IBM WW Z Security Conference

October 6-9, 2020

Protecting your Data in-use with IBM Secure Execution for Linux

- Reinhard Buendgen
- *Product owner security for Linux on Z and LinuxONE*
- *buendgen@de.ibm.com*

Trademarks

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

CICS*
Db2*
GDPS*
HiperSockets
IBM*
IBM (logo)*
IBM Cloud Pak
ibm.com
IBM Z*
LinuxONE
WebSphere*
z14*
z15
z/OS*
z/VM*
z/VSE*

*** Registered trademarks of IBM Corporation**

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.
IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.
ITIL is a Registered Trade Mark of AXELOS Limited.
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 other countries.
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.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Kubernetes and Container Initiative™ are registered trademark of The Linux Foundation.
Red Hat and Red Hat OpenShift are registered trademarks of Red Hat, Inc. Open
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
UNIX is a registered trademark of The Open Group in the United States and other countries.
VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Would you move sensitive workloads to the (public) cloud?

The IT strategy of many customers aims to deploy their workloads in the cloud.

- Cloud computing is "in"
- IBM propagates cloud computing (hybrid, multi)

But would you like

- the IRS do their tax processing
- your hospital and health insurance process your medical records
- your company to process financial strategies and results
- your R&D department design new products based on your intellectual property

in the public cloud

I guess not and there were good reasons for your hesitation …

Security Issues of Hosted Workloads

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image used to deploy my workload the one I provided?
- Are the secrets in my image kept confidential?

Guest protection

- Is the data in my running guest protected?

Data at-rest protection

- Is the data stored on disks protected

Workload placement

- Is the system that hosts my workload the system I think it is hosted on?

Guest isolation

- Is my running guest isolated from guest of other tenants?

Data data-in flight protection

- Is the data communicated to other systems protected

Security concerns when deploying workloads in the cloud?

Image protection

- Is th
  dep
  one
- Are
  imag

Guest protection

Data at-rest protection

on

Wo

- Is th
  my
  I thi

ght

other
systems protected

Image protection

- Is the image used to deploy may workload the one I provided?
  - Has the image been replaced by something that just looks like my image?
  - Has my boot image been corrupted? E.g. to leak data?
- Are the secrets in my image kept confidential?
  - Your image typically contains secrets.
    - E.g. passwords, SSH & TLS certificates, dm-crypt keys, …
  - Does the provider or any other party watching the transfer of the image get access to those secrets?

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image used to de or
- Ar im

Guest protection

- Is the data in my running

Data at-rest protection

Is the data stored on

Workload placement

- Is the system that hosts my workload the system I think it is hosted on?
- Did the image get intercepted and redirected to a different system?
- Is the host system managed by a provider I trust?
- Is it placed on the kind of system I pay the QoS for?
- Is it placed in a geography that is compatible with the legal requirements of my workload.

W

- Is my I think it is hosted on?

other tenants?

communicated to other systems protected

Security concerns when deploying workloads in the cloud?

Guest protection

- Is the sensitive data in my running guest protected?
  - with respect to both integrity and confidentiality
- Can the operators of the provider access/modify my data?
  - e.g. if the operator is rouge or threatened by an outsider
  - from the HW console, from the hypervisor
- Is the cloud environment secure?
  - no vulnerabilities, no zero days?
  - trustworthy SW?
  - well maintained?
    - good access control in place, security fixes applied, …
- Can an intruder who breaks into the cloud infrastructure access/modify my data?

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image
deploy may
one I provide
- Are the secr
image kept c

Guest protection

Data at-rest protection

stored on
cted

Guest isolation

- Is my running guest isolated from guest of other tenants?
- Can neighboring guests access data of my guest?
- Typically virtualization FW prevents memory accesses from on virtual machine to another, but a malicious guest may try to break into its hypervisor and from there …
- This boils down to can the data of my running guest be accessed from the hypervisor?

Workload p

- Is the systen
my workload
I think it is ho

a-in flight
ction

ted to other
systems protected

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image used to deploy may workload the one I provided?
- Are the secrets in my image kept confidential?

Guest protection

- Is the data in my running guest protected?

Data at-rest protec n

- Is the data stor protect

Pervasive
Encryption

Workload placement

- Is the system that hosts my workload the system I think it is hosted on?

Guest isolation

- Is my running guest isolated from guest of other tenants?

Data data-in fligh protection

- data i other sy ted

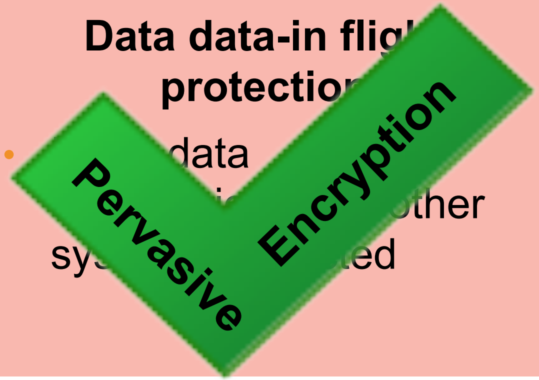

Trusting Cloud Deployments

Given all these security issues …

- Should you use the cloud at all?
- Why is IBM propagating cloud computing?

- Well, it is a matter of
  - trust in the cloud environment
    - provider
    - HW, SW, network infrastructure
    - policies,
    - contracts
  - the value of your sensitive data

Traditional Cloud Trust Model: Trust the CEOs

"Traditional cloud trust model"

- customer trusts good CEO
- good CEO enforces good policies on employees

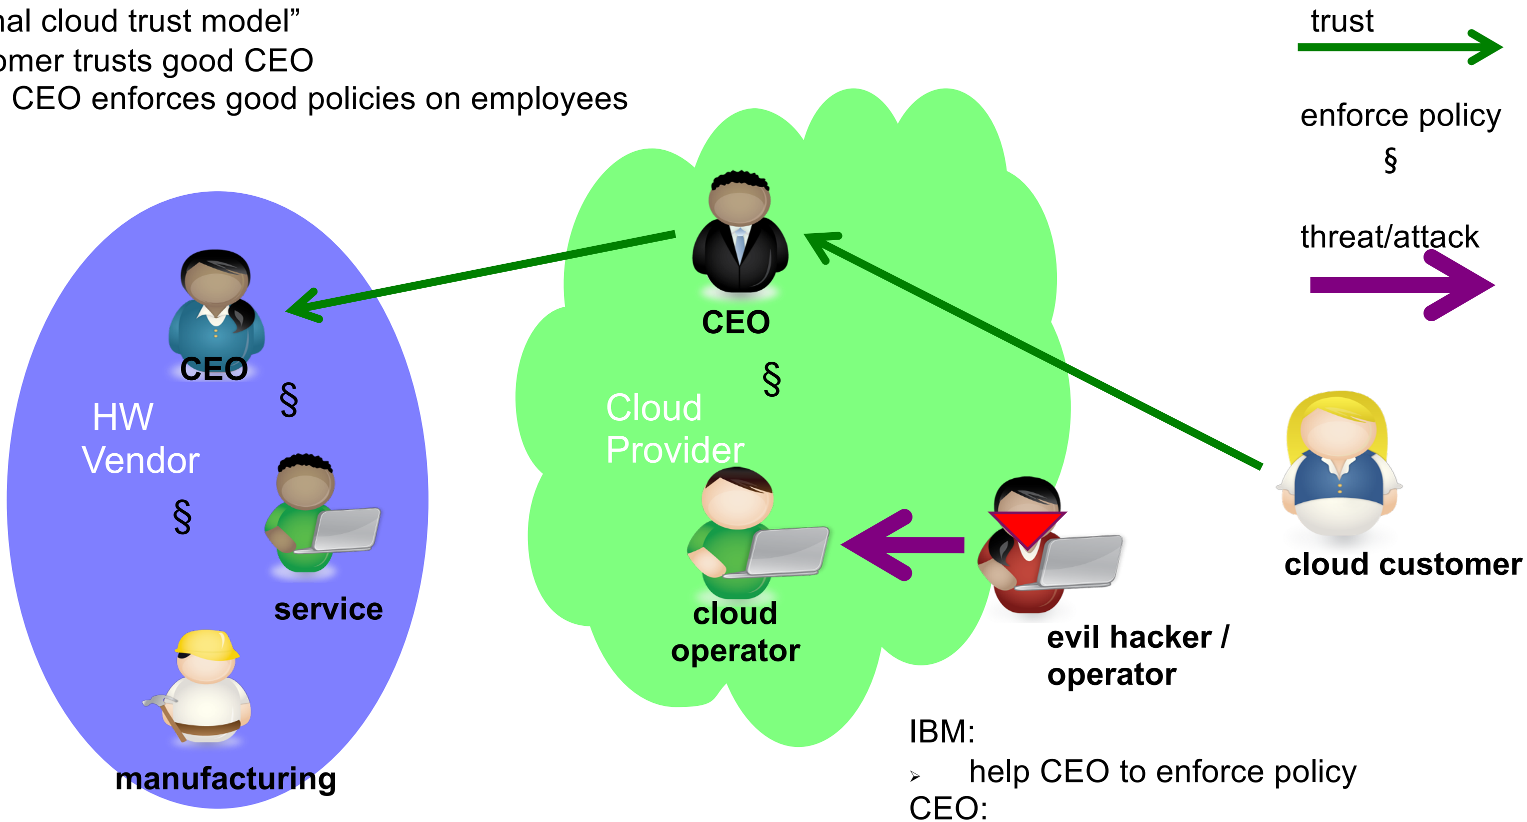

IBM:

- help CEO to enforce policy

CEO:

- publishes policy and means of enforcements

IBM Secure Execution for Linux

- A feature of IBM Z15 and LinuxONE III: FC 115*

- Allows to securely deploy a Linux workload in a KVM guest such that the computations inside the guest cannot be inspected by the hosting environment:
- HW console
- Hypervisor (here Linux/KVM)
- Cloud management infrastructure

- This leads to a new trust model for computing in the cloud

*) Feature code 115 is free of charge but needs to be ordered for availability

Secure Execution trust model: Trust HW vendor

"SE Trust Model"

