypt information
- Data encrypted by the private key can be encrypted by anyone with the public key
  - Only **Alice** has **Alice's** private key; if we can decrypt this message, it's from Alice.
  - If we encrypt the response with **Alice's** public key, only **Alice** will be able to read it.



- Mathematically more intensive than symmetric (and therefore much slower)
- **Question:** what if someone drops a bit? What happens to the message?

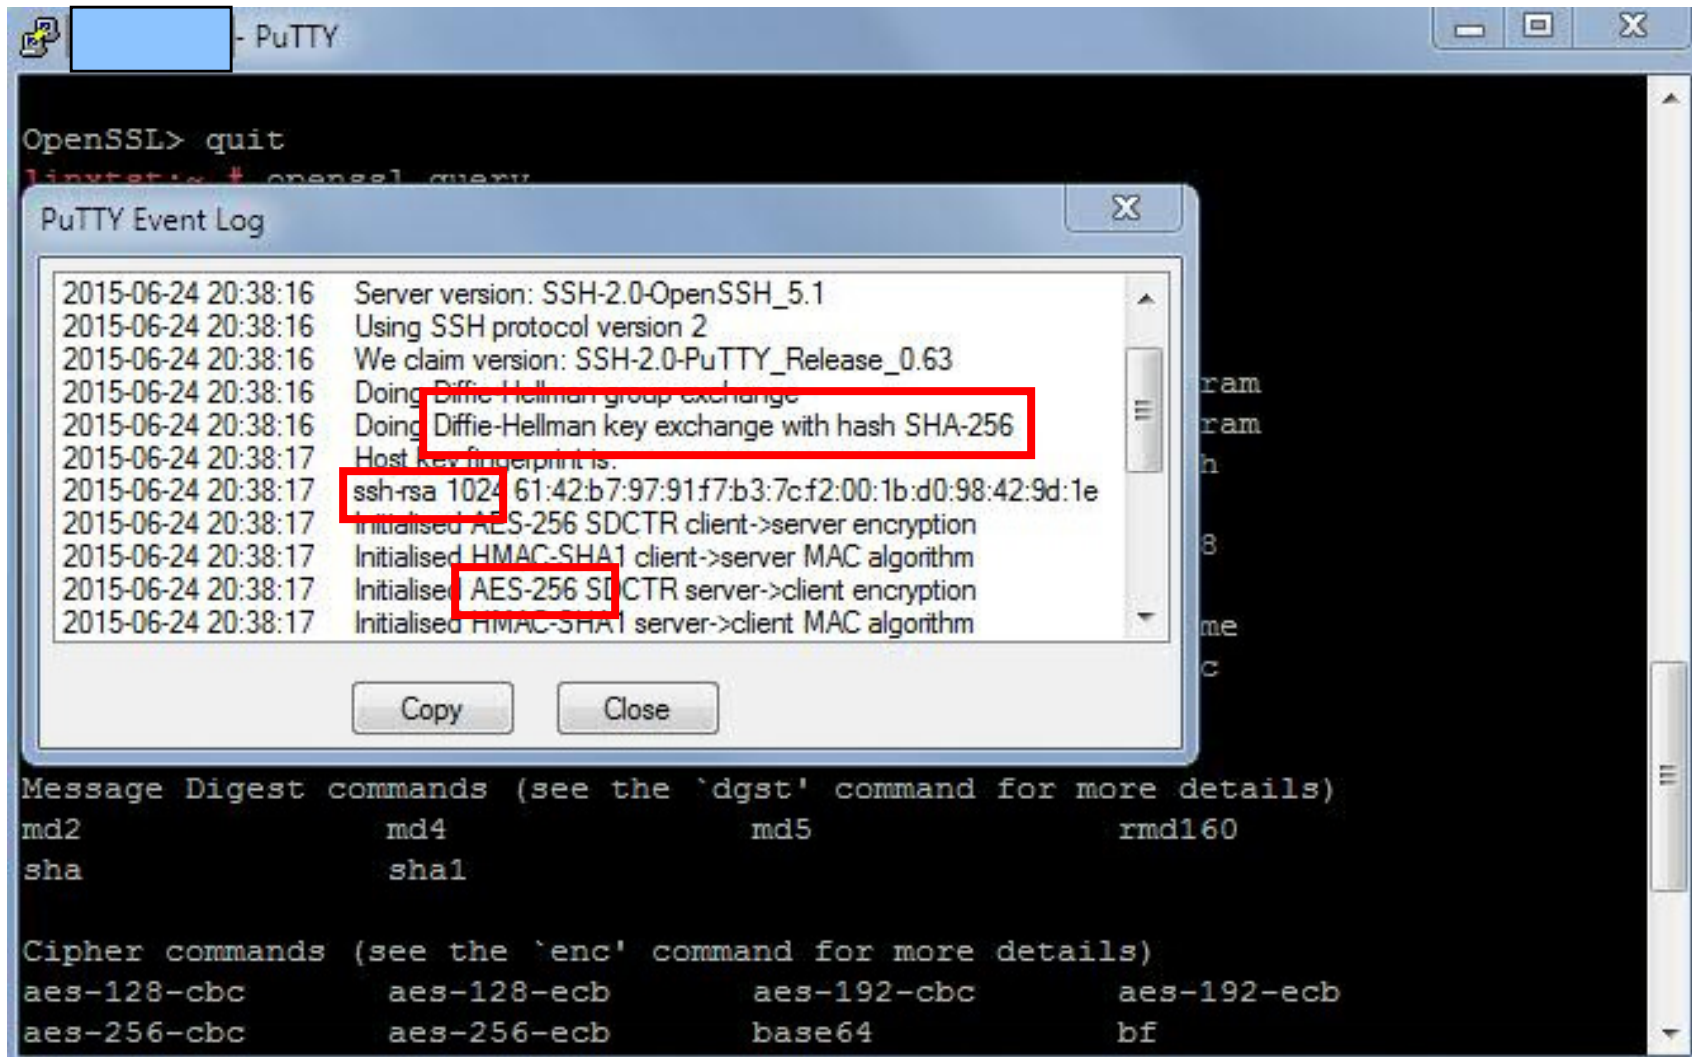
# Hashing

*(Examples: MD5, SHA-1, SHA-256, SHA-512, SHA-3)*

- Computes a "message digest" based on a set of data
- Used to ensure data integrity
  - Checksum computation
  - Message Authentication Codes (MACs)
  - Makes sure your data is the same at the destination as it was at the source



# It looks complicated, but it happens quickly.




## ... so let's apply our newfound knowledge!

- SSH connections and TLS connections use all three
  - Asymmetric key exchange to establish a connection
  - Symmetric keys to encrypt bulk traffic
  - Hashing to validate content between source and target
- That's a lot of math ... and it's processing power that adds up
  - The bigger (more secure) the keys, the longer it takes
  - Costs time, money




## IBM z Systems Cryptographic Features

- **IBM z Systems provide two flavors for offloading and accelerating cryptographic operations which help you to**
  - Move cryptographic workload away from central processors
  - Heighten your security level by protecting and securing keys
  - Accelerate encryption and decryption
  
- *CP Assist for Cryptographic Function (CPACF)*
  - Support for **symmetric** and hashing algorithms included in every CP and IFL
  - Pseudo-random number generator
  
- *Crypto Express features*
  - **Asymmetric** and hashing algorithm offload
  - Host master-key storage
  - Hardware RNG
  - PKCS #11 cryptographic support



## How do these features help?

- **Security** – encrypts your data at rest and in flight
- **Cost** – Saves on MIPS since it offloads from CPU
- **Capability** – modern algorithms aren't always implemented in the software libraries
- **Speed** – hardware is faster than software
- **Compliance** – meets regulations and complies with business or government standards



