

Enabling Intel SGX on Lenovo ThinkSystem Servers

Planning / Implementation

Intel Software Guard Extensions (Intel SGX) is an extension to Intel architecture that provides new CPU instructions and platform enhancements to allow applications to create private areas to protect sensitive information. Sensitive information is protected even when attackers have full control of the platform. Protection is achieved through the use of private regions of memory called enclaves.

Intel SGX is a feature that was first implemented in Intel client platforms and single-socket Intel Xeon E3 processors, and has been implemented in the 3rd Gen and 4th Generation Intel Xeon Scalable processors.

Intel SGX is available in the following ThinkSystem servers:

- ThinkSystem servers with Intel Xeon E3-2100, E3-2200 and E3-2300 processors (high-end processor SKUs only):
 - ThinkSystem ST50
 - ThinkSystem ST250
 - ThinkSystem SR150
 - ThinkSystem SR250
 - ThinkSystem ST50 V2
 - ThinkSystem ST250 V2
 - ThinkSystem SR250 V2
- ThinkSystem V2 servers with 3rd Gen Intel Xeon Scalable processors (all processor SKUs):
 - ThinkSystem ST650 V2
 - ThinkSystem SR630 V2
 - ThinkSystem SR650 V2
 - ThinkSystem SR670 V2
 - ThinkSystem SD630 V2
 - ThinkSystem SD650 V2
 - ThinkSystem SD650-N V2
 - ThinkSystem SN550 V2
- ThinkSystem V3 servers with 4th Gen Intel Xeon Scalable processors (all processor SKUs):
 - ThinkSystem ST650 V3
 - ThinkSystem SD650 V3
 - ThinkSystem SD650-I V3
 - ThinkSystem SD650-N V3
 - ThinkSystem SR630 V3
 - ThinkSystem SR650 V3
 - ThinkSystem SR850 V3
 - ThinkSystem SR860 V3
 - ThinkSystem SR950 V3

This paper covers the implementation of Intel SGX on the ThinkSystem V2 servers with 3rd Gen Intel Xeon Scalable processors, and ThinkSystem V3 servers with 4th Gen Intel Xeon Scalable processors.

Figure 1 shows the key components of Intel SGX. The ThinkSystem server UEFI provides the SGX initialization and reserves a contiguous memory region for enclaves. Application can use enclaves for critical data protection.

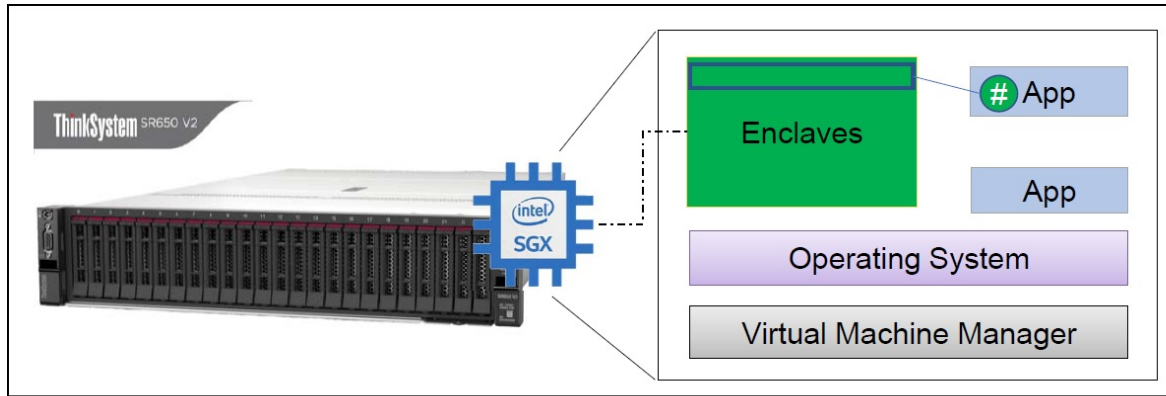


Figure 1. Applications using an SGX Enclave

Reserved Memory for SGX

The enclave region is allocated from special memory region in order to support SGX. ThinkSystem BIOS will reserve a contiguous memory region called the Processor Reserved Memory (PRM) in Processor Reserved Memory Range Registers (PRMRR).

An enclave is section of an application created from virtual address space of an application but located in secure part of physical memory referred to as Enclave Page Cache (EPC). EPC is part of PRM which is protected by Intel CPU via PRMRRs.

