Configuring Virtual Intel Software Guard Extensions (vSGX) in VMware ESXi on Lenovo ThinkSystem Servers

Introduces the features of Intel Software Guard Extensions (SGX)

Provides instructions on how to configure vSGX in VMware ESXi

Describes how to verify the vSGX feature in a virtual machine

Explains the prerequisites and limitations for using vSGX

Alpus Chen

LENOVO PRESS

Abstract

Intel Software Guard Extensions (Intel SGX) is a new processor-specific technology for application developers who are seeking to protect selected code and data from disclosure or modification. VMware vSphere 7.0 supports virtualizing Intel Secure Guard Extensions (vSGX) to virtual machines on selected Intel processors and provides additional security to the workloads.

This document provides a brief overview of the Intel SGX and describes how to configure and use Intel SGX in VMware ESXi 7.0 on Lenovo® ThinkSystem™ servers. This document is intended for IT specialists and IT managers who are familiar with VMware vSphere products.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you’re there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction

Intel Software Guard Extensions (Intel SGX) is a new processor-specific technology for application developers who are seeking to protect selected code and data from disclosure or modification. Intel SGX allows user-level code to define private region of memory which contains in an UEFI-reserved contiguous memory region called *Processor Reserved Memory* (PRM). An *enclave* is the section of an application created from virtual address space, located in secure part of physical memory referred to as *Enclave Page Cache* (EPC).

Unlike pages used for regular memory, enclave can be accessed using new special Instruction Set Architecture (ISA) commands that jump into per enclave predefined addresses. Data within an enclave can only be accessed from that same enclave code. The enclave contents are protected such that code running outside the enclave cannot access the enclave contents, including the operating system and the hypervisor, as shows in Figure 1.

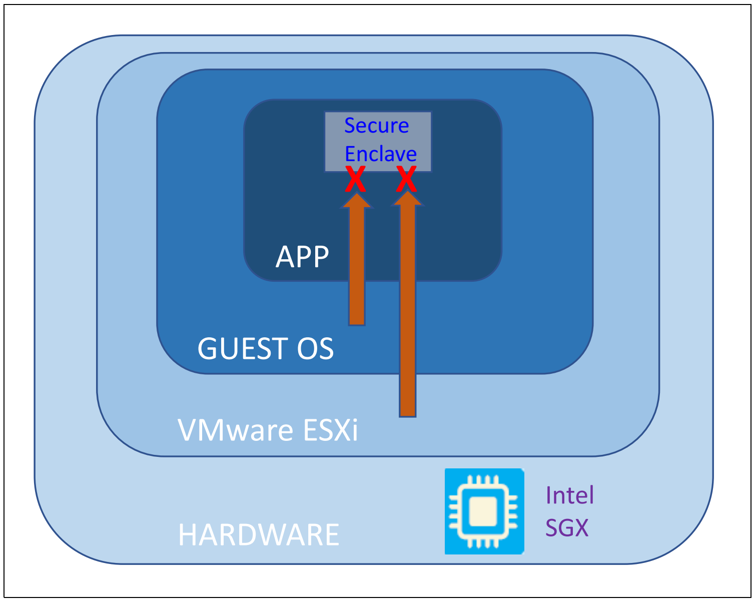

Figure 1 Intel SGX Model

Intel SGX relies on the system UEFI and OS for initial provisioning, resource allocation, and management. However, after an Intel SGX enclave starts execution, it runs on a cryptographically isolated environment separated from UEFI and OS, removes the privileged software (OS, virtual machine manager, device driver) and unprivileged software (ring 3 application, VM, container) from the trust boundary of the code running inside the enclave and enhances the security of sensitive application code and data.

For more details about the Intel SGX, visit:

https://www.intel.com/content/www/us/en/developer/tools/software-guard-extensions/overview.html

Note that some operations and features are not supported in a virtual machine when vSGX is enabled:

- vMotion/DRS migration
- Virtual machine suspend and resume
- Virtual machine snapshots (snapshots are supported if you do not snapshot
- the virtual machine’s memory)
- Fault tolerance
- Guest Integrity (GI, platform foundation for VMware AppDefense 1.0)

Intel SGX supported in Lenovo ThinkSystem servers

Intel SGX is a feature that was first implemented in Intel client platforms and Intel Xeon E3 single socket processors, and now has been implemented in the 3rd generation Intel Xeon Scalable processors.