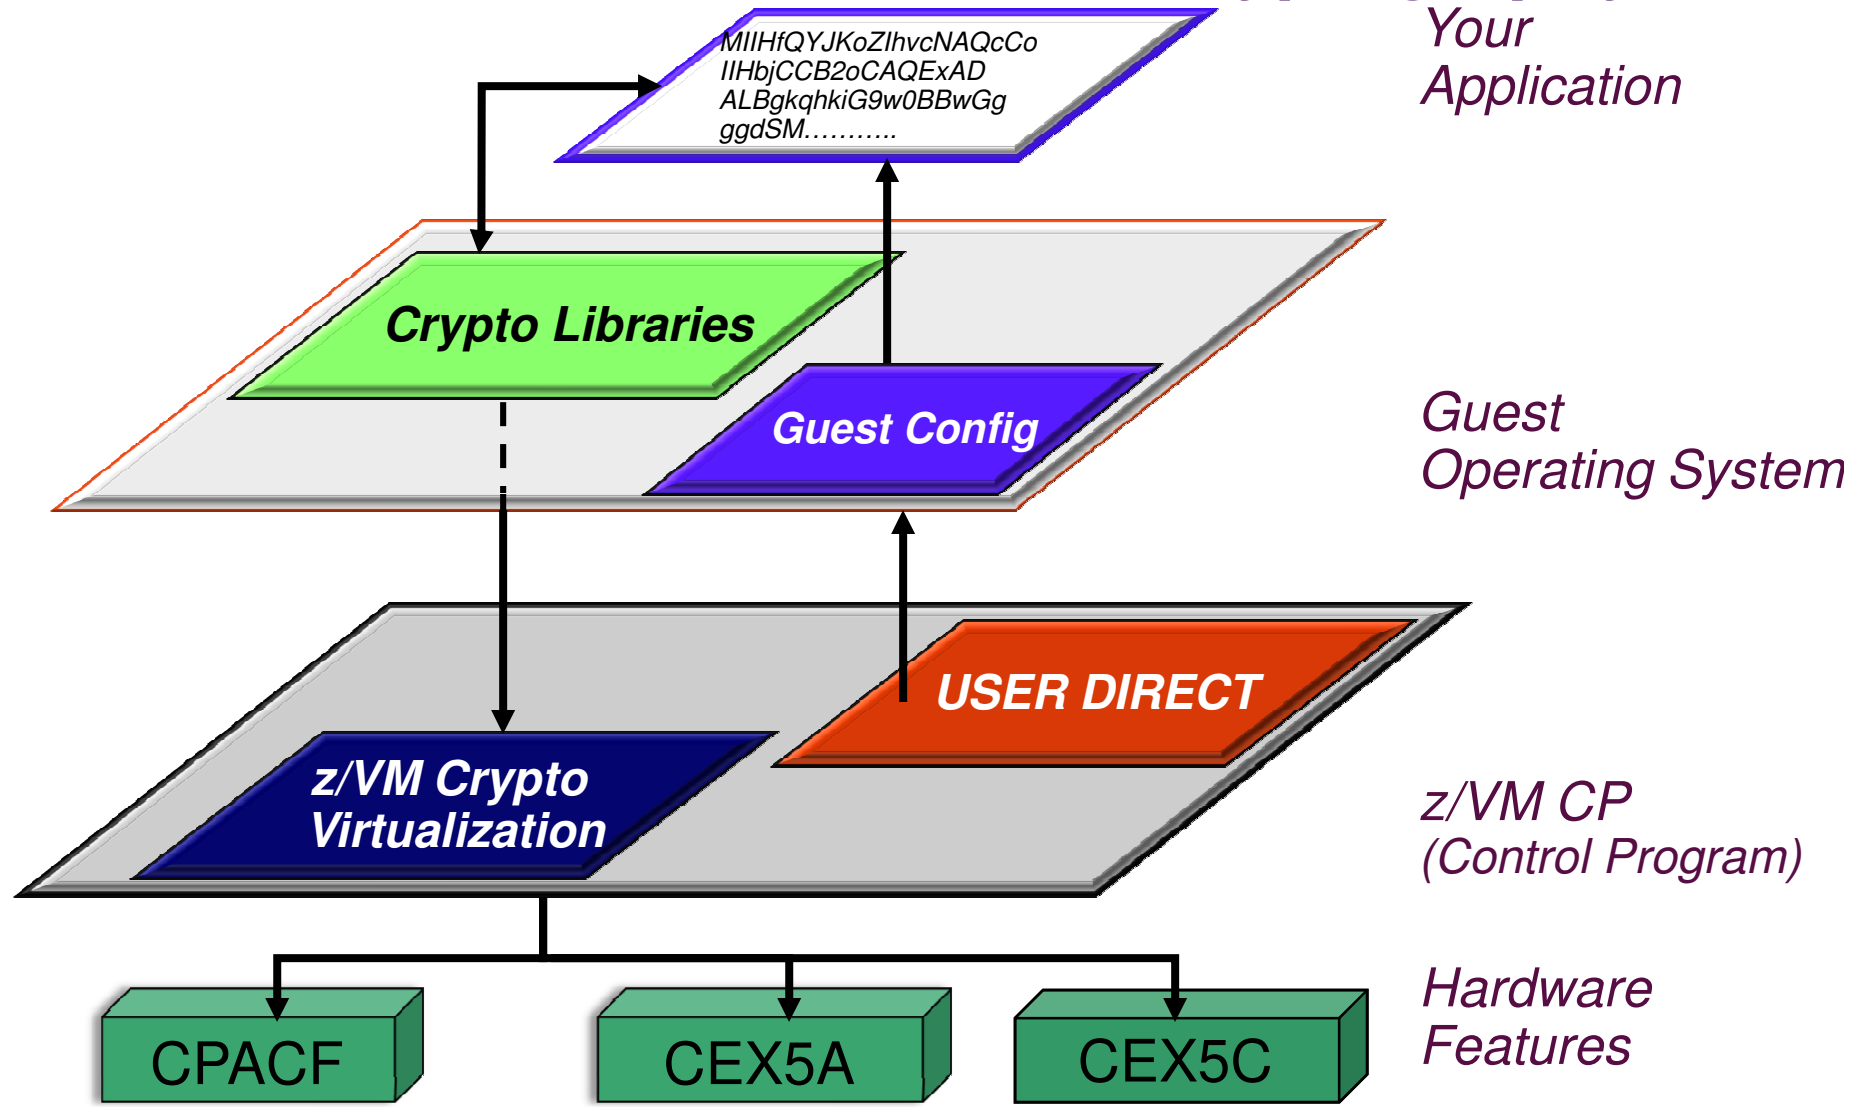




But that's just the hardware,  
and you're probably not running  
a single guest on an entire z13.

So let's take a look at how this ties  
into the rest of the z Systems  
virtual ecosystem.

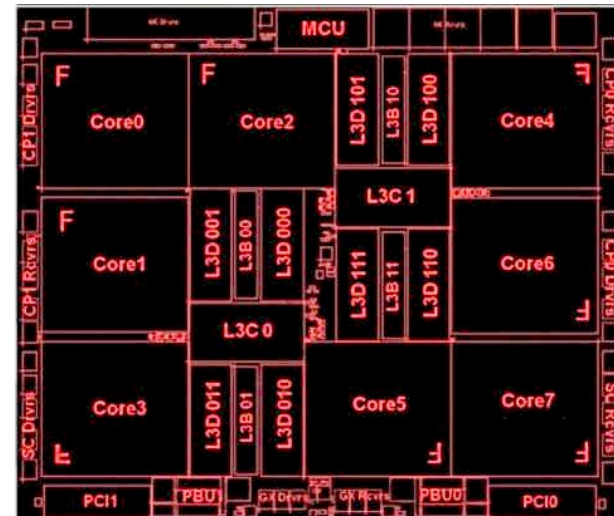
# z/VM Virtualization of Hardware Cryptography



# CP-Assisted Cryptographic Facility (CPACF)

## CPACF Support (Feature 3863)

- Available on all modern z Systems hardware but it must be explicitly enabled
- Provides on-CPU cryptographic processing *at a higher throughput*
- Supports the following algorithms:
  - DES
  - TDES
  - AES-128
  - AES-256 (z10 onward)
  
  - SHA-1
  - SHA-224 and SHA-256
  - SHA-384 and SHA-512 (z10 onward)
  
  - Single-length key MAC
  - Double-length key MAC



# CP-Assisted Cryptographic Facility (CPACF)

SCZP401 Details - SCZP401
i

Instance Information	Product Information	Acceptable CP/PCHID Status	STP Information	zBX Information	Energy Management
Ensemble name:	ITSO Ensemble	Ensemble HMC:	SCZHMxCB		
CP status:	Operating	Activation profile:	DEFAULT		
PCHID status:	Exceptions	Last profile used:	SCZP401		
zBX Blade status:	Not Operating	Service state:	false		
Group:	CPC	Number of CPs:	19		
IOCDS identifier:	A0	Number of ICFs:	8		
IOCDS name:	IODF78	Number of zAAPs:	6		
System mode:	Logically Partitioned	Number of IFLs:	4		
Alternate SE status:	Operating	Number of zIIPs:	6		
Lock out disruptive tasks:	<input type="radio"/> Yes <input checked="" type="radio"/> No	Dual AC power maintenance:	Fully Redundant		
		CP Assist for Crypto functions:	Installed		

OK Apply Change Options... Cancel Help

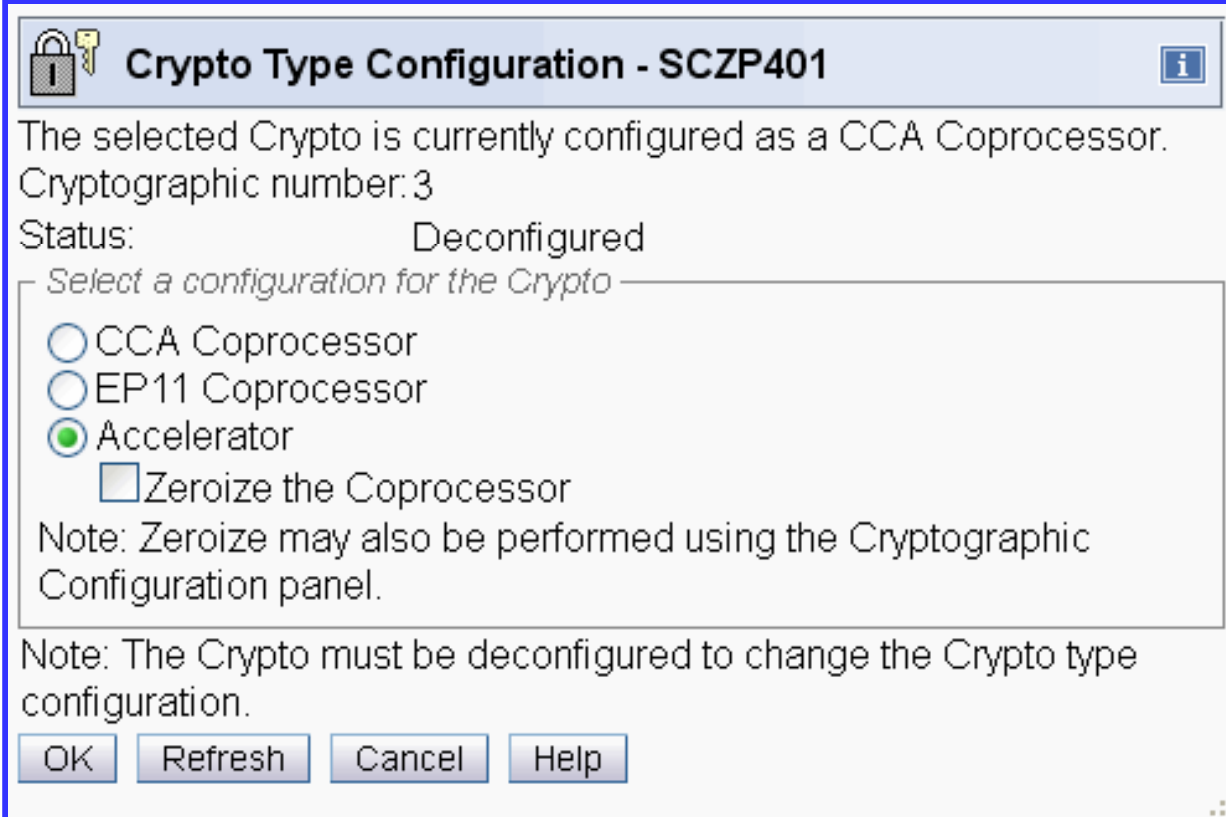
# IBM z Systems Crypto Express Features



## Crypto Express Support comes in three flavors

- IBM Common Cryptographic Architecture (CCA):
  - **CCA Accelerator mode**: meant for offload and acceleration of CPU intensive public/private key operations. Pertinent to workloads such as TLS and SSH, where secure handshaking factors heavily.
  - **CCA Coprocessor mode**: Accelerates public/private key operations and also supports secure key operations for encryption and decryption.
    - Coprocessor mode is the more cryptographically interesting of the two
    - Host master keys would be stored in Coprocessor domains
- **EP11 Mode (Coprocessor Only)**
  - Available on the Crypto Express4S
  - Enables PKCS#11 operations
  - *Must be set for the entire feature*