An application can create an enclave for its protected portion. Before the enclave is built, the enclave code and data are free for inspection and analysis. The protected portion is loaded into an enclave where its code and data is measured. Once the application's protected portion of the code and data are loaded into an enclave, memory access controls are in place to restrict access by external software. An enclave can prove its identity to a remote party and provide the necessary building-blocks for secure provisioning of keys and credentials. The application can also request an enclave-specific and platform-specific key that it can use to protect keys and data that it wishes to store outside the enclave.

The size of the SGX enclave is fixed but is different depending on the processor model. Sizes range from 8 GB to 512 GB per processor. For a 2-socket ThinkSystem server, if enough DDR memory is installed, the system BIOS can reserve between 16GB and 1TB based on processor model installed.

For the enclave size supported for each Intel processor model, see the Intel Xeon Scalable Processor Reference:

<https://lenovopress.com/lp1262-intel-xeon-sp-processor-reference#term=SGX>

Instructions of SGX

SGX feature need a new set of CPU instructions and mechanisms for memory accesses added to Intel Architecture processors. Intel SGX can encompass two collections of instruction extensions, referred to as SGX1 and SGX2.

The SGX1 extensions allow an application to instantiate a protected container, referred to as an enclave. The SGX2 extensions allow additional flexibility in runtime management of enclave resources and thread execution within an enclave. The 3rd and 4th Gen Intel Xeon Scalable processor family supports both SGX1 and SGX2 instruction extensions.

The enclave instructions available with Intel SGX are organized as leaf functions under three instruction mnemonics: ENCLS (ring 0), ENCLU (ring 3), and ENCLV (VT root mode). Table 1 provides a summary of the Supervisor (ring 0) and User (ring 3) instruction leaves that are available in the initial implementation of Intel SGX.

Table 1. Supervisor and User Mode Enclave Instruction Leaf Functions of SGX (source: [Intel](#))

Supervisor Instruction	Description	User Instruction	Description
ENCLS[EADD]	Add an EPC page to an enclave.	ENCLU[EENTER]	Enter an enclave.
ENCLS[EBLOCK]	Block an EPC page.	ENCLU[EEXIT]	Exit an enclave.
ENCLS[ECREATE]	Create an enclave.	ENCLU[EGETKEY]	Create a cryptographic key.
ENCLS[EDBGDR]	Read data from a debug enclave by debugger.	ENCLU[EREPORT]	Create a cryptographic report.
ENCLS[EDBGWR]	Write data into a debug enclave by debugger.	ENCLU[ERESUME]	Re-enter an enclave.
ENCLS[EEXTEND]	Extend EPC page measurement.		
ENCLS[EINIT]	Initialize an enclave.		
ENCLS[ELDB]	Load an EPC page in blocked state.		
ENCLS[ELDU]	Load an EPC page in unblocked state.		
ENCLS[EPA]	Add an EPC page to create a version array.		
ENCLS[EREMOVE]	Remove an EPC page from an enclave.		
ENCLS[ETRACK]	Activate EBLOCK checks.		
ENCLS[EWB]	Write back/invalidate an EPC page.		

Based on these instructions, OS/driver/application can create an enclave and initialize it and use for protected portions.

For more information, including details of these instructions, see Chapter 36 “Introduction to Intel Software Guard Extensions” of Intel 64 and IA-32 Architectures Software Developer’s Manual Volume 3:

<https://software.intel.com/content/www/us/en/develop/download/intel-64-and-ia-32-architectures-sdm-combined-volumes-1-2a-2b-2c-2d-3a-3b-3c-3d-and-4.html>

Enabling SGX on multi-socket servers

On single-socket platforms such as ThinkSystem ST50, SR150, ST250, SR250, ST50 V2, ST250 V2 and SR250 V2, as well as client platforms using the Xeon E processors, Intel SGX functionality and security properties are provided completely by the one socket. Each 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. Unique signing keys can also be derived. These keys are common across all logical processors on a single socket, thus, a key request will result in the same answer on all the virtual cores (threads) in the socket. Operating system schedulers rely on this coherency.

In order to keep the operating system side a simple design, Intel has designed a single coherent software environment on multi-socket platforms for enabling SGX. The server UEFI will perform Platform Establishment – the processor socket contains unique, per-part hardware keys, called the Intel SGX root keys, the UEFI will derive the common “platform” root keys, and store these platform keys in persistent storage (flash memory) to make sure these keys are still accessible after the platform is reset.