Supported ThinkSystem servers with Intel Xeon E3-2300, E3-2200 and E3-2100 processors include the following:

- Lenovo ThinkSystem ST250 V2
- Lenovo ThinkSystem SR250 V2
- Lenovo ThinkSystem ST250
- Lenovo ThinkSystem SR250
- Lenovo ThinkSystem SR150
- Lenovo ThinkSystem ST50

Supported ThinkSystem servers with third-generation Intel Xeon Scalable processor (Ice Lake) include the following:

- Lenovo ThinkSystem ST650 V2
- Lenovo ThinkSystem SR630 V2
- Lenovo ThinkSystem SR650 V2
- Lenovo ThinkSystem SR670 V2
- Lenovo ThinkSystem SD630 V2
- Lenovo ThinkSystem SD650 V2
- Lenovo ThinkSystem SD650-N V2
- Lenovo ThinkSystem SN550 V2

The size of the SGX enclave is fixed but is different depending on the processor model. For the enclave size supported for each Intel processor model, please refer to the Product Guide for the server, available from the Lenovo Press web site:

https://lenovopress.lenovo.com/

For the specific memory population rules to enable Intel SGX on a specific ThinkSystem server, please refer to the Setup Guide for the server on ThinkSystem documentation Information Center:

https://thinksystem.lenovofiles.com/help/index.jsp

How to configure vSGX in VMware ESXi

VMware vSphere 7.0 (and subsequent releases) supports virtualized Intel Secure Guard Extensions (vSGX) to virtual machines on selected Intel processors and provide additional security to the workloads. vSGX enables virtual machines to use Intel SGX technology if available on the hardware.

To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the UEFI of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine.

In order to use vSGX, the vSphere environment must meet these requirements:

- VMware ESXi running in the host is ESXi 7.0 or later.
- The virtual machine uses EFI firmware.
- The virtual machine hardware version is 17 or later.
- The guest VM is running either a major Linux distribution, Windows Server 2016 (64-bit) or later, or Windows 10 (64-bit) or later

On the single-socket platform such as ThinkSystem ST250 V2, Intel SGX functionality and security properties are provided by the one socket ships with per-part unique hardware keys built into the processor. Intel SGX instructions allow enclaves to access keys derived from these hardware keys to help protect secrets or securely communicate between enclaves.

The hardware configuration is listed in Table 1. The host OS is installed on the M.2 SATA SSD.

Table 1 ThinkSystem ST250 V2 Server HW configuration

Component	Configuration
System	ThinkSystem ST250 V2 Server
CPU	Intel Xeon E-2388G CPU 3.20GHz
Memory	4x TruDDR4™ 3200MHz 16GB UDIMM
M.2 SATA	800GB M.2 SATA SSD
Host hypervisor OS	ESXi 7.0 U3 Custom Image for Lenovo ThinkSystem
Guest VM OS	RHEL 8.4 / Windows 10

The following steps describe the procedures to configure Intel SGX in vSphere 7.0 U3 on a Lenovo ThinkSystem ST250 V2 server.

> **Hyperthreading:** You might need to turn off hyperthreading on certain CPUs to enable SGX on the ESXi host. For more information, see the VMware KB article at `https://kb.vmware.com/s/article/71367`

1. Enable Intel SGX function in the UEFI setup menu.
   a. Power on the system and enter UEFI setup menu by press F1 button.

b. Select **System Settings** → **Processors**. Find the Software Guard Extensions (SGX) setting and make sure the option is set to **Enabled** as shown in Figure 2.

