

Windows Secured-Core PC

Enterprise Evaluation Guide

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT. This document is provided “as-is.” Information and views expressed in this document, including URL and other Internet website references, may change without notice. You bear the risk of using it. Copyright © 2021 Microsoft Corporation. All rights reserved. Please refer to Microsoft Trademarks (<https://aka.ms/MSTrademarks>) for a list of trademarked products. The names of actual companies and products mentioned herein may be the trademarks of their respective owners

Contents

Windows Secured-Core PC	1
Securing your enterprise with Windows Secured-Core PC.....	5
How to use this reviewers guide.....	5
About Windows Secured-Core PC.....	5
Threats and complexities.....	5
Criminal groups are evolving their techniques	5
Nation-state actors are shifting their targets	6
Ransomware continues to grow as a major threat.....	6
How Windows 11 enables Zero Trust protection	6
Windows Secured-Core Benefits	7
Protecting identities from external threats	7
Securing the operating system from malware.....	7
Defending against hardware and firmware attacks.....	8
Who makes Windows secured-core Personal Computers?.....	9
Evaluation considerations.....	9
Decreasing management costs and improving user experience	9
Security Baselines	9
End user device experience	10
Deploying secured-core PC's.....	10
Testing.....	10
How to determine if your device is running as a secured-core PC.....	11
Success metrics	13
Management guides for Windows security services	13
Windows Defender Application Guard	13
Windows Defender Credential Guard.....	13
Windows Defender System Guard.....	13
Diving deeper into the technologies.....	14
Virtualization Based Security	14
Windows Hello	14
Windows Defender Credential Guard.....	15
Secure Boot	15
Early Launch Anti-Malware	16

Measured Boot	16
Windows Defender System Guard.....	17
Protecting peripherals	17
Differences between a Windows Secured-Core PC and Windows 11 PC.....	18
Some real-world examples of attacks.....	18
RobbinHood Ransomware	19
Thunderspy DMA Peripheral Attack	19
Resources (blog posts on malware protection by SCPC, Docs, etc.).....	21

Securing your enterprise with Windows Secured-Core PC

Thank you for reading the Windows secured-core PC Enterprise Evaluation guide.

How to use this reviewers guide

This document will help you understand the benefits provided by a Windows Secured-Core PC.

The guide describes how to understand these features and provides you with guidance and links to start exploring our protection benefits.

About Windows Secured-Core PC

Windows secured-core PCs are designed to provide an “**On-By-Default**” security experience for businesses and users. What this means is that from the moment a user takes the device out of the box and turns it on, businesses can be confident that the machine has protections in place that mitigate security risks and simplify the end user experience in configuring the device.

Microsoft works closely with OEM partners to help ensure that all certified Windows systems deliver a secure operating environment. Windows interoperates closely with the hardware to deliver protections that take advantage of available hardware capabilities.

Secured-core PC uses a variety of approaches to protect your device (and your data on that device) from malware, physical possession issues (like if it is lost, stolen, or confiscated) and access attacks such as Thunderspy. If access to your device is compromised, a Secured-core PC device helps ensure that your data remains secure and inaccessible to attack.

By using technologies like virtualization-based security (VBS), hypervisor-protected code integrity (HVCI), Trusted Platform Module (TPM) integrated identity protection and BitLocker Drive Encryption (a Windows feature that can be enabled), these devices help protect user data, identity, and the device itself from attack. We will go into more detail on this later in this document.

Threats and complexities

In the modern world threat actors have rapidly increased in sophistication over the past years, using techniques that make them harder to spot and that threaten even the savviest targets. For example, nation-state actors are engaging in new reconnaissance techniques that increase their chances of compromising high-value targets, criminal groups targeting businesses have moved their infrastructure to the cloud to hide among legitimate services, and attackers have developed new ways to scour the internet for systems vulnerable to ransomware.

In addition to attacks becoming more sophisticated, threat actors are showing clear preferences for certain techniques, with notable shifts towards credential harvesting and ransomware, as well as an increasing focus on exploiting hardware attacks on devices.

Criminal groups are evolving their techniques

Criminal groups are skilled and relentless. They have become adept at evolving their techniques to increase success rates, whether by experimenting with different phishing lures, adjusting the types of attacks they execute or finding new ways to hide their work.

Over the past year cybercriminals are playing their well-established tactics and malware against human curiosity and need for information. Attackers are opportunistic and will switch lure themes daily to align with news cycles, as seen in their use of the COVID-19 pandemic to stoke fear and trick people into doing things they should not. While the overall volume of malware has been relatively consistent over time, these campaigns have been used for broadly targeting consumers, as well as specifically targeting essential industry sectors such as health care.

