IBM Secure Execution for Linux: Hands-on Session

—

Marc Hartmayer
Software Engineer | Linux on Z & Virtualization Development
mhartmay@de.ibm.com

Viktor Mihajlovski
Product Owner KVM on IBM Z
mihajlov@de.ibm.com

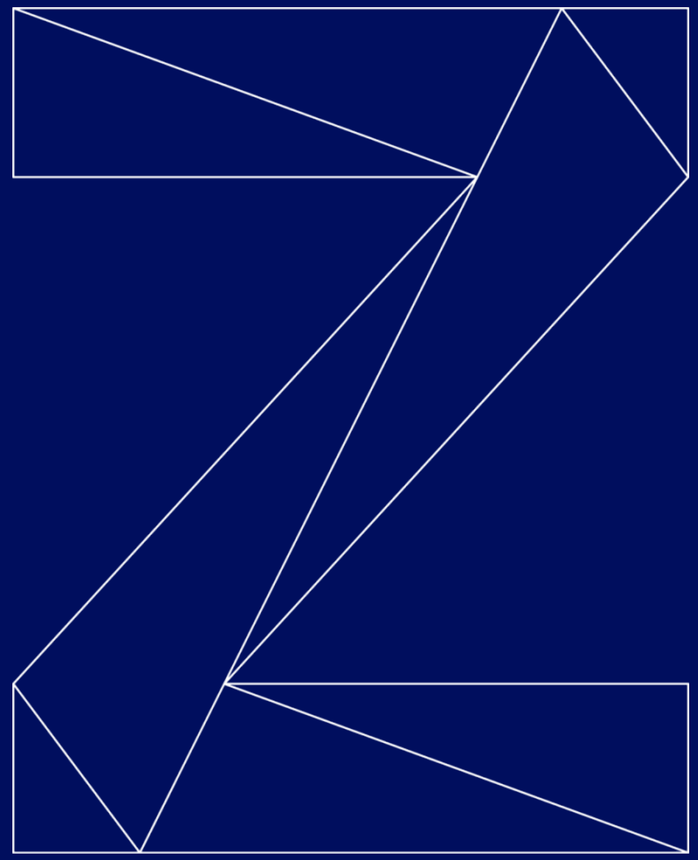

Trademarks & Disclaimer

Trademarks & Disclaimer #2

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 are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any

Notice Regarding Specialty Engines

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 (“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.

Outline

Assumptions
What you will learn
Tasks overview
Host owner tasks
Guest owner tasks
Demo

Assumptions

- You've heard of IBM Secure Execution for Linux
- You're used to Linux on IBM Z, especially
  - `zipl`, `HMC/SE`, ...
  - Disk encryption, e.g. LUKS2
- You're used to managing KVM guests
  - Libvirt: `virsh`, ...

Source: https://dam.ibm.com/l/4a8d8d20096ea1af/

What you won’t learn?

- Explanation how IBM Secure Execution works

Source: https://dam.ibm.com/l/5d3520157e88953c/

What you will learn?

- How to prepare an IBM Secure Execution host
- How to prepare an IBM Secure Execution guest

