TECHNOLOGY GUIDE

Intel Corporation

Protecting AI/ML Workloads: A Hardening Guide Using Hardware-Backed Virtualization

White Paper

Authors

Eshe N Pickett

Joshua Silverio

Zerene Sangma

1 Introduction

As the consumer service industry moves towards a future with contactless interactions, self-service technology is bound to play a key role in delivering a quality user experience. Artificial Intelligence (AI) and Machine Learning (ML) enable contactless interaction via gesture recognition, facial recognition, speech recognition, language processing, audience analytics, and so on. When AI/ML models are deployed on Kiosks and Intelligent Edge devices, they run alongside multiple business-critical applications while presenting a new set of user privacy data and intellectual property (IP), all of which must be secured.

This white paper delves into key security points for customers developing a solution that uses AI and ML running on Intel Architecture (IA) for usages such as access control, payment, and authentication at interactive kiosks and other customer-facing terminals.

The use case focus for this paper is biometric authentication using facial recognition. Biometric authentication is an important foundational component for delivering on the promise of a more secure, more convenient, and touchless experience[i]. The paper assumes that businesses planning to adopt biometric authentication have reviewed the appropriate laws and regulations applicable to their respective industry, jurisdictions, and use-cases. This is an evolving landscape of the regulatory landscape that presents an opportunity to deploy hardened security solutions to promote the protection of sensitive data.[1]

At a minimum, systems using biometric authentication must account for data privacy, data breach notification, and data security among others. For additional security, the use of multi-factor authentication is recommended[ii].

Depending on the facial authentication usage there may be additional compliance requirements, streaming from sector-specific standards such as Fast ID Online (FIDO - a standard)[iii].

While laws and regulations are designed to encourage correct use of biometric technologies according to standard security and privacy practices, they must be implemented correctly to be effective.

[1] No product or component can be absolutely secure.

Table of Contents

Figures

Tables

1.1 Terminology

ABBREVIATION	DESCRIPTION
AI	Artificial Intelligence
IoT	Internet of Things
ML	Machine Learning
POC	Proof of Concept
VMM	Virtual Machine Monitor
VM	Virtual Machine
SOC	System-On-Chip
TCB	Trusted Computing Base
Workload	An application or service running on the compute device

1.2 Reference Documentation

DOCUMENT	DOCUMENT NUMBER /LOCATION
Intel Security Initiatives	https://newsroom.intel.com/press-kits/intel-security-initiatives/#gs.nn7mj3
Intel® Virtualization Technology	https://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html?wapkw=virtualization
Intel® AES-NI and Intel® Secure Key Instructions	Introduction to Intel® AES-NI and Intel® Secure Key Instructions
Microsoft Virtualization-based security	https://www.linkedin.com/learning/microsoft-cybersecurity-stack-advanced-identity-and-endpoint-protection/what-is-virtualization-based-security
Microsoft* Windows TPM Recommendations	https://docs.microsoft.com/en-us/windows/security/information-protection/tpm/tpm-recommendations
[MS-CSPP] Credential Security Support Provider (CredSPP) Protocol	https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-cssp/85f57821-40bb-46aa-bfcb-ba9590b8fc30
Enabling Arbitrary Code Guard	https://medium.com/@benoit.sevens/arbitrary-code-guard-cd74c30f8dfe
Intel(R) Control-flow Enforcement Technology	https://software.intel.com/content/www/us/en/develop/articles/technical-look-control-flow-enforcement-technology.html
Intel Architecture, 32-bit (IA-32)	https://software.intel.com/content/www/us/en/develop/articles/ia-32-intelr-64-ia-64-architecture-mean.html
Instruction Set Architecture (ISA)	https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-instruction-set-reference-manual-325383.pdf
Intel Converged Security and Management Engine (CSME)	https://www.intel.com/content/dam/www/public/us/en/security-advisory/documents/intel-csme-security-white-paper.pdf
Ai for Social Good (AI4SocialGood)	https://www.intel.com/content/www/us/en/artificial-intelligence/ai4socialgood.html

1.3 Purpose and Scope

The white paper proposes an architecture to incorporate security on Intel® IoT Platforms. Structured as a hardening guide, the architecture uses facial authentication as a sample usage to create a system to demonstrate how various components can be secured using hardware technologies available on Intel-based platforms, with software provided by the Windows* 10 operating system. The defined architecture uses a layered security approach, addressing each layer in accordance with the Defense in Depth (DiD) strategy[iv]. It identifies key assets to protect and secure and proposes technologies to help mitigate security threats.

2 Facial Authentication at Self-Service Interactive Kiosks

This section describes the architecture for facial authentication at a kiosk for access control or loyalty. A typical end-to-end flow can be described as:

1. Customers go to the kiosk to access their loyalty account or bank account
2. They are given the option to enroll into the loyalty program using their face as the ID
3. If they agree to enroll:
   - A photo of their face is taken at the kiosk
   - A biometric template (BT) is created from the photograph
   - The BT is associated with the user's loyalty or membership account.
4. On the next visit, the kiosk authenticates the customer, following these steps:
   - The camera stream captures user in front of the device
   - A temporary BT is created
   - The temporary BT is matched to the user's stored BT. If there is no match, the temporary BT is discarded
5. Once authenticated, the user can access their account information and perform related tasks.

This model can be applied to additional forms of biometric authentication such as voice, fingerprint, and iris recognition. The key principles and techniques introduced in this model can also be applied to other architectures requiring secure artificial intelligence (AI) and machine learning (ML) workloads.

Facial authentication is a form of facial recognition technology that confirms the identity of a user. The underlying technology for facial authentication is deep learning inference. In this process, the user's facial features are mapped and matched against a known set of facial feature vectors in a database.

Figure 1: Interactive banking kiosks and teller machines

The key functions in facial authentication are:

1. **Facial detection and feature extraction**

   On the computing device, such as a kiosk, an application with an inference engine receives an incoming camera stream containing the user's face. The inference engine employs a deep learning model, such as face detection, to detect and extract relevant points on the face and create a biometric template specific to the user's face. The biometric template is then stored in a secure location to reference for later matching.

2. **Face matching**

   A facial identification model is used to perform face matching by comparing the biometric template generated at the kiosk with the stored template. The face matching inference can be done in the cloud or an on-premises server.