Nation-state actors are shifting their targets

Nation-states have shifted their targets to align with the evolving political goals in the countries where they originate.

In recent years there has been an important focus on vulnerabilities in critical infrastructure and have now expanded to other types of organizations. In fact, 90% of nation-state notifications in the past year have been to organizations that do not operate critical infrastructure. Common targets have included nongovernmental organizations (NGOs), advocacy groups, human rights organizations and think tanks focused on public policy, international affairs, or security. This trend may suggest nation-state actors have been targeting those involved in public policy and geopolitics, especially those who might help shape official government policies.

Ransomware continues to grow as a major threat

Governmental security organizations have continued to warn about ransomware, especially its potential use to disrupt government operations. Encrypted and lost files and threatening ransom notes have now become the top-of-mind fear for most executive teams. Attack patterns demonstrate that cybercriminals know when there will be change freezes, such as holidays, that will impact an organization's ability to make changes (such as patching) to harden their networks. They are aware of when there are business needs that will make organizations more willing to pay ransoms than incur downtime, such as during billing cycles in the health, finance, and legal industries.

Additionally, human-operated ransomware gangs are performing massive, wide-ranging sweeps of the internet, searching for vulnerable entry points, as they "bank" access – waiting for a time that is advantageous to their purpose.

How Windows 11 enables Zero Trust protection

The [Zero Trust Principles](#) are threefold. First, **verify explicitly**. That means always authenticate and authorize based on all available data points, including user identity, location, device health, service or workload, data classification, and anomalies. Second, use **least-privileged access**. Limit user access with just-in-time and just-enough-access (JIT/JEA), risk-based adaptive policies, and data protection to help secure both data and productivity. And lastly, **assume breach**. Operate in a manner that minimizes blast radius and segment access. Verify end-to-end encryption and use analytics to get visibility, drive threat detection, and improve defenses.

Windows 11 provides chip to cloud security, giving IT administrators the attestation and measurements to determine whether a device meets requirements and can be trusted. And Windows 11 works out of the box with Microsoft Intune and Azure Active Directory, so access decisions and enforcement are seamless.

Windows Secured-Core Benefits

Windows Secured-Core PCs provide several primary benefits to organizations and users. The most important of those benefits are that the protections are **turned on by default** so users will know that their device is protected right out of the box. Additionally, they provide these 3 core pillars of protection:

Protecting identities from external threats

Passwords alone are not enough to protect system data and identities. Secured-core PC ensures that user identities and credentials are protected against theft, compromise, and phishing attacks.

People are using the same password repeatedly and are more exposed than ever. Microsoft wanted to help improve the security of identity and partnered with OEMs to help address this.

Windows Hello prevents phishing and credential-based attacks through a combination of biometric sensors and hardware-based credential storage. Using your face, fingerprint, secure FIDO2 key, or PIN, Windows Hello allows you to sign-in password-free and gives you the fastest, most secure way to unlock your Secured-core PC devices.

Windows Defender Credential Guard, an optional service that can be enabled, leverages Virtualization-Based Security (VBS) to prevent identity attack techniques such as Pass-the-Hash and Pass-the-Ticket and isolates company confidential information so that only privileged system software can access it.

Credential Guard helps prevent attacks through virtualization-based security. When Credential Manager domain credentials, NTLM, and Kerberos derived credentials are protected using virtualization-based security, the credential theft attack techniques and tools used in many targeted attacks are blocked. Secured-core PCs are built on the principles of assume breach and defense in depth. These devices have virtualization-based security (VBS) turned on by default meaning that even if an attacker gains administrative privileges through malware, authentication tokens are better protected in an isolated environment. To use Windows Hello with biometrics specialized hardware, including fingerprint reader, illuminated IR sensor, or other biometric sensors is required. Hardware based protection of the Windows Hello credential/keys requires TPM 2.0 or greater.

Securing the operating system from malware

Along with ensuring a small, trusted computing base by establishing a hardware root of trust, Secured-core PC ensures that code running within that trusted computing base runs with integrity and is not subject to outside exploit or attack.

Secured-core PCs use policies enabled with Hypervisor Enforced Code Integrity (HVCI) to check system software before it is loaded, and only start executables that are signed by known, approved authorities. HVCI runs in the Virtualization Based Security (VBS), which protects it from outside attack. Kernel mode code integrity checks all kernel mode drivers and binaries before they are started and prevents unsigned drivers or system files from being loaded into system memory. This ensures that only code from trusted

and verified sources can run on the system, providing strong protections against kernel viruses and malware.