The implementation of Intel SGX is more complex with two-socket or four-socket servers. Figure 2 shows Intel SGX data protection model when two sockets are used. UEFI will initialize a platform manifest containing the platform root keys along with the information on the sockets that participated in establishing these platform root keys. Each socket encrypts its own copy of the platform keys using its hardware keys, and this ensures that if any socket fails, the remaining sockets can still access the platform keys. Each processor can handle the its own enclaves in Processor Reserved Memory (PRM).

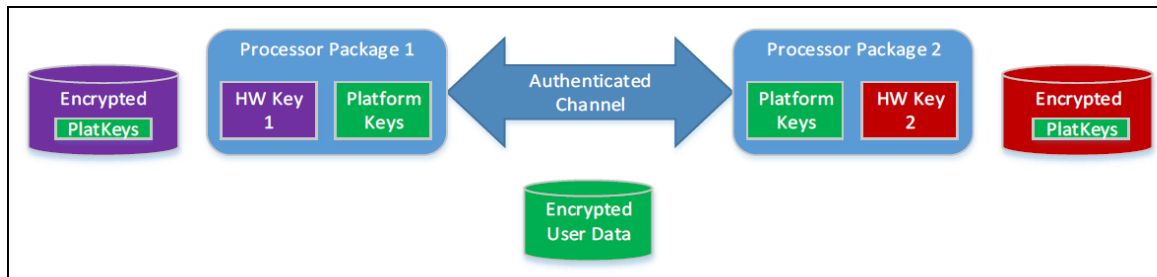


Figure 2. Multi-socket Intel SGX Keys

Figure 3 shows one example of ThinkSystem SR650 V2, where each socket has 16x 16GB DIMMs for a total 512GB of memory with two processors, but the whole memory physical address range will be from 0 to 514GB because there is one 2GB hole in the system address space below 4GB. The processors each have a 64GB enclave size.

Intel SGX requires system memory to be configured as NUMA (Non-Uniform Memory Access) memory, thus each processor has its own consecutive memory space. SGX allocates its Processor Reserved Memory (PRM) which is configured into Processor Reserved Memory Range Registers (PRMRR).

