

Innovative Technology for CPU Based Attestation and Sealing

Ittai Anati Shay Gueron Simon P Johnson Vincent R Scarlata
Intel Corporation

{ittai.anati, shay.gueron, simon.p.johnson, vincent.r.scarlata}@intel.com

ABSTRACT

Intel is developing the Intel® Software Guard Extensions (Intel® SGX) technology, an extension to Intel® Architecture for generating protected software containers. The container is referred to as an enclave. Inside the enclave, software's code, data, and stack are protected by hardware enforced access control policies that prevent attacks against the enclave's content. In an era where software and services are deployed over the Internet, it is critical to be able to securely provision enclaves remotely, over the wire or air, to know with confidence that the secrets are protected and to be able to save secrets in non-volatile memory for future use.

This paper describes the technology components that allow provisioning of secrets to an enclave. These components include a method to generate a hardware based attestation of the software running inside an enclave and a means for enclave software to seal secrets and export them outside of the enclave (for example store them in non-volatile memory) such that only the same enclave software would be able un-seal them back to their original form.

General Terms

Measurement, Security

Keywords

Enclave, Measurement, Attestation, Local Attestation, Remote Attestation, Sealing

1 INTRODUCTION

In an era where software and services are deployed over the Internet, Intel® Software Guard Extensions (Intel® SGX), and extension to Intel® Architecture enables service providers to provision applications over the wire or air with sensitive content, and to know with confidence that their secrets are properly protected. In order to do so, the service providers must be able to know with certainty what software is running on the remote platform and in which environment it is executing.

1.1 Software Lifecycle

Intel® SGX-enabled software does not ship with sensitive data. After the software is installed, it contacts the service provider to have data remotely provisioned to the enclave. The software then encrypts the data and stores it for future use. Figure 1 illustrates the steps software takes to complete this processes.

1. Enclave Launch – The untrusted application launches the enclave environment to protect the service provider's software. While the enclave is built, a secure log is recorded reflecting the contents of the enclave and how it was loaded. This secure log is the enclave's "Measurement."




Figure 1: Software Lifecycle

2. Attestation – The enclave contacts the service provider to have its sensitive data provisioned to the enclave. The platform produces a secure assertion that identifies the hardware environment and the enclave.
3. Provisioning – The service provider assesses the trustworthiness of the enclave. It uses the attestation to establish secure communication and provision sensitive data to the enclave. Using the secure channel, the service provider sends the data to the enclave.
4. Sealing/Unsealing – The enclave uses a persistent hardware-based encryption key to securely encrypt and store its sensitive data in a way that ensures the data can be retrieved only when the trusted environment is restored.
5. Software Upgrade – Enclave software updates might be required by the service provider. To streamline the migration of data from an older software version to the newer version, the software can request seal keys from older versions to unseal the data and request the new version's seal so that the sealed data won't be available to previous versions of the software.

Finally, when platform owners plan to transfer ownership of the platform, the secrets available during their ownership period should be made inaccessible. Intel® SGX includes a user owned special persistent value that, when changed, changes all keys available to software.

1.2 Security Model

A service provider that wishes to provision secrets to a remote platform must know, in advance, that the protection policies of the remote platform meet the protection requirements of the secrets to be deployed. According to Intel® SGX's security model [1], the trusted computing base (TCB) in charge of protecting the secrets includes the processor's firmware and hardware and only the software inside the enclave.

An enclave writer can use Intel® SGX's dedicated sealing and attestation components that uphold the security model

assertions below to prove to the service provider that the secrets will be protected according to the intended security level:

- Intel® SGX provides an enclave instance with the means to request a secure assertion from the platform of the enclave’s identity.
 - Intel® SGX also allows the enclave to bind enclave ephemeral data to the assertion.
- Intel® SGX provides the means for an enclave instance to verify assertions originating from other enclave instances on the same platform.
- Intel® SGX provides the means for remote entity to verify assertions from an enclave instance.
- Intel® SGX allows an enclave instance to obtain keys that are bound to the platform and the enclave.
 - Intel® SGX prevents software access to keys of other Enclave Identities

1.3 Intel® SGX Instructions

The Intel® SGX Architecture [1] provides the hardware instructions, *EREPOR* and *EGETKEY*, to support attestation and sealing. Secret owners that accept SGX’s security model can rely on these instructions to report back the TCB in charge of the secrets.