The hypervisor, the most privileged level of system software, sets and enforces page permissions across all system memory. Pages are only made executable after code integrity checks inside the secure region have passed, and executable pages are not writable. With Kernel Memory Protection turned on by default, Windows secured-core PCs better defend against malware attempting to modify kernel memory and kernel code.

Defending against hardware and firmware attacks

At the heart of the Secured-core PC promise is a hardware-enforced root of trust that means a Secured-core PC runs in a clean, trusted state, regardless of the state of any firmware code that runs before the device boots.

To do so, Secured-core PC leverages the Trusted Platform Module 2.0 (TPM) and a modern, capable CPU with dynamic root of trust measurement (DRTM) capability to boot securely and minimize the impact of firmware vulnerabilities.

To protect critical resources such as the Windows authentication stack, single sign-on tokens, the Windows Hello biometric stack, and the Virtual Trusted Platform Module, a system's firmware and hardware must be trustworthy.

Windows Defender System Guard reorganizes the existing Windows 11 system integrity features under one roof and sets up the next set of investments in Windows security. It's designed to make these security promises:

- Protect and maintain the integrity of the system as it starts up
- Validate that system integrity has truly been maintained through local and remote attestation

While Windows Defender System Guard provides advanced protection that will help protect and maintain the integrity of the platform during boot and at run time, the reality is that we must apply an "assume breach" mentality to even our most sophisticated security technologies. We should be able to trust that the technologies are successfully doing their jobs, but we also need the ability to verify that they were successful in achieving their goals. When it comes to platform integrity, we can't just trust the platform, which potentially could be compromised, to self-attest to its security state. So Windows Defender System Guard includes a series of technologies that enable remote analysis of the device's integrity.

PCI devices are DMA-capable, which allows them to read and write to system memory at will, without having to engage the system processor in these operations. The DMA capability is what makes PCI devices the highest performing devices available today. These devices have historically existed only inside the PC chassis, either connected as a card or soldered on the motherboard. Access to these devices required the user to turn off power to the system and disassemble the chassis.

This is no longer the case with hot plug PCIe ports (e.g., Thunderbolt™ and CFXpress). Hot plug PCIe ports such as Thunderbolt™ technology have provided modern PCs with extensibility that was not available before for PCs. It allows users to attach new classes of external peripherals, such as graphics cards or other PCI devices, to their PCs with a hot plug experience identical to USB. Having PCI hot plug ports externally and easily accessible makes PCs susceptible to drive-by DMA attacks.

Drive-by DMA attacks are attacks that occur while the owner of the system is not present and usually take less than 10 minutes, with simple to moderate attacking tools (affordable, off-the-shelf hardware and software) that do not require the disassembly of the PC. A simple example would be a PC owner leaves the PC for a quick coffee break, and within the break, an attacker steps in, plugs in a USB-like device and walks away with all the secrets on the machine, or injects a malware that allows them to have full control over the PC remotely.

Windows leverages the system Input/Output Memory Management Unit (IOMMU) to block external peripherals from starting and performing DMA unless the drivers for these peripherals support memory isolation (such as DMA-remapping). Peripherals with DMA Remapping compatible drivers will be automatically enumerated, started, and allowed to perform DMA to their assigned memory regions.

By default, peripherals with DMA Remapping incompatible drivers will be blocked from starting and performing DMA until an authorized user signs into the system or unlocks the screen.

Who makes Windows secured-core Personal Computers?

The current list of Windows secured-core PC's is available here:

<https://www.microsoft.com/en-us/windowsforbusiness/view-all-devices?col=secured-core-pcs>

Evaluation considerations

Decreasing management costs and improving user experience

When businesses deploy new machines to users, they typically approach it through one of two ways:

1. Deploy a custom corporate image to the device which contains their security settings and software, or
2. Have the user join the device to their Active Directory or Azure AD instance.

In the first case the business expends resources to create and maintain the custom image and deployment scripts for their environment. When a major new release of Windows occurs, these images must be updated and tested to ensure that they still perform as expected.

With the second approach the settings are pushed to the device as part of the normal system management process. This can require the user to perform multiple restarts of the machine as services and hardware devices such as the Trusted Platform Module are provisioned and configured. It also can delay key policies from appearing on the device immediately.

A Windows secured-core PC helps address these issues and can minimize the amount of effort required to ensure that the business environment is configured correctly. Evaluating a Windows secured-core PC should include an audit of the existing deployment experience in your business.

Security Baselines

Every organization faces security threats. However, the types of security threats that are of most concern to one organization can be completely different from another organization. For example, an e-commerce company may focus on protecting its Internet-facing web apps, while a hospital may focus on protecting confidential patient information. The one thing that all organizations have in common is a

need to keep their apps and devices secure. These devices must be compliant with the security standards (or security baselines) defined by the organization.