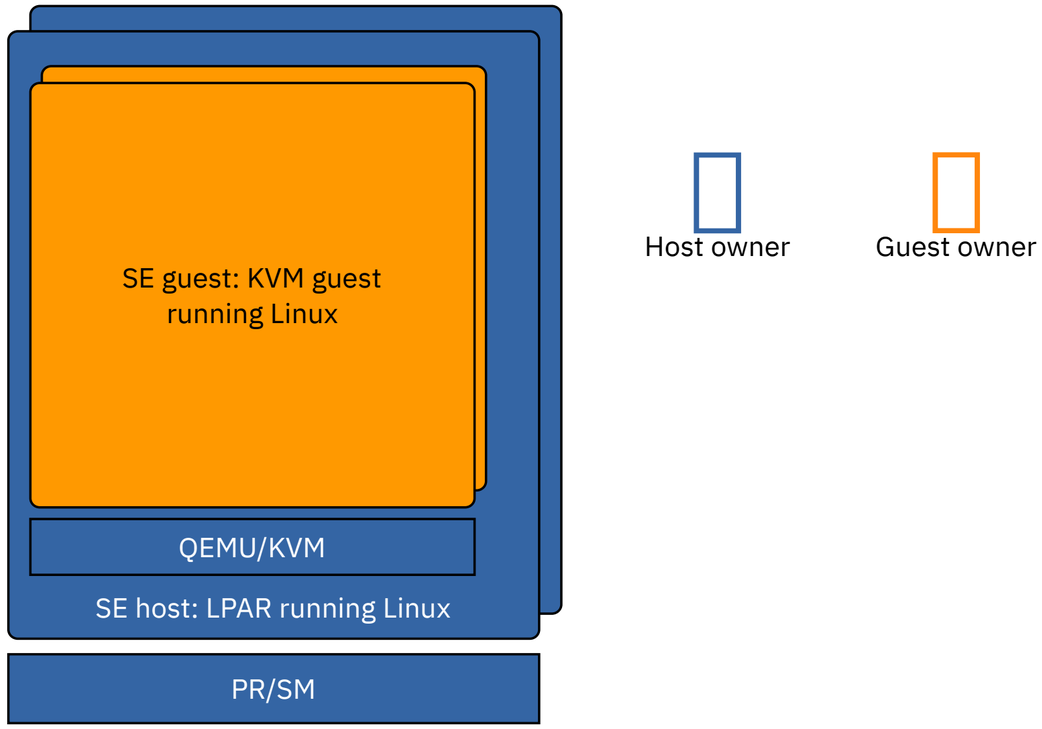

Tasks overview

1. Prepare CEC

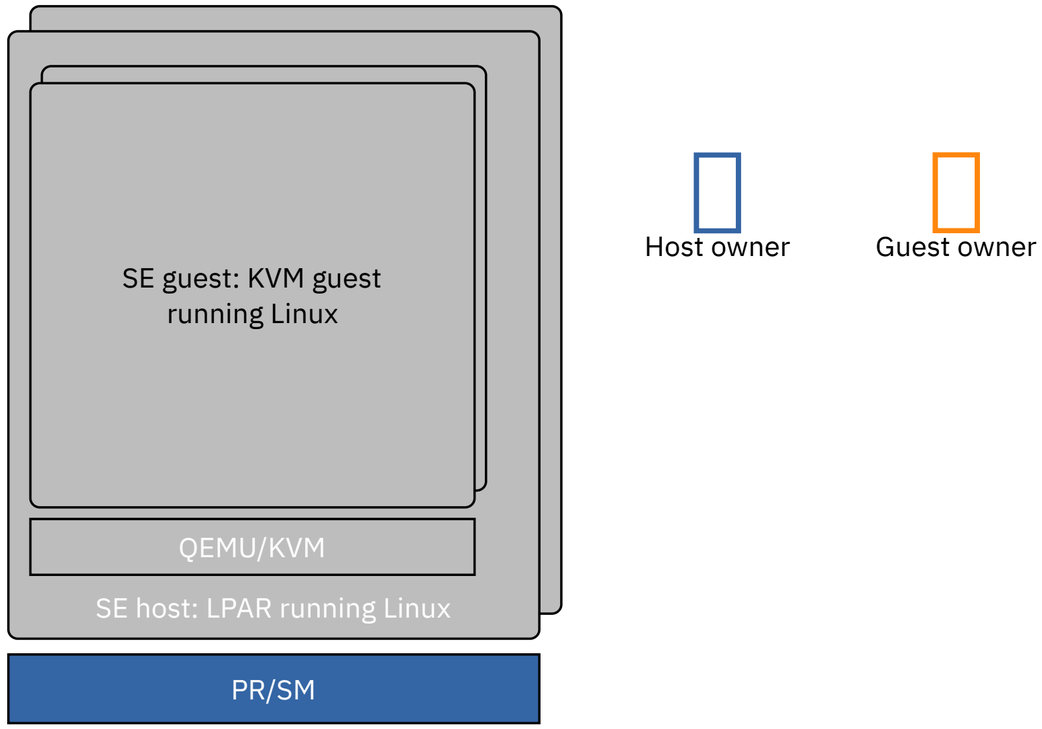

Tasks overview

1. Prepare CEC
2. Prepare KVM host

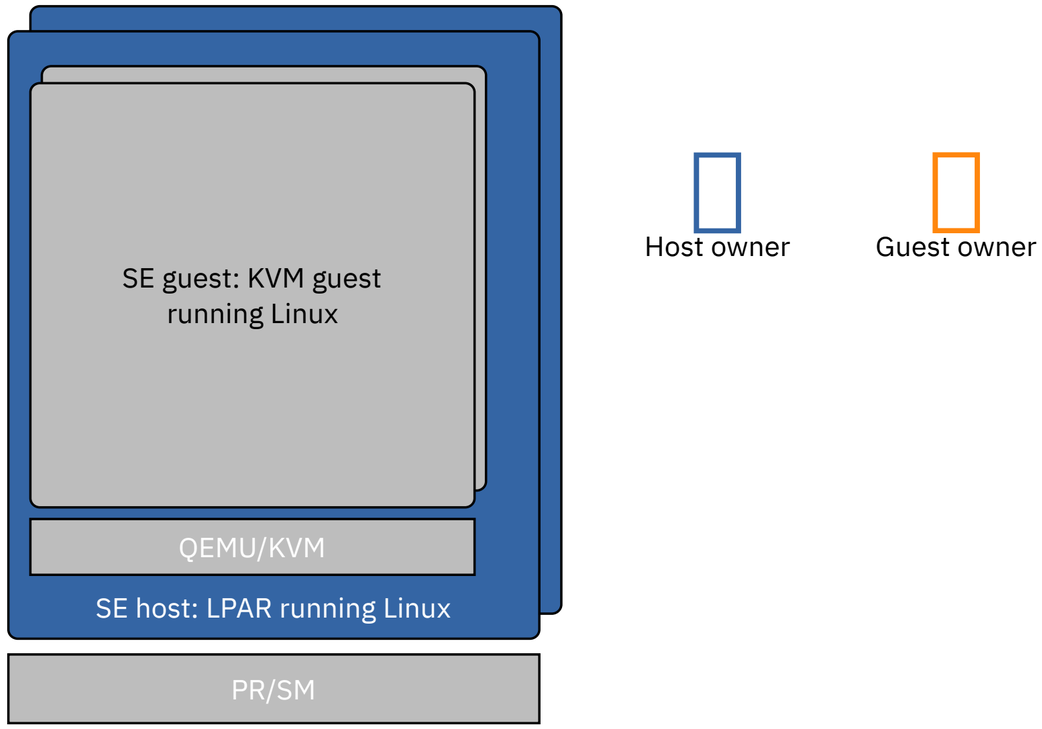

Tasks overview

1. Prepare CEC
2. Prepare KVM host
3. Prepare KVM virtual server resources:
   libvirt XML domain definition

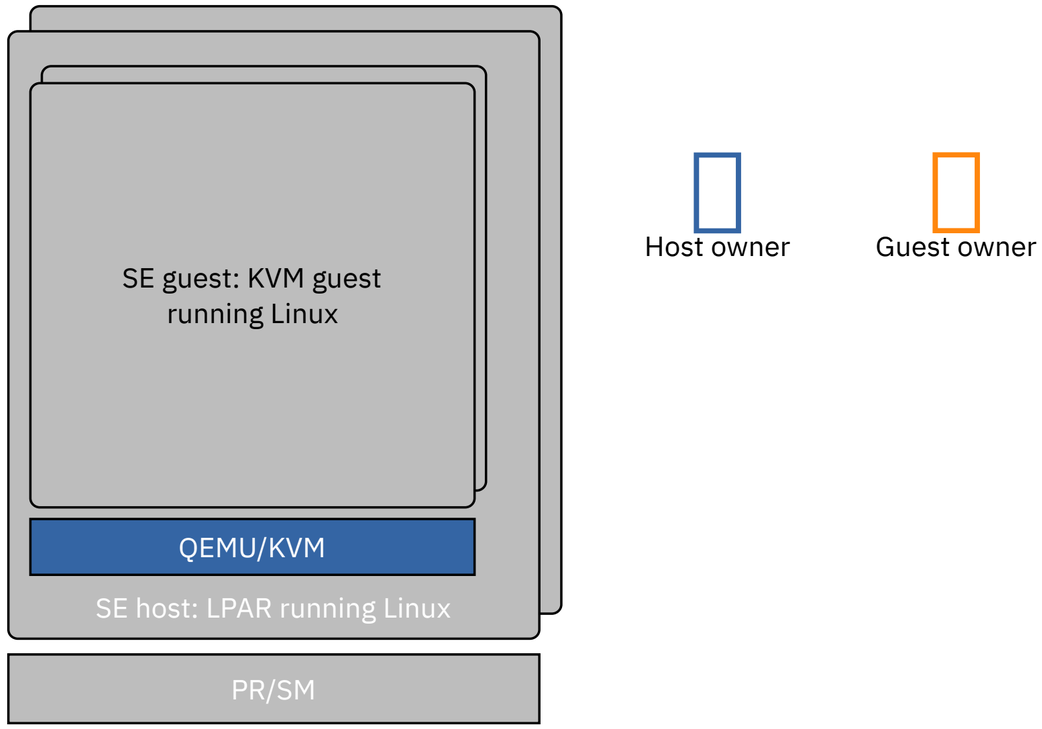