To create the enclave environment, untrusted software uses Intel® SGX instructions. These instructions also compute the cryptographic measurement of the launched environment. These processes are further described in section 2 of this paper.

To enable attestation and sealing, the hardware provides two additional instructions *EREPOR* and *EGETKEY*. The *EREPOR* instruction provides an evidence structure that is cryptographically bound to the hardware for consumption by attestation verifiers. *EGETKEY* provides enclave software with access to the “Report” and “Seal” keys used in the attestation and sealing process. The use of these instructions to provide an attestation of the enclave is discussed in section 3, and to protect the secrets delivered to the enclave in section 4.

In section 5 we briefly review related work in the field of establishing remote trust in platforms.

2 MEASUREMENT

The Intel® SGX architecture is responsible for establishing identities for attestation and sealing. For each enclave it provides two measurement registers, MRENCLAVE and MRSIGNER; MRENCLAVE provides an identity of the enclave code and data as it’s constructed and MRSIGNER provides an identity of an authority over the enclave. These values are recorded while the enclave is built, and are finalized before enclave execution commences. Only the TCB has access to write to these registers in order to ensure an accurate reflection of the identities is available when attesting and sealing.

2.1 MRENCLAVE - Enclave Identity

The “Enclave Identity” is the value of MRENCLAVE, which is a SHA-256[2] digest of an internal log that records all the activity done while the enclave is built [1]. The log consists of the following information:

- The contents of the pages (code, data, stack, heap).
- The relative position of the pages in the enclave.
- Any security flags associated with the pages.

Once enclave initialization is complete, through the *EINIT* instruction, no more updates are made to MRENCLAVE. The final value of MRENCLAVE is a SHA-256 digest that identifies, cryptographically, the code, data, and stack placed inside the

enclave, the order and position in which the enclave’s pages were placed, and the security properties of each page. Any change to any of these variables would lead to a different value in MRENCLAVE.

2.2 MRSIGNER - Sealing Identity

The enclave has a second identity used for data protection called the “Sealing Identity.” The Sealing Identity includes a “Sealing Authority,” a product ID and a version number. The Sealing Authority is an entity that signs the enclave prior to distribution, typically the enclave builder. The enclave builder presents the hardware with an RSA signed enclave certificate (SIGSTRUCT) that contains the expected value of the Enclave Identity, MRENCLAVE, and the public key of the Sealing Authority. The hardware checks the signature on the certificate, using the public key contained within, and then it compares the value of the measured MRENCLAVE against the signed version. If these checks pass, a hash of the public key of the Sealing Authority is stored in the MRSIGNER register. It is important to note that if multiple enclaves are signed by the same Sealing Authority, they will all have the same MRSIGNER value. As shown in Section 4, the value of Sealing Identity can be used for sealing data in a way that enclaves from the same Sealing Authority (e.g., different versions of the same enclave) can share and migrate their sealed data.

3 ATTESTATION

Attestation is the process of demonstrating that a piece of software has been properly instantiated on the platform. In Intel® SGX it is the mechanism by which another party can gain confidence that the correct software is securely running within an enclave on an enabled platform. In order to do this the Intel® SGX architecture produces an attestation assertion (shown in Figure 2) that conveys the following information:

- The identities of the software environment being attested
- Details of any non-measurable state (e.g. the mode the software environment may be running in)
- Data which the software environment wishes to associated with itself
- A cryptographic binding to the platform TCB making the assertion

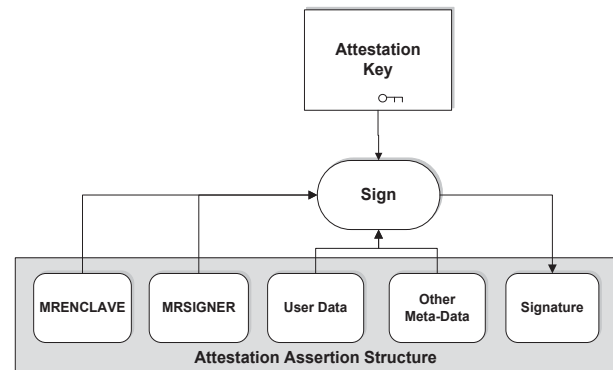


Figure 2: Attestation Assertion Creation

The Intel® SGX architecture provides one mechanism for creating an authenticated assertion between two enclaves running on the same platform (local attestation), and another mechanism for extending local attestation to provide assertions to 3rd parties outside the platform (remote attestation).

