

Windows Server Security Guide

August 2017

Contents

| | |
|---|----------|
| Windows Server 2016 Security Guide | 3 |
| Why is Windows Server 2016 security important?..... | 3 |
| How does Windows Server 2016 help prevent and detect compromise? | 4 |
| Additional resources | 5 |
| Build a secure foundation | 5 |
| Stay current on Windows Server security updates..... | 5 |
| Configure Windows Server security settings | 6 |
| The high-level process for obtaining and deploying the security baselines can be found in the Microsoft Security Compliance Toolkit 1.0. You can find out more about current Microsoft security guidance at Microsoft Security Guidance blog..... | 7 |
| Back up your information and systems | 7 |
| Management and monitoring using Operations Management Suite..... | 7 |
| Protect privileged identities..... | 8 |
| How do privileged identities get compromised?..... | 8 |
| How to prevent attackers from gaining access to privileged identities | 9 |
| Harden Windows Server | 12 |
| Improve threat detection..... | 15 |
| Harden Hyper-V environments | 15 |
| Why harden a virtualization environment?..... | 15 |
| How to harden Hyper-V environments..... | 15 |
| Appendix | 18 |

Windows Server 2016 Security Guide

Windows Server® 2016 is the most secure version of Windows Server developed to date. However, just as with every previous version of Windows Server, Windows Server 2016 needs to be secured and hardened to your specific apps and environment.

This guide will help you secure Windows Server 2016 and previous versions of Windows Server for your environment. It provides additional resources that contain step-by-step instructions you can use to implement the guide's security recommendations.

Why is Windows Server 2016 security important?

Security affects everyone in your organization from upper-level management (such as CEO-level) to the information worker. A lack of security is a real risk for organizations; a security breach can potentially disrupt all normal business and bring your organization to a halt. Recent studies from McKinsey, the Ponemon Institute, and Verizon show that cyber security has a \$3 trillion impact each year in terms of lost productivity and growth, with the average security breach costing \$3.5 million. It is imperative for organizations to detect and prevent security breaches.

Note Although this guide focuses on Windows Server, you need to have a comprehensive security plan that encompasses your clients and network infrastructure, which is beyond the scope of this guide. For additional Microsoft® security resources, see <http://www.microsoft.com/en-us/security/default.aspx>.

Much like any other crime, the sooner that you can detect a potential attack, the more that you can mitigate any compromise in security. Typically, an attacker starts by researching an environment's weak points and then proceeds to performing the attack. After an attacker breaches an environment (through phishing or vulnerable entry points), they proceed to escalate their privileges through lateral movement within the environment until they take control over the organization within a short period, typically 24 to 48 hours from the first compromise (as shown in the following figure). Your goal is to detect and respond to such attacks as fast as possible.

To do that, you need to extend the time it takes an attacker to take control to weeks or even months by blocking their lateral movements and hardening your systems. Then you can detect the attack by improving the various warning signals and respond by removing compromised identities and systems.

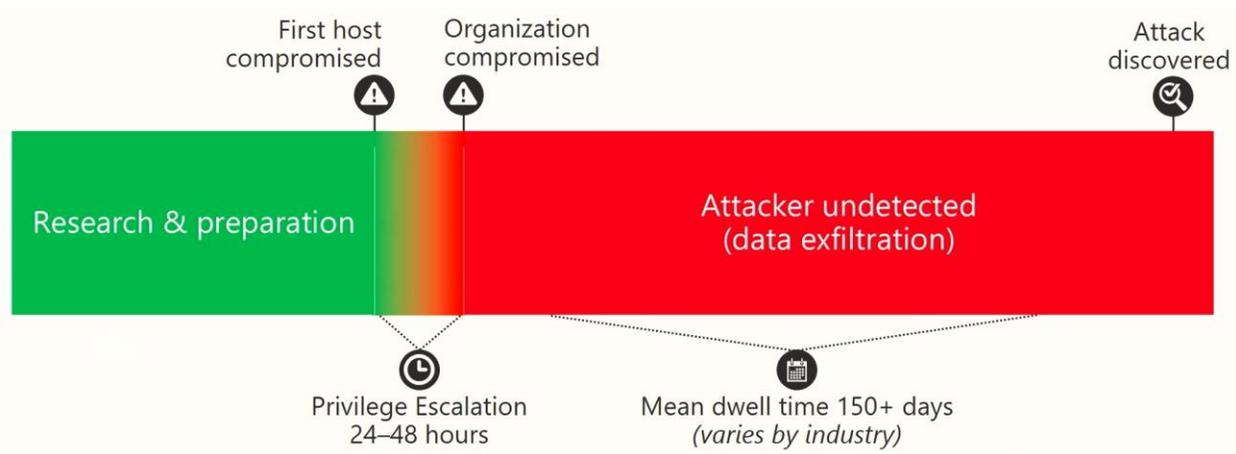


Figure 1. Timeline for typical attack scenario

The following is a typical attack scenario:

1. The attacker does some research and preparation about an organization (such as by using Facebook, Linked In, search engines, or other social networking services).
2. The attacker determines the best method for initiating an attack (such as a phishing email or probing edge-of-network services).
3. The attacker initiates an attack to gain a foothold into the organization's network and services.
4. The attacker gains access and then, using one or more compromised identities, attempts to escalate their privileges.
5. The attacker gains escalated privileges and continues to compromise services and servers within the organization, compromising data and/or causing denial of service.