3. **Liveliness detection**

   A deep learning model is trained to detect whether the incoming camera stream is a real live person or a static picture/video. One way to do this detection is to use a 3D camera to capture facial depth such as Intel® RealSense™.

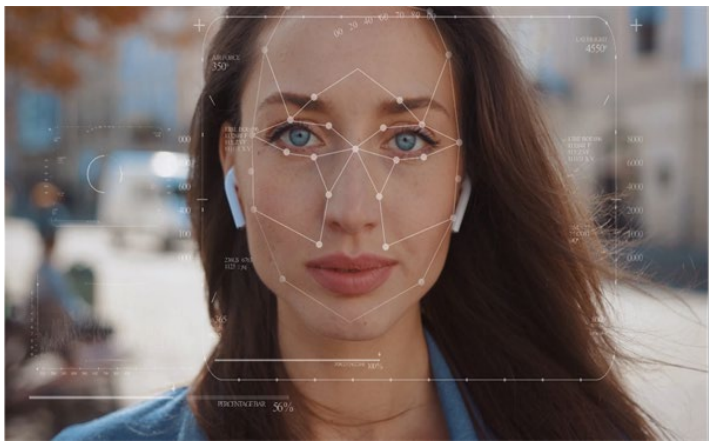

Figure 2: Facial Detection and Feature Extraction

An end-to-end architecture for such a system is depicted in Figure 3. The face detection, feature extraction, and liveliness detection models are deployed at the kiosk client end. The kiosk sends the biometric template to an on-premises or cloud server, where a face-matching deep learning model compares the template against the stored biometric templates. When a match is found, the deep learning model returns the matched user to the client as authenticated.

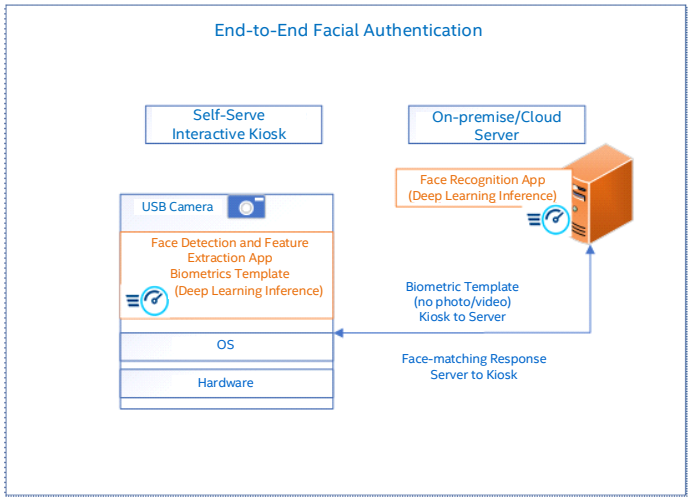

Figure 3: End-to-end Facial Authentication Architecture

3 Secure Facial Authentication

Successful deployment and usage of facial authentication or any other biometric authentication requires user data and privacy protection throughout the system as data flows from one point to another. In almost all cases, authentication data, biometrics template, or the IP responsible for generating biometric templates are the target of attackers and are considered as assets[v]. An asset must be protected at rest, in transit and during execution. Additional assets that need to be protected are the trained models deployed on the system, which are the intellectual property (IP) of the model provider.

This section analyzes security threats for the architecture introduced in section 2, it then covers modifications to the architecture by DiD mechanisms to introduce security capabilities provided by Intel platforms. These modifications can be enabled to provide an end-to-end secure facial authentication system.

While the security features introduced in this section provide ingredient-level capability, the features must be utilized by application developers and solution providers to deliver enterprise-grade security and privacy. Implemented security measures must be applied to protect all identified assets on the system relevant to the usage and environment.

3.1 Threat Analysis

This section focuses on analyzing the threats in the architecture proposed in section 2, using the CIA security principles. Attack surfaces and threats for IoT Edge devices and AI/ML models include, but are not limited to the following events[vi,vii]:

- Data theft
- Malicious code injection
- System boot process vulnerabilities
- Integrity attack (False Positive induction, Spoofing)
- Privacy violation
- Replay attack

3.1.1 Threat Analysis Scope

The platform must mitigate threats against the IoT edge device, using hardware and software capabilities. The technology capabilities in this document address the device, sensor, and application layers. Social engineering and behavioral attacks are not included in this analysis.

Providers must caution users against a malicious actor observing data input, masquerading as an employee, or otherwise exploiting user behavior. Users of the system remain a vulnerability and must be properly educated. This includes, but is not limited to:

- Credentials guarding
- Environmental awareness
  - Conscious of malicious actors while interacting with the device
- Exercising caution when sharing personal information

Physical device tampering device, such as replacing I/O components, opening the kiosk cabinet, or attaching sniffers must be addressed by the device manufacturer. AI models provided by third-party vendor are out of scope. The model provider must secure the training, test, and inference phases to prevent accuracy, impersonation, and poisoning attacks.

3.1.2 Assets

Table 1 lists the assets and the CIA principles[viii] associated with them. A threat analysis of these assets is required to build a robust facial authentication architecture. The assets listed are of value to the system and require protection. The Owner in the table refers to the party responsible for securing the asset.

Table 1: Threat Analysis for Facial Authentication System

ASSET	OWNER	DESCRIPTION	CIA PRINCIPLE
Application AI model	3rd Party Software Vendor	AI model for facial detection and biometrics template generation is secure and free of vulnerabilities.	Confidentiality, Integrity, Availability
Application secrets	Solution Architect	Application secret like Encryption keys, configuration data, and so on.	Confidentiality, Integrity, Availability
Backup Credential	Solution Architect	For use in case of facial recognition failure, this credential is used.	Confidentiality, Integrity
Biometrics application	3rd Party Software Vendor	Application is free of vulnerabilities and secure.	Confidentiality, Integrity, Availability
Biometrics template	3rd Party Software Vendor	In motion, sent from client to the cloud for identification.	Confidentiality, Integrity, Availability
Data from camera	Solution Architect	In motion, raw or encoded camera stream via USB data bus.	Confidentiality, Integrity, Availability