Tasks overview

1. Prepare CEC
2. Prepare KVM host
3. Prepare KVM virtual server resources:
   libvirt XML domain definition
4. Prepare KVM guest image and provide it to the
   host owner

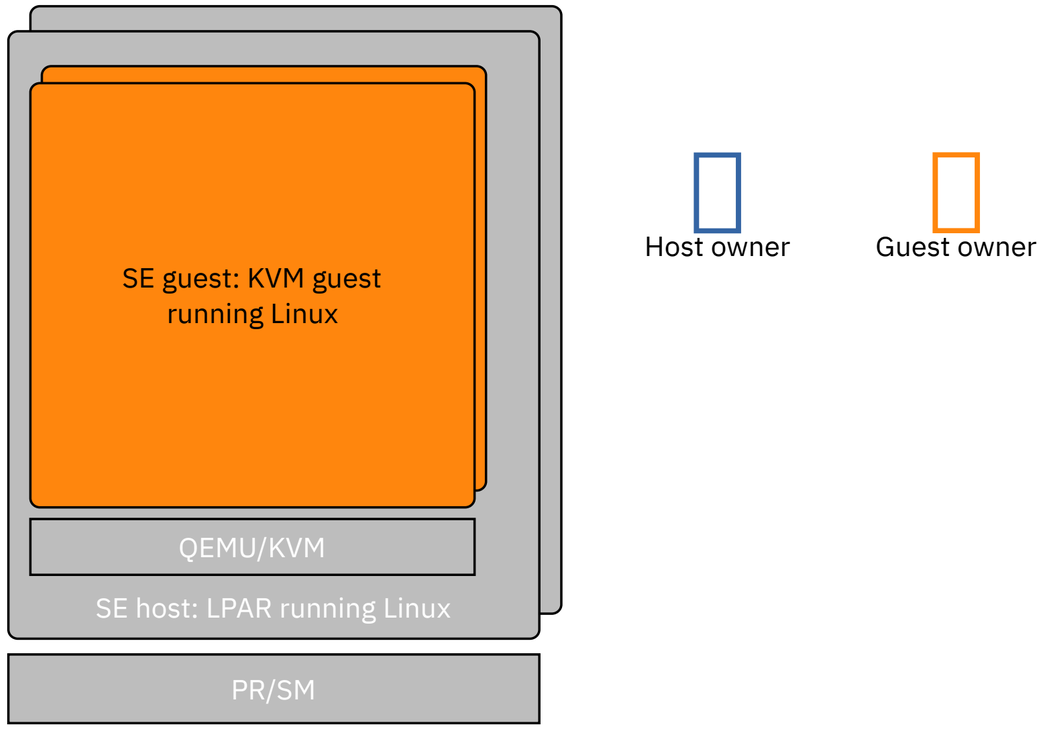

Host owner tasks

Prepare CEC

- IBM z15 or LinuxONE III with the IBM Secure Execution for Linux feature enabled
- Install IBM provided key bundles [1]