- customer trusts HW vendor
  - CEO + manufacturing process

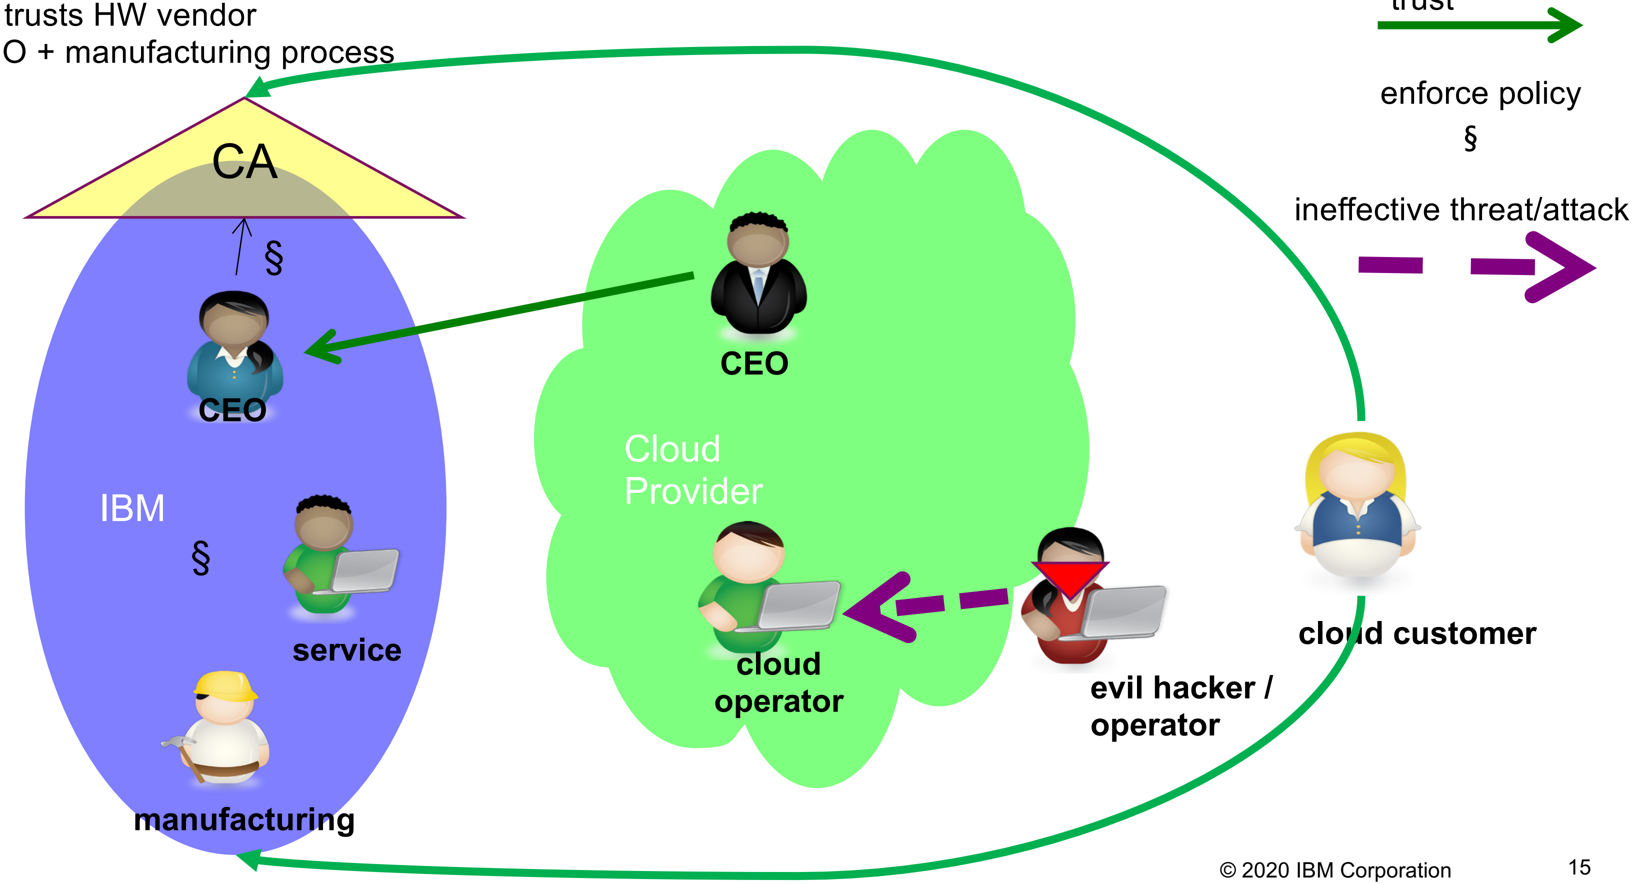

Eliminating the hypervisor's liability.

The problem

- Guest owners must trust:
  - hypervisor code
    - No malicious HV
    - No security vulnerabilities
  - HW / hypervisor management
    - Must be trustworthy
    - Non negligent
    - Implement secure access control policies
    - apply all security fixes
- Cloud providers / admins
  - cannot repudiate having done an attack
  - may be liable for a breach into a guest

The solution:

- Protect guest content from hypervisor
  - even if the hypervisor is not trustworthy or under control of an untrusted person
- The hypervisor may not see
  - the code and data that is loaded into a guest
  - the code and data inside the guest memory while the guest is running
  - any guest state in CPUs or guest control structures
- The hypervisor runs a guest as black box

Technical Background

How does it work?

- Each IBM Z15 or LinuxONE III system is associated with a **host key pair**, of which the private key is only accessible to IBM Z / LinuxONE hardware and firmware.
- A client can prepare an encrypted and integrity protected Linux image of which the secret keys are securely communicated with the help of the public host key
  - The encrypted image can only be executed in a virtual machine on the host(s) it has been prepared for
  - The image cannot be decrypted outside of the designated host(s)
  - The secure guest owner must make sure that disk and network data is encrypted (e.g., dm-crypt, TLS)
- Z hardware and firmware (ultravisor UV) ensure that unencrypted memory or processor state of a running secure execution guest cannot be accessed by the Linux KVM hypervisor (or Support Element / HMC).
- Z hardware and firmware (ultravisor UV) will detect whenever integrity of the memory of a running secure execution guest is violated (from outside the guest).

Running a Secure Execution guest in a nutshell (1)

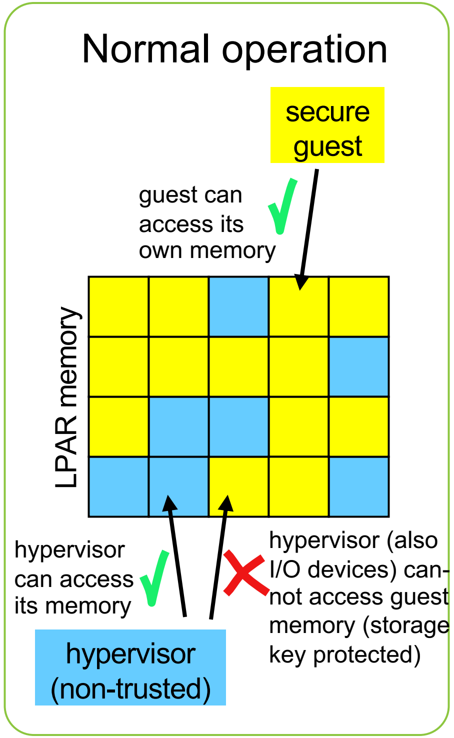

Normal operation

secure guest

guest can access its own memory ✓

LPAR memory

hypervisor can access its memory ✓

hypervisor (also I/O devices) can-not access guest memory (storage key protected) ✗

hypervisor (non-trusted)

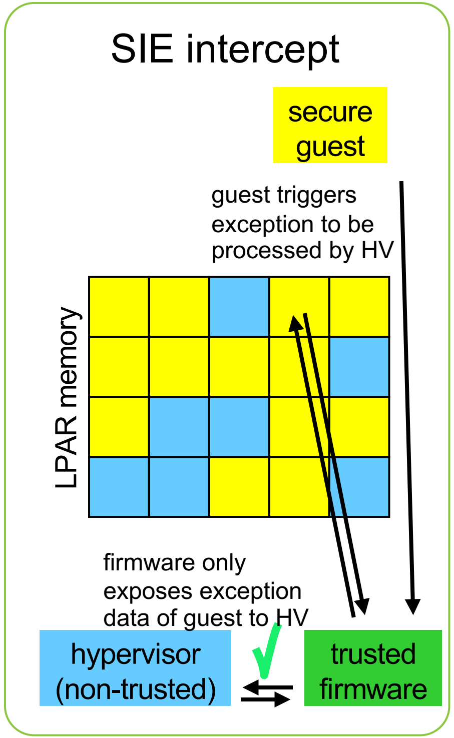

SIE intercept

secure guest

guest triggers exception to be processed by HV

LPAR memory

firmware only exposes exception data of guest to HV ✓

hypervisor (non-trusted)

trusted firmware

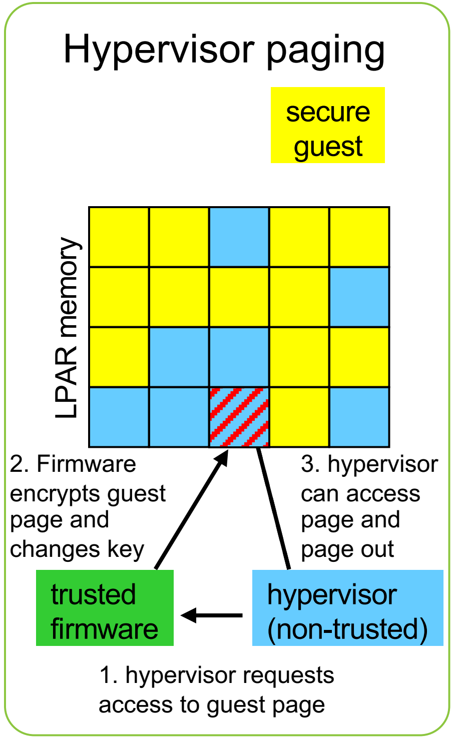

Hypervisor paging

secure guest

LPAR memory

2. Firmware encrypts guest page and changes key

3. hypervisor can access page and page out

trusted firmware

hypervisor (non-trusted)

1. hypervisor requests access to guest page

Yellow: memory of confidential guest

Blue: hypervisor memory

HV memory with encrypted contents

Running a Secure Execution guest in a nutshell (2)