# Setting Personality for a Crypto Express Feature

- Hardware configuration for the feature is done on the Support Element
  - Step 1: turn on CPACF.
  - Step 2: Select feature, then choose personality type



 **Crypto Type Configuration - SCZP401** 

The selected Crypto is currently configured as a CCA Coprocessor.  
Cryptographic number: 3  
Status: Deconfigured

Select a configuration for the Crypto

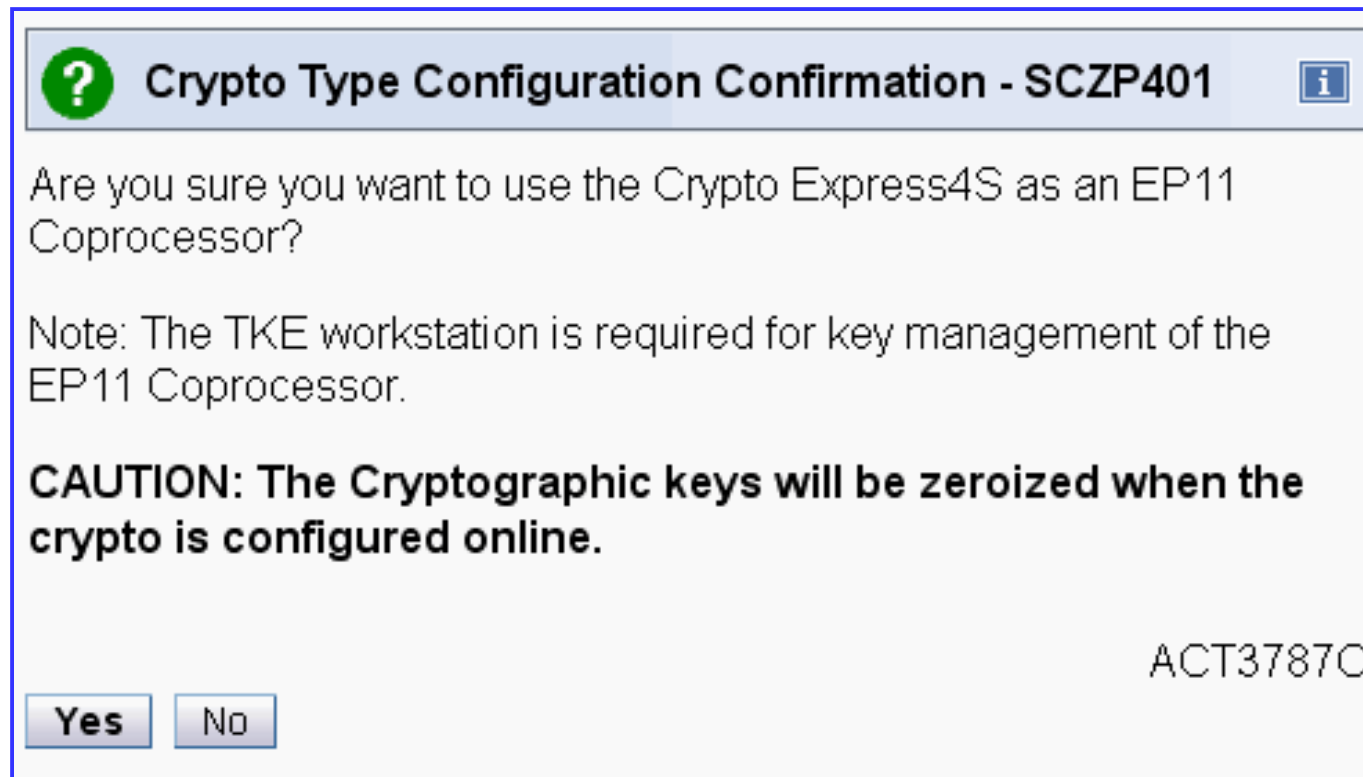
CCA Coprocessor  
 EP11 Coprocessor  
 Accelerator  
 Zeroize the Coprocessor

Note: Zeroize may also be performed using the Cryptographic Configuration panel.

Note: The Crypto must be deconfigured to change the Crypto type configuration.

# Setting Personality for a Crypto Express Feature

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



# Activating a Crypto Express Feature

- Hardware activation is done from the Support Element
- Select pertinent feature, "Configure On/Off"

**Support Element**

System Management > SCZP401 > Partitions > A01 > **Cryptos**

Cryptos Topology

Filter

Tasks Views

Select	Crypto ID	PCHID	Status	Channel Operations	Crypto Express
<input type="checkbox"/>	00	05FC	Operating		
<input type="checkbox"/>	01	05C0	Operating		
<input checked="" type="checkbox"/>	02	05BC	Stopped	Standby	Crypto Express4S CCA Coprocessor
<input checked="" type="checkbox"/>	03	0584	Stopped	Standby	Crypto Express4S EP11 Coprocessor
<input type="checkbox"/>	04	057C	Operating	Online	Crypto Express4S EP11 Coprocessor
<input type="checkbox"/>	05	0540	Operating	Online	Crypto Express4S CCA Coprocessor
<input checked="" type="checkbox"/>	06	053C	Stopped	Standby	Crypto Express4S Accelerator
<input type="checkbox"/>	07	0504	Operating	Online	Crypto Express4S CCA Coprocessor

Max Page Size: 500 Total: 8 Filtered: 8 Selected: 3

Context menu for selected row (02):

- Crypto Details
- Channel Operations
- Crypto Service Operations
  - Channel Problem Determination
  - Configure On/Off



# Attaching a Crypto Express Domain to an LPAR

- LPAR assignment is done from the HMC (building an activation profile)
  - **Candidate list:** domains on this AP which are **eligible to be accessed** by this partition
  - **Online List:** processors automatically brought online at LPAR startup.
  - **Usage Domain:** bundles domains together inside a common cryptographic boundary
  - **Control Domain:** identifies domain index pertinent to TKE control of the LPAR. Must also contain Usage Domain.
- **z/VM will only detect those cards and domains assigned to the LPAR**

Index	Control Domain	Usage Domain	Crypto Number	Cryptographic Candidate List	Cryptographic Online List
0	<input type="checkbox"/>	<input type="checkbox"/>	0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
1	<input type="checkbox"/>	<input type="checkbox"/>	1	<input type="checkbox"/>	<input type="checkbox"/>
2	<input type="checkbox"/>	<input type="checkbox"/>	2	<input type="checkbox"/>	<input type="checkbox"/>
3	<input type="checkbox"/>	<input type="checkbox"/>	3	<input type="checkbox"/>	<input type="checkbox"/>
4	<input type="checkbox"/>	<input type="checkbox"/>	4	<input type="checkbox"/>	<input type="checkbox"/>
5	<input type="checkbox"/>	<input type="checkbox"/>	5	<input type="checkbox"/>	<input type="checkbox"/>
6	<input type="checkbox"/>	<input type="checkbox"/>	6	<input type="checkbox"/>	<input type="checkbox"/>
7	<input type="checkbox"/>	<input type="checkbox"/>	7	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
8	<input type="checkbox"/>	<input type="checkbox"/>	8	<input type="checkbox"/>	<input type="checkbox"/>
9	<input type="checkbox"/>	<input type="checkbox"/>	9	<input type="checkbox"/>	<input type="checkbox"/>
10	<input type="checkbox"/>	<input type="checkbox"/>	10	<input type="checkbox"/>	<input type="checkbox"/>
11	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	11	<input type="checkbox"/>	<input type="checkbox"/>
12	<input type="checkbox"/>	<input type="checkbox"/>	12	<input type="checkbox"/>	<input type="checkbox"/>
13	<input type="checkbox"/>	<input type="checkbox"/>	13	<input type="checkbox"/>	<input type="checkbox"/>
14	<input type="checkbox"/>	<input type="checkbox"/>	14	<input type="checkbox"/>	<input type="checkbox"/>
15	<input type="checkbox"/>	<input type="checkbox"/>	15	<input type="checkbox"/>	<input type="checkbox"/>

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.

## Getting Keys into z Systems

- **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 z System
  - Card reader for crypto secret storage
  - Generates new secrets, stores data in Crypto Express domains
  
- **z/OS Integrated Cryptographic Services Facility (ICSF)** – a base component which allows interaction with Crypto Express features. (Requires MVS.)
  