Finally, in order to gain maximum trustworthiness in the system, the attestation key should only be bound to a specific

platform TCB. Should the platform TCB change, say through a microcode update, the platform attestation key should be replaced in order to properly represent the trustworthiness of the TCB.

3.1 Intra-Platform Enclave Attestation

Application developers may wish to write enclaves which can cooperate with one another to perform some higher-level function. In order to do this they need a mechanism for the enclaves to authenticate one another. For this purpose Intel® SGX architecture provides the *EREPORT* instruction.

When invoked by an enclave, *EREPORT* creates a signed structure, known as a REPORT. The REPORT structure contains the two identities of the enclave, the attributes associated with the enclave (attributes identify modes and other properties established during *ECREATE*), the trustworthiness of the hardware TCB, and additional information the enclave developer wishes to pass on to the target enclave, and a message authentication code (MAC) tag. The target enclave is the enclave which will verify the MAC over the REPORT allowing it to determine that the enclave that created the REPORT has run on the same platform.

The MAC is produced with a key called the “Report Key”. As shown in Table 1, Report Key is known only to the target enclave and to the *EREPORT* instruction. The validating (target) enclave can retrieve its own Report Key using the *EGETKEY* instruction. *EGETKEY* provides enclaves with keys, among them the Report Key, usable for symmetric encryption and authentication. The target enclave uses the Report Key to recompute the MAC over the REPORT data structure, and verify that the REPORT was produced by the attesting (reporting) enclave. The Intel® SGX architecture uses AES128-CMAC [3] as the MAC algorithm.

Table 1: Access to Report Key and Seal Key

Key	Instruction	Key association
Seal Key	EGETKEY	Current enclave
Report Key	EREPORT	Target enclave
	EGETKEY	Current enclave

Each REPORT structure also includes a 256-bit field for User Data. This field binds data that is inside the enclave to the identity of the enclave (as expressed by the REPORT). This field can be used to extend the REPORT with auxiliary data by populating it with a hash digest of the auxiliary data, which is then provided alongside the REPORT. The use of the User Data field enables an enclave to build a higher level protocol to form a secure channel between itself and another entity.

For example, by exchanging REPORTs that authenticate public Diffie-Hellman keys, that were randomly generated inside the enclave using mutually agreed parameters, the enclaves can generate an authenticated shared secret and use it to protect further communications between themselves. The Intel® Architecture supports the generation of true random values through the RDRAND instruction [4] available for use by enclave software.

The figure below shows an example flow of how two enclaves on the same platform would authenticate each other and verify that the other party is running inside an enclave on the same platform and therefore meets Intel® SGX’s security model




Figure 3: Intra-Platform Attestation Example

1. After communication paths between enclave A and B have been established, enclave A obtains enclave B’s MRENCLAVE value. Note that the communication path in this step doesn’t have to be secure.
2. Enclave A invokes the *EREPORT* instruction together with enclave B’s MRENCLAVE to create a signed REPORT destined for enclave B. Enclave A transmits its REPORT to enclave B via the untrusted communication path.
3. After receiving the REPORT from enclave A:
 - Enclave B calls *EGETKEY* to retrieve its Report Key, recomputes the MAC over the REPORT structure, and compares the result with the MAC accompanying the REPORT. A match in the MAC value affirms that A is indeed an enclave that is running on the same platform as enclave B and as such that A is running in an environment that abides to Intel® SGX’s security model.
 - Once the firmware and hardware components of the TCB have been verified, Enclave B can then examine Enclave A’s REPORT to verify the software components of the TCB:
 - MRENCLAVE reflects the contents of the software image running inside the enclave.
 - MRSIGNER reflects the sealer’s identity
 - Enclave B can then reciprocate by creating a REPORT for enclave A, by using the MRENCLAVE value from the REPORT it just received.
 - Enclave B transmits its REPORT to enclave A. Enclave A can then verify the report in a similar manner to enclave B confirming that enclave B exists on the same platform as enclave A.