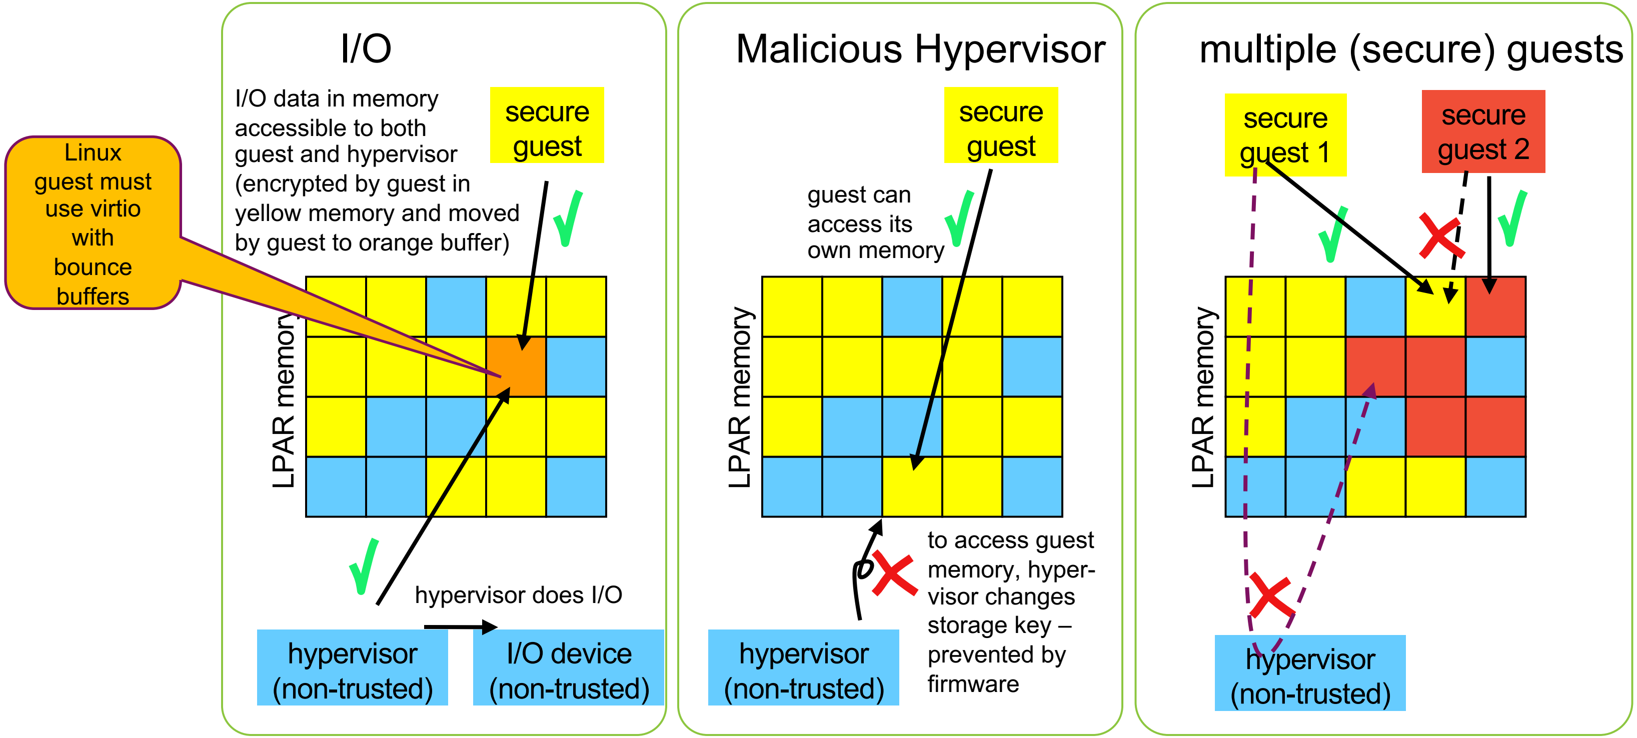

I/O

I/O data in memory accessible to both guest and hypervisor (encrypted by guest in yellow memory and moved by guest to orange buffer)

Linux guest must use virtio with bounce buffers

Malicious Hypervisor

guest can access its own memory

to access guest memory, hyper-visor changes storage key – prevented by firmware

multiple (secure) guests

Yellow: memory of secure guest (1)

Blue: hypervisor memory

Red: memory of secure guest 2

Memory accessible to host and guest

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image used to deploy may workload the one I provided?
- Are the secrets in my image kept confidential?

Guest protecti

- Is the data in r ng t prote

Secure Execution

Data at-rest protec n

- Is the data stor protect

Pervasive Encryption

Workload placement

- Is the system that hosts my workload the system I think it is hosted on?