- **Panel.exe Utility for Linux** – a Linux package installed as part of the IBM .rpms which allows for key management function.
  - `/opt/IBM/CEX3C/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>

# Securing the Keys Once They're Installed

Three different types of **key protection** in the **IBM Crypto hardware**:

- **Clear keys:**

- The security of keys is provided by operating procedures.
- This means keys may appear in the clear in the environment somewhere

- **Secure keys: (FIPS 140-2 Level 4 certified)**

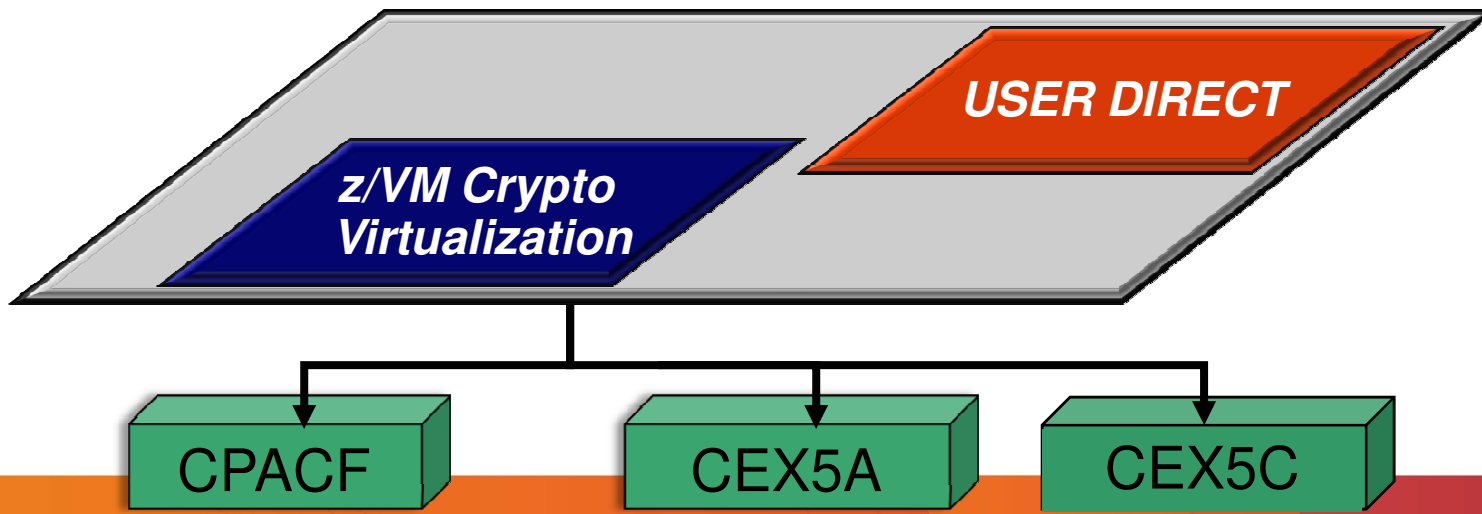
- Secure keys are protected by another key (the master key) stored in hardware
- When a secure key must leave the hardware, the key encrypted under the master key ... so the value of the secure key is never exposed to the operating system



- **Protected keys (CCA only):**

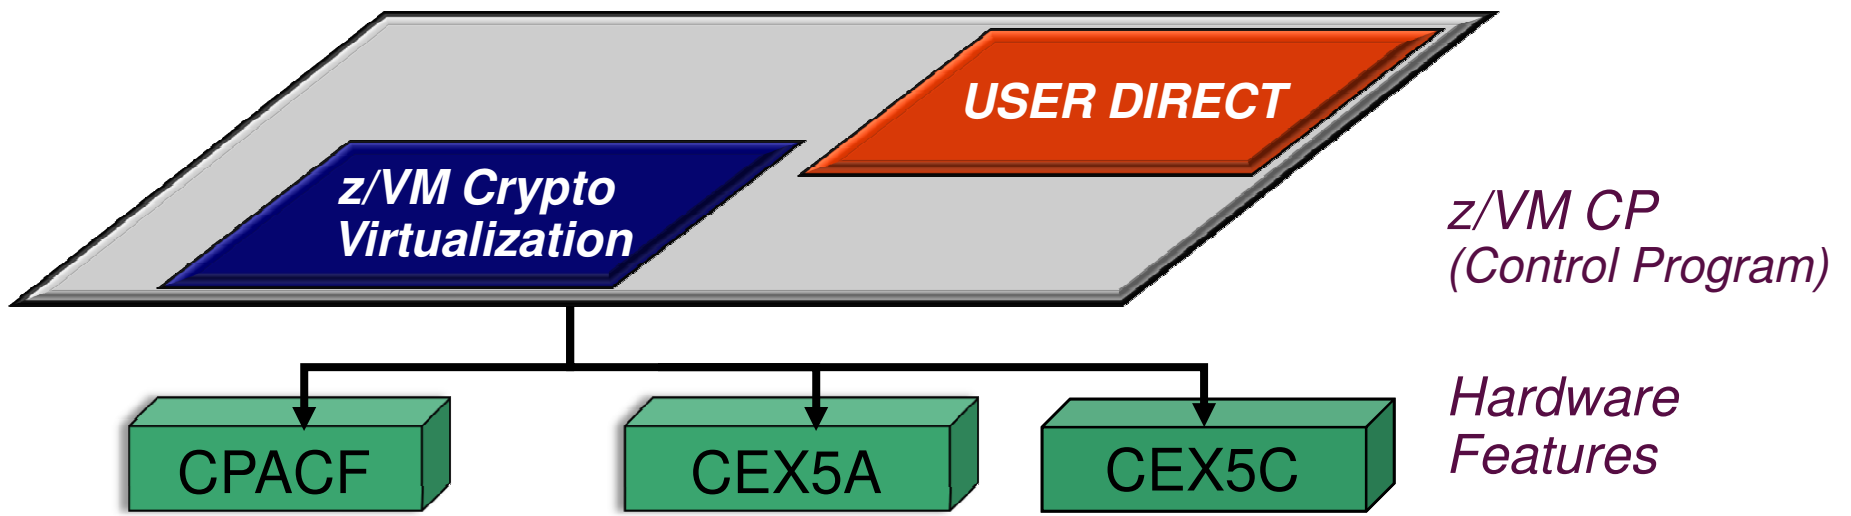
- Protected keys are encrypted under a Wrapping Key uniquely created for each LPAR
- Cryptographic operations using protected keys can benefit from **CPACF** performance

## z/VM Configuration for z Systems Cryptography




# z/VM Virtualization of Hardware Cryptography (stack view)

- Once domains are added to an LPAR running z/VM, they become available for guest use



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



# z/VM Virtualization of Hardware Cryptography

The **CRYPTO User Directory statement** grants a z/VM userid access to cryptographic features associated with the hardware:

```

                                v-----+
CRYPTo  +- DOMAIN ---+domains +- APDEDicated +- aps --+----><
        |
        +- APVIRTual-----+-----^
    
```