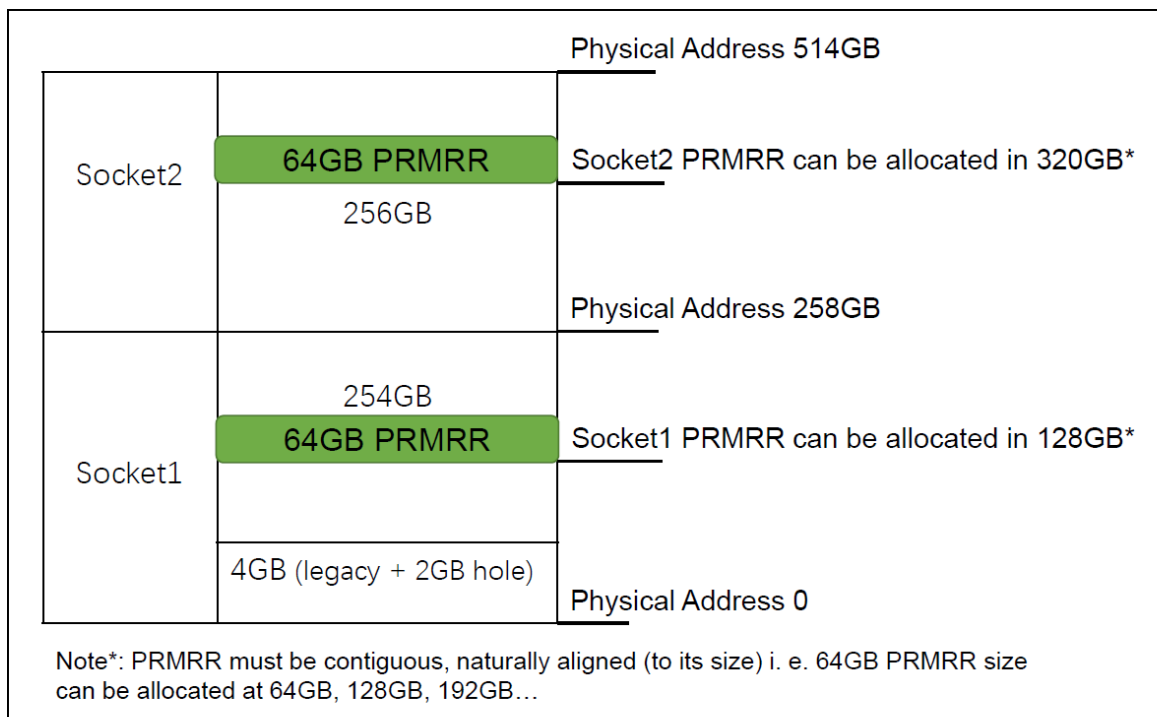


Figure 3. PRMRR Allocation Example in ThinkSystem

DIMM population requirement on multi-socket servers

3rd and 4th Gen Intel Xeon Scalable processors have four Integrated Memory Controllers (iMC), and each iMC has two DDR channels and each channel supports two DIMMs, so one processor can have a maximum of 16 DIMMs. 3rd and 4th Gen Scalable processors only support SGX feature with specific DIMM configurations.

DIMM population requirements are as follows:

- Only configurations of 8 DIMMs, 12 DIMMs and 16 DIMMs are supported.
- Figure 4 and Figure 5 show the DIMM installation (that is, the exact DDR channels and slots of each processor) required to enable SGX on ThinkSystem V2 / V3 Server.
- If different DIMMs are populated in the system, then the populated DIMMs must be symmetric between (iMC0, iMC1) and (iMC2, iMC3), and the populated DIMMs must be identical between socket 1 and socket 2 if two processors are installed.
- Memory Mirroring is not supported and must be disabled.

The following figure shows the memory configuration for 3rd Gen Intel Xeon Scalable processors.

Figure 4. Memory configuration for 3rd Gen Intel Xeon Scalable processors

iMC#	iMC0				iMC1				iMC2				iMC3			
	Chan 0 (A)		Chan 1 (B)		Chan 0 (C)		Chan 1 (D)		Chan 0 (E)		Chan 1 (F)		Chan 0 (G)		Chan 1 (H)	
	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1
8	DIMM		DIMM		DIMM		DIMM		DIMM		DIMM		DIMM		DIMM	
12	DIMM	DIMM	DIMM		DIMM	DIMM	DIMM		DIMM	DIMM	DIMM		DIMM	DIMM	DIMM	
16	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM

The following figure shows the memory configuration for 4th Gen Intel Xeon Scalable processors.

Figure 5. Memory configuration for 4th Gen Intel Xeon Scalable processors

iMC#	iMC0				iMC1				iMC2				iMC3			
	Chan 0 (A)		Chan 1 (B)		Chan 0 (C)		Chan 1 (D)		Chan 0 (E)		Chan 1 (F)		Chan 0 (G)		Chan 1 (H)	
	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1	Slot0	Slot1
8	DIMM		DIMM		DIMM		DIMM		DIMM		DIMM		DIMM		DIMM	
16	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM

For the specific memory population rules to enable SGX on a ThinkSystem server, see DIMM installation rules in “Server hardware setup” of the related ThinkSystem server page on “Lenovo Documentation” web:

<https://pubs.lenovo.com/>

Enabling SGX in System Setup on multi-socket servers

By default, SGX is disabled on ThinkSystem servers. To enable SGX, perform the following steps:

1. Boot the server and press F1 when prompted to enter System
2. Enable NUMA (disable UMA) for 3rd Gen Scalable processors: Select **System Settings > Processors > UMA-Based Clustering** and set it to **Disabled** for ThinkSystem V2 server.
3. Enable Intel Total Memory Encryption (TME). Both 3rd and 4th Gen Scalable processors require TME enabled for the use of SGX. Select **System Settings > Processors > Total Memory Encryption (TME)** and set it to **Enabled**.
4. Enable SGX by selecting **System Settings > Processors > SW Guard Extension (SGX)** and set it to **Enabled**.
5. Configure PRMRR size if needed, the options base on the available consecutive system memory space and processor's capability: **System Settings > Processors > PRMRR Size**.
Figure 6 shows the settings on the Processors page in ThinkSystem SR650 V2 Setup after SGX is enabled.