Guest isolatio

- Is my running d from of

Secure Execution

Data data-in fligh protection

- data other sy ed

Pervasive Encryption

Deploying a Secure Guest

Deploying a secure guest
steps 1 & 2

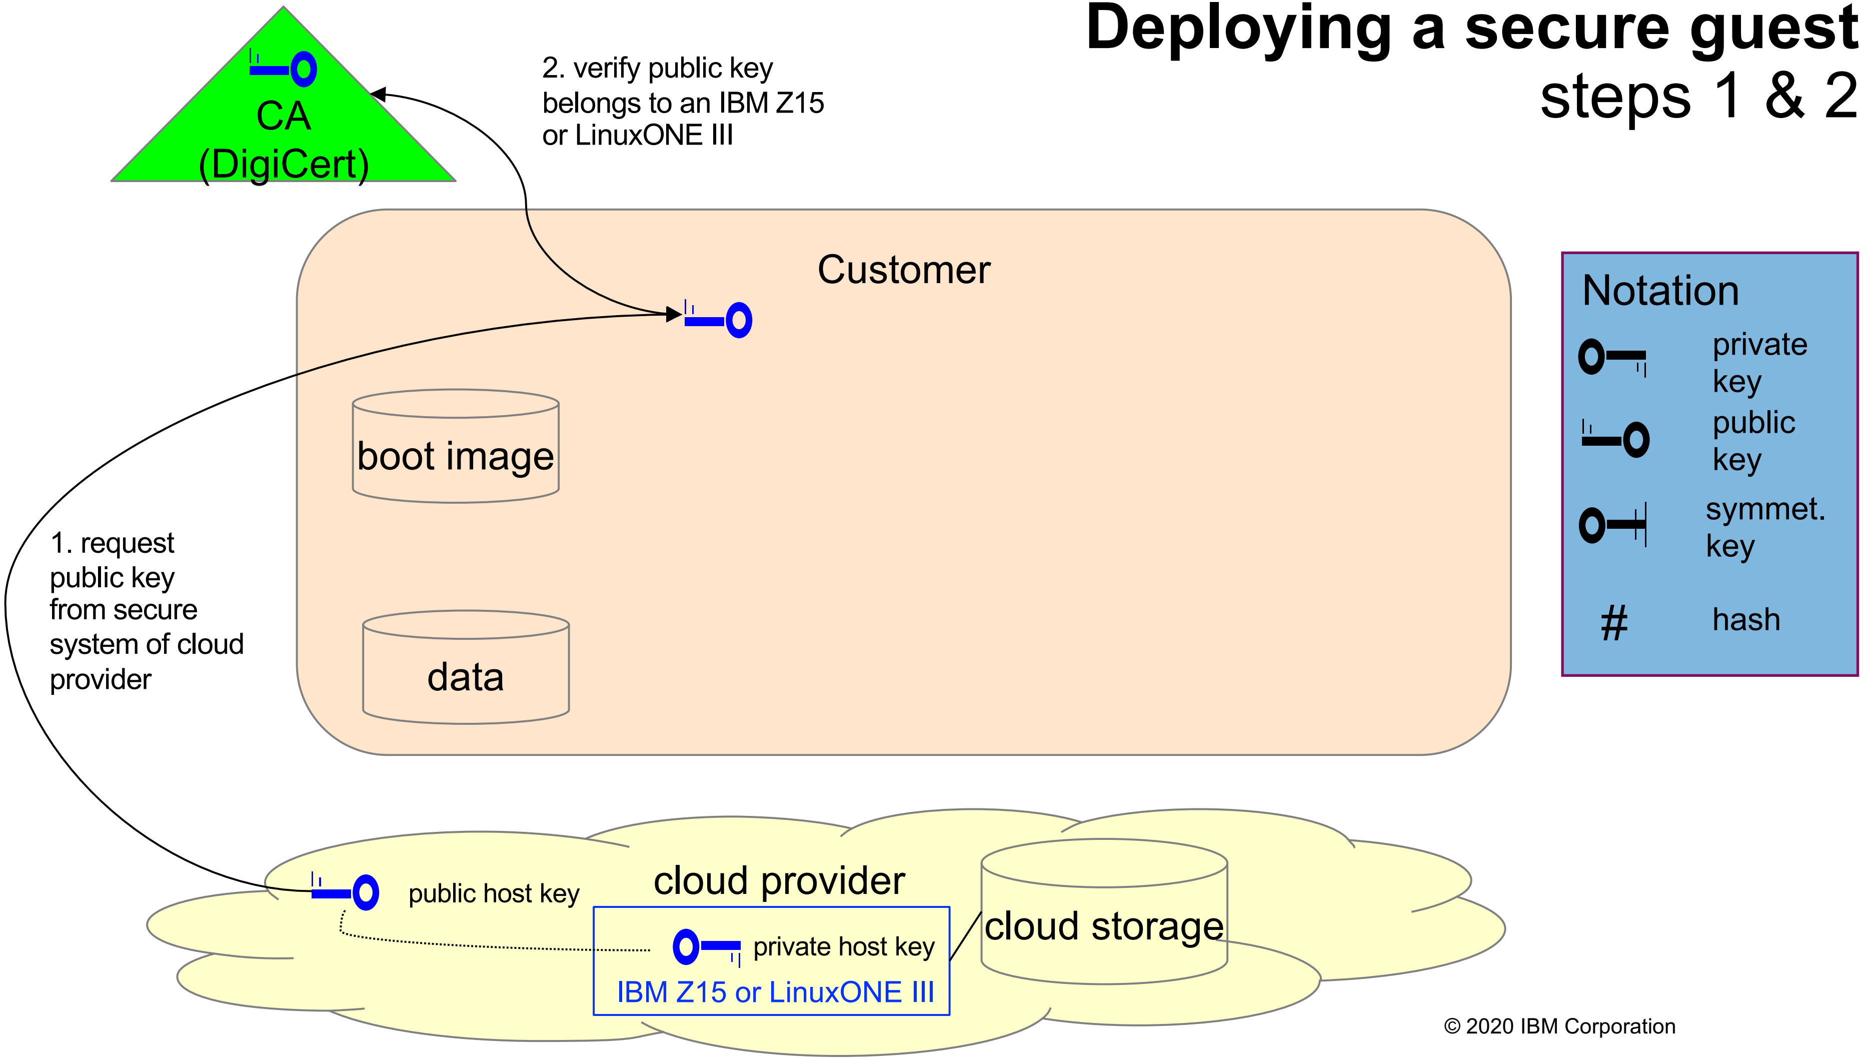

Deploying a secure guest
steps 3 & 4

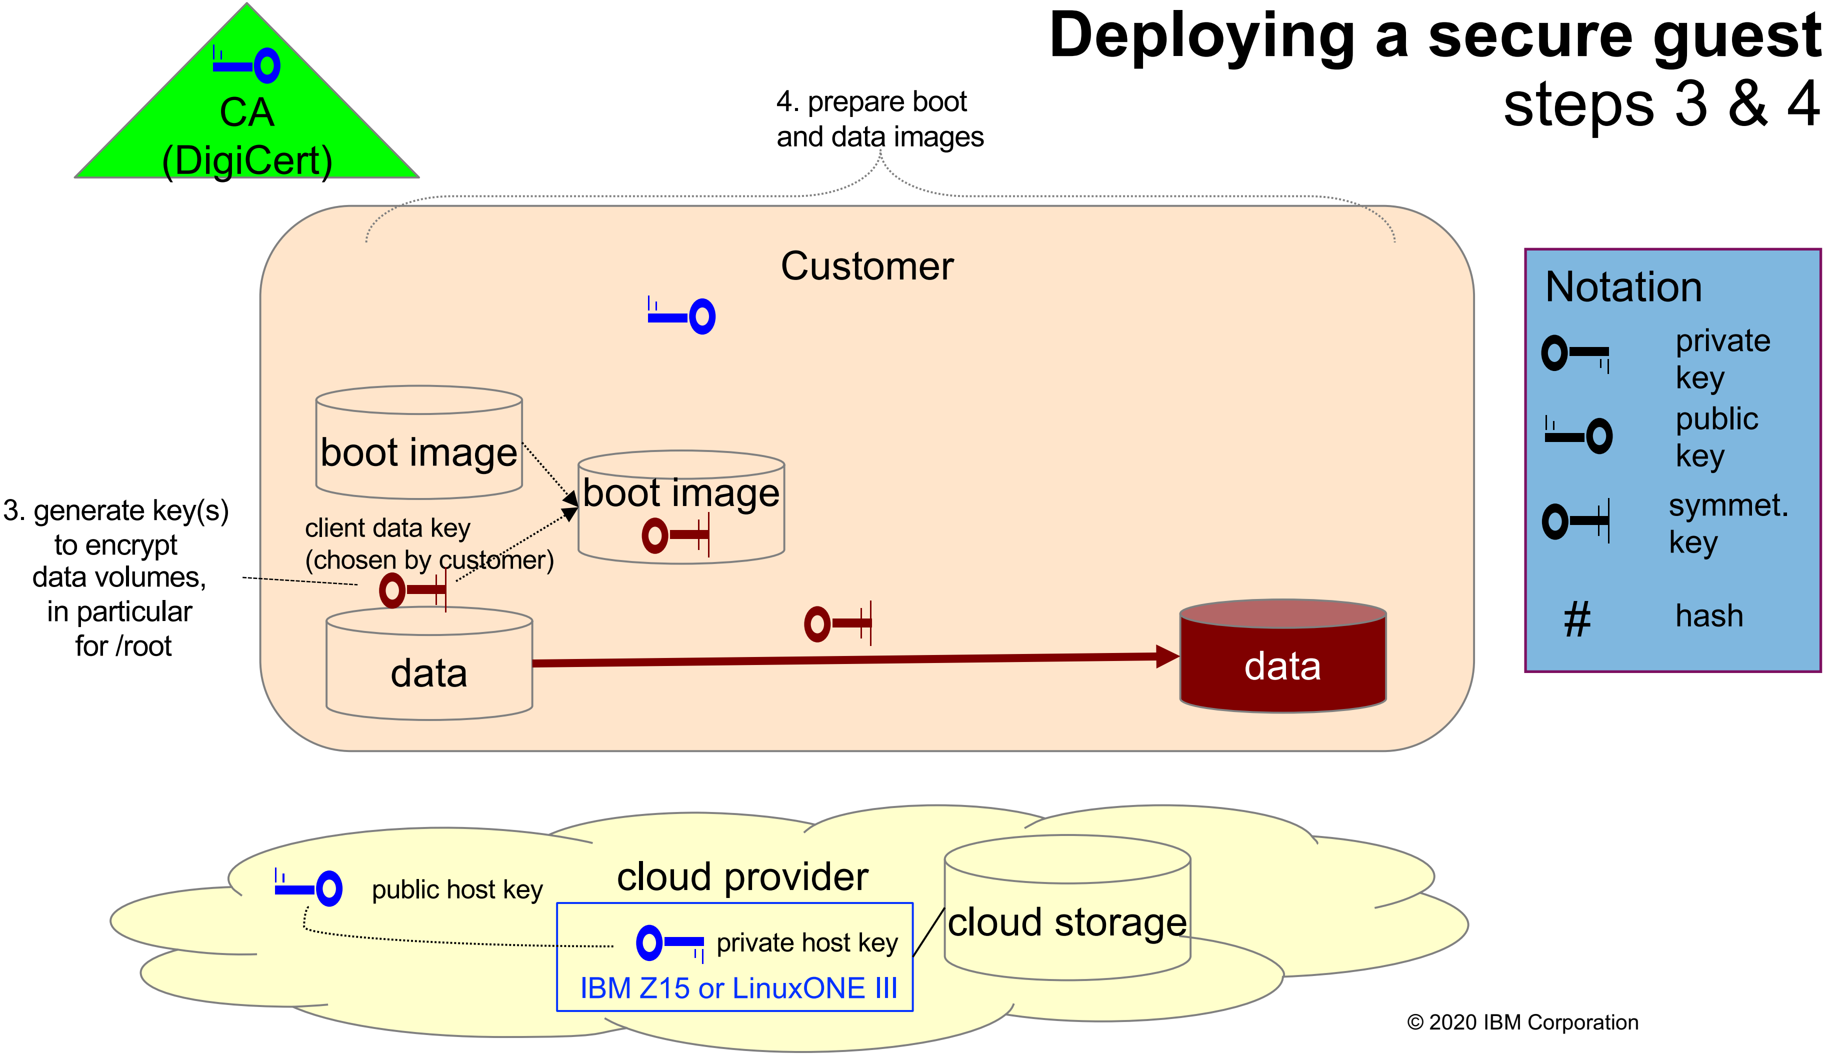

Deploying a secure guest
step 5 & 6

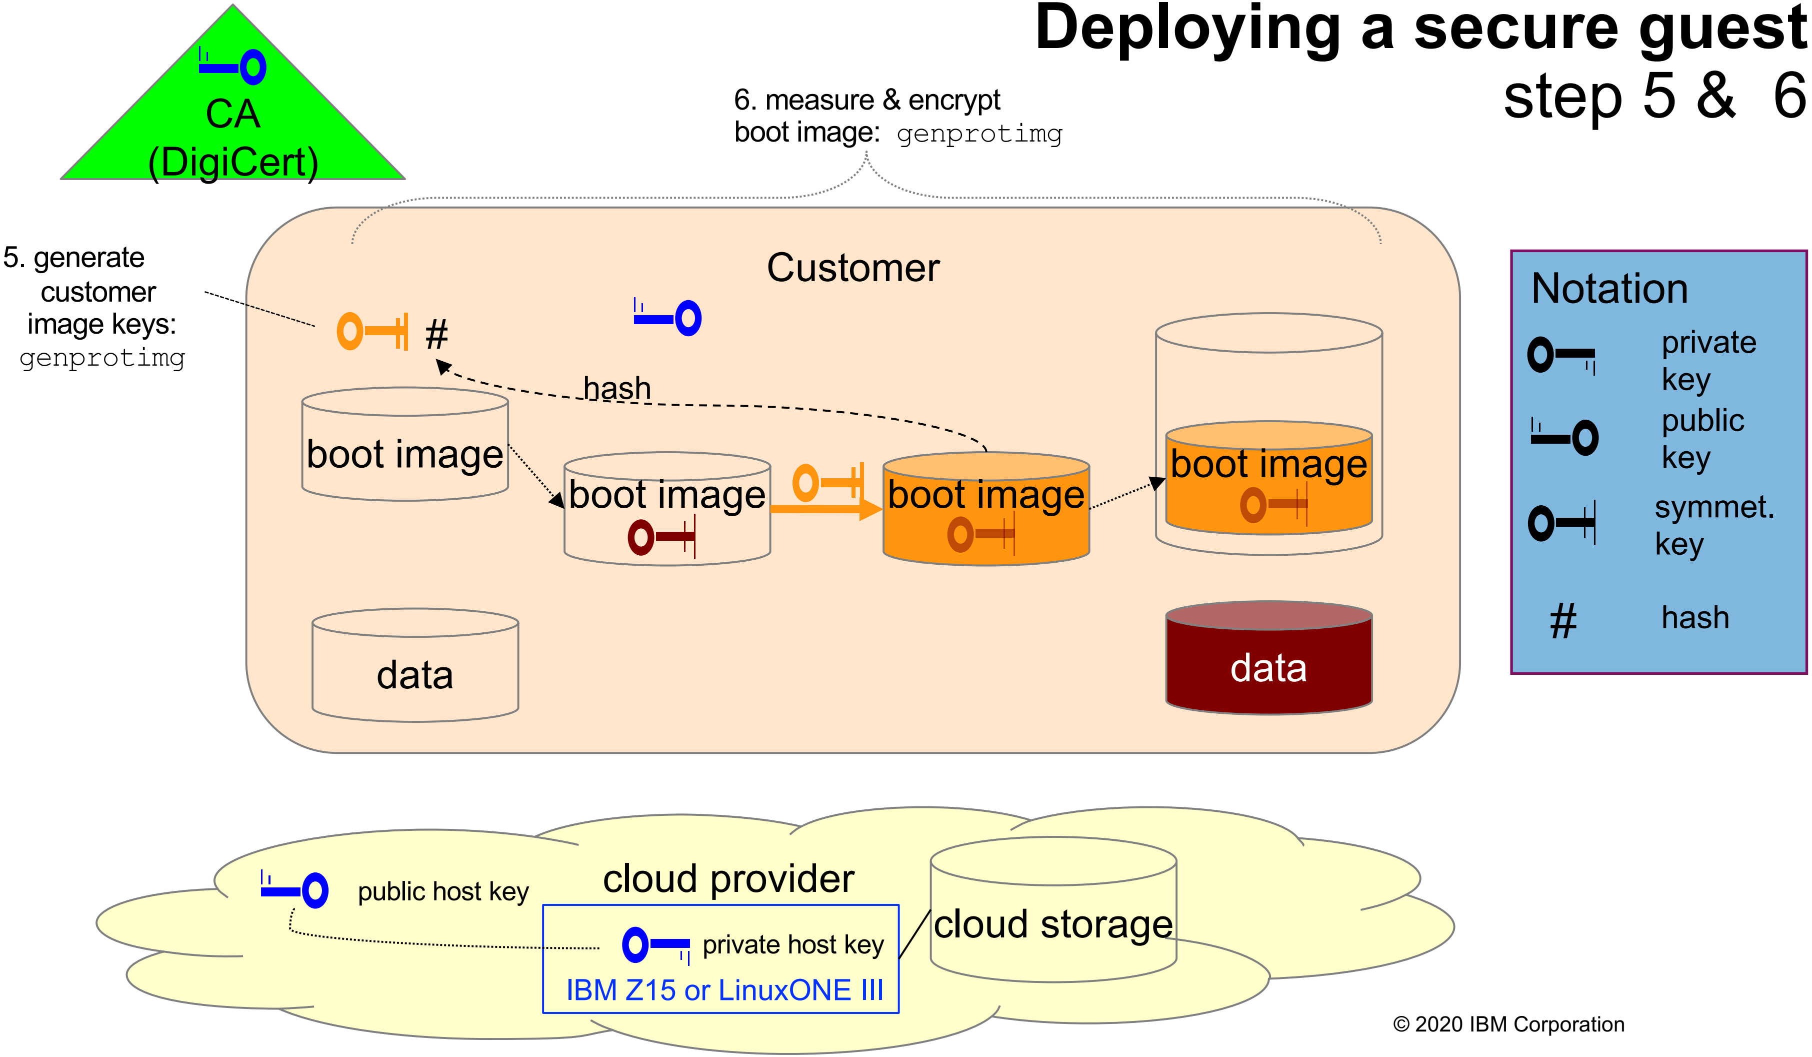

Deploying a secure guest step 7

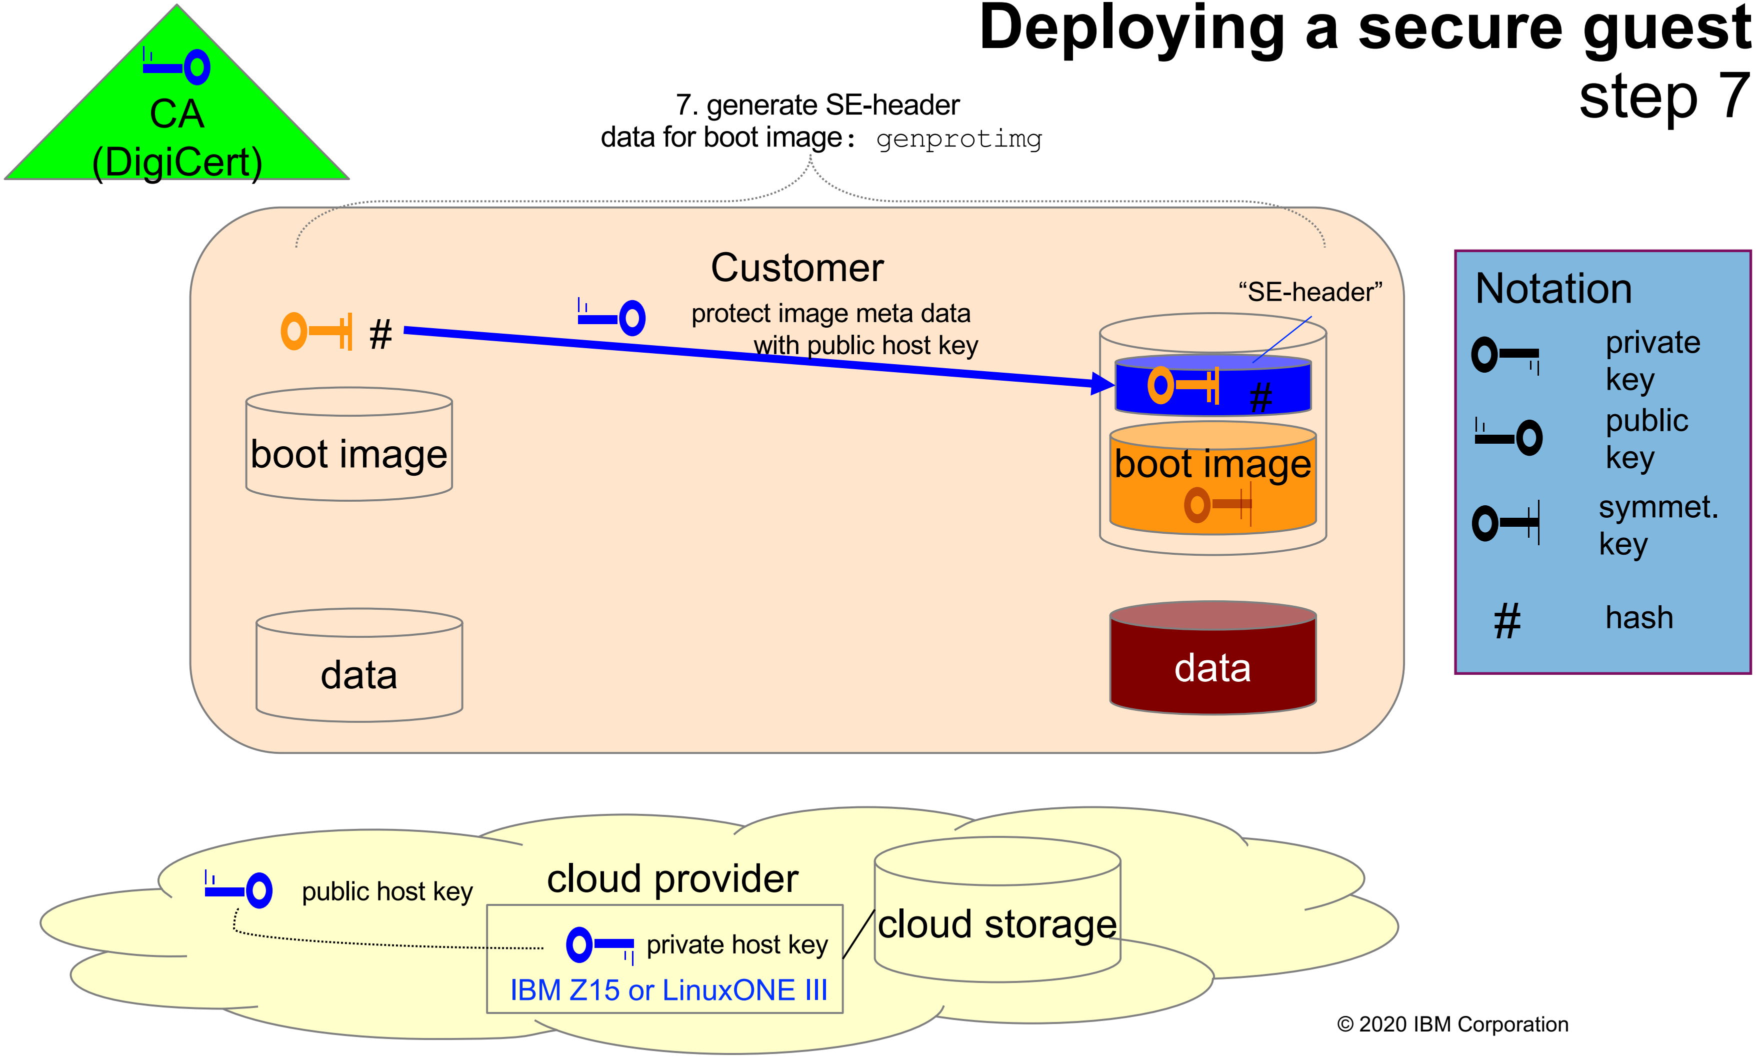

Deploying a secure guest
step 8

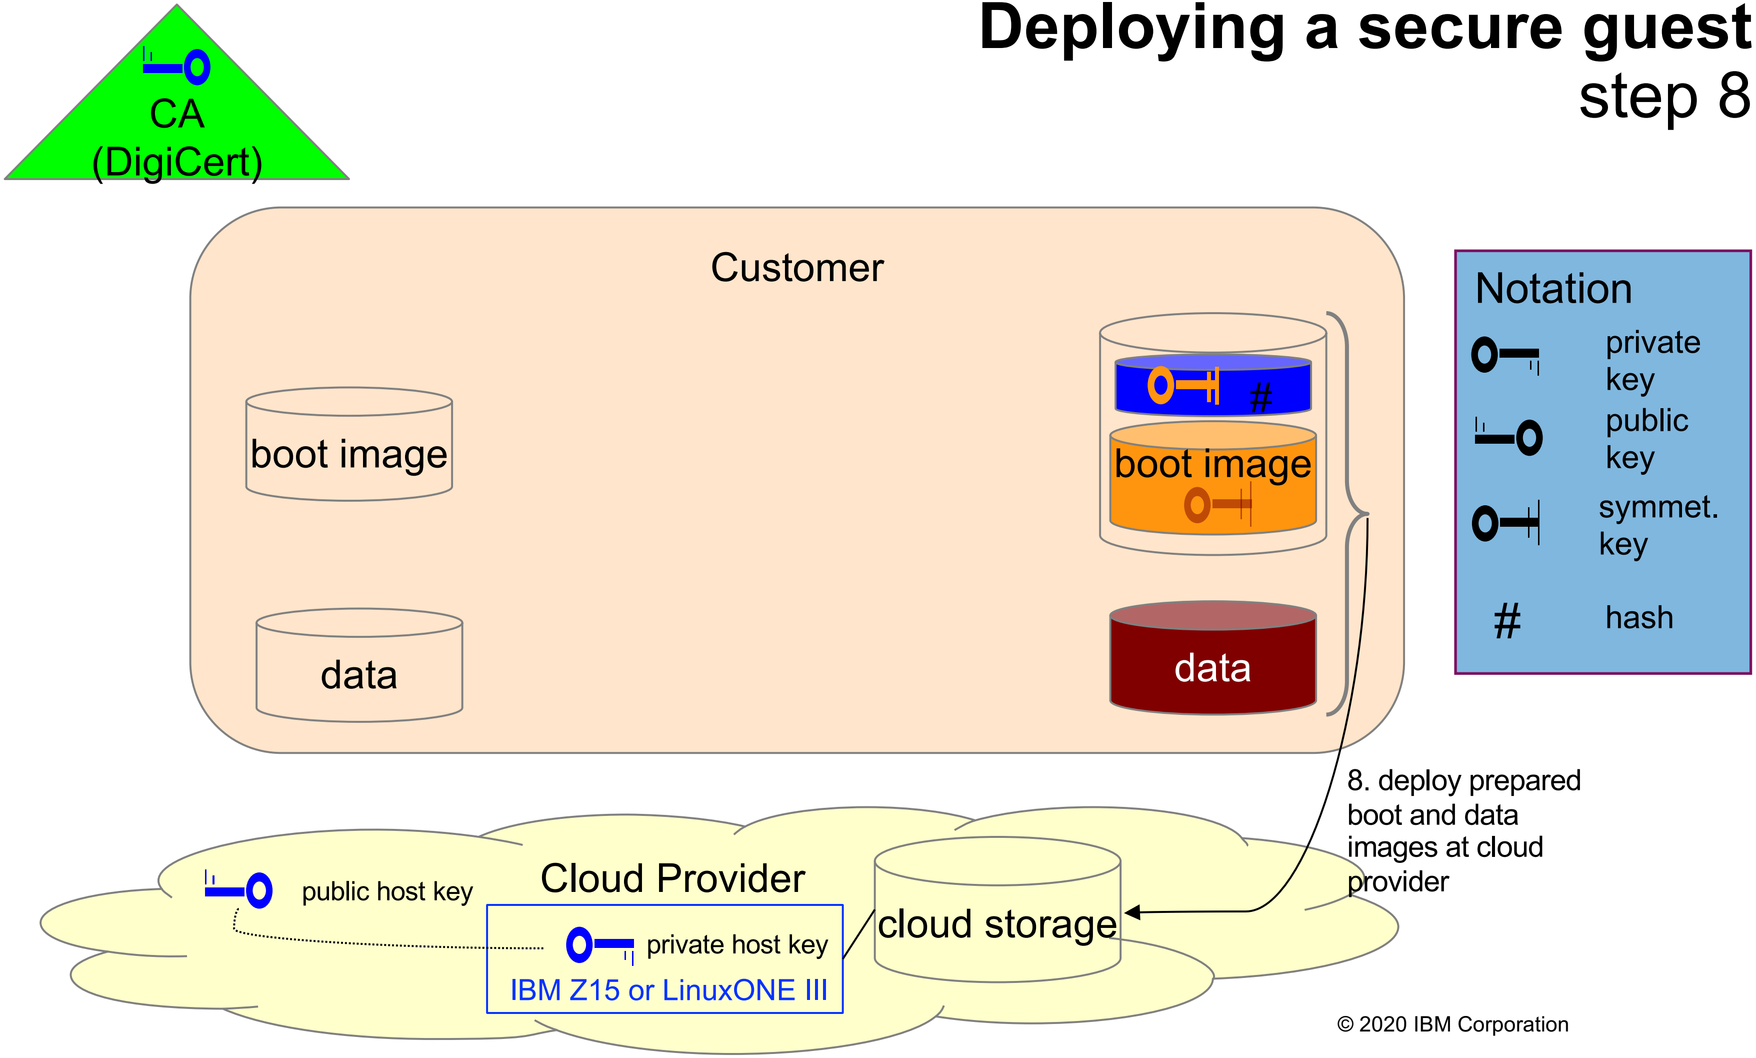

8. deploy prepared boot and data images at cloud provider

Host Keys

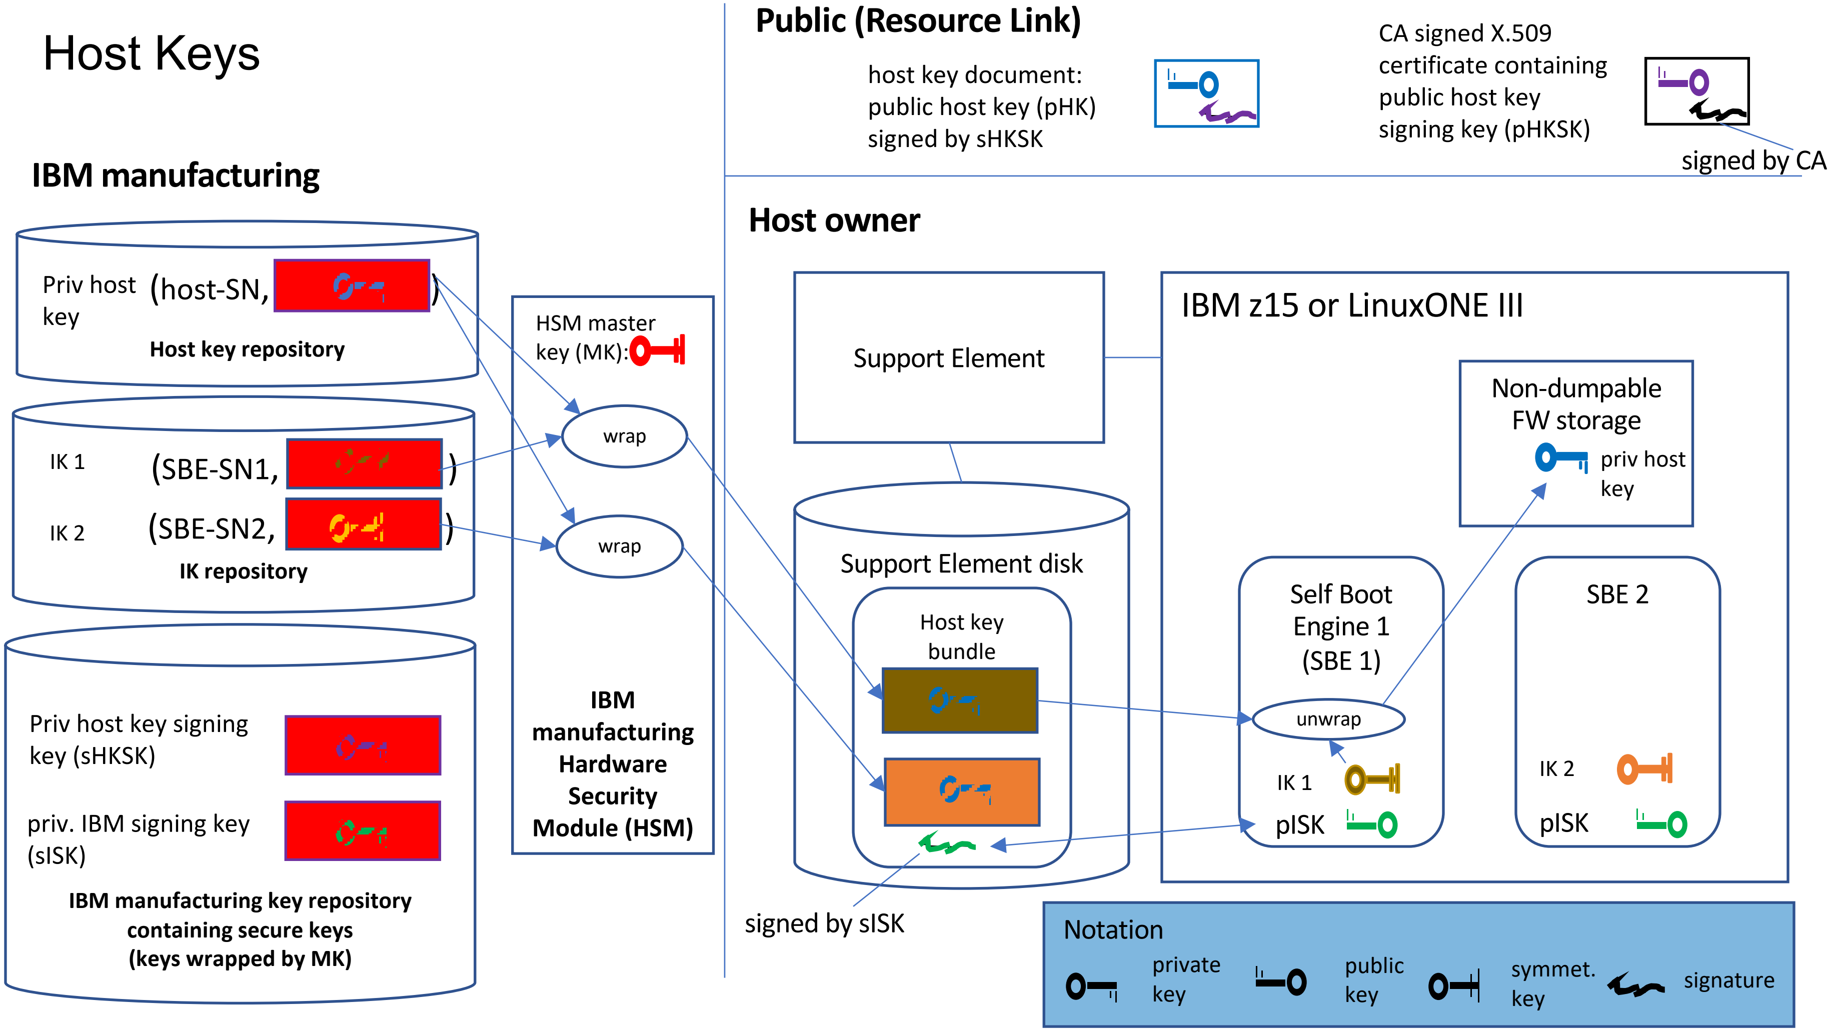

Public (Resource Link)

host key document:
public host key (pHK)
signed by sHKSK

CA signed X.509
certificate containing
public host key
signing key (pHKSK)