[1] See “Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721” for details

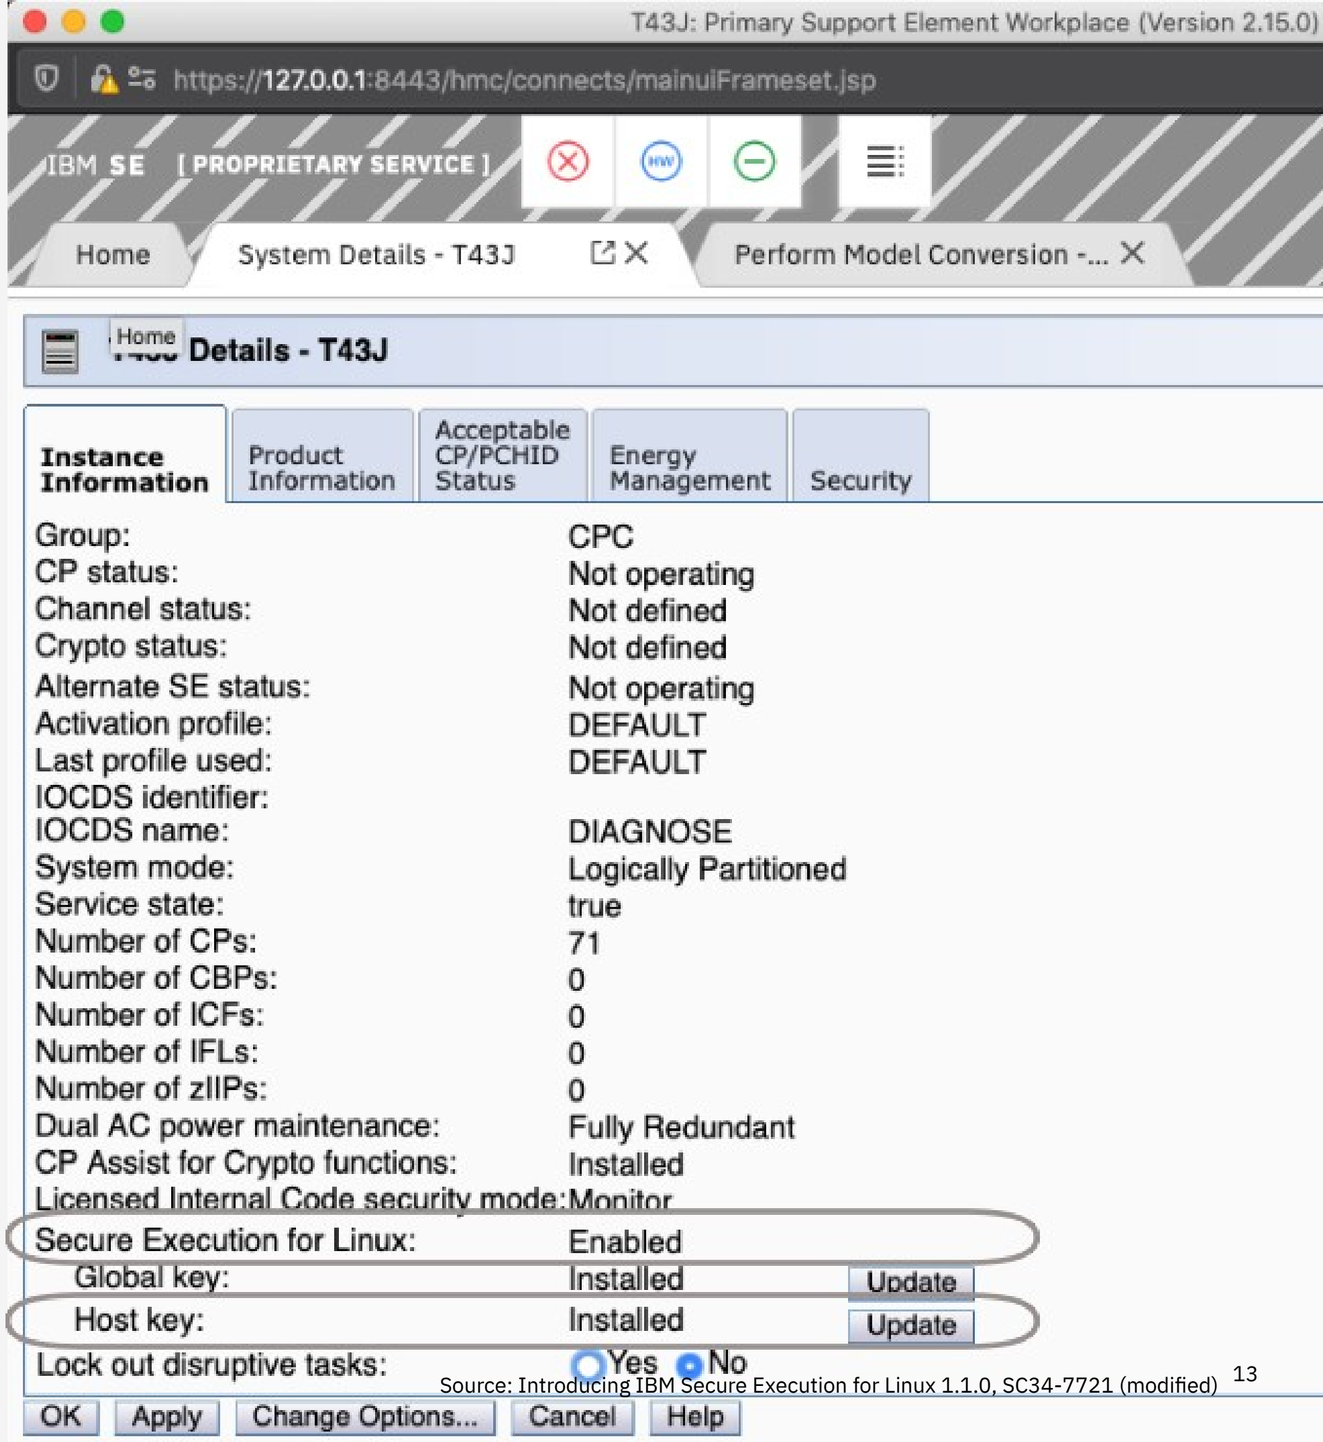

Source: Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721 (modified)

Prepare KVM host

Install OS, QEMU, and libvirt

Install Hypervisor with IBM Secure Execution host support

- Linux OS (with KVM)
- QEMU
- libvirt

OS with IBM Secure Execution host support
Ubuntu 20.04
SLES 15 SP2
RHEL 8.3

For development:

- Upstream Linux kernel >= 5.7
- Upcoming version of upstream QEMU (probably 5.1)

Prepare KVM host

Enable IBM Secure Execution

Enable `prot_virt=1` Linux kernel option

- e.g. edit `zipl.conf`
- Run `zipl`

Reboot and verify that the opt-in was successful

```
# dmesg
...
[0.311322] prot_virt: Reserving 322MB as ultravisor base storage
...
```

Prepare the libvirt domain definition [1]

- Currently supported devices: `sclp`, `virtio-blk`, `virtio-scsi`, `virtio-net`, and `virtio-serial`
- Enable bounce buffers for virtio devices by using the option `iommu='on'`
- Use host CPU model
- **Pitfalls:**
  - Special handling: e.g. for `virtio-serial` and `virtio-scsi`

    Enable `iommu='on'` on the associated controller
  - `virtio-memballoon` is unsupported => disable it