ASSET	OWNER	DESCRIPTION	CIA PRINCIPLE
Device Identity Key	Solution Architect	Used for attestation to remote services.	Confidentiality, Integrity, Availability
Face-matching result	3rd Party Software Vendor	Result of matched user for authentication.	Integrity, Availability

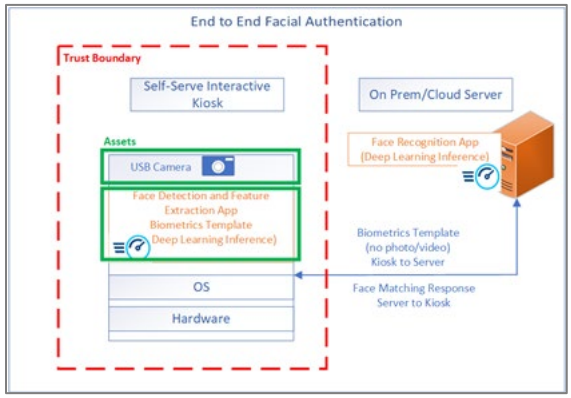

Figure 4: Threat Analysis

3.1.3 Adversaries

Table 2 lists relevant adversaries, who are potential attackers of the system, their capabilities, and the assets under threat from the specified attacker.

Table 2: Relevant Adversaries, Capabilities, Assets Under Threat

ADVERSARY	CAPABILITIES	ASSETS UNDER THREAT
Unprivileged Software Attacker	An adversary typically known as a "ring-3" attacker, whose capabilities are limited by IA-32 to those granted by the system software.	Biometrics template Biometrics application Application AI model Application secrets Face matching result Credentials
System Software	An adversary with full control over the operating system, or virtual machine monitor that can manipulate the IA-32 ISA in any manner allowed by the ISA specification.	Assets from unprivileged software attacker Data from camera
Network Attacker	An adversary that has access to and may have control over various network fabrics that are used to connect the platform to other platforms, intranet or extranet resources.	Device identity Face matching result

The attackers in the adversary list in Table 2 can execute attacks at the host-level to facilitate:

- Node capture
- Eavesdropping and inference
- Booting vulnerabilities
- Service interruption

For the VM where the AI workload is running, the attacks can result in:

- False data injection
- Data theft
- Access control exploitation

3.2 Security Capabilities

This section introduces various security capabilities available on Intel platforms that can be used to help protect the assets in Table 2. In the next section, the capabilities described here are used as building blocks to modify the facial authentication architecture at various stages to help improve the security posture. The key principles (on which these security capabilities are based) include:

- Defense-in-depth
- Workload and I/O isolation
- Data encryption
- Attestation (device and workload)
- Secure storage

The security capabilities described in this section are available (where indicated) on the specified Intel platforms and implemented in Microsoft's Windows* 10 operating system.

3.2.1 AI Model

The working assumption is that all AI models employed for face detection, biometrics template extraction, and face matching are accurate, and are constantly retrained to improve accuracy. The measures employed to improve accuracy are outside the scope of this paper, but they are necessary to be implemented. Without them, a wrong user could be authenticated, thereby nullifying all security measures on the device.

Well-known accuracy tests can be used as references to measure accuracy for facial authentication. The National Institute of Standards and Technology (NIST) Face Recognition Vendor Test (FRVT) is one such accuracy test. Solution providers may also do their own extensive testing to measure the accuracy based on location, demographics, and usage of the solution (refer to Appendix B).

For models to continuously improve and maintain accuracy, they must be *retrained.* The user's facial characteristics may also change due to weight, age, and facial hair etc. in which case the reference biometrics template may need to be updated. The model provider must comprehend such changes and accordingly offer a secure channel to update models or the reference template. The mechanism for which is outside the scope of this paper.

3.2.2 System-Level Security

Securing an IoT device requires that it is protected throughout its lifecycle, from boot to connectivity. The securing includes protection of data sources, communication, data storage, and applications[ix]. Verification of the hardware and software requires the establishment of a trust relationship between the device and its applications.

Intel® Security Essentials provide Root of Trust capabilities for devices, the network, and the cloud. Device-level capabilities that enable trusted execution include the technologies listed in Table 3 to establish a *chain of trust* rooted in silicon that makes the device more trustworthy and secure.

Table 3: Technologies with Root of Trust Capabilities

TECHNOLOGY	DESCRIPTION	PLATFORM INTERCEPT[2]
Intel® Boot Guard	Hardware root of trust verifies static early boot UEFI components provided in OEM firmware (Initial Boot Block code) as host CPU comes out of reset and transfers control to the OEM firmware.	Intel® 4th generation core processors onwards

[2] The introduction platform where the capability begins to have support.

TECHNOLOGY	DESCRIPTION	PLATFORM INTERCEPT[2]
Intel® Trusted Execution Technology	Hardware root of trust verifies the launched environment using Dynamic Root of Trust for Measurement (DRTM).	Intel® 3rd generation core processors onwards
UEFI Secure Boot	Part of Unified Extensible Firmware Interface (UEFI) specification that verifies integrity of boot software, pre-boot applications, and the operating system. The technique of measuring static early boot Unified Extensible Firmware Interface components is called the Static Root of Trust for Measurement (SRTM). UEFI Secure boot assumes that the system firmware is a trusted entity, ideally verified by a hardware root of trust like the Boot Guard technology. Secure Boot is also available for Virtual Machines (VMs) and provides an added layer of protection when enabled.	Intel® 4th generation core processors
Trusted Platform Module (TPM) 2.0	TPM (Trusted Platform Module) is a microchip installed on the motherboard of a computing device. It provides basic security functions including secure storage of credentials and encryption keys as well as storage for platform measurements to help ensure platform remains trustworthy. Intel platforms support two types of TPM solutions following the TPM 2.0 specifications] by Trusted Computing Group (TCG): 1. **Integrated TPM solution:** Intel® Platform Trust Technology (Intel® PTT) 2. **Discreet TPM solution:** provided by 3rd party hardware vendors ***Note:*** Different parts of the world may have a TPM specific to their regions' requirements.	Intel® 4th generation core processors
Intel® Total Memory Encryption	Full memory encryption ensuring all memory accessed from CPU is encrypted, on the external memory bus. Protects system memory against hardware attacks. ***Note:*** Next-generation Alder Lake platforms (2022) will allow Multiple Key TME and allow applications and guest VM to encrypt their data with their own keys, generated by the VMM.	Intel® 4th generation core processors
Intel Secure Device Onboard (FDO)	A "zero-touch" automated service that enables a device to be dynamically and securely provisioned at power-on. Eliminates the need for shipping default passwords.	