A security baseline is a group of Microsoft-recommended configuration settings that explains their security impact. These settings are based on feedback from Microsoft security engineering teams, product groups, partners, and customers.

For example, there are over 3,000 Group Policy settings for Windows 11, which does not include over 1,800 Internet Explorer 11 settings. Of these 4,800 settings, only some are security related. Although Microsoft provides extensive guidance on different security features, exploring each one can take a long time. You would have to determine the security impact of each setting on your own. Then, you would still need to determine the appropriate value for each setting.

A Windows secured-core PC provides a pre-configured, on by default implementation of the Microsoft Security Baseline to help simplify and ensure that devices are protected correctly.

End user device experience

When a user receives a new device, they often want to start using it immediately. This desire can run into the reality of deploying and provisioning security settings. Many users will try and limit the number of times that they restart a machine early in the acquisition or deployment cycle so that they can transfer personal data or settings. Deploying Windows and the security settings to a business device either as a custom image or through group policy, several of the settings and services can require Windows to restart to ensure that provisioning and configuration is performed correctly.

For example, the configuration of Windows Defender Credential Guard and the provisioning of the underlying TPM device may require that the user restart the device to ensure that the service is running. Until that restart is performed the device may be vulnerable to identity attacks.

The best experience for both users and the business is to ensure that the required services and hardware are provisioned and operational before the first user interaction with the device. A secured-core PC arrives pre-configured with these on by default and thus provides the best possible experience for the end user while still achieving the business goals of ensuring that the device is secured out of the box.

Deploying secured-core PC's

Very few organizations can plan to deploy devices from scratch. Often new devices are deployed either as part of a device refresh cycle or to new employees as they are onboarded. When evaluating Windows secured-core PC's you should consider how you currently provide new devices to employees and the non-asset cost to do so.

Understanding the applications, settings and usage patterns of your devices will assist you in both evaluating Windows secured-core PC's and in the deployment process.

Testing

Protecting user's identity, operating system, hardware peripherals and firmware can have impacts on your applications and peripheral devices. A robust testing plan that is designed to cover the applications, use cases and peripheral devices can help to ensure that the deployment is conducted as easily as possible.

Examples of the areas that need to be tested are:

- Authentication breaks applications – Applications that utilize NTLM v1, Kerberos DES encryption or Kerberos unconstrained delegation will not work correctly on a secured-core PC. If you have applications that use these technologies then you will need to consider how to mitigate them (update, transition, etc.) to successfully deploy a secured-core PC.
- Authentication exposure – Applications that make use of MS-CHAP v2, credential delegation or digest authentication will work but there is increased risk as these applications will cause a system identity prompt that may expose the credentials. As above you need to consider the mitigation approach.
- Credential Performance impacts – Applications that expect to be able to hook the credential authentication process will have decreased performance due to the isolated nature of the Windows Defender Credential Guard (WDCG) process.
- General Performance impacts – Virtualization Based Security (VBS) can have an impact on applications. Profiling your applications running both with and without VBS can provide insight into the application architecture that impacts performance. Again, a mitigation process needs to be considered.
- Peripheral driver limitations – When using Hypervisor Code Integrity (HVCI) you will need drivers that support HVCI and are signed using an Extended Validation (EV) certificate. This includes both external and internally developed drivers. To mitigate this, you will either need to modify and sign your own drivers or update external devices and drivers to support HVCI.


Additional testing around how you expect devices to be used can help discover issue prior to deployment.


How to determine if your device is running as a secured-core PC

Windows provides an application that can help you determine the status of a device and whether or not it is compatible with the policies and settings of a secured-core PC.

The Windows Security application is installed as part of your normal Windows 11 installation and provides information regarding the status of the underlying hardware and whether Windows is configured to use those features.

There are 3 primary systems that the application will provide information on. The first is “Core Isolation”, this refers to whether your system can support virtualization-based security (VBS). The text presented indicates whether the hardware supports it but is not configured, the hardware supports it and is configured for VBS or the hardware supports it, and the system is configured for VBS and Hypervisor Code Integrity (HVCI).





The second area is the “Security Processor”. This refers to the Trusted Platform Module (TPM) which is used to store secrets like your biometric logon information or encryption keys. Again, this section will identify if you have a TPM and whether it is configured for use by Windows.

The final area refers to “Secure Boot”. This section provides you with information concerning the

status of the secure boot in the device UEFI settings. If all these items indicate that Windows is configured to utilize the hardware, then you can be assured that you device is running the full secured-core PC experience. If these items do not show that the underlying hardware exists, then you will not be able to configure the device to use the full secured-core PC settings.