    `<memballoon model='none'/>`

[1] See https://libvirt.org/kbase/s390_protected_virt.html for details

```
<domain type="kvm">
  <name>secguest1</name>
  ...
  <cpu mode='host-model'/>
  <devices>
    <disk type="file" device="disk">
      <driver name="qemu" type="qcow2" iommu="on"/>
      <source file="/var/lib/libvirt/images/secg1.qcow2"/>
      <target dev="vda" bus="virtio"/>
    </disk>
    ...
    <memballoon model='none'/>
  </devices>
</domain>
```

Guest owner tasks 

Guest preparation

Requirements

- Trusted (s390x) system to build the disk image: LPAR, KVM guest, …
  - This doesn’t have to be a IBM z15
- Host key document(s) for the CEC(s) on which the prepared guest should run:
  - This is either provided by the host owner
  - Or download it: https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/HKDSearch
- IBM Z Host key signing certificate, intermediate DigiCert CA, and IBM Z host key revocation list: https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/IBM-Secure-Execution-for-Linux

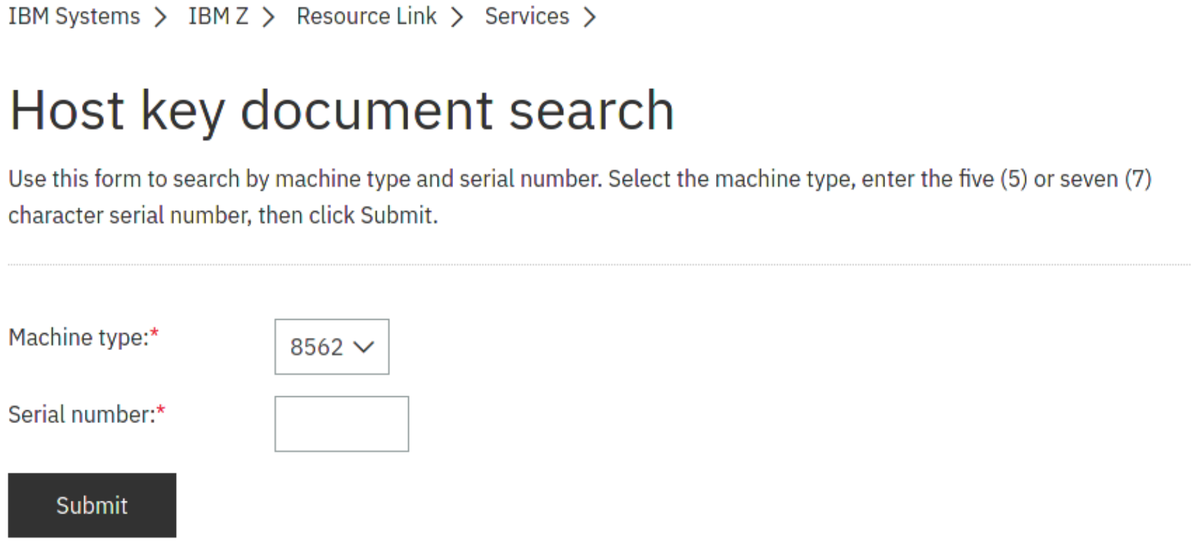

Visited https://www-01.ibm.com/servers/resourcelink/hom03010.nsf/pages/HKDSearch on 02.07.2020

Guest preparation

Prepare guest OS

1. Prepare a libvirt KVM guest definition
   - mind the target

2. Install a OS with IBM Secure Execution guest support in the KVM guest and encrypt all partitions except `/boot`

3. Deploy your workload on the encrypted partition(s) in the guest

OS with IBM Secure Execution guest support
RHEL 7.8
RHEL 8.2
SLES12 SP5
SLES15 SP2
Ubuntu 20.04

For development:

- Linux upstream kernel >= 5.3 and `CONFIG_PROTECTED_VIRTUALIZATION_GUEST=y`

Prepare guest OS

In the KVM guest:

- Enforce secure remote login only
  - set up SSHD and the SSH keys
  - Disable login on kernel consoles

    e.g. by disabling serial and virtual TTYs

```
# cat /etc/systemd/system/serial-getty@.service.d/disable.conf
[Unit]
ConditionKernelCommandLine=allowlocallogin

# cat /etc/systemd/system/autovt@.service.d/disable.conf
[Unit]
ConditionKernelCommandLine=allowlocallogin
```

Guest preparation

Prepare guest OS

In the KVM guest:

- Disable debug shell in `initramfs`

  e.g. `panic=...`

- Disable debug, emergency, and rescue shells

  e.g. for systemd

```
# systemctl mask emergency.service
# systemctl mask emergency.target

# systemctl mask rescue.service
# systemctl mask rescue.target
```

- Remove information leaks on the kernel console

  e.g. `loglevel=0 systemd.show_status=no`

- Install `genprotimg` (s390-tools)

```
...
Begin: Waiting for root file system ...
Begin: Running /scripts/local-block ... Not enough
available memory to open a keyslot.
cryptsetup: ERROR: vda6_crypt: cryptsetup failed, bad
password or options?
cryptsetup: ERROR: vda6_crypt: maximum number of tries
exceeded
  Volume group "vgubuntu" not found
  Cannot process volume group vgubuntu
done.
done.
Gave up waiting for root file system device.  Common
problems:
 - Boot args (cat /proc/cmdline)
   - Check rootdelay= (did the system wait long
enough?)
 - Missing modules (cat /proc/modules; ls /dev)
ALERT!  /dev/mapper/vgubuntu-root does not exist.
Dropping to a shell!
Rebooting automatically due to panic= boot argument
```