3.2.3 Workload Isolation

Isolating the AI workloads on a device by virtualizing the operating environment with its dedicated I/O resources enhances security by:

- Providing a dedicated guest virtual machine for each AI workload, which prevents other applications from sharing resources, thereby helping limit the attack surface
- Isolating the I/O hardware for the guest virtual machine such as USB device, memory, and network access
- Helping ensure that only signed driver, applications, and users can access the guest virtual machine
- Replacing a compromised virtual machine with a cloned VM that is in a known good state

Workload isolation by virtualization is the first step toward mitigating threats to the compute environment where the AI workload executes. Windows* 10 Hyper-V with Virtualization-Based Security (VBS) and Intel® Virtualization Technology

(Intel® VT) add additional capabilities to help secure the compute environment. Hyper-V is a standard feature for devices running Windows* 10 Enterprise. The Hyper-V (hypervisor) is backed by Intel® VT technology and must be enabled on the operating system. Virtual machines running a Windows* 10 Guest OS require a separate license.

3.2.3.1 Intel® Virtualization Technology (Intel® VT)

This technology provides hardware assist to the virtualization software, reducing its size, cost, and complexity. Optimizations are implemented to reduce the virtualization overhead occurring in cache, I/O, and memory. Refer to Intel Virtualization Technology for more information on various virtualization capabilities. Table 4 summarizes Intel® VT features enabled on Windows* 10 Hyper-V.

Table 4: Intel® VT Features in Windows* 10 Hyper-V

FEATURE/TECHNOLOGY	DESCRIPTION	WINDOWS* 10 HYPER-V	PLATFORM AVAILABILITY
Intel® Virtualization Technology Intel® 64 and Intel® Architecture (Intel® VT-x)	CPU Virtualization: Hardware support in the Intel processor to improve the Virtualization performance and robustness. Use of hardware transitions in the VMM strengthens the isolation of VMs and adds more security.	VM Monitor Mode Extension	All
Extended Page Tables (Part of Intel® VT-x)	Hardware-assisted page table virtualization.	Second Level Address Translation (SLAT)	All
Execute Disable Bit	Additional Memory check that marks memory locations as non-executable to prevent malicious code from exploiting buffer-overrun vulnerabilities.	Hardware Data Execution Prevention (DEP)	All
Intel® Virtualization Technology for Directed I/O (Intel® VT-d)	Provides hardware-assisted DMA-remapping for I/O devices securing direct memory access (DMA) of devices.	IOMMU (I/O Memory Management)	All
Hypervisor-Managed Linear-Address Translation Support	Allows OS SW to create a read-only block of memory designed to prevent viruses from accessing memory area thereby protecting OS kernel.	Hypervisor Managed	Feature available with Alder Lake (2022)

3.2.3.2 Windows* 10 Hyper-V with Virtualization-Based Security (VBS)

Windows* 10 and Windows* Server 2016+ implement Virtualization-Based Security (VBS) to enhance Windows* system security. VBS creates and isolates a secure region of memory from the normal operating system. Windows uses this *virtual secure mode* to host several security solutions that are available for host and guest virtual machines (VMs). Table 5 summarizes Windows* 10 Hyper-V security features that are used to mitigate threats mentioned in this paper. VBS is enabled by default on enterprise devices.

Table 5: Windows* 10 VBS Features

FEATURE NAME	SECURITY BENEFIT	STATUS	HARDWARE-ENFORCED
Windows Defender System Guard	Protects and maintains integrity of the system at start-up (including authentication stack, virtual TPM, system firmware, hardware, and so on). Validates system integrity maintained for remote attestation.	Enabled	Intel® TXT (DRTM), Intel® VT-X, Intel® VT-d
Hypervisor-Enforced Code Integrity (HVCI) (Memory Integrity)	Performs Code Integrity verification to prevent unsigned drivers or system files from being loaded into system memory.	On for host and guest OS	Intel® VT-X, Intel® VT-d

FEATURE NAME	SECURITY BENEFIT	STATUS	HARDWARE-ENFORCED
Windows Defender Application Control (WDAC)	Restricts the VM and device to run only authorized apps, protect the enforcement mechanism with HVCI.	Configured & Deployed for host and guest OS	Intel® VT-X, Intel® VT-d (when protected using HVCI)
Windows Defender Credential Guard	Protects NTLM, Kerberos, and credential manager secrets for host and VM.	Enabled	Intel® VT-X, Intel® VT-d

3.2.4 Application Security

In the use case for this paper, the application is provided by a third-party vendor. When considering any workload for inclusion in the end solution, the provider must develop with a *security first* mindset. The solution provider must confirm that the software vendor is following well-known processes from the Security Development Lifecycle (SDL). Following SDL ensures that mitigation measures for security threats are incorporated early in the software and hardware development lifecycles.

Note: Application security is outside the scope of this paper and is the responsibility of the application developer(s).

Application secrets such as credentials, and biometric template must be protected when at rest. One approach to protecting the application secrets is to store encryption keys locally in secure storage such as Trusted Platform Module (TPM), more details on this is covered in section 3.2.5.

All patches and updates for security bulletins on the system must be applied for both the host *and* the guest OSes. Just updating the host OS is not sufficient as each VM on the system is a separate installation. Patch management tools can ease the support burden and are advisable to use to reduce the likelihood of threat exposure[x].

Additional capabilities in Windows* 10 that are recommended to enhance the security of the host and the guest VM are shown in table below.

Table 6: Windows* 10 Application Security Features

FEATURE NAME	SECURITY BENEFIT	HARDWARE-ENFORCED
Arbitrary Code Guard and Control Flow Guard	Generates arbitrary code and enables control flow hijacking protection.	Intel® Control-Flow Enforcement Technology (Intel® CET)
Windows Defender Application Control (WDAC)	Mitigates risk of unauthorized applications running on the system by blocking unsigned scripts and MSIs from running. This protection, when enabled, helps to address file-based malware.	Protect WDAC enforcement mechanism with HVCI

3.2.5 Data Encryption

In order to protect assets - such as application secrets, biometric templates, and AI models, relevant code, and data must be encrypted at various stages. Data must be encrypted in two states to protect it (recommended):

- At rest: when data is stored on a hard drive or other storage device.
- In transit: when data is transported/moved between components or devices (locations).