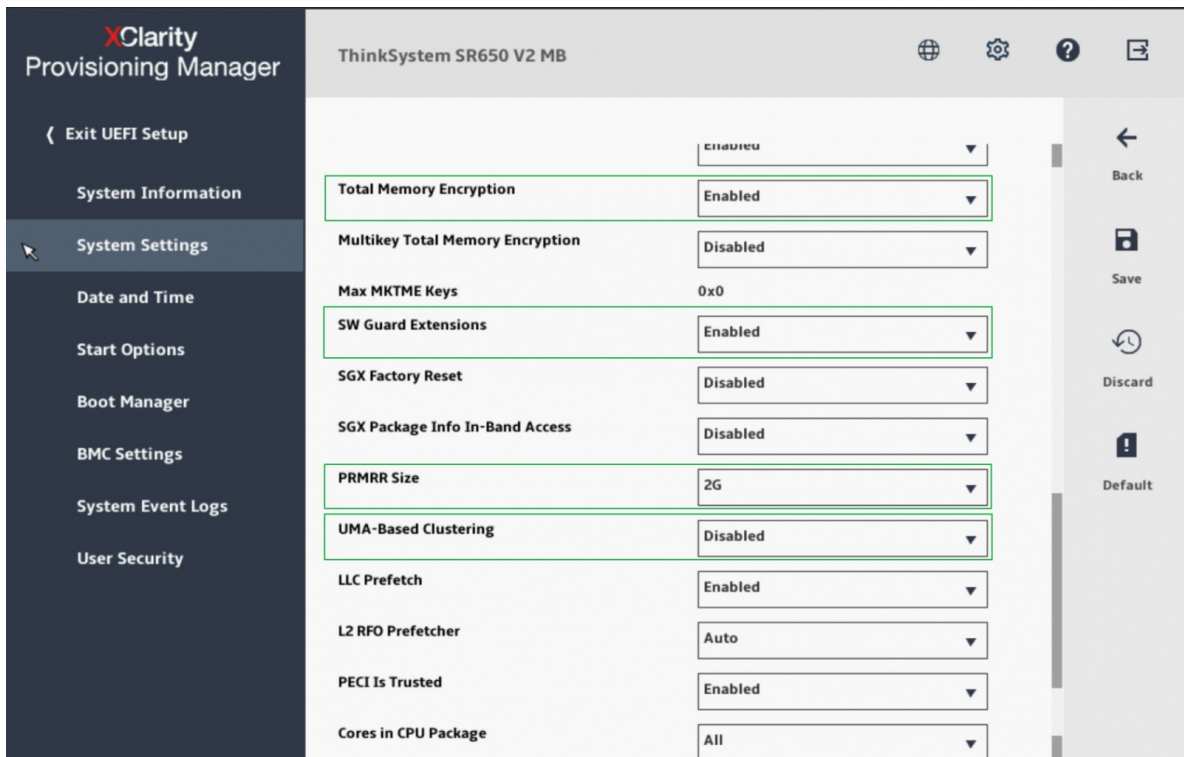


Figure 6. UEFI Setup Page with SGX Enabled on ThinkSystem SR650 V2 Server

6. Save the settings and reboot the
7. During the system reboot, ThinkSystem UEFI will perform an SGX initialization and enable the

You can now use SGX feature in a supported operating system.

Enabling SGX in System Setup on Xeon E3 servers

On Intel Xeon E3 based ThinkSystem ST50, SR150, ST250, SR250, ST50 V2, ST250 V2 and SR250 V2, Intel SGX functionality and security properties are provided completely by the one socket, and the related configuration is simple. As E3 processor only has one iMC, there is no special DIMM population requirement but just follow the DIMM installation rules of the server, please check more details in “Server hardware setup” of the ThinkSystem server page on “Lenovo Documentation” web: <https://pubs.lenovo.com/>

Intel Xeon E3 based ThinkSystem is one socket server, thus it does not have UMA capability. And E3 processor does not have TME feature either. By default, SGX is disabled on E3 based ThinkSystem servers. To enable SGX, just perform the following steps:

1. Enable SGX by selecting **System Settings > Processors > Software Guard Extensions (SGX)** and set it to **Enabled**.
Note: there is another option “Software Controlled” for Xeon E3 based system, if your OS is UEFI version which supports UEFI variable services, and you want to have SGX platform software to enable/disable SGX in OS directly, then you can select “Software Controlled” option; but the SGX configuration in OS still need reboot system to let UEFI enable/disable SGX and then SGX platform software can get SGX status in OS.
2. Configure PRMRR size if needed, the options base on the available consecutive system memory space and processor’s capability: **System Settings > Processors > PRMRR Size**. Figure 7 shows the settings on the Processors page in ThinkSystem ST250 V2 Setup after SGX is enabled.
3. Save the settings and reboot the system.
4. During the system reboot, ThinkSystem UEFI will perform an SGX initialization and enable the feature.