Prepare IBM Secure Execution boot image overview

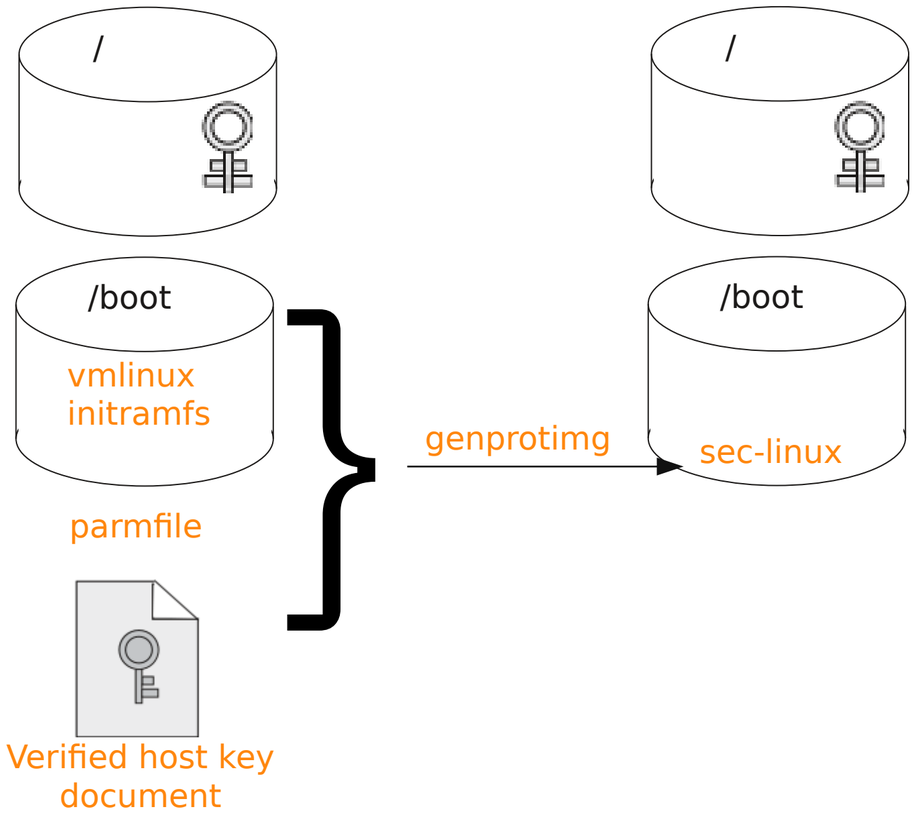

Source: Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721 (modified)

Guest preparation

Prepare `initramfs`

Prepare `initramfs` so the disk encryption keys are included [1]

- Save references to keys (plain format) or pass phrases (LUKS/LUKS2) for each volume in the `/etc/crypttab` configuration file
- Include the `/etc/crypttab` configuration file in the initial RAM file system
- Set `KEYFILE_PATTERN` in `/etc/cryptsetup-initramfs/conf-hook`

Because the initial RAM file system will be encrypted, it can hold keys and pass phrases without compromising security

[1] See https://cryptsetup-team.pages.debian.net/cryptsetup/encrypted-boot.html#avoiding-the-extra-password-prompt

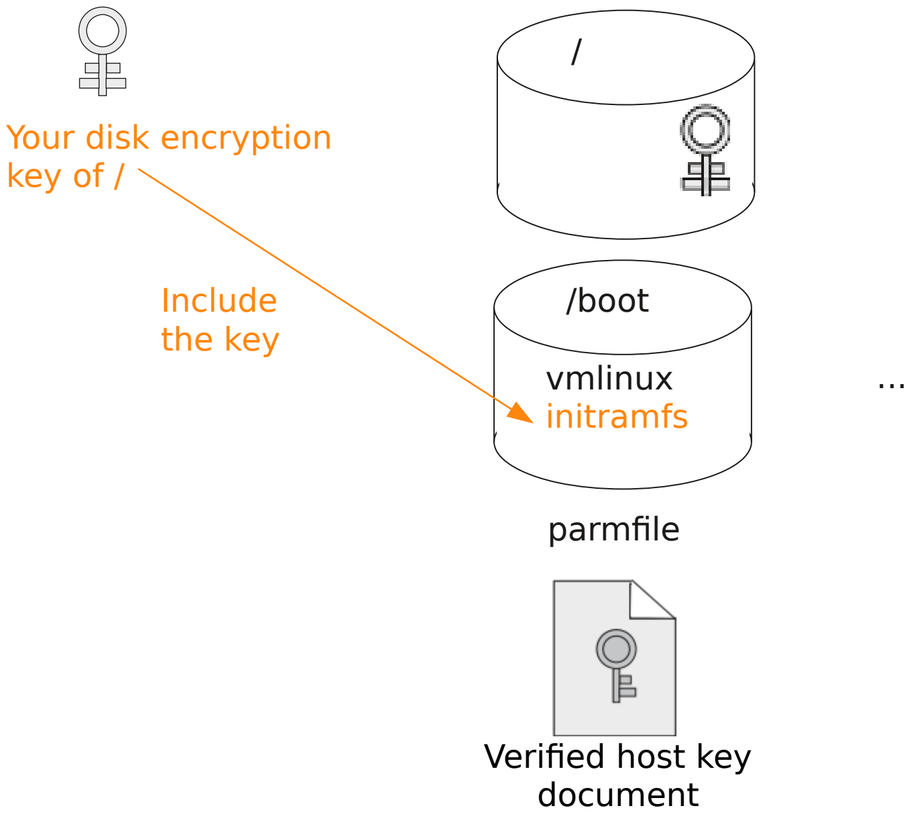

Source: Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721 (modified)

Prepare `parmfile`

Take the guests kernel command line[1] (e.g. from `zipl.conf`) and:

- Set recommended buffer size for bounce buffering

  `swiotlb=262144`

- Disable debug shell, e.g. in the `initramfs`

  e.g. `panic=...`

- Remove kernel console information leaks

  e.g. `loglevel=0 systemd.show_status=no`

- **Optional:**

  Increase crashkernel size by using `crashkernel=...` option for `kdump`

[1] See https://www.kernel.org/doc/html/latest/admin-guide/kernel-parameters.html for details

Guest preparation

Host key document verification

Verify the chain of trust with `check_hostkeydoc[1]`:

Host key document

IBM signing key certificate

```
# check_hostkeydoc host-key.crt \
      ibm-z-host-key-signing.crt \
      -c DigiCertCA.crt -r ibm-z-host-key.crl
```

CA intermediate CA

IBM host key revocation list

In a later version of `genprotimg` this function will be performed by the tooling itself.