signed by CA

IBM manufacturing

Priv host key (host-SN,)

Host key repository

IK 1 (SBE-SN1,)

IK 2 (SBE-SN2,)

IK repository

Priv host key signing key (sHKSK)

priv. IBM signing key (sISK)

IBM manufacturing key repository containing secure keys (keys wrapped by MK)

HSM master key (MK):

wrap

wrap

IBM manufacturing Hardware Security Module (HSM)

Host owner

Support Element

Support Element disk

Host key bundle

signed by sISK

IBM z15 or LinuxONE III

Non-dumpable FW storage

priv host key

Self Boot Engine 1 (SBE 1)

unwrap

IK 1

pISK

SBE 2

IK 2

pISK

Notation

private key

public key

symmet. key

signature

How to verify a public host key?

Important task: make sure the public key you got is really a valid public host key of an IBM z15 or LinuxONE III

get host key document which includes the public host key of a target host from cloud provider or from resource link (for a given host serial number)

download list of revoked host keys from resource link: ibm-z-host-key.crl

download host key signing key certificate from resource link: ibm-z-host-key-signing.crt

download DigiCert CA certificate from resource link: DigiCertCA.crt

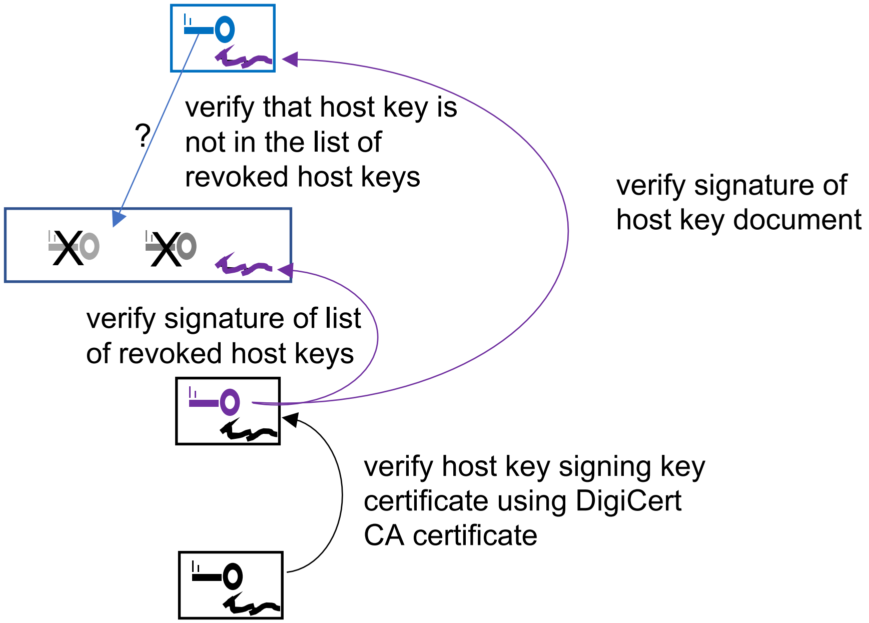

?
verify that host key is not in the list of revoked host keys
verify signature of host key document
verify signature of list of revoked host keys
verify host key signing key certificate using DigiCert CA certificate

documentation:	https://www.ibm.com/support/knowledgecenter/linuxonibm/com.ibm.linux.z.lxse/lxse_t_verify.html
sample verification script:	https://github.com/ibm-s390-tools/s390-tools/blob/master/genprotimg/samples/check_hostkeydoc