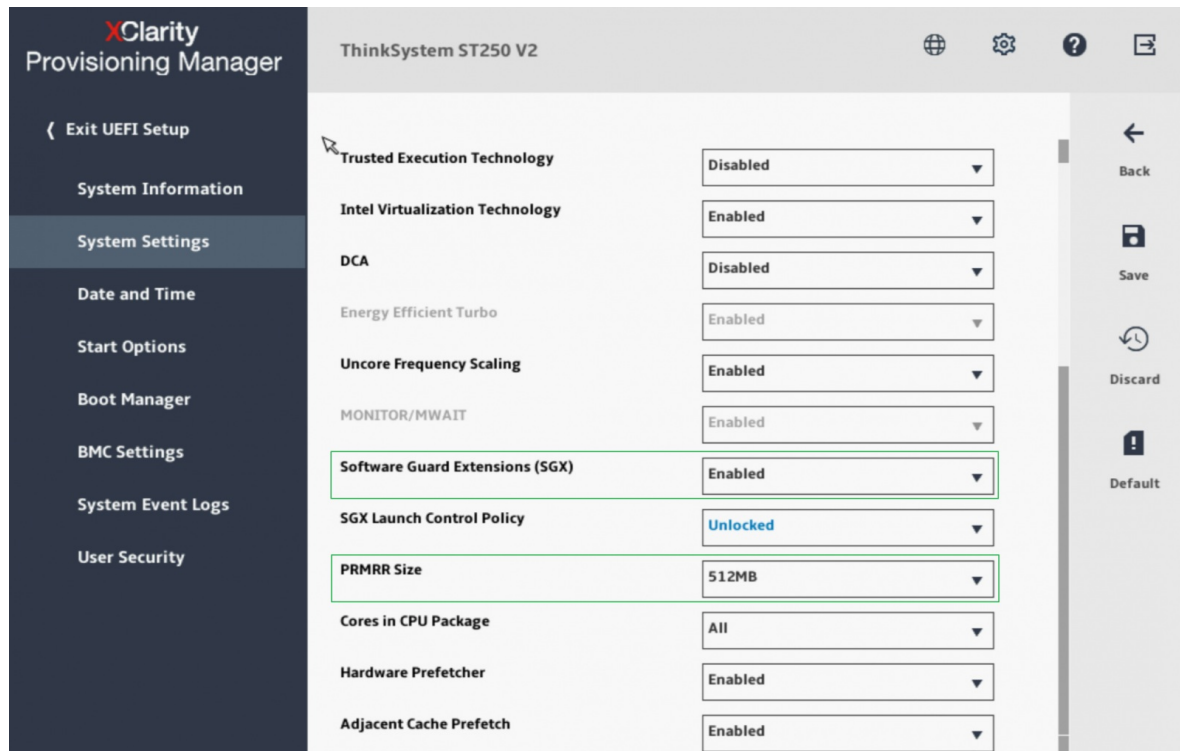


Figure 7. UEFI Setup Page with SGX Enabled on ThinkSystem ST250 V2 Server

References

For more information, see the following documents:

- Intel SGX public web has more introduction:
<https://software.intel.com/content/www/us/en/develop/topics/software-guard-extensions.html>
- Chapter 36 “Introduction to Intel® Software Guard Extensions” of Intel® 64 and IA-32 Architectures Software Developer’s Manual Volume 3:
<https://software.intel.com/content/www/us/en/develop/download/intel-64-and-ia-32-architectures-sdm-combined-volumes-1-2a-2b-2c-2d-3a-3b-3c-3d-and-4.html>
- Supporting Intel® SGX on Multi-Socket Platforms
<https://www.intel.com/content/www/us/en/architecture-and-technology/software-guard-extensions/supporting-sgx-on-multi-socket-platforms.html>

Authors

Jason Liu is a Senior UEFI Architect at Lenovo where he provides high-level infrastructure design support for Lenovo ThinkSystem UEFI firmware and leads the enabling, customization and innovation of new technologies into UEFI firmware. Jason also leads Reliability, Availability and Serviceability (RAS) architecture design for ThinkSystem, both UEFI firmware and BMC firmware.

Blake Liu is a Senior Engineer at Lenovo Infrastructure Solutions Group. Blake is a technical leader in UEFI Secure domain, responsible for secure feature design and implementation, and also leads PSIRT issue tracking and fixing for ThinkSystem UEFI firmware.

Related product families

Product families related to this document are the following:

- [Processors](#)

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.
8001 Development Drive
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.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1471, was created or updated on December 18, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1471>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1471>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.