[1] If the tool is not provided by the distribution you can download it https://raw.githubusercontent.com/ibm-s390-tools/s390-tools/master/genprotimg/samples/check_hostkeydoc

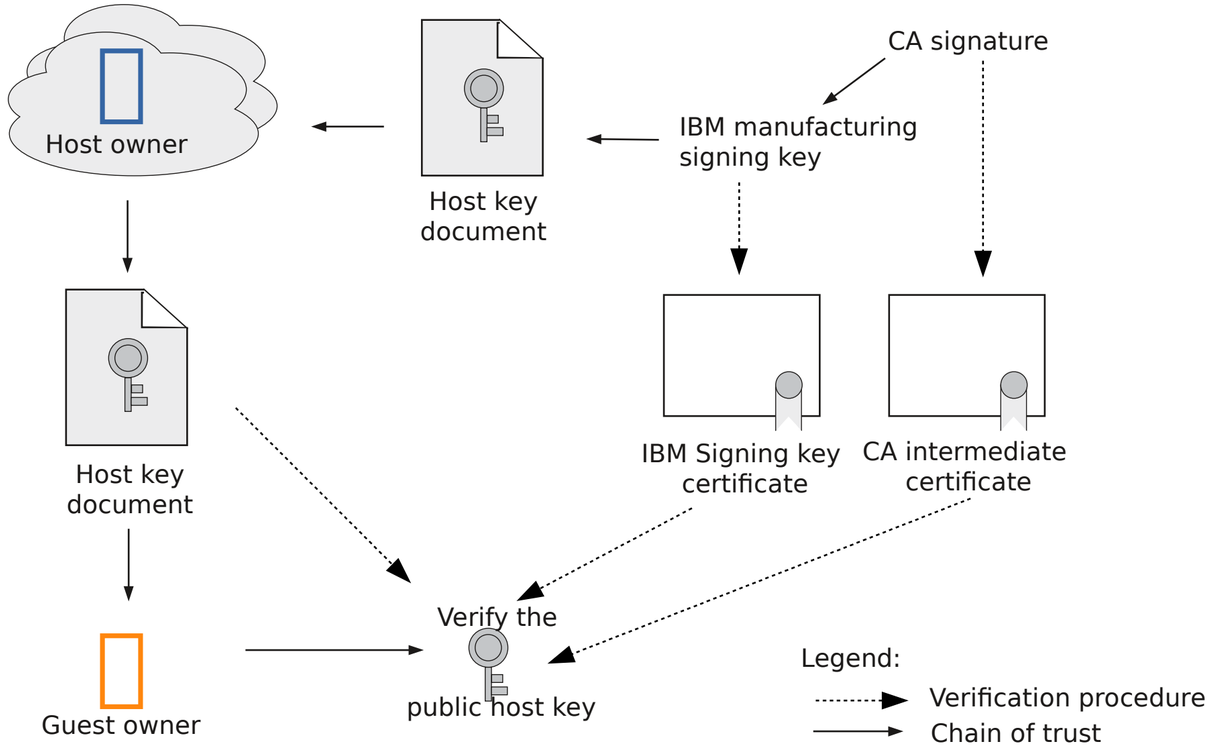

Source: Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721 (modified)

Guest preparation

Create boot image

Basic usage:


```
# genprotimg -k host-key.crt -i vmlinuz -p parmfile -r initramfs.img -o /boot/sec-linux --no-verify
```

- Use the option `--no-verify` only if the host key document has been verified!

- The output file `/boot/sec-linux` can be zipl'ed and used for QEMU direct kernel boot

Remaining steps in the guest

1. Make sure there are no secrets lying around on an unencrypted disk! If there are, use a secure deletion tool to delete these files (e.g. `srm`)
2. Create a zipl entry for the created image
3. Remove all other “unsecured” zipl entries
4. Run `zipl`

```
# vim zipl.conf

...

[secure]

target=/boot

image=/boot/sec-linux

...
```

Guest preparation

Provide the prepared workload to the host owner

Provide the prepared workload (e.g. `QCOW2` disk image) to the host owner.

DEMO TIME

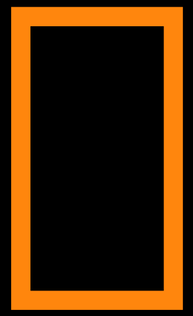

Learn more about IBM Secure Execution

- **Knowledge Center:**
  https://ibm.com/support/knowledgecenter/linuxonibm/com.ibm.linux.z.lxse/lxse_t_secureexecution.html
- **Where to get host key documents?**
  https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/HKDSearch
- **Where to get IBM Z signing key document, IBM Z host key revocation list, and DigiCert Intermediate CA?**
  https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/IBM-Secure-Execution-for-Linux
- **Technical overview blog:**
  https://developer.ibm.com/blogs/technical-overview-of-secure-execution-for-linux-on-ibm-z/
- **One pager:**
  https://ibm.com/downloads/cas/GPLNZLE2
- **FAQ:**
  https://ibm.com/downloads/cas/G1WLJDAY
- **Libvirt documentation for IBM Secure Execution:**
  https://libvirt.org/kbase/s390_protected_virt.html

Thank you

Marc Hartmayer
Software Engineer | Linux on Z & Virtualization Development
—
mhartmay@de.ibm.com
+49-7031-16-1944