The SE-header

meta data needed to start a secure execution image

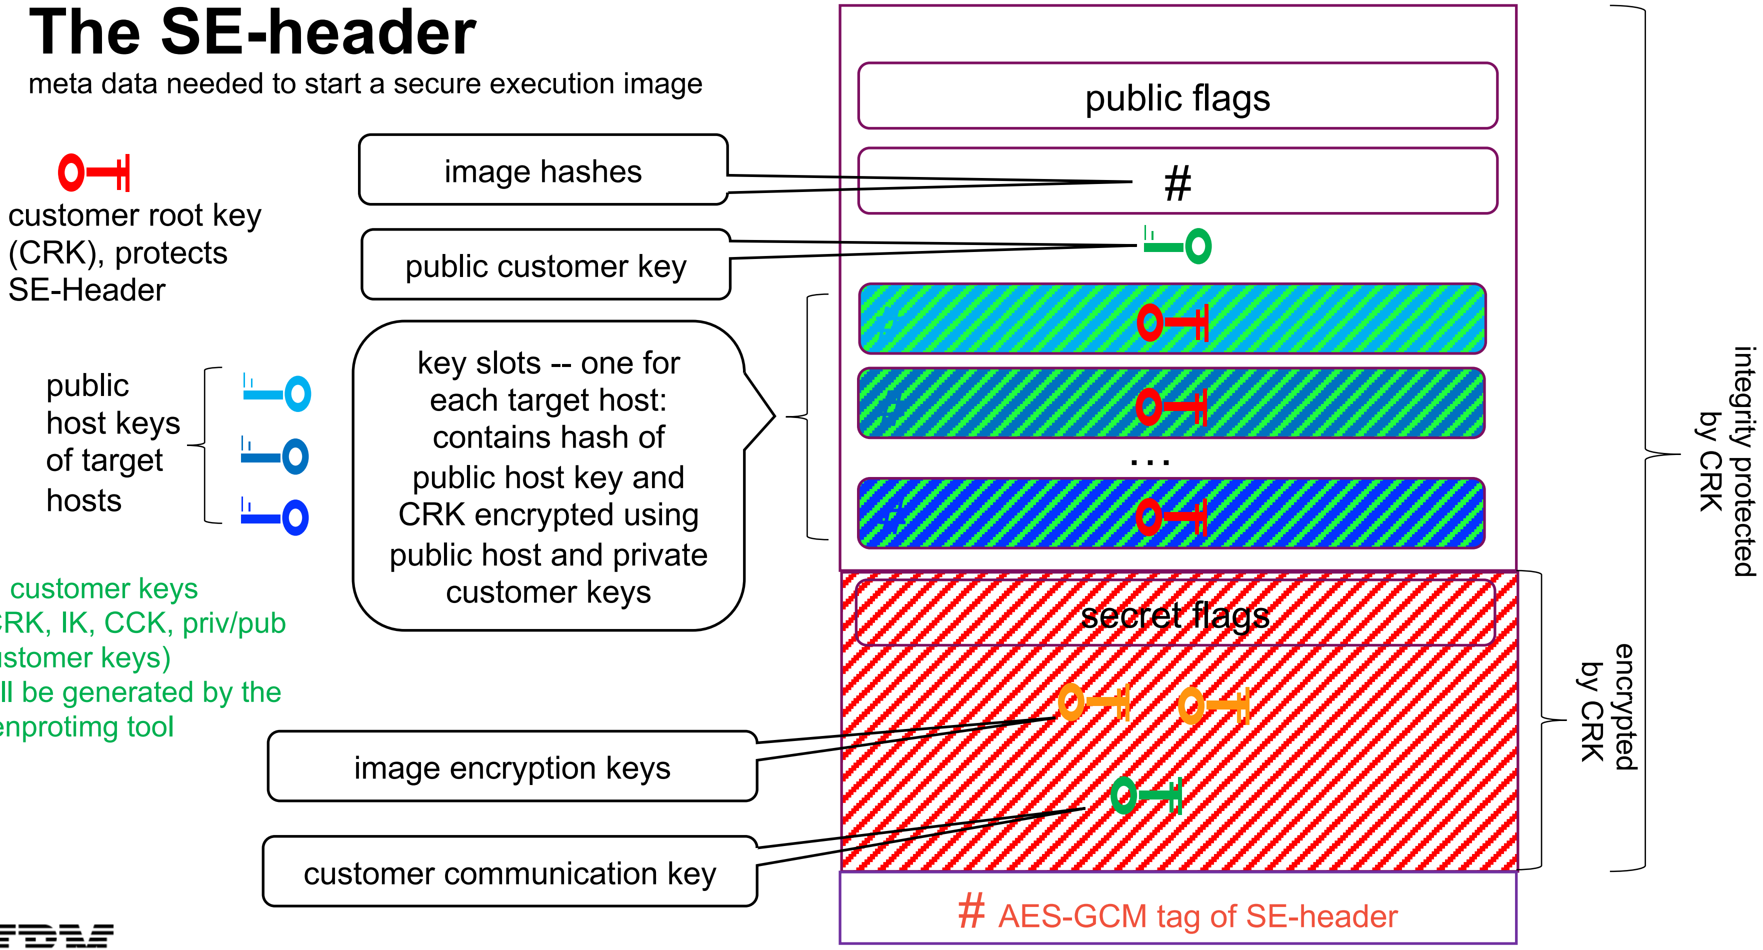

customer root key (CRK), protects SE-Header

public host keys of target hosts

image hashes

public customer key

key slots -- one for each target host: contains hash of public host key and CRK encrypted using public host and private customer keys

image encryption keys

customer communication key

public flags

#

. . .

secret flags

# AES-GCM tag of SE-header

integrity protected by CRK

encrypted by CRK

all customer keys (CRK, IK, CCK, priv/pub customer keys) will be generated by the genprotimg tool

The SE-header

meta data needed to start a secure execution image

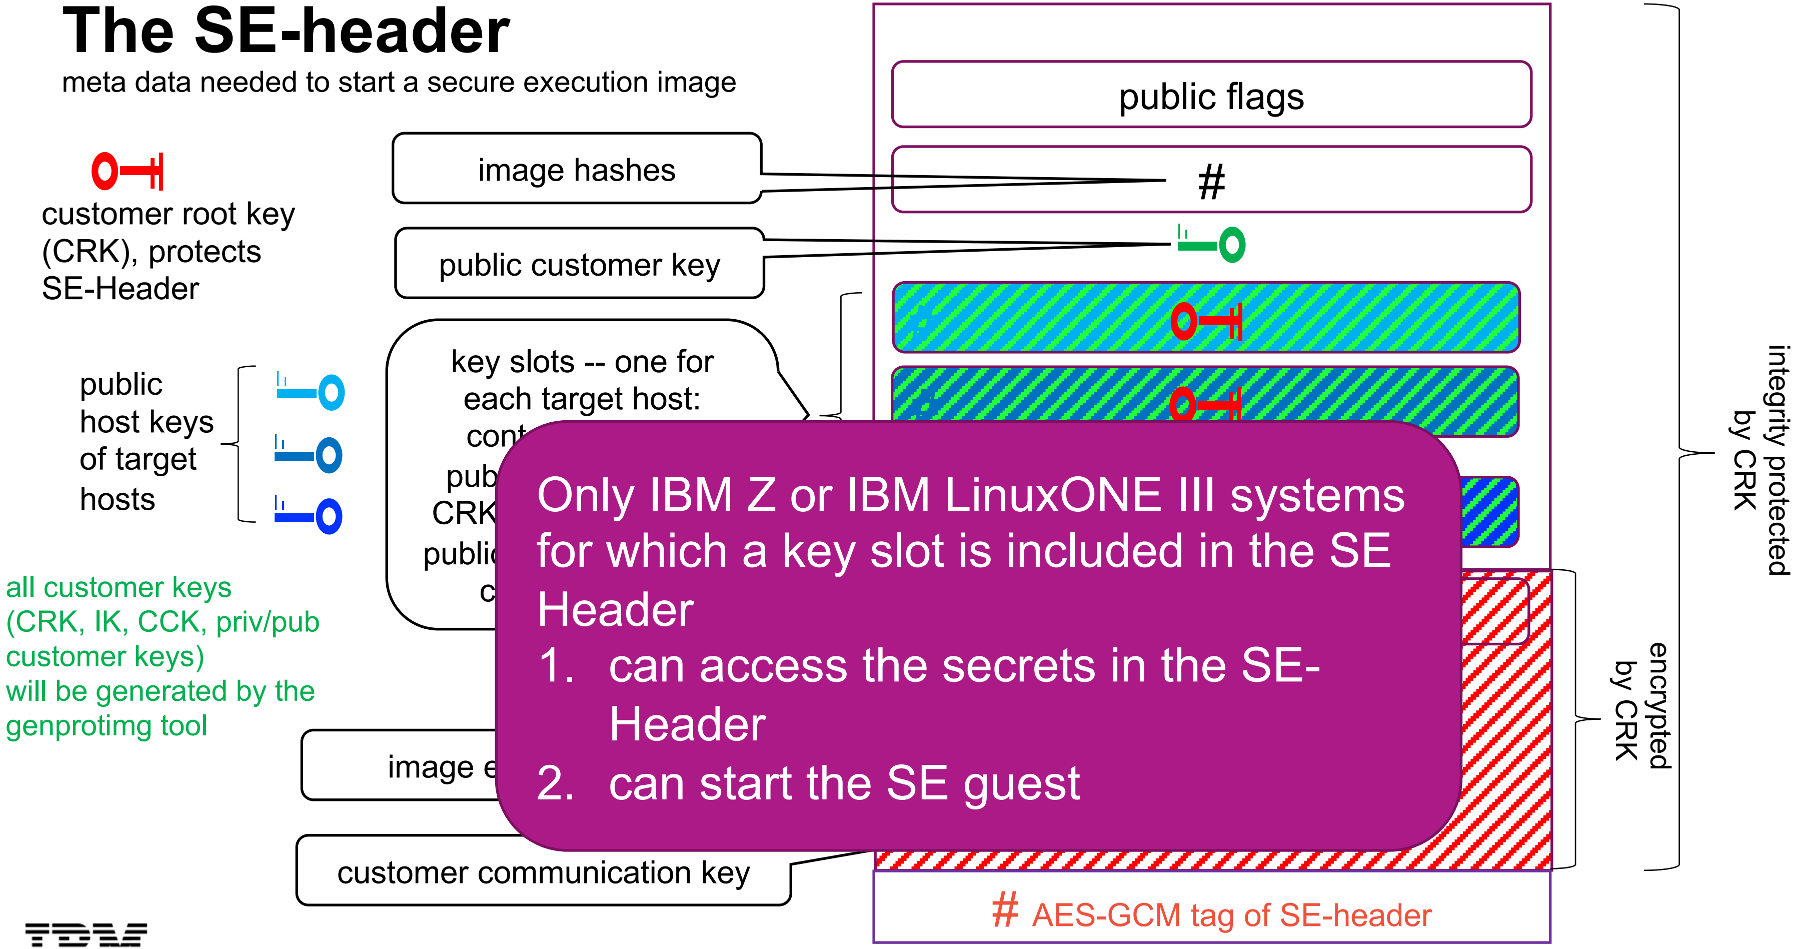

customer root key (CRK), protects SE-Header

image hashes

public customer key

public host keys of target hosts

key slots -- one for each target host:
cont
pub
CRK
public
c

all customer keys (CRK, IK, CCK, priv/pub customer keys) will be generated by the genprotimg tool

image e

customer communication key

Only IBM Z or IBM LinuxONE III systems for which a key slot is included in the SE Header

1. can access the secrets in the SE-Header
2. can start the SE guest

Security concerns when deploying workloads in the cloud?

Image protection

- Is the image used to deploy may workload the one I provided?
- Are the secrets in my image kept confidential?