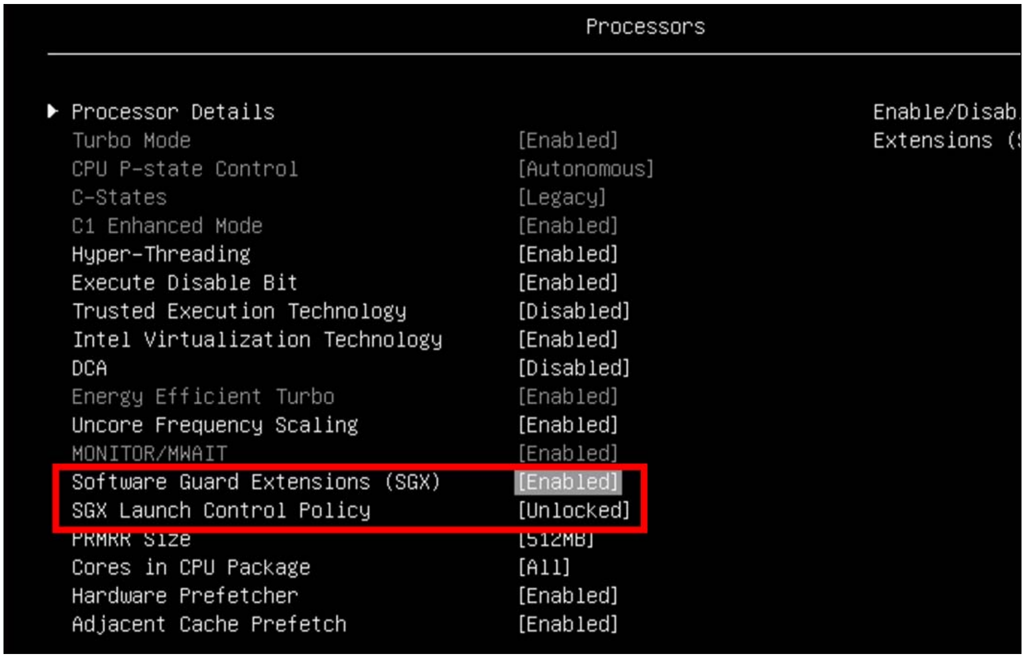

```
                                    Processors
 ───────────────────────────────────────────────────────────────────────────────

▶ Processor Details                                              Enable/Disab
  Turbo Mode                          [Enabled]                  Extensions (
  CPU P-state Control                 [Autonomous]
  C-States                            [Legacy]
  C1 Enhanced Mode                    [Enabled]
  Hyper-Threading                     [Enabled]
  Execute Disable Bit                 [Enabled]
  Trusted Execution Technology        [Disabled]
  Intel Virtualization Technology     [Enabled]
  DCA                                 [Disabled]
  Energy Efficient Turbo              [Enabled]
  Uncore Frequency Scaling            [Enabled]
  MONITOR/MWAIT                       [Enabled]
  Software Guard Extensions (SGX)     [Enabled]
  SGX Launch Control Policy           [Unlocked]
  PRMRR Size                          [512MB]
  Cores in CPU Package                [All]
  Hardware Prefetcher                 [Enabled]
  Adjacent Cache Prefetch             [Enabled]
```

Figure 2 Intel SGX option set to Enabled

c. Ensure SGX Launch Control Policy option is set to **Unlocked** (this option enables the launch enclave configuration of OS/VMM),

d. Save the settings and exit the UEFI setup menu.

2. Install VMware ESXi 7.0 U3 on ThinkSystem ST250 V2 system. When OS Installation completed, login to the ESXi shell as root user.

3. In the ESXi shell, obtain SGX global information by using the following command:

```
~# vsish -e get /hardware/cpu/sgxInfo | more
```

The output is shown in Figure 3 where you can see that Intel SGX feature is enabled in the system and the total EPC size is 96768 pages, which is totally 378MB (1 page = 4K bytes) available for use.