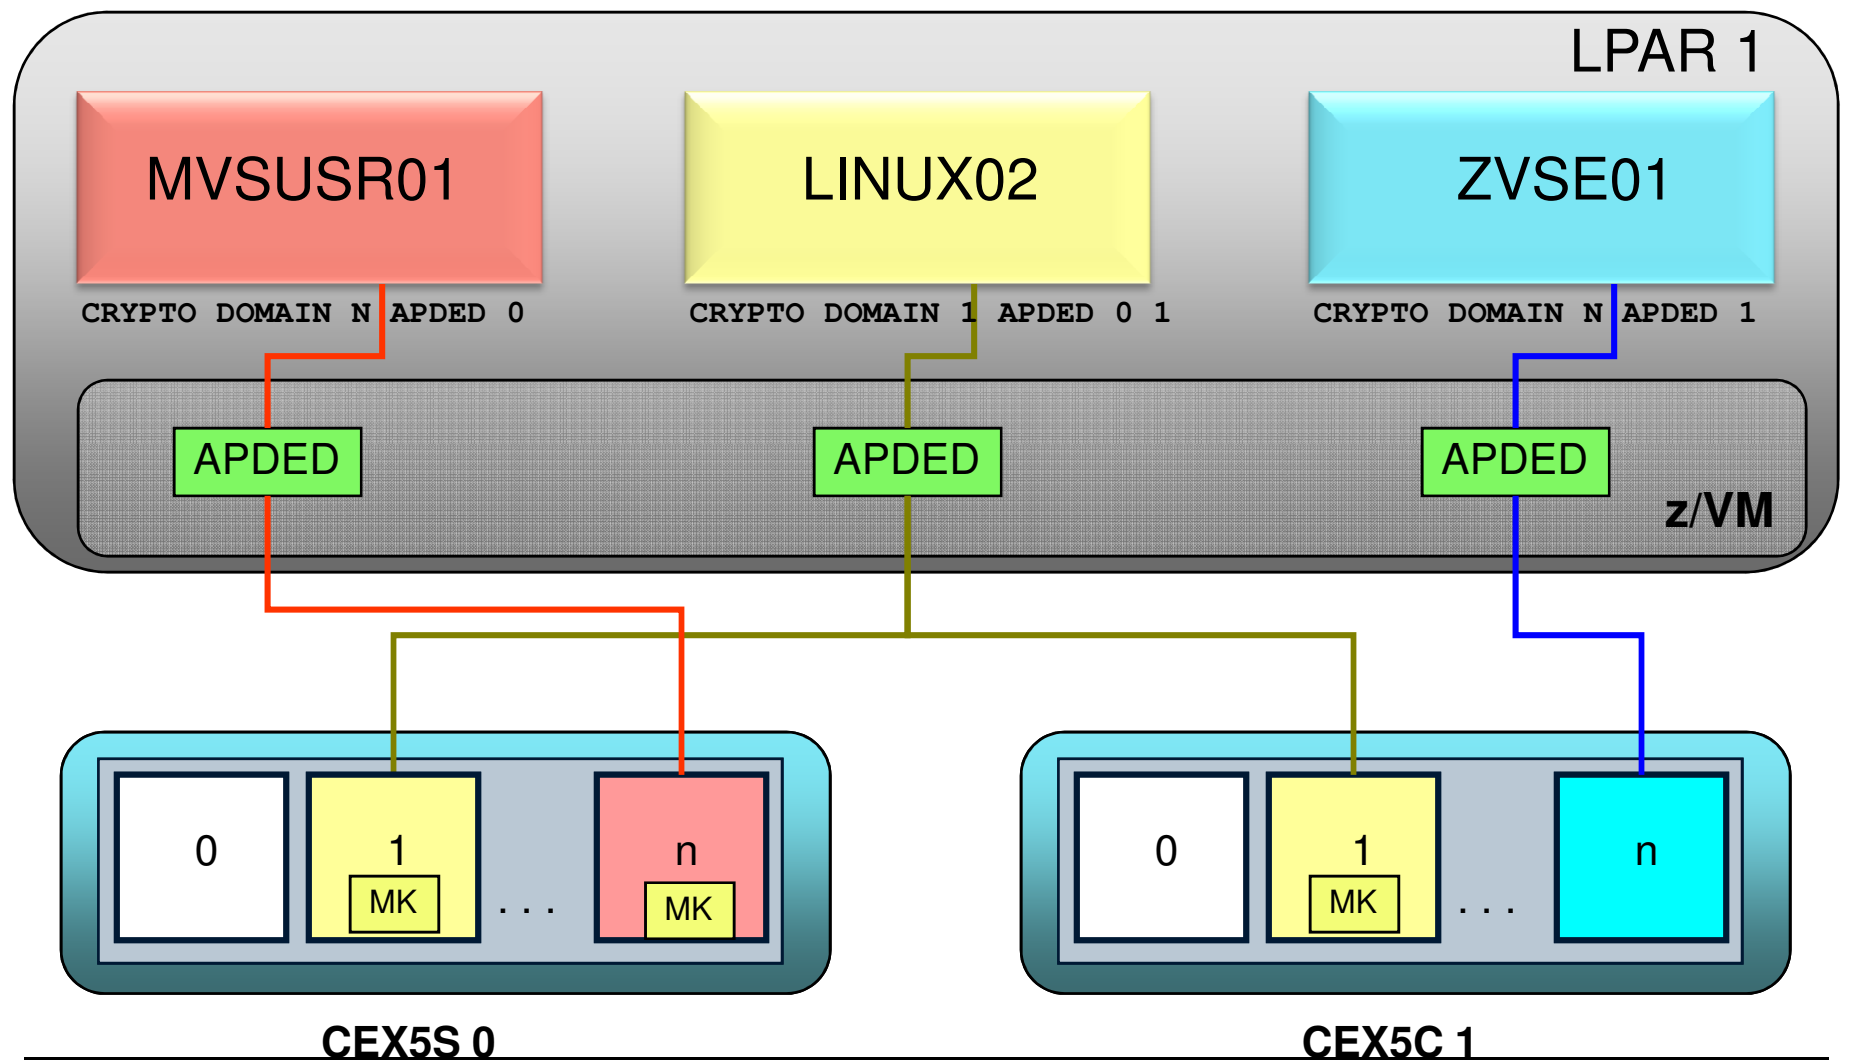
## APDED

Dedicates a particular AP domain (or set of domains) to this virtual machine. Domains granted in the directory are “reserved for dedication”; they are not actually in-use until the virtual machine logs on.

## APVIRT

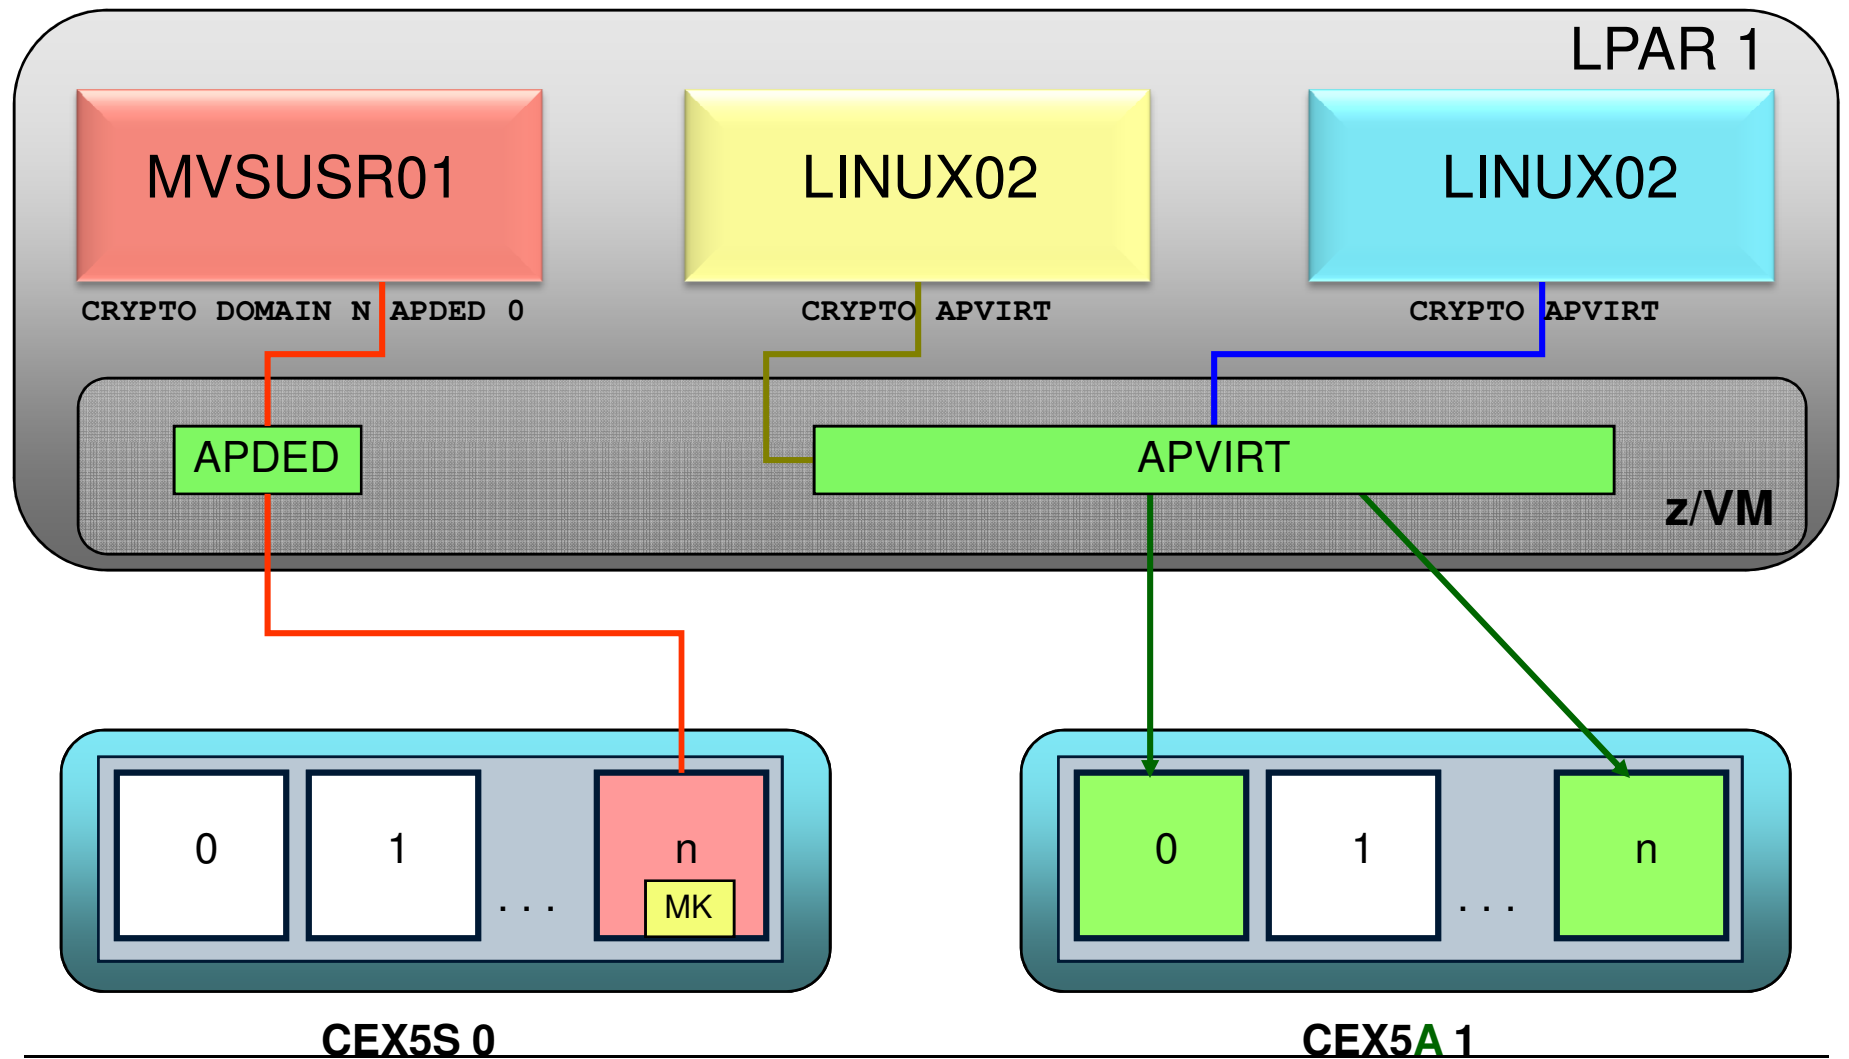
Virtual machine can access a collection of domains controlled by the system.

# Assigning AP Domains to z/VM Guests



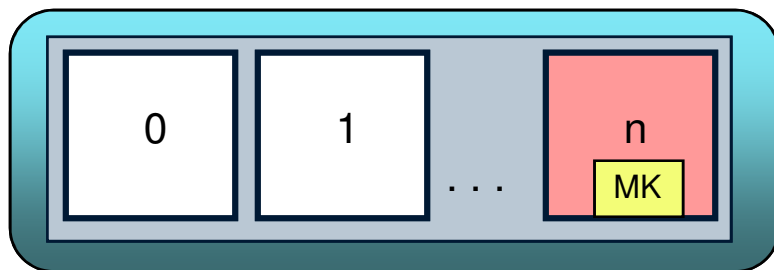


# Assigning AP Domains to z/VM Guests




## Notes on APVIRT Domain Selection

- Any domain in APVIRT will behave as an accelerator
  - Whether it is or not it is configured as one or not
  - CP will discard coprocessor operations sent to an APVIRT domain
  - This is done for security context reasons (and why APVIRT is meant for clear-key)
- APVIRT domains are selected by mode and release level
  - Accelerator is chosen before coprocessor
  - CEX5S is chosen before CEX4S before CEX3 ...
- EP11 domains cannot be used for APVIRT



**CEX5S 0**



**CEX5A 1**

## Assigning Domains to APVIRT (APAR VM65577)

- z/VM supports a new [System Configuration statement](#) for z/VM V6 which allows a system administrator to [assign APVIRT domains](#) for use by CP:

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

- Usage Notes:
  - z/VM will designate the first available domain in its list as the type
  - Any other available domains in SYSTEM CONFIG that are [also of that type](#) are designated for APVIRT usage
  - Domains that do not meet criteria are ignored.
- If this statement is not present in the System Configuration file, z/VM will use default APVIRT domain selection behavior

## Assigning Domains to APVIRT (APAR VM65577)

- Given the following System Configuration:

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


... z/VM V6 will check domains in the following order:

```
AP 1 DOMAIN 7          /* CEX5A */
AP 1 DOMAIN 8          /* CEX5A */
AP 2 DOMAIN 7          /* CEX4A */
AP 2 DOMAIN 8          /* CEX4A */
AP 4 DOMAIN 9          /* CEX5C */
```

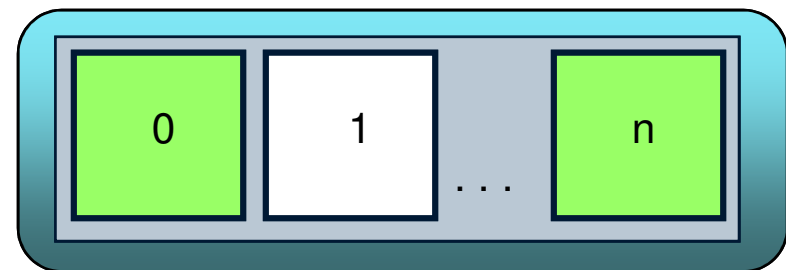
- If **AP 1 DOMAIN 7** is available at system initialization, it will be APVIRT.
  - APVIRT must use type CEX5A
  - Only AP 1 DOMAIN 8, with a matching type, is set as APVIRT
  - If a guest lists AP 1 DOMAIN 7 as **APDED**, the guest will be denied access

## Support for Crypto Express5S (APAR VM65577)

- Expanded domain selection for dedicated domains
  - z/VM supports architected limits for CryptoExpress domains
  - CryptoExpress5S supports 85 domains per feature, maximum of 16 features
  
- APDED really does mean dedicated; no collision is permitted
  - In a race, the first guest to LOGON has all requests fulfilled
  - Collisions void the latter guest's domain claims for an entire AP




CEX5C 0



CEX5A 1

## Example: Assigning Domains for z/VM

- **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 ...

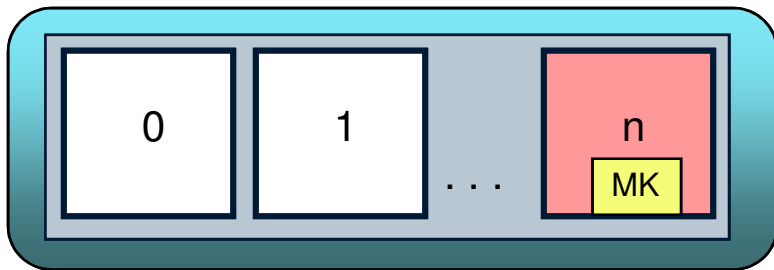


# z/VM Virtualization of Hardware Cryptography

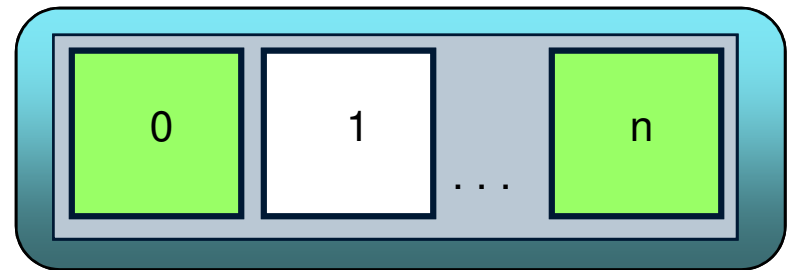
## QUERY CRYPTO

(Class A, B, C, or E) will display which domains/APs are available. Note that this list will be limited to devices available to a z/VM instance.

```
>>-Query--CRYPTo--+-----+-----+-----><
                '-DOMains--+-----+-'
                    '-Users-'
```



**CEX5S 0**



**CEX5A 1**

# z/VM Virtualization of Hardware Cryptography

```
QUERY CRYPTO DOMAINS USERS
```

<u>AP</u>	<u>device</u>	<u>Domain nn</u>	<u>device status</u>	<u>system usage</u>	<u>planned usage</u>
01: AP 02	CEX3C	Domain 08	available	free	unspecified
01: AP 03	CEX3A	Domain 06	available	dedicated to BWHUGEN	dedication
01: AP 03	CEX3A	Domain 07	available	free	unspecified
01: AP 03	CEX3A	Domain 08	available	shared	shared
01: AP 04	CEX4C	Domain 06	available	free	dedication
01: AP 04	CEX4C	Domain 07	available	free	dedication
01: AP 04	CEX4C	Domain 08	available	free	unspecified

Ready;



# z/VM Virtualization of Hardware Cryptography

## QUERY VIRTUAL CRYPTO

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

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

```

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

```

## QUERY VIRTUAL CRYPTO

```


AP 03 CEX3A Domain 06 dedicated
Ready;

```

# Assigning AP Domains to z/VM Guests

- **The Big Question: Which type of domain 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?
  
- Different guests will have different needs, based upon their drivers and configuration requirements ...

\*Note: some restrictions apply. Consult the *CP Planning and Administration Guide* or *Getting Started With Linux* manuals.



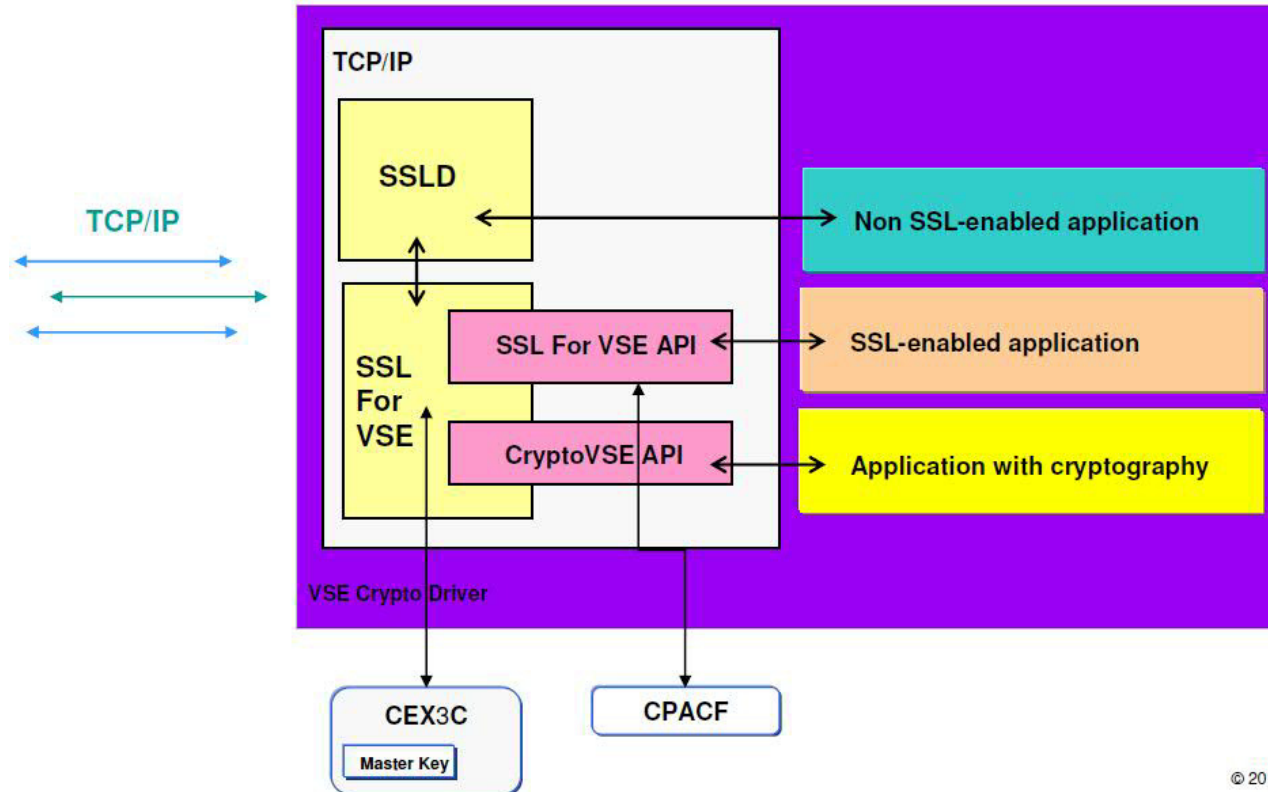
## Guest Usage of the z Systems Cryptographic Features