Guest protecti

- Is the data in r ng
  t prote

Secure Execution

Data at-rest protec n

- Is the data stor
  protect

Pervasive Encryption

Workload placem t

- Is the system t s
  rkloa stem
  I on?

Secure Execution

Guest isolatio

- Is my running
  d fro f

Secure Execution

Data data-in fligh protectio

- data
  other
  sy ted

Pervasive Encryption

IBM

Starting a secure guest

- When booting a secure guest QEMU/KVM passes the
  - SE-header and
  - the encrypted image
- to the firmware (ultravisor).

- The ultravisor then
  - searches in the SE-header for a key slot matching its host key
  - checks the integrity of the key slot
  - checks the integrity of the SE-header
  - checks integrity of the encrypted image
  - decrypts the image
- only starts secure guest if all integrity checks were passed.

Security concerns when deploying workloads in the cloud?

Image protectio

- Is the image us may the
- Are in my image confidential?

Secure Execution

Guest protectio

- Is the data in m g protec

Secure Execution

Data at-rest protec n

- Is the data stor protect

Pervasive Encryption

Workload placem

- Is the system t s rkloa tem on?

Secure Execution

Guest isolatio

- Is my running c d from r

Secure Execution

Data data-in flig protection

- data other sy ted

Pervasive Encryption

Linux support for IBM Secure Execution for Linux

Linux Support for IBM Secure Execution

- hypervisor: Linux/KVM
  - upstream: Linux kernel 5.7, upcoming QEMU version 5.1
  - supporting host distributions: Ubuntu 20.04, SLES 15 SP2
- secure guests
  - supporting guest distributions: RHEL 7.8, 8.1 SLES 12 SP5, 15 SP2, Ubuntu 20.04
  - supported devices: sclp, virtio-blk, virtio-scsi, virtio-net, virtio-serial
    - with bounce buffers enabled (`iommu = ‘on’`)
  - to establish security, all data communicated to or from the secure guest must be explicitly encrypted inside the guest
- image preparation tools: s390 tools 2.13
  - genprotimg
  - check_hostkeydoc

What does IBM Secure Execution for Linux protect?

What Secure Execution shall protect against?

Guest data theft or corruption due to

- bad operation of the HW Console (Support Element / HMC) by rogue HW admins
- bad operation of HV by rogue/negligent HV admins
- hacked HVs (e.g. from a neighbor guest)
- corrupt HVs

What shall Secure Execution <u>not</u> protect against?

- damage due to inappropriate physical operations (e.g. inspecting HW with oscilloscope)
- stealing memory (and inspecting its contents)
- denial of service attacks
- bad operation of the guest by guest admins
- hacking the guest through guest I/O channels
- loading infected code (viruses, worms, key loggers, ransom ware, …) into guest via network

Goal: If you operate a hosted workload according to best security practices, its data is protected by secure execution as if it is run on your own premises.

Summary

- IBM Secure Execution for Linux allows you to protect your data from access of cloud (HW or KVM) operators while being used within a running guest.
- You need to prepare your image to start it as secure guest
  - no changes to the applications required
- Then together with Pervasive Encryption your data is protected everywhere:
  - in-use
  - in-flight
  - at-rest

Backup

Documentation

https://www.ibm.com/support/knowledgecenter/linuxonibm/com.ibm.linux.z.lxse/lxse_t_secureexecution.html

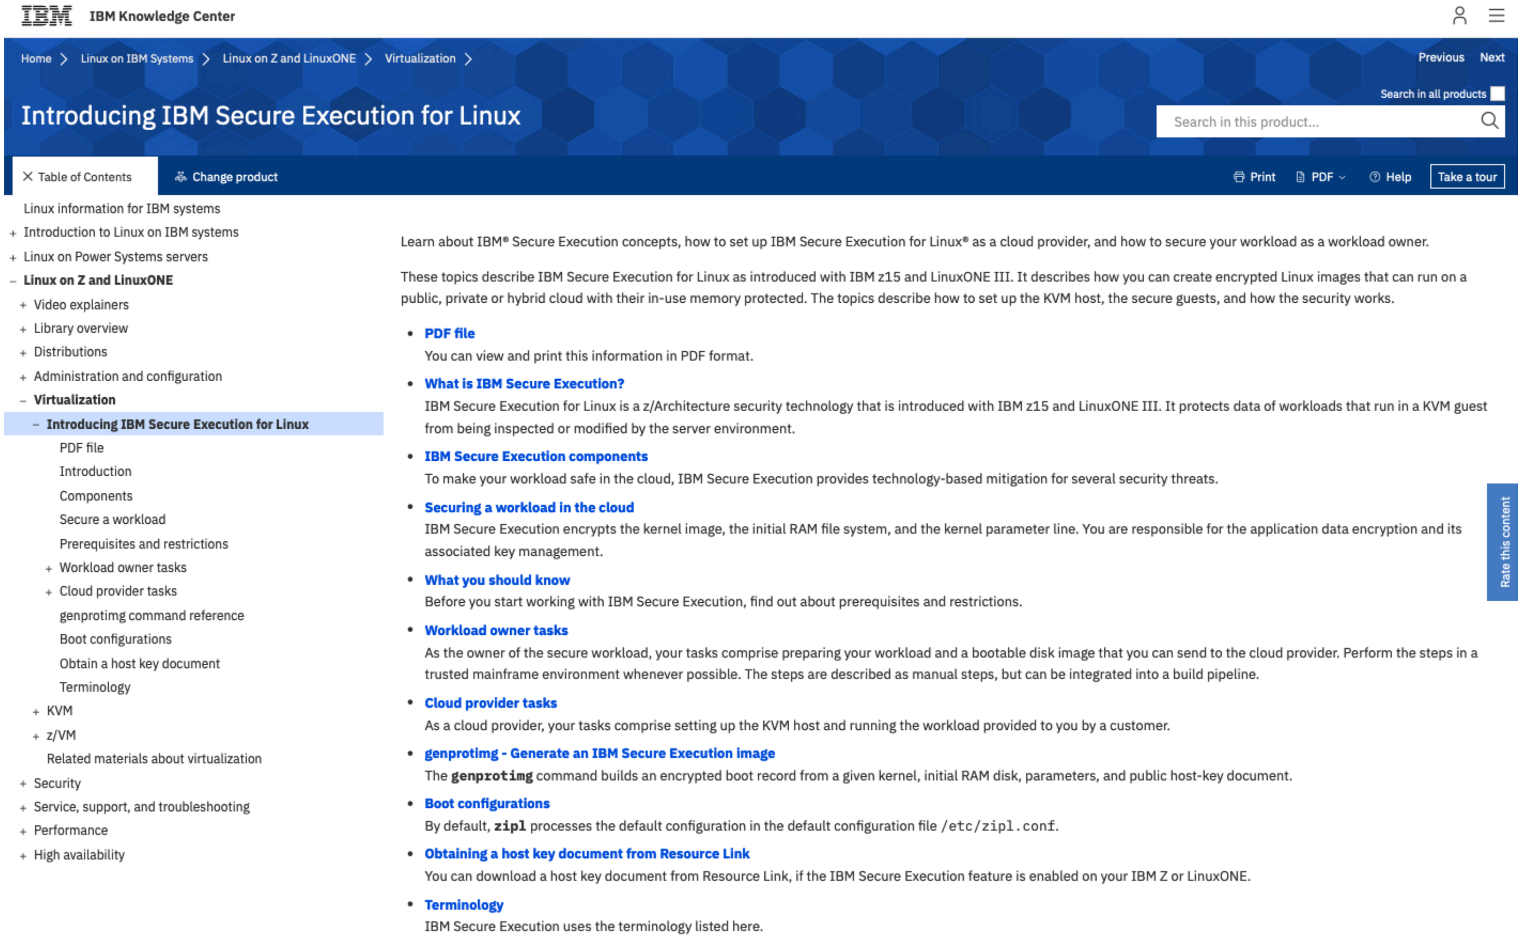

Introducing IBM Secure Execution for Linux

Linux information for IBM systems
+ Introduction to Linux on IBM systems
+ Linux on Power Systems servers
− **Linux on Z and LinuxONE**
  + Video explainers
  + Library overview
  + Distributions
  + Administration and configuration
  − **Virtualization**
    − **Introducing IBM Secure Execution for Linux**
      PDF file
      Introduction
      Components
      Secure a workload
      Prerequisites and restrictions
      + Workload owner tasks
      + Cloud provider tasks
      genprotimg command reference
      Boot configurations
      Obtain a host key document
      Terminology
    + KVM
    + z/VM
    Related materials about virtualization
  + Security
  + Service, support, and troubleshooting
  + Performance
  + High availability

Learn about IBM® Secure Execution concepts, how to set up IBM Secure Execution for Linux® as a cloud provider, and how to secure your workload as a workload owner.

These topics describe IBM Secure Execution for Linux as introduced with IBM z15 and LinuxONE III. It describes how you can create encrypted Linux images that can run on a public, private or hybrid cloud with their in-use memory protected. The topics describe how to set up the KVM host, the secure guests, and how the security works.

- **PDF file**
  You can view and print this information in PDF format.
- **What is IBM Secure Execution?**
  IBM Secure Execution for Linux is a z/Architecture security technology that is introduced with IBM z15 and LinuxONE III. It protects data of workloads that run in a KVM guest from being inspected or modified by the server environment.
- **IBM Secure Execution components**
  To make your workload safe in the cloud, IBM Secure Execution provides technology-based mitigation for several security threats.
- **Securing a workload in the cloud**
  IBM Secure Execution encrypts the kernel image, the initial RAM file system, and the kernel parameter line. You are responsible for the application data encryption and its associated key management.
- **What you should know**
  Before you start working with IBM Secure Execution, find out about prerequisites and restrictions.
- **Workload owner tasks**
  As the owner of the secure workload, your tasks comprise preparing your workload and a bootable disk image that you can send to the cloud provider. Perform the steps in a trusted mainframe environment whenever possible. The steps are described as manual steps, but can be integrated into a build pipeline.
- **Cloud provider tasks**
  As a cloud provider, your tasks comprise setting up the KVM host and running the workload provided to you by a customer.
- **genprotimg - Generate an IBM Secure Execution image**
  The **`genprotimg`** command builds an encrypted boot record from a given kernel, initial RAM disk, parameters, and public host-key document.
- **Boot configurations**
  By default, **`zipl`** processes the default configuration in the default configuration file `/etc/zipl.conf`.
- **Obtaining a host key document from Resource Link**
  You can download a host key document from Resource Link, if the IBM Secure Execution feature is enabled on your IBM Z or LinuxONE.
- **Terminology**
  IBM Secure Execution uses the terminology listed here.

Restrictions

what does not work (yet)

- life guest migration
  - offline migration works if SE-header has a key slot for the target host
- save guest to disk & restore guest from disk
- hypervisor initiated memory dump
- using huge memory pages on the host for backing guest memory.
- memory ballooning through a virtio-balloon device.
- pass-through of host devices, for example PCI, CCW, and AP (CryptoExpress).
  - note, with Secure Execution, it is OK to use clear key crypto for dm-crypt because the key or passphrase of /root can be stored in the boot partition that belongs to the encrypted image