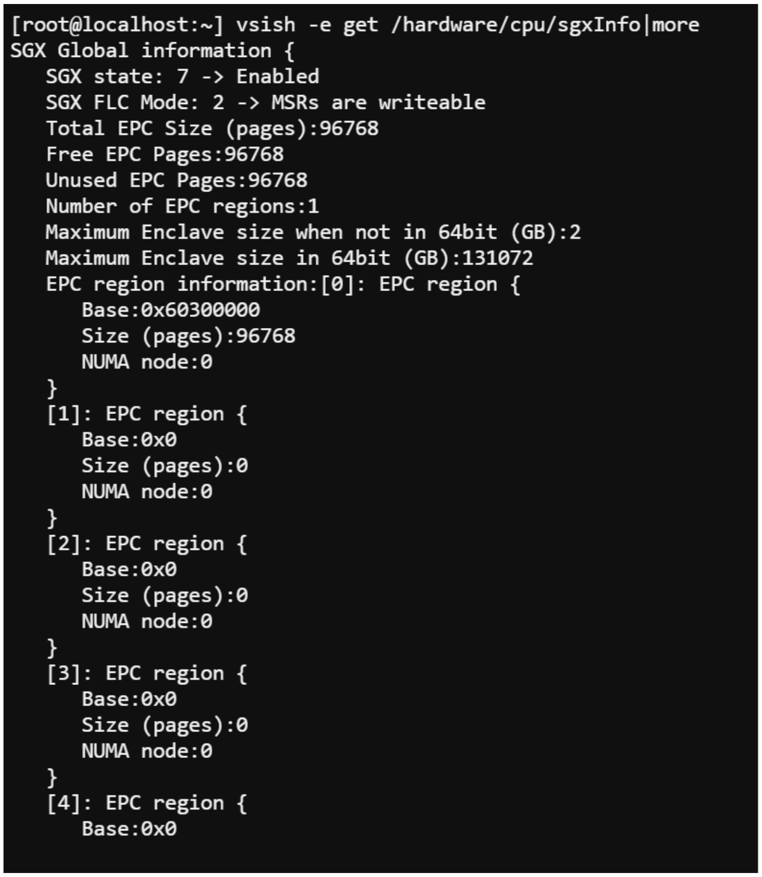

```
[root@localhost:~] vsish -e get /hardware/cpu/sgxInfo|more
SGX Global information {
   SGX state: 7 -> Enabled
   SGX FLC Mode: 2 -> MSRs are writeable
   Total EPC Size (pages):96768
   Free EPC Pages:96768
   Unused EPC Pages:96768
   Number of EPC regions:1
   Maximum Enclave size when not in 64bit (GB):2
   Maximum Enclave size in 64bit (GB):131072
   EPC region information:[0]: EPC region {
      Base:0x60300000
      Size (pages):96768
      NUMA node:0
   }
   [1]: EPC region {
      Base:0x0
      Size (pages):0
      NUMA node:0
   }
   [2]: EPC region {
      Base:0x0
      Size (pages):0
      NUMA node:0
   }
   [3]: EPC region {
      Base:0x0
      Size (pages):0
      NUMA node:0
   }
   [4]: EPC region {
      Base:0x0
```

Figure 3   SGX information in VMware OS

4. Create a virtual machine and install a guest OS that supports Intel SGX (for example RHEL 8.4 or Windows 10).
5. Setup up the virtual machine.
   a. Before power on the VM, click the Edit button to modify virtual machine settings.
   b. In the Security devices category, ensure that SGX is selected and input the Enclave page cache size of your choice (the size must be multiple of 2MB, the maximum size is 378MB for the test machine)

c. Launch control configuration option select Unlocked, as shows in Figure 4.