# z/VM Virtualization of Hardware Cryptography



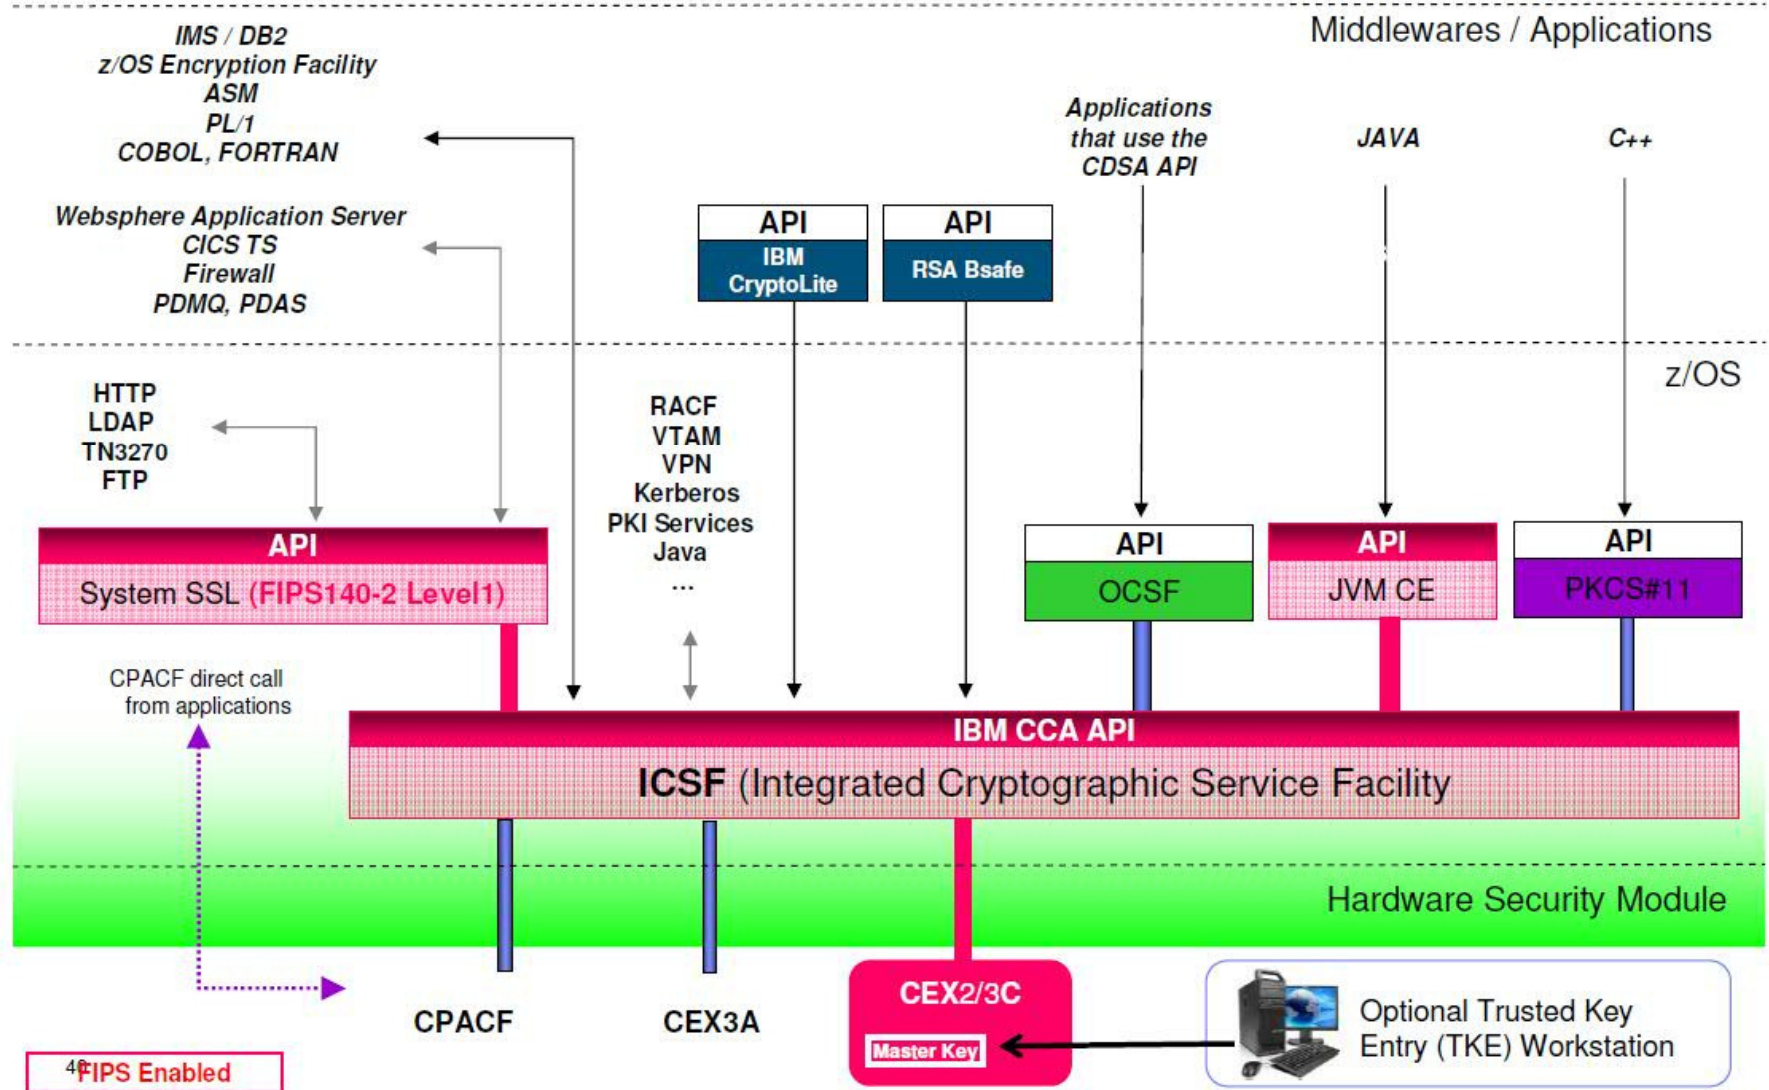
- Cryptographic libraries will vary from operating system to operating system
- Some may require specific configuration to make use of certain features
- Consult pertinent local documentation

# z/VSE Cryptographic Infrastructure

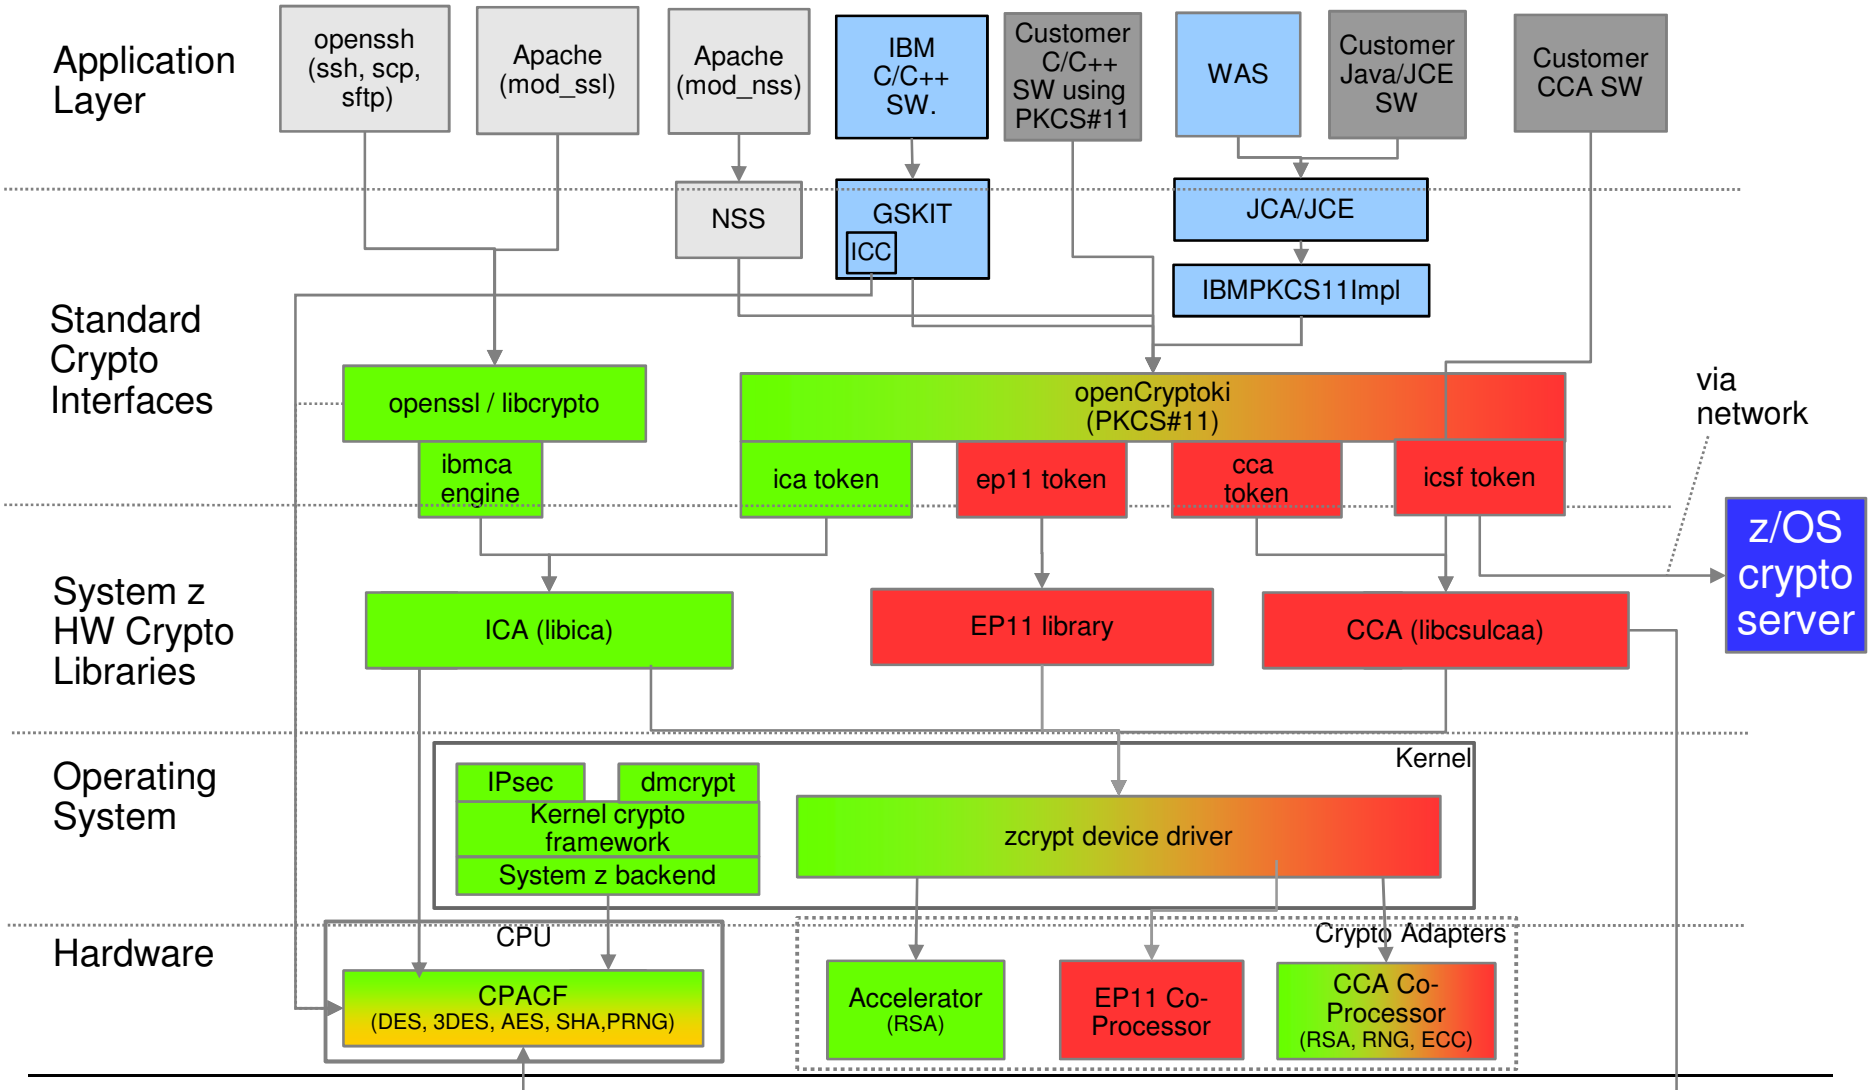


- z/VSE automatically detects any Crypto Express features dedicated to (or shared with) the virtual machine in which it's running