3.2 Inter-Platform Enclave Attestation

The authentication mechanism used for Intra-platform enclave attestation uses a symmetric key system, where only the enclave verifying the REPORT structure and the *EREPORT* instruction that creates the REPORT have access to the authentication key. Creating an attestation that can be verified outside the platform requires using asymmetric cryptography. The Intel® SGX enables a special enclave, called the Quoting Enclave, which is devoted to remote attestation. The Quoting Enclave verifies REPORTs from other enclaves on the platform using the Intra-platform enclave attestation method described above, and then replaces the MAC over these REPORTs with a signature created with a device specific (private) asymmetric key. The output of this process is called a QUOTE.

3.2.1 Intel® Enhanced Privacy ID (EPID)

Attestation using standard asymmetric signing schemes has drawn privacy concerns when a small number of keys are used across the life of the platform. To overcome this problem Intel has introduced an extension to the Direct Anonymous Attestation scheme used by the TPM [5] & [6] called Intel® Enhanced Privacy ID (EPID)[7] that is used by the Quoting Enclave to sign enclave attestations.

EPID is a group signature scheme that allows a platform to sign objects without uniquely identifying the platform or linking different signatures. Instead, each signer belongs to a “group”, and verifiers use the group’s public key to verify signatures. EPID supports two modes of signatures. In the fully anonymous mode of EPID a verifier cannot associate a given signature with a particular member of the group. In Pseudonymous mode an EPID verifier has the ability to determine whether it has verified the platform previously.

3.2.2 The Quoting Enclave

The Quoting Enclave creates the EPID key used for signing platform attestations which is then certified by an EPID backend infrastructure. The EPID key represents not only the platform but the trustworthiness of the underlying hardware.

Only the Quoting Enclave has access to the EPID key when the enclave system is operational, and the EPID key is bound to the version of the processor’s firmware. Therefore, a QUOTE can be seen to be issued by the processor itself.

3.2.3 Example Remote Attestation Process

Figure 4 shows an example of how an application with a secure processing element on the user platform could provide an attestation to a challenging service provider in order to receive some value added service from the provider. Note that many usages will use this process infrequently (e.g. at enrolment time) to provision the enclave with a communication key that will then be used directly in subsequent connections.

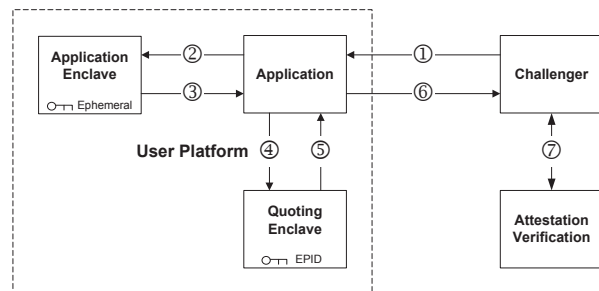


Figure 4: Remote Attestation Example

- Initially, the application needs service from outside the platform, and establishes communication with the service providing system. The service provider issues a challenge to the application to demonstrate that it is indeed running the necessary components inside one or more enclaves. The challenge itself contains a nonce for liveness purposes.
- The application is provided with the Quoting Enclave’s Enclave Identity and passes it along with the provider’s challenge to the application’s enclave.
- The enclave generates a manifest that includes a response to the challenge and an ephemerally generated public key to be used by the challenger for communicating secrets back to the enclave. It then generates a hash digest of the manifest and includes it as User Data for the *EREPORT* instruction that will generate a REPORT that binds the manifest to the enclave, as described in section 3.2. The enclave then sends the REPORT to the application.
- The application forwards the REPORT to the Quoting Enclave for signing.
- The Quoting Enclave retrieves its Report Key using the *EGETKEY* instruction and verifies the REPORT. The Quoting enclave creates the QUOTE structure and signs it with its EPID key. The Quoting Enclave returns the QUOTE structure to the application.
- The application sends the QUOTE structure and any

associated manifest of supporting data to the service challenger.

- The challenger uses an EPID public key certificate and revocation information or an attestation verification service to validate the signature over the Quote. It then verifies the integrity of the manifest using USERDATA and checks the manifest for the response to the challenge it sent in step 1.

4 SEALING

When an enclave is instantiated, the hardware provides protections (confidentiality and integrity) to its data, when it is maintained within the boundary of the enclave. However, when the enclave process exits, the enclave will be destroyed and any data that is secured within the enclave will be lost. If the data is meant to be re-used later, the enclave must make special arrangements to store the data outside the enclave.

Table 1 above shows that *EGETKEY* provides access to persistent Sealing Keys that enclave software can use to encrypt and integrity-protect data. Intel® SGX has no restrictions on the encryption scheme used by the enclave. Replay protection for the data is also possible when combined with other service offered by the platform, such as monotonic counters.