Overview

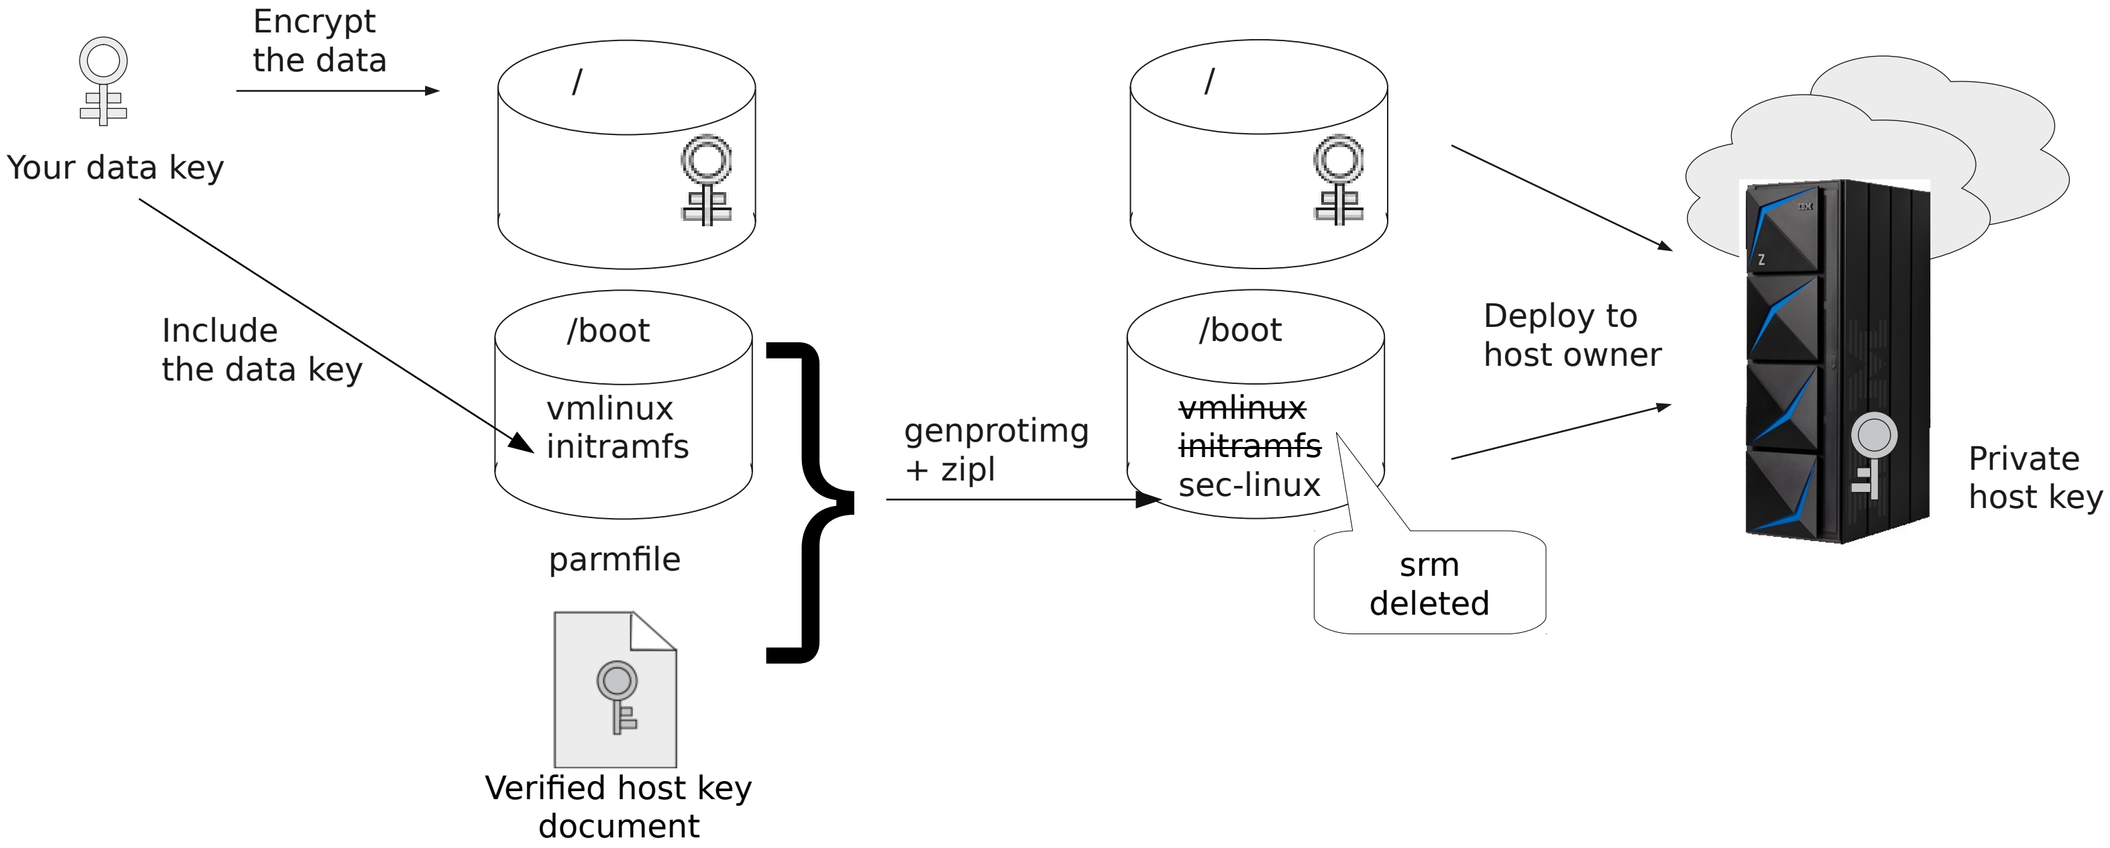

Source: Introducing IBM Secure Execution for Linux 1.1.0, SC34-7721 (modified)

Backup: FAQ

- **How do I know if the Linux host supports IBM Secure Execution feature?**

  When using libvirt >= 6.5.0 you can run

  ```
  $ virt-host-validate
  ```

  or check for the CPU facility 158

  ```
  $ grep facilities /proc/cpuinfo | grep 158
  ```

- **How do I know if QEMU supports IBM Secure Execution feature?**

  When using libvirt >= 6.5.0 you can run

  ```
  $ virsh domcapabilities | grep unpack
  <feature policy='require' name='unpack'/>
  ```

Backup:

There are also some experimental options available:

```
$ genprotimg --help-all
```

Example usage:

```
$ genprotimg --x-comm-key ~/comm.key --x-comp-key ~/comp.key --x-header-key ~/header.key ...
```