The images above are examples of what you will see. Depending on the version of Windows in use you may see the text in these locations as indicated by the table below:

Hardware capabilities	Current Windows release (20H2)	Future releases
No DRTM, HVCI-capable, TPM 2.0, Secure Boot ON	Your device meets the requirements for standard hardware security.	Your device meets the requirements for standard hardware security.
No DRTM, HVCI-enabled, TPM 2.0, Secure Boot ON	Your device meets the requirements for enhanced hardware security.	Your device meets the requirements for enhanced hardware security.
DRTM + SMM 10 (PPAM), HVCI-enabled, TPM 2.0, Secure Boot ON	Your device has all Secured-core PC features enabled. [Core Isolation] Your device meets firmware protection version one	Your device has all Secured-core PC features enabled. [Core Isolation] Your device meets firmware protection version one
DRTM + SMM 20 (DGR), HVCI-enabled, TPM 2.0, Secure Boot ON	Your device has all Secured-core PC features enabled.	Your device has all Secured-core PC features enabled.

	[Core Isolation] Your device meets firmware protection version two	[Core Isolation] Your device meets firmware protection version two
DRTM + SMM 30, HVCI-enabled, TPM 2.0, Secure Boot ON	Your device has all Secured-core PC features enabled. [Core Isolation] Your device meets firmware protection version three	Your device has all Secured-core PC features enabled. [Core Isolation] Your device meets firmware protection version three

Success metrics

It is important to create success metrics for a deployment and to gather a baseline for comparison. For example, conducting surveys into the out of box experience (OOBE) for your existing deployment process can provide a baseline to enable you to evaluate the success of the new OOBE experience with secured-core PC's.

Success metrics should not only include satisfaction information but also the non-asset cost of providing a new device (pre-delivery cost, image maintenance, etc.) along with the change in security incidents for users. Adopting metrics that reflect the entire cost of providing and managing a device can highlight the success of your new deployment.

Management guides for Windows security services

Windows Defender Application Guard

[Configure the Group Policy settings for Microsoft Defender Application Guard \(Windows 10\) - Windows security | Microsoft Docs](#)

[Virtualization-based Security \(VBS\) | Microsoft Docs](#)

[Enable virtualization-based protection of code integrity - Windows security | Microsoft Docs](#)

Windows Defender Credential Guard

[Manage Windows Defender Credential Guard \(Windows 10\) - Microsoft 365 Security | Microsoft Docs](#)

Windows Defender System Guard


[System Guard Secure Launch and SMM protection \(Windows 10\) - Windows security | Microsoft Docs](#)

Diving deeper into the technologies

Virtualization Based Security

Virtualization-based security, or VBS, uses hardware virtualization features to create a secured environment, which can host several security features. Running these security applications inside VBS offers greatly increased protection from vulnerabilities in the operating system and helps prevent the use of malicious OS exploits that attempt to defeat protections.

Virtualization Based Security




VBS uses the Windows hypervisor to create this virtual secure mode, and to enforce restrictions which protect vital system and operating system resources, or to protect security assets such as authenticated user credentials. With the increased protections offered by VBS, even if malware gains access to the OS kernel, the possible exploits can be greatly limited and contained because the hypervisor can prevent the malware from executing code or accessing platform secrets.

Hypervisor Code Integrity

Hypervisor-Protected Code Integrity can use hardware technology and virtualization to isolate the Code Integrity (CI) decision-making function from the rest of the Windows operating system. When using virtualization-based security to isolate Code Integrity, the only way kernel memory can become executable is through a Code Integrity verification. As a result, a Windows secured-core PC can better defend against attacks even if the attack manages to gain access to the privilege level required to run kernel code.

Windows Hello

In Windows 11, Windows Hello for Business replaces passwords with strong two-factor authentication on PCs and mobile devices. This authentication consists of a different type of user credential that is tied to a device and uses a biometric or PIN.



Windows Hello addresses the following problems with passwords:

- Strong passwords can be difficult to remember, and users often reuse passwords on multiple sites.
- Server breaches can expose symmetric network credentials (passwords).
- Passwords are subject to [replay attacks](#).
- Users can inadvertently expose their passwords due to [phishing attacks](#).

You may wonder how a PIN can help protect a device better than a password. Passwords are shared secrets; they are entered on a device and transmitted over the network to the server. An intercepted account name and password can be used by anyone, anywhere. Because they are stored on the server, a server breach can reveal those stored credentials.