4.1 Intel® SGX supported Sealing Policies

When invoking *EGETKEY*, the enclave selects criteria, or a policy, for which enclaves may access this sealing key. These policies are useful for controlling the accessibility of sensitive data to future versions of the enclave.

Intel® SGX supports two policies for Seal Keys:

- Sealing to the Enclave Identity
- Sealing to the Sealing Identity

Sealing to the Enclave’s Identity produces a key that is available to any instance of this exact enclave. This does not allow future software to access the secrets of this enclave. Sealing to the enclave’s Sealing Identity produces a key that is available to some other enclaves signed by the same Sealing Authority. This can be used to allow newer enclaves to access data stored by previous versions.

Only a subsequent instantiation of an enclave, executing *EGETKEY* with the same policy specification, will be able to retrieve the Sealing Key and decrypt data that was sealed using that key by a previous instantiation.

4.1.1 Sealing to the Enclave Identity

When sealing to the enclave’s Enclave Identity, *EGETKEY* bases the key on the value of the enclave’s MRENCLAVE. Any change that impacts the enclave’s measurement will yield a different key. This results in a different key for each enclave, providing full isolation between enclaves. A byproduct of using this policy is that different versions of the same enclave will also have different seal keys, preventing offline data migration.

This policy is useful for usages where the old data should not be used after a vulnerability is found. For example, if the data is an authentication credential, the service provider may revoke those credentials and provision new ones. Access to the old credential might be harmful.

4.1.2 Sealing to the Sealing Identity

When sealing to the enclave’s Sealing Identity, *EGETKEY* bases the key on the value of the enclave’s MRSIGNER, and the enclave’s version. MRSIGNER reflects the key/identity of the Sealing Authority that signed the enclave’s certificate.

The advantage of sealing to the Sealing Authority over sealing to the Enclave Identity is that it allows offline migration of sealed data between enclave versions. The Sealing Authority may

sign multiple enclaves and enable them to retrieve the same seal key. These enclaves can transparently access data that was sealed by the other.

When sealing to a Sealing Authority, the older software should not be allowed to access data created by the newer software. This is true when the reason for releasing new software is to fix security issues. To facilitate this, the Sealing Authority has the option to prescribe a Security Version Number (SVN) as part of the Sealing Identity. *EGETKEY* allows the enclave to specify which SVN to use when producing the Seal Key. It will only allow the enclave to specify SVNs for its Sealing Identity or previous ones. When the enclave seals data, it has a choice to set the minimum SVN value of the enclave permitted to access that Sealing Key. This protects future secrets from access by old vulnerable software, but still enables a seamless upgrade transition where all previous secrets are available after the upgrade.

The SVN is not the same as a product version number. A product may have several versions, with different functionalities, but with the same SVN. It is the responsibility of the enclave writer to communicate to their customers (if necessary), which product versions have the same security equivalence.

4.2 Removing Secrets from a Platform

The architecture proposes a mechanism, known as OwnerEpoch, which enables the platform owner to change all the keys in the system by changing a single value. As OwnerEpoch is automatically included when requesting keys through the *EGETKEY* instruction, data objects which were sealed with a particular OwnerEpoch value can only be unsealed if the OwnerEpoch is set to the same value.

The primary purpose of this mechanism is to allow a platform owner to deny access to all sealed secrets on the platform, in a simple and recoverable step, before transferring the platform to someone else (permanently, or temporarily). Prior to transferring the platform, the platform owner can change the OwnerEpoch to a different value by using the hooks provided by the OEM. In case of a temporary transfer (such as for platform maintenance), the platform owner can restore OwnerEpoch to its original value once the platform is returned and can restore access to the sealed secrets.

5 RELATED WORK

Previous work to define hardware additions to a micro-processor chip for remote attestation [8] was designed for a single authority in hold of multiple platforms and was not easily scalable to support an unbound number of trusted authorities. In contrast, the Intel® SGX architecture allows for a plurality of distrusting enclaves, with complete separation between the enclaves and each one capable of securely communicating with one or more trusted authorities.