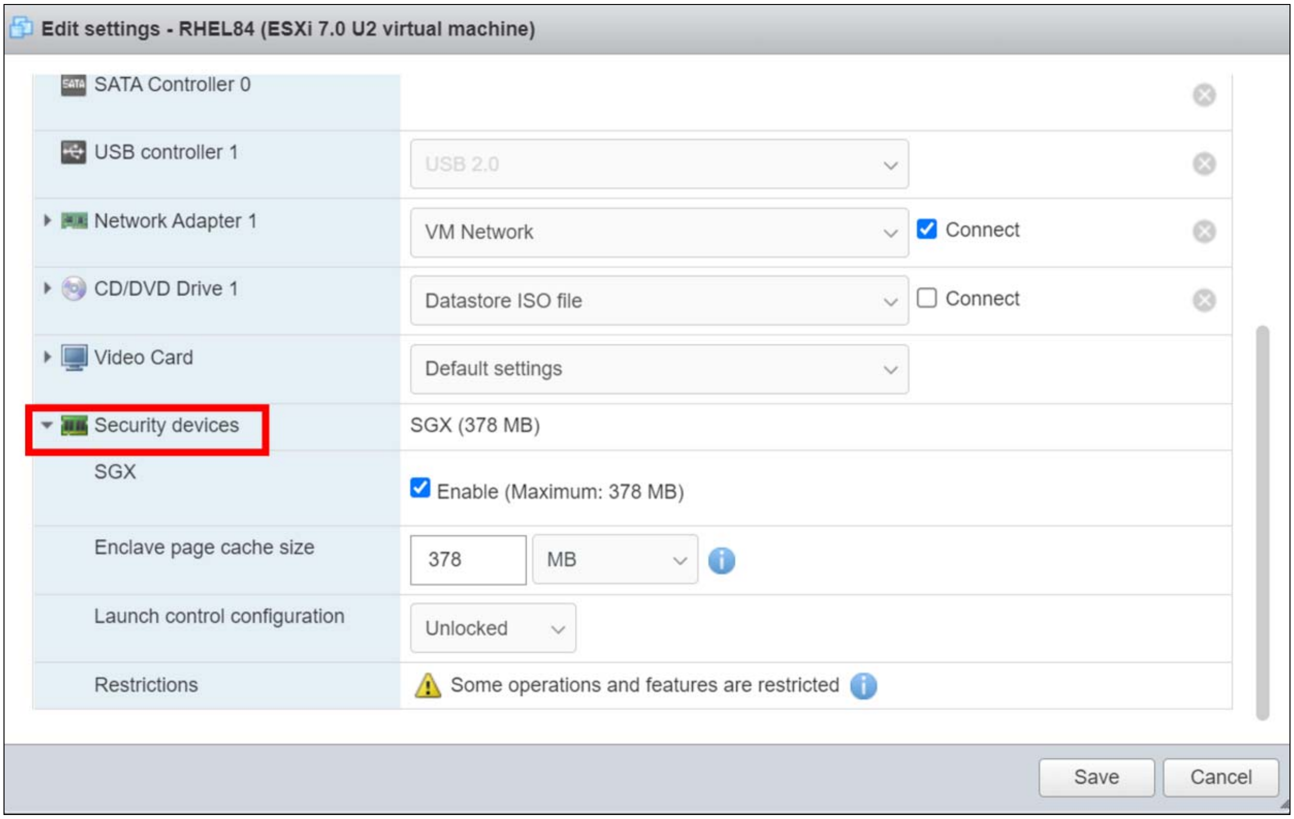

Edit settings - RHEL84 (ESXi 7.0 U2 virtual machine)

SATA Controller 0

USB controller 1: USB 2.0

Network Adapter 1: VM Network ☑ Connect

CD/DVD Drive 1: Datastore ISO file ☐ Connect

Video Card: Default settings

Security devices: SGX (378 MB)

SGX: ☑ Enable (Maximum: 378 MB)

Enclave page cache size: 378 MB

Launch control configuration: Unlocked

Restrictions: Some operations and features are restricted

Figure 4 Enable SGX in virtual machine setting

d. Go to VM Options -> Boot Options -> Firmware, ensure that EFI is selected, as shown in Figure 5.