# z/OS Cryptographic Infrastructure



# Linux on z Systems Crypto 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 z Systems 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 Systems Encryption (dm-crypt)

- **dm-crypt** (transparent disk encryption subsystem)
  - Inserts layer of crypto between block device & accessing file systems or apps
  - Positioned between file system and device mapper
- Administration done through **cryptsetup**
  - Uses LUKS (Linux Unified Key Setup)
  - Choose cipher/hashing algorithms from **/proc/crypto/procfs**
  - HW crypto (AES-CBC, XTS-AES)
- Can also set up encrypted filesystems during init
  - **/etc/crypttab** (referenced before **/etc/fstab** )
  - Bear in mind, though, interactive password prompts will still wait for you



# Linux Support for Crypto Express5S

## Today: Toleration Support

- Linux kernel recognizes CEX5S adapter and treats it as CEX4S adapter
- New **sysfs** attribute shows its real identity under `/sys/bus/ap/raw_hwtype`
- Supported Distributions
  - SLES 11 SP3 + maintenance
  - SLES 12 + maintenance
  - RHEL 7.1
  - RHEL 6.6 + maintenance
  - RHLE 5.11
- Some Restrictions Apply
  - [http://www.ibm.com/developerworks/linux/linux390/distribution\\_hints\\_z13.html](http://www.ibm.com/developerworks/linux/linux390/distribution_hints_z13.html)

## Validating Linux and z/VM Configuration

```
certlxb:~ # 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

```
certlxb:~ # sudo vmcp QUERY VIRTUAL CRYPTO  
AP 01 CEX4A Queue 01 shared
```

- 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.



## Frequently Asked Questions

# Frequently Asked Questions

Do these crypto features meet any particular industry standards?



**Answer:** The Crypto Express cards are certified to the Federal Information Processing Standard (FIPS) 140-2 at Level 4. The **secure-key protection** not only meets HSM requirements, but is confirmed as zeroizing Master Keys in case of physical tampering, x-rays, power-supply interruption ...



+ Master Key zeroization in case of tampering attempt



# Frequently Asked Questions



- **Terminology question** – is it a *domain*? or a *queue*? or an *AP*?
- **Answer:** In this context, “domain” and “queue” are mostly synonymous.

z/VM’s QUERY CRYPTO command (as of z/VM 6.2) documents the sub-structures associated with the Crypto Express features as “domains.” APQS (short for ‘Adjunct Processor Queues’) is still accepted as an operand, and the terminology of ‘queues’ may still appear in documentation related to other IBM products.

The ‘AP’ in abbreviations like ‘APDED’ and ‘APVIRT’ refers to ‘Adjunct Processor’ ... which is another term of the Crypto Express features (CEX2 and onward).

# Frequently Asked Questions



- **CPACF for the z/VM TLS Server? Cool!**  
**But does z/VM TLS use the Crypto Express Cards?**

- **Answer:** No. While SSLSERV and LDAPSRV use **CPACF** if it is enabled, z/VM only virtualizes Crypto Express support for hosted operating systems.

z/VM's CMS-based servers will not utilize the Crypto Express features.

# Frequently Asked Questions



- **What happens if two z/VM guests have the same domain DEDicated to them on the CRYPTO statement?**
- **Answer:** The domain is considered “Reserved for Dedication” until one of the guests IPLs. At that time, the domain is considered dedicated. If the second guest IPLs at that time, the virtual machine will not receive that domain for use.
- **Update for z13:** Not only will the second guest not receive the conflicting domain, but it will not be able to access any of the domains it's reserved on that entire AP.
- **Final Answer:** *Be careful in your domain assignments.* Your guests should not swap dedicated domains!

# Frequently Asked Questions



- **Bonus Question!** Explain the following statement:

CRYPTO DOMAIN 0 1 APDED 14 15

- **Answer:** The guest receives dedicated access to the following domains:

[0, 14]

[0, 15]

[1, 14]

[1, 15]

- Domain assignation is a **union** of the AP queues and specific domains listed; be careful about assigning too many domains when configuring your z/VM virtual machines.

# 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



- **Question:** I just overhauled my USER DIRECT, and suddenly my guests can't use their shared crypto domains. What happened??
- **Answer:** On z/VM 5.4 (and on z/VM V6 before z13 support), there is no way for the system administrator to assign APVIRT domains specifically for system use. Instead, APVIRT domains are assigned at system IPL and **are managed by CP.**

If you've rearranged your User Directory and reserved a previously shared domain for dedicated use, you may see errors related to availability. You may need to restart your z/VM LPAR to regain specific domains.

**Note:** This will continue to be the default behavior for z/VM V6 for any system where a CRYPTO APVIRT statement is not specified in your System Configuration file.

# Frequently Asked Questions



- **Question:** what are those restrictions on migrating guests with crypto domains assigned to them?

- **Answer:** For APVIRT, the target system must have Crypto Express domains available for APVIRT which match the same mode as what was available on the source system. So, if **SYSTEMA** is using CEX3A for APVIRT, then CEX3A must similarly be available on **SYSTEMB**.

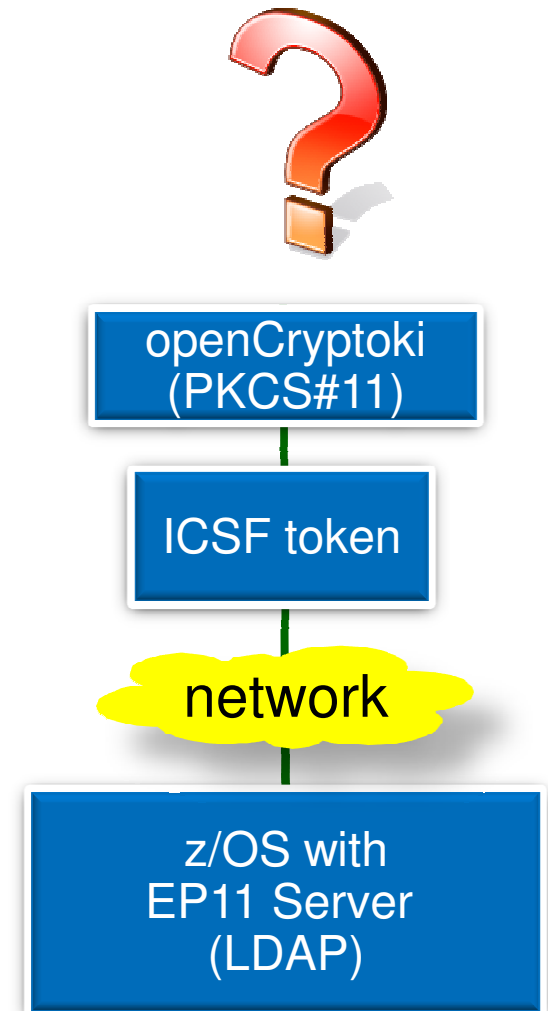
Additionally, the domain on the target system must provide the same level of *function*.

Relocation of a Linux guest with dedicated use of a domain is not permitted.

- \***Reminder:** Consult the *CP Planning and Administration Guide* or *Getting Started With Linux* manuals for more details!

# 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`







Questions?


## Summary

- z Systems **hardware** cryptography accelerates the hard math of crypto
  - Saves **time**, saves CPU processing **power**, saves MIPS **cost**
  - Secure Key operations are FIPS 140-2 Level 4 certified
  
- z/VM **virtualizes** z Systems hardware cryptography
  - Architectural fidelity in all things z
  - A "shared" flavor as well as dedicated domain use
  
- **Guests** that understand cryptography can utilize z Systems cryptography
  - May require configuration of the guest to exploit
  - Different guests provide different options
  
- Don't let cryptography (or its terminology) scare you away
  - Security is meant to enhance business, not impede it
  - Cryptography protects your data, whether at rest or in flight

## For More Information ...

- z/VM Security: <http://www.VM.ibm.com/security>
- z Systems Security: <http://www.ibm.com/systems/z/advantages/security/>
- IBM zEC12 Configuration Setup (SG24-8034), IBM RedBooks  
<http://www.redbooks.ibm.com/abstracts/sg248034.html?Open>
- **Security for Linux on System z** (SG24-7728), IBM RedBooks  
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247728.pdf>
- *z/VM Secure Configuration Guide*: <http://publibz.boulder.ibm.com/epubs/pdf/hcss0b30.pdf>
- IBM z13: <http://www-03.ibm.com/systems/z/hardware/z13.html>
- IBM z Systems Crypto Express Features:  
<http://www-03.ibm.com/security/cryptocards/pciecc/overview.shtml>

### *Contact Information:*

[Brian W. Hugenbruch](#), CISSP  
IBM z Systems Virtualization Security  
[bwhugen at us dot ibm dot com](mailto:bwhugen@us.ibm.com)  
 [@Bwhugen](#)

**Dank u**

Dutch

**Merci**

French

**Спасибо**

Russian

**Gracias**

Spanish

شكراً

Arabic

감사합니다

Korean

**Tack så mycket**

Swedish

धन्यवाद

Hindi

תודה רבה

Hebrew

**Obrigado**

Brazilian  
Portuguese

谢谢

Chinese

**Dankon**

Esperanto

**Thank You**

ありがとうございます

Japanese

**Trugarez**

Breton

**Danke**

German

**Tak**

Danish

**Grazie**

Italian

நன்றி

Tamil

děkuji

Czech

ขอบคุณ

Thai

go raibh maith agat

Gaelic