For instance, in the facial authentication usage model, the biometrics template is moved from the kiosk to the cloud, where the biometrics identification application resides. The data and disk encryption must be extended to both host and guest VMs.

Intel provides various hardware-based capabilities for data encryption.

Table 7: Hardware-based Capabilities for Data Encryption

OFFERING	DESCRIPTION	PLATFORM AVAILABILITY
Intel® AES New Instructions (Intel® AES-NI)	Advanced instruction sets for application acceleration of data encryption.	All Intel® Core Processor family

OFFERING	DESCRIPTION	PLATFORM AVAILABILITY
Trusted Platform Module (TPM) 2.0	Secure Encryption Key storage on the device.	Intel® 4th generation core processors
Intel® Secure Key	Hardware-based cryptographic keys using Digital Random Number Generator (DRNG), Hardware	3rd Generation Intel® Core Processor family
Intel® Total Memory Encryption (Intel® TME)	Full memory encryption ensuring all memory accessed from CPU is encrypted, on the external memory bus at rest, at runtime and in transit. Protects system memory against hardware attacks. ***Note:*** Next-generation Alder Lake platforms (2022) will allow Multiple Key TME for Windows* systems and allow applications and guest VM to encrypt their own memory data with the key assigned to the VM or application by VMM.	Intel® 11th generation core processors

3.2.5.1 SecureKey Storage on Device

To store cryptographic keys securely, there are two well-known hardware modules—Trusted Platform Module (TPM) and Hardware Security Module (HSM).

A TPM is a microchip installed on the motherboard of a computing device to provide basic security functions. Key-storage of application secrets is one such function. The TPM communicates with the system by using a Low Pin Count (LPC) hardware bus. Multiple TPM implementation options are available.

- **A discrete TPM solution:** using a TPM chip deployed as a separate component in its own semiconductor package.
- **An integrated TPM solution:** using dedicated hardware integrated into one or more semiconductor packages alongside, but logically separate from, other components.
- **A firmware TPM solution:** running the TPM in the firmware in a Trusted Execution mode of a general-purpose computational unit.

Intel platforms support two types of TPM solutions [following the TPM 2.0 specifications]:

1. **Integrated TPM solution:** Intel® Platform Trust Technology (Intel® PTT)
2. **Discreet TPM solution:** provided by 3rd party hardware vendors.

Note: To use discrete TPM, check with the Kiosk hardware provider.

3.2.5.1.1 Intel® Platform Trust Technology (Intel® PTT)

Intel® PTT is a TPM 2.0 implementation integrated with the Intel® Chipset and SOC. Intel® PTT is supported by Windows* 10 Hyper-V for key management and storage. The integrated solution provides the benefit of BOM (bill of materials) cost reduction and real estate savings on the board.

Note: While the hardware implementation makes Intel® PTT resistant to software attacks, there may be region or compliance-specific requirements requiring discrete TPM

3.2.5.1.2 Windows* 10 Hyper-V Virtual TPM

Application secrets—such as configuration data, credentials, tokens, SSH keys, encryption keys, and so on, as well as the biometric template generated as an output by the application—must be encrypted to prevent them from being exposed or stolen. The required encryption keys can be stored in the TPM.

Windows* 10 Hyper-V offers Gen 2 virtual machines with a virtual TPM. Virtual TPM can be a physically discrete TPM or an integrated TPM, both of which can be easily enabled using Hyper-V. Microsoft provides a key storage API for application developers to use the TPM for key storage or use TSS (TCG Software Stack) APIs via TPM Base Services (TBS). TPM is also used by Hyper-V to protect data and the state of the VM. Table 8 describes the encryption features in Hyper-V.

Table 8: Hyper-V Encryption Features for VM Data Protection

ENCRYPTION FEATURE IN HYPER-V	DESCRIPTION
Enable TPM	Encrypts VM hard drive using BitLocker and TPM (disk encryption).
VM State and VM Migration Traffic	Encrypts virtual machine saved state and live migration traffic.
Shielded VM	Enables virtual TPM, encrypts VM state and Migration traffic using BitLocker* to prevent against compromised host.

3.2.6 Network Security

The assets requiring protection in transit are the biometric template and the identification token. These assets are transferred, from the edge device (kiosk) to the cloud, for recognition. Applications typically use RESTful APIs to send data across the network, which is the standard practice to use secured protocols for the data transfer. Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols must be used to exchange data across the network.

Windows* 10 provides the Windows Defender Firewall with Advanced Security to restrict traffic flow over unexpected ports. Common communication ports for REST APIs utilizing HTTPS over TCP include ports 443 and 80.

The VM network traffic can be isolated by software method by creating a dedicated Hyper-V virtual switch and assigning it dedicated VLAN (virtual LAN) to isolate the network traffic from the management host operating system. Since VM network traffic is still visible to the host, the host OS must be trusted and is part of the trusted computing base (TCB). Additional security considerations on the host OS are recommended in section 3.2.7.

Note: True network isolation is more secure and can be achieved with single-root I/O virtualization (SR-IOV), which allows the virtual machine network traffic to bypass the virtual switch and go directly to the physical network interface card (NIC). The capability is provided via Intel® Virtualization for directed I/O (Intel® VT-d). SR-IOV for network is not available on Windows* 10 Hyper-V and for usages requiring enhanced network security/ isolation. Windows* Server Hyper-V and VMware* Workstation are other VMMs that offer hardware virtualization via the Discrete Device Assignment feature and VMDirectPath I/O feature, respectively.

3.2.7 I/O Security

The camera stream input is a potential point of intrusion from an attacker. The assets requiring protection include the physical camera-to-device connection and the camera stream itself. Both these assets must be protected to ensure that the camera stream is not compromised. Windows* 10 provides the RemoteFX USB redirection feature for device such as webcam redirection to VMs, available under the Windows* RDP technologies.

Once the camera is redirected via RemoteFX USB redirection, it is no longer available to the host. In addition to securing the host with VBS and network security, additional group policy settings must be applied to control which devices can be restricted and made available for redirection, based on device ID and classes.