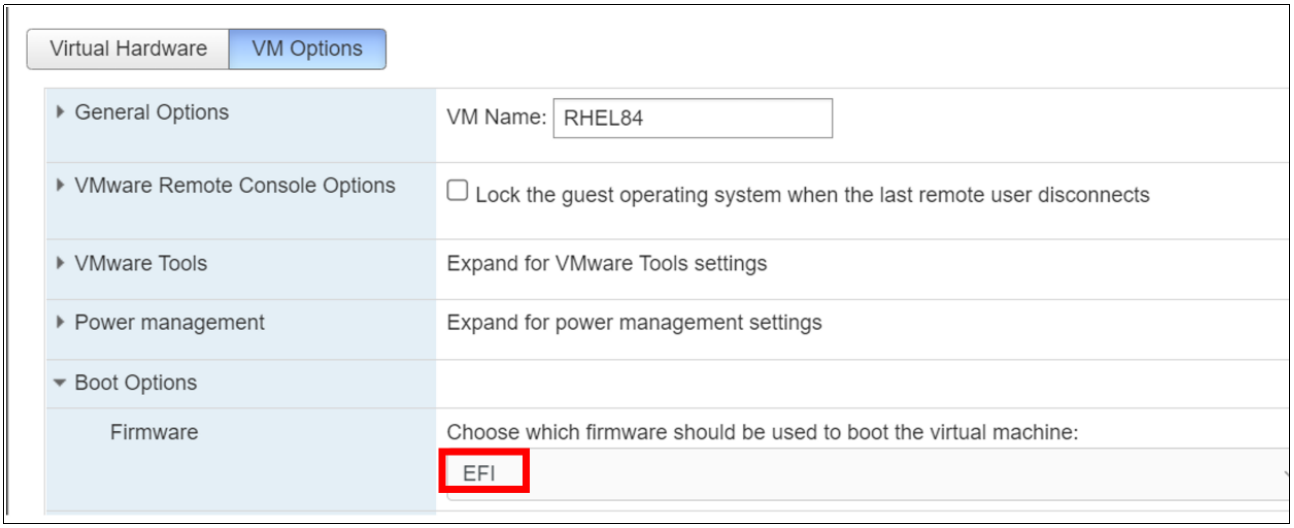

General Options: VM Name: RHEL84

VMware Remote Console Options: ☐ Lock the guest operating system when the last remote user disconnects

VMware Tools: Expand for VMware Tools settings

Power management: Expand for power management settings

Boot Options

Firmware: Choose which firmware should be used to boot the virtual machine: EFI

Figure 5 Setup EFI mode in VM options

e. Click Save to finish the VM settings and perform guest OS installation.

Test Intel SGX in Guest OS

This section demonstrates how to verify Intel SGX features in a guest OS. We will use RHEL 8.4 and Windows 10 as example guest OSes.

RHEL 8.4

We can verify vSGX functionality in RHEL 8.4 as a guest OS as follows:

1. Boot up the VM and check the OS dmesg with the following command to get the EPC section size.

   ```
   ~# dmesg | grep -i sgx
   ```

   The output is shown in Figure 6.

   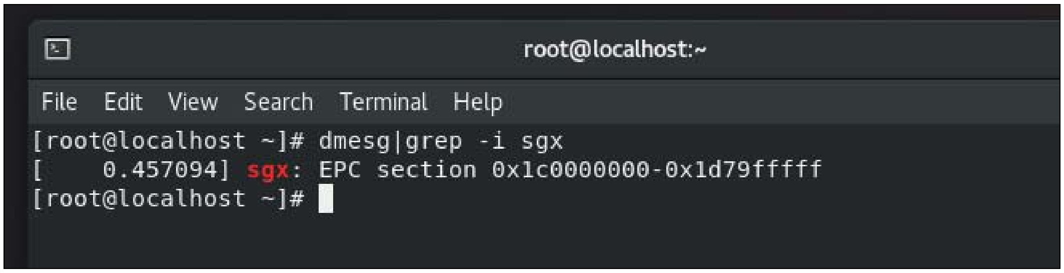

   ```
   [root@localhost ~]# dmesg|grep -i sgx
   [    0.457094] sgx: EPC section 0x1c0000000-0x1d79fffff
   [root@localhost ~]#
   ```

   Figure 6 Get EPC size in RHEL8 VM

   As the output shows, the EPC section size is `0x17A00000` (`0x1d79fffff - 0x1c0000000 + 1`) in hexadecimal. Convert the number to decimal and we get the size is 378MB, which matches the number of the EPC size configured in the virtual machine setting in Figure 4 on page 8.

2. Check that SGX kernel devices exists by using the following command:

   ```
   ~ # ls /dev/sgx*
   ```

   The output is shown in Figure 7.