In Windows 11, Windows Hello replaces passwords. When the identity provider supports keys, the Windows Hello provisioning process creates a cryptographic key pair bound to the Trusted Platform Module (TPM), if a device has a TPM 2.0, or in software. Access to these keys and obtaining a signature to validate user possession of the private key is enabled only by the PIN or biometric gesture. The two-step verification that takes place during Windows Hello enrollment creates a trusted relationship between the identity provider and the user when the public portion of the public/private key pair is sent to an identity provider and associated with a user account. When a user enters the gesture on the device, the identity provider knows from the combination of Hello keys and gesture that this is a verified identity and provides an authentication token that allows Windows 11 to access resources and services.


Windows Defender Credential Guard

Introduced in Windows 11 Enterprise and Windows Server 2016, Windows Defender Credential Guard uses virtualization-based security to isolate secrets so that only privileged system software can access them.

By enabling Windows Defender Credential Guard, the following features and solutions are provided:

- Hardware security NTLM, Kerberos, and Credential Manager take advantage of platform security features, including Secure Boot and virtualization, to protect credentials.
- Virtualization-based security Windows NTLM and Kerberos derived credentials and other secrets run in a protected environment that is isolated from the running operating system.
- Better protection against advanced persistent threats:
 - NTLM, and Kerberos derived credentials are protected using virtualization-based security
 - Malware running in the operating system with administrative privileges cannot extract secrets that are protected by virtualization-based security.

Windows Defender Credential Guard



Secure Boot

When a PC starts, it first finds the operating system bootloader. PCs without Secure Boot simply run whatever bootloader is on the PC's hard drive. There is no way for the PC to tell whether it is a trusted operating system or a rootkit.

When a PC equipped with UEFI starts, the PC first verifies that the firmware is digitally signed, reducing the risk of firmware rootkits. If Secure Boot is enabled, the firmware examines the bootloader's digital

signature to verify that it has not been modified. If the bootloader is intact, the firmware starts the bootloader only if one of the following conditions is true:

- **The bootloader was signed using a trusted certificate.** In the case of PCs certified for Windows 11, the Microsoft® certificate is trusted.
- **The user has manually approved the bootloader's digital signature.** This allows the user to load non-Microsoft operating systems.

These requirements help protect you from rootkits while allowing you to run any operating system you want.

Early Launch Anti-Malware

Because Secure Boot has protected the bootloader and Trusted Boot has protected the Windows kernel, the next opportunity for malware to start is by infecting a non-Microsoft boot driver. Traditional anti-malware apps do not start until after the boot drivers have been loaded, giving a rootkit disguised as a driver the opportunity to work.

Early Launch Anti-Malware (ELAM) can load a Microsoft or non-Microsoft anti-malware driver before all non-Microsoft boot drivers and applications, thus continuing the chain of trust established by Secure Boot and Trusted Boot. Because the operating system has not started yet, and because Windows needs to boot as quickly as possible, ELAM has a simple task: examine every boot driver and determine whether it is on the list of trusted drivers. If it is not trusted, Windows will not load it.


Measured Boot

If a PC in your organization does become infected with a rootkit, you need to know about it. Enterprise anti-malware apps can report malware infections to the IT department, but that does not work with rootkits that hide their presence. In other words, you cannot trust the client to tell you whether it is healthy.

As a result, PCs infected with rootkits appear to be healthy, even with anti-malware running. Infected PCs continue to connect to the enterprise network, giving the rootkit access to vast amounts of confidential data and potentially allowing the rootkit to spread across the internal network.

Working with the TPM and non-Microsoft software, Measured Boot in Windows 11 allows a trusted server on the network to verify the integrity of the Windows startup process. Measured Boot uses the following process:

- The PC's UEFI firmware stores in the TPM a hash of the firmware, bootloader, boot drivers, and everything that will be loaded before the anti-malware app.
- At the end of the startup process, Windows starts the non-Microsoft remote attestation client. The trusted attestation server sends the client a unique key.
- The TPM uses the unique key to digitally sign the log recorded by the UEFI.
- The client sends the log to the server, possibly with other security information.



Depending on the implementation and configuration, the server can now determine whether the client is healthy and grant the client access to either a limited quarantine network or to the full network.

We explain the details of how this system works to deliver Device Health Attestation at the following link:

<https://docs.microsoft.com/en-us/windows-server/security/device-health-attestation>

Secure Boot, and Measured Boot create an architecture that is fundamentally resistant to bootkits and rootkits. In Windows 11, these features have the potential to eliminate kernel-level malware from your network. This is the most ground-breaking anti-malware solution that Windows has ever had and helps improve your trust in the integrity of the operating system.