Note: For solution requiring additional security via I/O virtualization for I/O devices like camera, Windows* Server Hyper-V and VMware* Workstation are other VMMs that offer hardware virtualization via the Discrete Device Assignment feature and VMDirectPath I/O feature respectively.

3.2.8 Securing infrastructure

Securing overall system must include both securing the guest VMs and the Host OS of the virtualized infrastructure. VBS capabilities such as HVCI, Control Flow Guard for OS kernel mitigations must be applied to the host OS, and an efficient patch and security update strategy must be employed.

Self-service kiosks are user facing, interactive devices that run a single application. Windows* 10 provides the Kiosk Mode to run an authorized kiosk application in full screen mode, preventing the user access to other system functions. The virtualized facial authentication workload can run in the background, providing authentication information to the kiosk application via TCP/IP. While the Kiosk Mode prevents threats from the user at the device, additional steps must be taken at the host OS to prevent remote threats (remote network adversary), as outlined in Plan for Hyper-V security.

Note: TCB can further be reduced by running the Kiosk application in its own VM, thereby reducing the host OS attack surface by limiting workloads running in it. This may incur additional licensing cost while providing additional security to infrastructure.

The kiosk application also receives the face matching result and must be free of vulnerabilities to prevent replay attack. Refer to section 3.2.4 for more information on kiosk application development.

In the cloud-based implementation, facial authentication matching happens on a public cloud. Alternative architectures, where the matching is done on-device or in a private cloud, are feasible as well. In the scenario(s) where the facial matching is not done on-device or on an on-premises cloud within a trusted network, the ability to establish a trust relationship must exist.

In a trusted computing model, private and public keys are used to establish a relationship between the remote service and the physical device and/or the hosted application services. The device must be secured throughout its lifecycle from onboarding to retirement. Table 9 details Intel capabilities that enable trust relationship and lifecycle management.

Table 9: Intel Capabilities that Enable Trust Relationship and Lifecycle Management

INTEL CAPABILITIES	DESCRIPTION
Intel® Secure Device Onboarding (Intel® FDO)	An open, secure, zero-touch device provisioning solution. Integrates to cloud providers for establishing trust to remote devices.
Intel® Active Management Technology (Intel® AMT)	A technology to remotely discover, repair, and protect networked computing assets.
Intel® Remote Secure Erase (Intel® RSE)	Remotely wipe Solid State Drives to erase an infected OS in response to an attack. Remote erase can also be used to retire or repurpose a device. After erasing an infected OS, use Intel® AMT One-Click Recovery to recover device with a clean OS.

3.3 Modified Architecture for Secure Facial Authentication Architecture

This section outlines the modifications made to the original architecture to secure facial authentication services using security capabilities discussed in section 3.2.

3.3.1 Mitigations

In the original architecture, facial detection, feature extraction application, and the inferencing model, were all running alongside the native applications on the host OS. Hardware security features were not enabled in the platform.

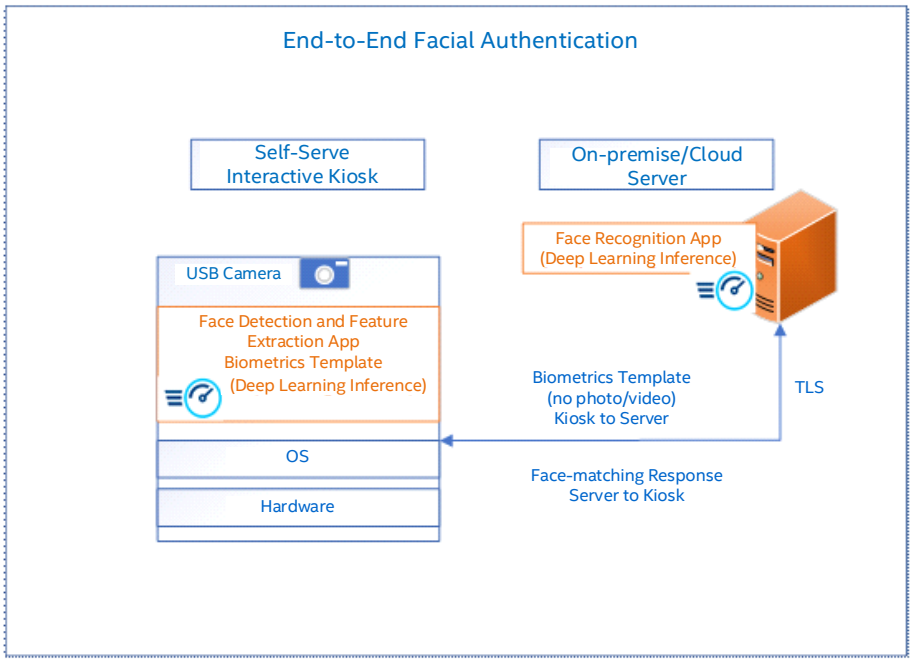

Figure 5: Original Architecture

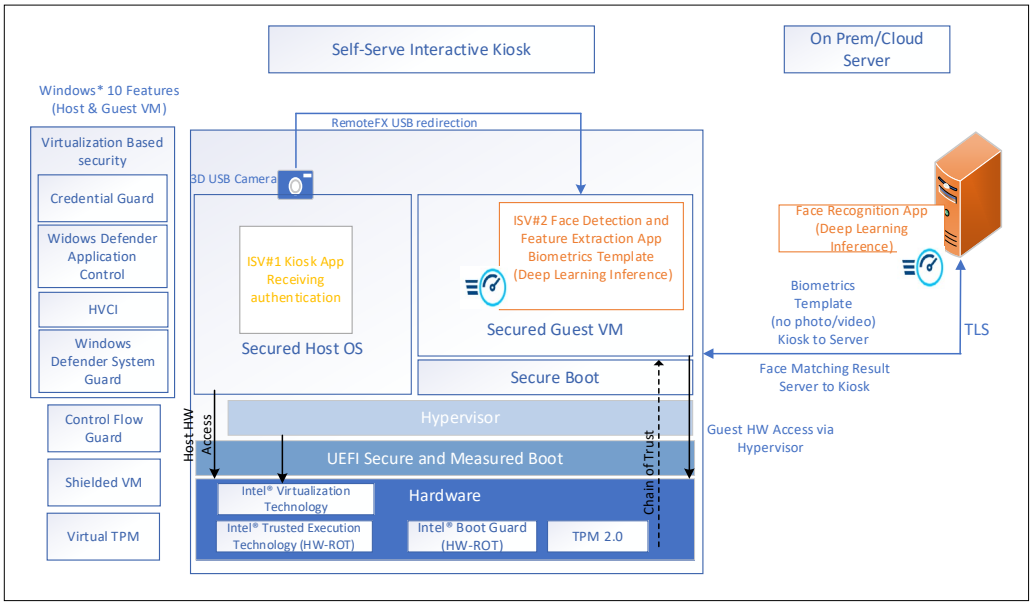

Figure 6: Modified Architecture

The modified architecture implements the security capabilities as discussed in the previous sections, which are summarized in Table 10.