   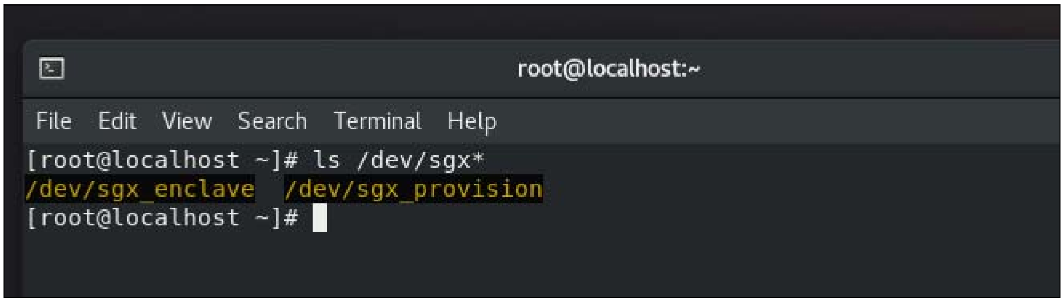

   ```
   [root@localhost ~]# ls /dev/sgx*
   /dev/sgx_enclave  /dev/sgx_provision
   [root@localhost ~]#
   ```

   Figure 7 Check SGX kernel device

   The output shows the two SGX devices /dev/sgx_enclave and /dev/sgx_provision exist. RHEL 8.4 has native driver support for SGX kernel devices. Mainline kernel release 5.11 or higher includes the SGX in-box driver.

3. Check sgx and sgx_lc flag in the CPU information by using the following command:

   ```
    ~# cat /proc/cpuinfo
   ```

The output is shown in Figure 8.

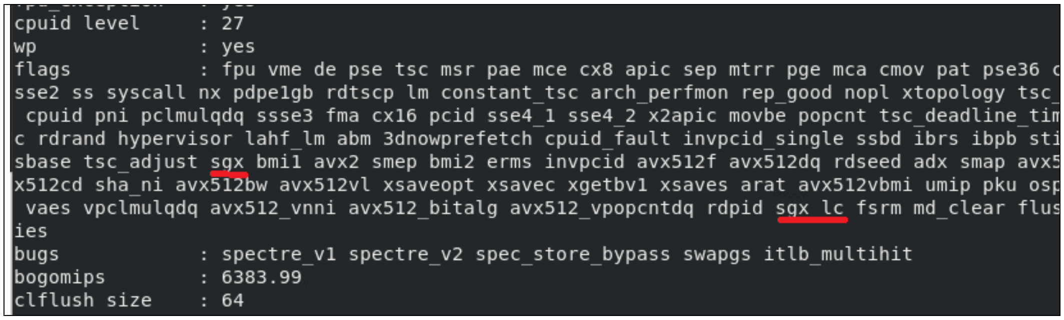

```
fpu_exception   : yes
cpuid level     : 27
wp              : yes
flags           : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 c
sse2 ss syscall nx pdpe1gb rdtscp lm constant_tsc arch_perfmon rep_good nopl xtopology tsc_
 cpuid pni pclmulqdq ssse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_deadline_tim
c rdrand hypervisor lahf_lm abm 3dnowprefetch cpuid_fault invpcid_single ssbd ibrs ibpb sti
sbase tsc_adjust sgx bmi1 avx2 smep bmi2 erms invpcid avx512f avx512dq rdseed adx smap avx5
x512cd sha_ni avx512bw avx512vl xsaveopt xsavec xgetbv1 xsaves arat avx512vbmi umip pku osp
 vaes vpclmulqdq avx512_vnni avx512_bitalg avx512_vpopcntdq rdpid sgx_lc fsrm md_clear flus
ies
bugs            : spectre_v1 spectre_v2 spec_store_bypass swapgs itlb_multihit
bogomips        : 6383.99
clflush size    : 64
```

Figure 8 Check CPU information in RHEL8

The output shows that the `sgx` and `sgx_lc` CPU flags are present and kernel supports it.

Windows 10

To verify vSGX is enabled in a Windows 10 guest OS, simply check for the SGX entries in the following sections in Windows Device Manager as shown in Figure 9.