Commercially, the Trusted Platform Module 1.2 [5] has seen widespread introduction onto PC platforms in both discrete and integrated form. It supports a variety of usage models that include attestation of and sealing to the pre-boot and OS environments. Attesting and sealing to programs at the application level requires secret owners to evaluate the whole platform stack. Previous attempts to address this scalability issue have included [9]. The Intel® SGX architecture significantly reduces the software TCB by including only the enclave software in it and provides for the direct measurement of application software. This significantly reduces the attestation evaluation work any secret owner has to perform.

The TPM relies on the platform providing a Root of Trust of Measurement. To gain maximum trust when issuing TPM Attestation Certificates a binding to a platform credential identifying the Root of Trust for Measurement is required. By

providing the Root of Trust for Measurement along with the Root of Trust for Reporting, the Intel® SGX architecture allows us to use a single key to endorse and attest at the same time, significantly reducing the eco-system overhead in evaluating the trustworthiness of platforms.

6 CONCLUSIONS

In this paper, a novel hardware assisted mechanism in the form of a new ISA extension to Intel® Architecture is proposed to allow secure attestation and sealing to application software executing in a secure environment (referred to as an enclave).

The ISA extension provides the means for the enclave software to prove to another party that it has been properly instantiated on the platform and by extension is the correct software and is securely running inside an enclave on an enabled platform. This attestation does not require the verifier to understand the platform software context in which the enclave is executing and is limited to trusting the enclave software.

The Intel® SGX architecture also provides a mechanism to obtain a persistent unique key that could be used by the software to seal secrets and unseal the secrets later. The sealing mechanism also provides the ability to seamlessly transition secrets when an enclave's software is upgraded.

The attestation and sealing mechanisms, are defined in a scalable way that supports a plurality of enclaves running simultaneously, each handling its own secrets, and securely communicating with remote parties and proving to them that they are in compliancy with their security policies.

7 ACKNOWLEDGMENTS

The authors of this paper wish to acknowledge the contributions of many hardware and software architects and designers who have worked in developing this innovative technology.

8 REFERENCES

- [1] F. Mckeen, I. Alexandrovich, A. Brenenzon, C. Rozas, H. Shafi, V. Shanbhogue and U. Savagaonkar, "Innovative Instructions and Software Model for Isolated Execution," in *Hardware and Architectural Support for Security and Privacy*, 2013.
- [2] NIST, "FIPS 180-4:Secure Hash Standard," NIST, Gaithersburg, 2012.
- [3] NIST, "Special Publication 800-38B: Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication," May 2005. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-38B/SP_800-38B.pdf.
- [4] Intel, "Intel(r) 64 and IA-32 Architectures Software Developers Reference Manual," June 2013. [Online]. Available: <http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>.
- [5] Trusted Computing Group, "Trusted Platform Module Main Specification (TPM1.0)," March 2011. [Online]. Available: http://www.trustedcomputinggroup.org/resources/tpm_main_specification.
- [6] Trusted Computing Group, "Trusted Platform Module Library Specification (TPM2.0)," March 2013. [Online]. Available: http://www.trustedcomputinggroup.org/resources/tpm_library_specification.

- [7] E. Brickell and J. Li, "Enhanced privacy ID from bilinear pairing for Hardware Authentication and Attestation," *International Journal for Information Privacy, Security and Integrity*, vol. 1, no. 1, pp. 3-33, 2011.
- [8] J. S. Dworkin and R. B. Lee, "Hardware-rooted Trust for Secure Key Management and Transient Trust," in *CCS '07 Proceedings of the 14th ACM conference on Computer and communications security*, New York, 2007.
- [9] J. M. McCune, B. Parno, A. Perrig, M. K. Reiter and H. Isozaki, "Flicker: An Execution Infrastructure for TCB Minimization," *Proceedings of the ACM European Conference on Computer Systems (EuroSys'08)*, March 2008.
- [10] M. Hoekstra, R. Lal, P. Pappachan, C. Rozas, V. Phegade and J. Del Cuvillo, "Using Innovative Instructions to Create Trustworthy Software Solutions," in *Hardware and Architectural Support for Security and Privacy*, 2013.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

No computer system can provide absolute security under all conditions. Built-in security features available on select Intel® processors may require additional software, hardware, services and/or an Internet connection. Results may vary depending upon configuration. Consult your system manufacturer for more details.

Intel®, the Intel® Logo, Intel® Inside, Intel® Core™, Intel® Atom™, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and/or other countries. Other names and brands may be claimed as the property of others.

Copyright © 2013 Intel® Corporation