Table 10: Summary of Security Capabilities Implemented by Modified Architecture

ASSET	THREAT	ADVERSARY	MITIGATION
Application AI Model	IP/ Data Theft	Network Attacker, Unprivileged Software Attacker, System Software Attacker	Virtualized OS (Isolate workload) VBS (HVCI, WDAC, Shielded VM, System Guard) Encrypted (vTPM) Multi-Key Total Memory Encryption (MK-TME)
AI Application and application secrets	Malicious Code Injection, System Boot Process Vulnerabilities. Integrity Attack	Network Attacker, Unprivileged Software Attacker, System Software Attacker	Virtualized OS (Isolate workload) VBS (HVCI, WDAC, Shielded VM, vTPM, Credential Guard, System Guard) Arbitrary Code Guard and Control Flow Guard Secrets and Encryption keys stored in vTPM (not in plain sight) Multi-Key Total Memory Encryption Isolated network traffic (via dedicated VLAN) Network virtualization via SR-IOV
Biometrics Template, Back up Credential	Data Theft, Integrity Attack (tampered)	Network Attacker, System Software Attacker	Encrypted (vTPM) TLS (in transit to Cloud) Isolated network traffic (via dedicated VLAN) Network virtualization via SR-IOV
Data from camera	Integrity Attack, Privacy Violation	System Software Attacker	VBS (WDAC, HVCI, System Guard, Credential Guard, applied to host) SR-IOV (I/O virtualization)
Device Identity Key	Integrity Attack	Network Attacker,	TLS (in transit between Cloud to Kiosk)
Face matching result supplied to the Kiosk application for authentication	Replay Attack, Integrity Attack	System Software Attacker, Network Attacker	Arbitrary Code Guard and Control Flow Guard (applied to Kiosk Application) Kiosk Mode

ASSET	THREAT	ADVERSARY	MITIGATION
			TLS (in transit from Cloud to Kiosk Application) VBS (WDAC, HVCI, System Guard, Credential Guard, applied to host or VM where the Kiosk application resides) Device Attestation Total Memory Encryption or MK-TME

4 Proof-of-Concept Implementation

A proof-of-concept (POC) implementation and performance impact of virtualization are described in this section.

4.1 System Setup

To fully test the technologies discussed previously, two different scenarios were implemented. The first scenario serves as a baseline system with no hardening strategies applied, while the second ties the various applications and data securities mentioned in Table 10 with no changes to the physical hardware, to create a system that is secured through hardware technologies available on Intel-based platforms along with software provided by the Windows* 10 operating system.

The facial recognition workload was provided by Yoonik and was deployed to both implementations while conducting performance metrics.

Table 11: POC Host Hardware Configuration

COMPONENT	DESCRIPTION
CPU	11th-Gen Intel® Core™ i5-1135G7
RAM	32 GB
Storage	500 GB
OS	Windows* 10 Enterprise Build 19041.928
Camera	Logitech Brio 4k and Intel® RealSense™

Table 12: POC Virtual Machine Hardware Configuration

COMPONENT	DESCRIPTION
vCPU	2
RAM	2 GB
Storage	150 GB
OS	Windows* 10 Enterprise 19041.928

4.2 Performance Summary

Table 13 details the performance summary of both POC implementations. The data was collected using Windows Performance Monitor. The Intel® RealSense camera was used for liveliness detection and a Logitech* Brio 4k camera was used for performance benchmarking.

The increase in system resources can be attributed to the greater overhead necessary to virtualize workloads through Hyper-V*. The modified architecture showed no discernible end-user degradation of the AI/ML workload even with this overhead.

Table 13: Implementation System Performance

IMPLEMENTATION	SYSTEM STATE	CPU USAGE	NETWORK USAGE	DISK USAGE (READ/WRITE PER SEC)	MEMORY USAGE
Usage #1	Idle	4%	< 1%	2	10%
Usage #1	Original Architecture (Native)	13%	< 1%	47	13.78%
Usage #2	Modified Architecture (Virtualized)	21.14%	< 1%	13	23%

5 Summary

When using technologies like facial authentication at kiosks, it is recommended that you are aware of laws and regulations of the locations where they are likely to be deployed. Laws are designed to encourage use of security technologies as they protect consumers when implemented correctly—more security, more convenience, and better user experience with a contactless interaction.

Sensitive AI/ML workloads—such as facial authentication at Self-Service Kiosks—can be secured by introducing virtualization at the kiosk, and isolating the sensitive workload along with its I/O and network components from other applications on the kiosk. Intel hardware-based technologies are available to enhance the performance of a virtualized system while securing the workload without impacting user experience. The hardware capabilities are deployed via OS technologies and available for application vendors to take advantage of in securing their applications as well as for Solution providers to secure the kiosk infrastructure.

In this paper, a Windows* enterprise-based Self-Service Kiosk Architecture was modified to isolate the facial authentication workload in a Hyper-V-based Windows* 10 VM. The Hyper-V VMM along with Virtualization Based Security is equipped with multiple security capabilities to secure the workload and the infrastructure, along with virtual TPM to store application secrets.

The solution, however, is dependent on a vulnerabilities-free host OS for USB camera and network packet redirection. This gap can be mitigated by minimizing the host OS attack surface, which can be achieved by minimizing the production workload on the host OS, such as a kiosk application that can be hosted on another guest VM. This may incur additional OS licenses. Another mitigation is to utilize SR-IOV (Single Root I/O virtualization) for the network provided by Intel® VT-d, which is available with many VMMs in the market today—such as Windows* Server Hyper-V and VMware* Workstation, among others.