- Software components
- System devices

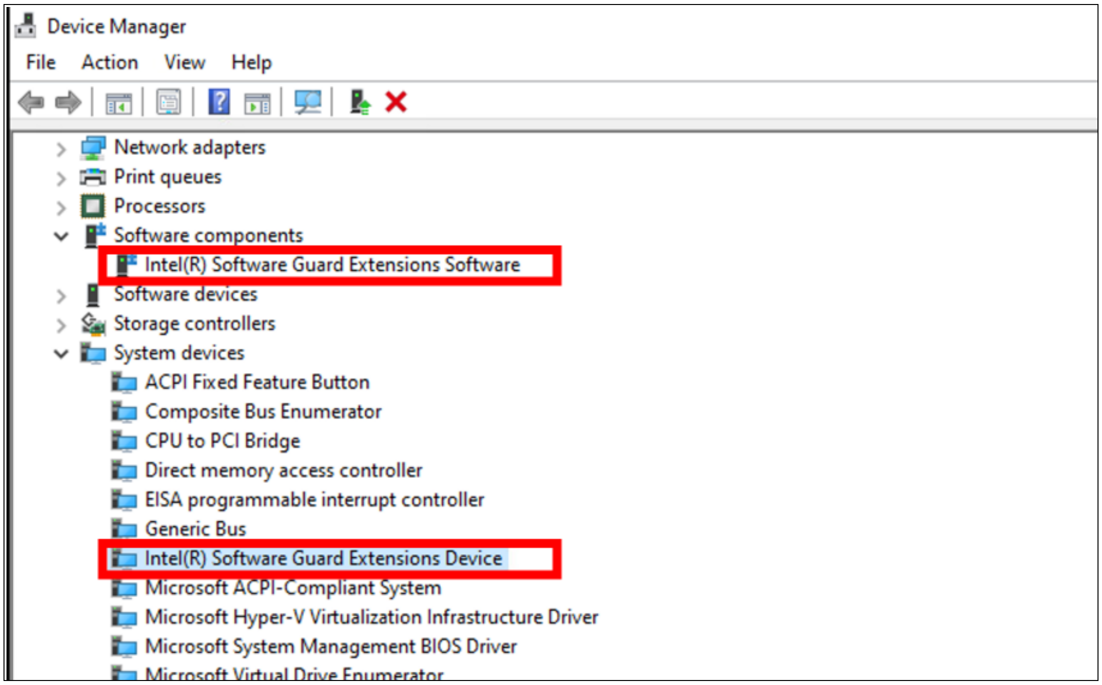

- Network adapters
- Print queues
- Processors
- Software components
  - Intel(R) Software Guard Extensions Software
- Software devices
- Storage controllers
- System devices
  - ACPI Fixed Feature Button
  - Composite Bus Enumerator
  - CPU to PCI Bridge
  - Direct memory access controller
  - EISA programmable interrupt controller
  - Generic Bus
  - Intel(R) Software Guard Extensions Device
  - Microsoft ACPI-Compliant System
  - Microsoft Hyper-V Virtualization Infrastructure Driver
  - Microsoft System Management BIOS Driver
  - Microsoft Virtual Drive Enumerator

Figure 9 Intel SGX supported in Windows Device Manger

Reference

For additional information, see these resources:

- Intel SGX

  https://www.intel.com/content/www/us/en/developer/tools/software-guard-extensions/overview.html
- Enabling Intel SGX on Lenovo ThinkSystem V2 Servers

  https://lenovopress.lenovo.com/lp1471-enabling-intel-sgx-on-thinksystem-v2
- VMware vSphere documentation, Securing Virtual Machines with Intel Software Guard Extensions

  https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-4CFB8CE3-4104-41B3-B7F9-D330392F0732.html
- Intel Software Guard Extension SDK for Linux

  https://01.org/intel-softwareguard-extensions

Author

Alpus Chen is an OS Engineer at the Lenovo Infrastructure Solutions Group in Taipei, Taiwan. As a specialist in Linux and VMware technical support for several years, he is interested in operating system operation and recently focuses on VMware OS.

Thanks to the following specialists for their contributions and suggestions:

- Chengcheng Peng, Lenovo VMware Engineer
- Gary Cudak, Lenovo OS Architect
- David Watts, Lenovo Press

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on August 22, 2022.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp1639

Trademarks