Windows Defender System Guard

Windows Defender System Guard reorganizes the existing Windows 11 system integrity features under one roof and sets up the next set of investments in Windows security. It is designed to make these security promises:

- Protect and maintain the integrity of the system as it starts up
- Validate that system integrity has truly been maintained through local and remote attestation

Windows Defender System Guard Secure Launch, first introduced in Windows 10 version 1809, aims to alleviate limitations with UEFI secure boot by leveraging a technology known as the Dynamic Root of Trust for Measurement (DRTM). DRTM lets the system freely boot into untrusted code initially, but shortly after launches the system into a trusted state by taking control of all CPUs and forcing them down a well-known and measured code path. This has the benefit of allowing untrusted early UEFI code to boot the system, but then being able to securely transition into a trusted and measured state. Essentially, this means that the large amount of code in the UEFI firmware is isolated from sensitive system assets protected by the hypervisor and VBS.

After the system boots:

- WDSG signs and seals these measurements using the TPM.
- Management systems like Intune or Microsoft Endpoint Configuration Manager can acquire them for remote analysis.
- System Guard also enabled SMM protection since SMM code can potentially access hypervisor memory

If these measurements indicate that the device lacks integrity, the management system can take a series of actions, such as denying the device access to resources.

Protecting peripherals

In Windows 10 version 1803, Microsoft introduced a new feature called Kernel DMA Protection to protect PCs against drive-by Direct Memory Access (DMA).

Drive-by DMA attacks can lead to disclosure of sensitive information residing on a PC, or even injection of malware that allows attackers to bypass the lock screen or control PCs remotely. A simple example would be a PC owner leaves the PC for a quick coffee break, and within the break, and attacker steps in, plugs in a USB-like device and walks away with all the secrets on the machine, or injects a malware that allows them to have full control over the PC remotely.

Windows leverages the system Input/Output Memory Management Unit (IOMMU) to block external peripherals from starting and performing DMA unless the drivers for these peripherals support memory isolation (such as DMA-remapping). Peripherals with [DMA Remapping compatible drivers](#) will be automatically enumerated, started and allowed to perform DMA to their assigned memory regions.

By default, peripherals with DMA Remapping incompatible drivers will be blocked from starting and performing DMA until an authorized user signs into the system or unlocks the screen. IT administrators can modify the default behavior applied to devices with DMA Remapping incompatible drivers using the [DmaGuard MDM policies](#).

Differences between a Windows Secured-Core PC and Windows 11 PC

The two key differences between a Windows Secured-Core PC and a normal Windows 11 PC are:

- You can be sure that the hardware in the device can **support all the security features** of Windows described in this document, and
- That all the services are configured and **turned on by default**.

What this means for organizations is that they can send a new device out to a user and be confident that when the user turns the device on that it will be secure and immediately be able to adopt specific corporate security policies.

To summarize:

What makes a secured-core PC?

Feature	Windows 10 Default	Secured-core PC
Trusted Platform Module 2.0 (TPM)	✓	✓
Hardware protection for peripherals		✓
Hardware protection for applications		✓
Hardware protection for boot	✓	✓
Hardware protection for identity		✓
BitLocker encryption	✓	✓

Some real-world examples of attacks

A Windows Secured-Core PC would have defended users and organizations against a couple of recent attacks so it is worthwhile to look at how it would have helped.

RobbinHood Ransomware

[RobbinHood](#) ransomware is distributed as a packed executable that contains multiple binaries. One of these files is a Gigabyte driver (*GDRV.sys*), which has a [vulnerability](#) that could allow elevation of privilege, enabling an adversary to gain kernel privileges. In RobbinHood campaigns, adversaries use these kernel privileges to disable kernel-mode signing to facilitate the loading of an unsigned driver. The unsigned malicious driver is then used to disable security products from the kernel.

Two of the security promises of Secured-core PCs are directly applicable to preventing RobbinHood attacks:

- Defending against vulnerable and malicious drivers
- Defending against unverified code execution

Defending against vulnerable and malicious drivers

Secured-core PCs are the latest hardware to provide driver control out of the box, with baseline configuration already set. Driver control is provided by a combination of HVCI & Windows Defender Application Control ([WDAC](#)) technologies.

Every driver loaded into the kernel is verified by HVCI before it can run. HVCI runs in a hardware-protected execution environment isolated from the kernel space and cannot be tampered with by other code running in the kernel, including drivers.

Defending against unverified code execution and kernel data corruption attacks

There are several unverified code execution mitigations built-in to Windows. These are readily available on Secured-core PCs. The RobbinHood attack utilized the vulnerable *GDRV.sys* driver to change a crucial variable within the system memory. Although HVCI already protects against this attack, other areas of memory may still be susceptible, and we need broader defense against kernel data corruption attacks.