While Intel provides the necessary building blocks to deploy more secure solutions, solution providers must proactively audit and maintain their components and applications throughout the device's lifecycle[3].

[3]No product or component can be absolutely secure.

Appendix A Independent Software Vendor ECO-System Partners

Examples of currently deployed Biometric solutions in the market:

1. Self-Service VR eKiosk allows biometric authentication for purchase of prepaid SIM cards: https://www.ntsretail.com/vr-ekiosk-innovative-self-service-solution-launches-23-locations-across-germany
2. X5 Retail Group uses algorithms from VisionLabs solutions to deploy face recognition payment solutions: https://www.electronicpaymentsinternational.com/news/x5-retail-debuts-facial-recognition-payment-system-in-russia/
3. San Diego restaurant taps facial recognition kiosk as part of COVID recovery: https://www.kioskmarketplace.com/articles/san-diego-restaurant-taps-facial-recognition-as-part-of-covid-recovery/
4. VeriTrans Provides Facial Recognition Payment Service for Shizuoka Prefecture's Facial Recognition System PoC with NEC: https://www.garage.co.jp/en/pr/release/2021/03/20210308/
5. Software used to develop the POC in this document is provided by YooniK, a SaaS platform that can authenticate customers on any IoT device to prove their identity: https://yoonik.me/

Appendix B Facial Recognition Accuracy Measures

Face Recognition Vendor Test (FRVT) Ongoing provides ongoing published reports that provide accuracy measures for various vendors. See nist.gov/programs-projects/face-recognition-vendor-test-frvt-ongoing.

Labeled Faces in the Wild (LFW) is a public benchmark for algorithms performing face verification, for example, www.cs.umass.edu/lfw/results.html.

MegaFace has one of the largest public datasets for facial recognition accuracy benchmarking on megaface.cs.washington.edu.

Dataset / Stats	MegaFace (this paper)	CASIA-WebFace	LFW	PIPA	FaceScrub	YouTube Faces	Parkhi et al.	CelebFaces	DeepFace (Facebook)	NTechLab	FaceNet (Google)	WebFaces Wang et al.	IJB-A IAPRA
#photos	1,027,060	494,414	13K	60K	100K	3425 videos	2.6M	202K	4.4M	18.4M	>500M	80M	25,813
#subjects	**690,572**	10,575	5K	2K	500	1595	2.6K	10K	4K	200K	>10M	N/A	500
Source of photos	**Flickr**	Celebrity search	Yahoo News	**Flickr**	Celebrity search	Celebrities on YouTube	Celebrity search	Celebrity search	Internal	Internal	Internal	Web crawling	Internal
Public/private dataset	**Public**	**Public**	**Public**	**Public**	**Public**	**Public**	Private	Private	Private	Private	Private	Private	**Public**

The table above is a representative sample of recent face recognition datasets (in addition to LFW). Current public datasets include up to 10K unique people, and a total of 500K photos.

Several companies have access to more photos and subjects; these, however, are subject to privacy constraints and are not public. Mega Face includes 1M photos of more than 690K unique subjects, collected from Flickr (from creative commons photos), and is available publicly; refer to, The Mega Face Benchmark: 1 Million Faces for Recognition at Scale.

Document Revision History

REVISION	DATE	DESCRIPTION
001	July 2021	Initial release.

6 Bibliography

[i] Global 360 Research Team at Frost & Sullivan. "*The Top Trends for 2020: Disruptive Technologies Go Mainstream"*. 2020; Gokulan, Dhanusha. *UAE: Your face is now your passport at Dubai airport*. Khaleej Times, 22 Feb. 2021, khaleejtimes.com/news/uae-your-face-is-now-your-passport-at-dubai-airport. Accessed 23 February 2021.

[ii] Federal Bureau of Investigation. "*Cyber Criminals Use Social Engineering and Technical Attacks to Circumvent Multi-Factor Authentication"*. Private Industry Notification. PIN Number 20190917-001. 2019.

[iii] FIDO Alliance. "*FIDO Alliance Specifications Overview".* fidoalliance.org/specifications. Accessed 24 February 2021.

[iv] Jenkins, Ira Ray. "*Defense in Depth of Resource-Constrained Devices"*: digitalcommons.dartmouth.edu/dissertations/59. 2020.

[v] N. Messe, V. Chiprianov, N. Belloir, J. El-Hachem, R. Fleurquin and S. Sadou, "*Asset-Oriented Threat Modeling"*, in IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), 2020, pp. 491-501, doi: 10.1109/TrustCom50675.2020.00073.

[vi] Q. Liu, P. Li, W. Zhao, W. Cai, S. Yu and V. C. M. Leung, "*A Survey on Security Threats and Defensive Techniques of Machine Learning: A Data Driven View*" in IEEE Access, vol. 6, pp. 12103-12117, 2018.; V. Hassija, V. Chamola, V. Saxena, D. Jain, P. Goyal and B. Sikdar, "*A Survey on IoT Security: Application Areas, Security Threats, and Solution Architectures*" in IEEE Access, vol. 7, pp. 82721-82743, 2019.

[vii] "*RIPTE: Runtime Integrity Protection Based on Trusted Execution for IoT Device*". hindawi.com/journals/scn/2020/8957641. Accessed 19 May 2021.

[viii] Kota, S. N. Swamy and S. R., Standards for Security Categorization of Federal Information and Information Systems, "*Security: Application Areas, Security Threats, and Solution Architectures*" in IEEE Access, vol. 7, pp. 82721-82743, 2019.

[viii] "*RIPTE: Runtime Integrity Protection Based on Trusted*" nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf. 2004. Accessed 9 June 2021.

[ix] Kota, S. N. Swamy and S. R., "*An Empirical Study on System Level Aspects of Internet of Things (IoT)"*, in IEEE Access, vol. 8, pp. 188082-188134, 2020.

[x] NIST. "*Guide to Enterprise Patch Management Technologies"*. www.nist.gov/publications/guide-enterprise-patch-management-technologies. 2013. Accessed 16 March 2021.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's Global Human Rights Principles. Intel's products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

0721/JJ/WIT/PDF 632038-001US