In addition to existing mitigations, Windows introduced a feature called [Kernel Data Protection](#) (KDP), which provides driver developers and software running in the Windows kernel (and the OS code itself) with the ability to mark some kernel memory containing sensitive information as read-only protected. The memory is protected through the second level address translation (SLAT) tables by the hypervisor, such that no software running in VTLO have access to the protected memory. KDP does not protect executable pages, as those are already protected with HVCI.

Secured-core PCs have KDP **enabled by default**.

Thunderspy DMA Peripheral Attack

Researchers at the Eindhoven University of Technology recently revealed information around "[Thunderspy](#)," an attack that relies on leveraging direct memory access (DMA) functionality to compromise devices. An attacker with physical access to a system can use Thunderspy to read and copy data even from systems that have encryption with password protection enabled.


[Secured-core PCs](#) provide customers with Windows 11 systems that come configured from OEMs with a set of hardware, firmware, and OS features enabled by default, mitigating Thunderspy and any similar attacks that rely on malicious DMA.



How Thunderspy Works

Like any other modern attack, Thunderspy relies on not one but multiple building blocks being chained together. Below is a summary of how Thunderspy can be used to access a system where the attacker does not have the password needed to sign in. A video from the Thunderspy research team showing the attack is available [here](#).

- **Step 1:** A serial peripheral interface (SPI) flash programmer called Bus Pirate is plugged into the SPI flash of the device being attacked. This gives access to the Thunderbolt controller firmware and allows an attacker to copy it over to the attacker's device
- **Step 2:** The Thunderbolt Controller Firmware Patcher (TCFP), which is developed as part of Thunderspy, is used to disable the security mode enforced in the Thunderbolt firmware copied with the Bus Pirate device in Step 1.
- **Step 3:** The modified insecure Thunderbolt firmware is written back to the SPI flash of the device being attacked
- **Step 4:** A Thunderbolt-based attack device is connected to the device being attacked, leveraging the PCILeech tool to load a kernel module that bypasses the Windows sign-in screen



The result is that an attacker can access a device without knowing the sign-in password for the device. This means that even if a device were powered off or locked by the user, someone that could get physical access to the device in the time it takes to run the Thunderspy process could sign in and exfiltrate data from the system or install malicious software.

Secured-core PC protections

In order to counteract these targeted, modern attacks, Secured-core PCs use a defense-in-depth strategy that leverage features like [Windows Defender System Guard](#) and virtualization-based security (VBS) to mitigate risk across multiple areas, delivering comprehensive protection against attacks like Thunderspy.

Secured-core PCs ship with hardware and firmware that support [Kernel DMA protection](#), which is enabled by default in the Windows OS. Kernel DMA protection relies on the Input/Output Memory Management Unit (IOMMU) to block external peripherals from starting and performing DMA unless an authorized user is signed in and the screen is unlocked. Watch this [video from the 2019 Microsoft Ignite](#) to see how Windows mitigates DMA attacks.

This means that even if an attacker were able to copy a malicious Thunderbolt firmware to a device, the Kernel DMA protection on a Secured-core PC would prevent any accesses over the Thunderbolt port unless the attacker gains the user's password in addition to being in physical possession of the device, significantly raising the degree of difficulty for the attacker.

Secured-core PCs ship with hypervisor protected code integrity (HVCI) enabled by default. HVCI utilizes the hypervisor to enable VBS and isolate the code integrity subsystem that verifies that all kernel code in Windows is signed from the normal kernel. In addition to isolating the checks, HVCI also helps ensure that kernel code cannot be both writable and executable, ensuring that unverified code does not execute.

HVCI helps to ensure that malicious kernel modules like the one used in Step 4 of the Thunderspy attack cannot execute easily as the kernel module would need to be validly signed, not revoked, and not rely on overwriting executable kernel code.

Resources (blog posts on malware protection by SCPC, Docs, etc.)

Protecting identity

- [Windows Hello](#)
- [Windows Hello for Business](#)
- [Windows Defender Credential Guard](#)

Securing the Operating System

- [Virtualization Based Security](#)
- [Hypervisor Code Integrity](#)
- [Windows Defender Application Guard](#)

Defending the Hardware and Firmware

- [Trusted Platform Module 2.0](#)
- [Secure Boot](#)
- [Roots of trust](#)
- [Measured Boot](#)
- [Windows Defender System Guard](#)

Blog posts

- [Secured-core PCs: A brief showcase of chip-to-cloud security against kernel attacks - Microsoft Security](#)
- [Security blog series - Microsoft Security](#)