It is important to note that the longer the attacker goes undetected, the more damage they can do and the harder it will be to expunge the attacker from the network. Again, your goal is to extend the time it takes to escalate privilege to weeks and months so that you can detect an attack and respond to it before the attacker can gain full control. The remainder of this guide focuses on how you can make it harder for an attacker to escalate privilege and move freely in your network, and how to detect attacks sooner.

How does Windows Server 2016 help prevent and detect compromise?

As the latest version of Windows Server, Windows Server 2016 has built-in security features to help better harden the operating system and detect malicious activity. The following bullet points identify the security features available in Windows Server, and they are discussed in more detail in the corresponding sections later in this guide:

- **Build a secure foundation.** This section discusses how to help ensure Windows Server is a secure foundation for running your apps and services by using Windows Server security updates, Group Policy settings, Local Script tools, and Microsoft Operations Management Suite (OMS).

- **Protect privileged identities.** This section discusses how to help protect your privileged identities (accounts with elevated privileges, such as members of Domain Admins) from theft by using Just Enough Administration (JEA), Just in Time Administration (JIT), Credential Guard, Remote Credential Guard, and Advanced Threat Analytics. Additional protections include the usage of Privileged Access Workstations, which is not covered in this document.
- **Harden Windows Server.** This section describes how to help protect the apps and services running on Windows Server by using Control Flow Guard (similar to /GS, DEP, and ASLR), Windows Defender, Device Guard, AppLocker®, and Microsoft OMS.
- **Improve threat detection.** This section describes how to help detect security threats faster by using improvements in Windows event log entries, Windows Server auditing, and Microsoft OMS.
- **Harden Hyper-V® environments.** This section describes how to help protect sensitive workloads running in Hyper-V environments by using Guarded fabric, TPM in Hyper-V, and the Datacenter Firewall in Software Defined Networking (SDN).

Additional resources

In addition to the resources listed in this guide, you can use the following resources to help you secure Windows Server 2016 in your environment:

- Security and assurance documentation (<https://technet.microsoft.com/en-us/library/mt130644.aspx>)
- Securing privileged access guidance (<http://aka.ms/privsec>)
- Privileged Access Workstation (<https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/privileged-access-workstations>)
- Microsoft Virtual Academy online courses (<https://mva.microsoft.com>)

Build a secure foundation

Windows Server is deployed in a secure configuration. To keep it secure, you need to ensure that Windows Server is current on security updates, make sure your data is backed up, and configure the Windows Server security settings based on Microsoft security recommendations and your organization's security standards.

Stay current on Windows Server security updates

Microsoft regularly releases updates for Windows operating systems, including Windows client and Windows Server. These updates include security updates to keep Windows Server secure as new threats and vulnerabilities are discovered as well as antimalware and antispymware definition updates for Windows Defender.

You can deploy these updates to the servers in your organization by using one of the methods listed in the following table.

Table 1. Methods for Deploying Windows Updates to Servers

| Method | When to select this method |
|---|--|
| Windows Update only | Use this method when you have a small number of servers that have direct access to the Internet and can download updates directly from Windows Update. A potential drawback to this method is that you cannot easily manage which updates are deployed and when they are deployed. |
| Windows Server Update Services (WSUS) | Use this method when you do not have System Center Configuration Manager, but desire a centralized method of downloading and managing updates. Windows Server Update Services downloads the desired updates locally and then distributes the updates to the servers on your network. You can select the updates to be deployed and control which groups of servers receive the updates. WSUS is a built-in Windows Server role. For more information, see Manage Windows updates using Windows Server Update Services (WSUS) . |
| System Center Configuration Manager | Use this method when you want to have even more precise control of the updates to be deployed and which servers will receive the updates. This method leverages Windows Server Update Services to download the updates, but then uses the deployment flexibility of the Software Update feature in System Center Configuration Manager to deploy the updates to servers on your network. For more information, see Introduction to software updates in System Center Configuration Manager . |
| Operations Management Suite Update Management service | Use this method for full scanning, monitoring and update orchestration capabilities. An Azure based orchestrated Update Management across any OS (Windows/Linux) and any cloud. Update Management solution in OMS |

Note All the methods for deploying and monitoring updates in this section are applicable to Windows Server 2008 R2 and later versions of Windows Server.

Configure Windows Server security settings

All Windows operating systems include security settings that you can use to help harden computer security profiles. Microsoft publishes security baselines that are based on Microsoft security recommendations, which are established from real-world security experience obtained through partnership with commercial organizations and the US government (such as the Department of Defense [DoD]).

These security baselines include recommended settings for Windows Firewall, Windows Defender, and other security settings. These security baselines are provided as Group Policy object (GPO) backups that you can import into Active Directory® Domain Services (AD DS) and then deploy to domain-joined servers. You can also use the Local Script tools to configure standalone (non domain-joined) servers.

The high-level process for obtaining and deploying the security baselines can be found in the Microsoft Security Compliance Toolkit 1.0. You can find out more about current Microsoft security guidance at [Microsoft Security Guidance blog](#).

Back up your information and systems

You should perform scheduled backups of the Windows Server operating system, including the applications and data stored on Windows Server. Doing so will help protect against ransomware attacks on Windows Server. You should perform backups frequently so that you can easily restore to a point-in-time prior to a ransomware attack.

You can perform backups on-premises by using solutions such as [System Center Data Protection Manager](#) or cloud-based backups by using [Microsoft Azure Backup Server](#). There are also a number of backup solutions available from Microsoft partners.

Management and monitoring using Operations Management Suite

[Microsoft Operations Management Suite](#) (OMS) is a cloud-based IT management solution that helps you manage and protect your on-premises and cloud infrastructure. OMS is implemented as a cloud-based service, and you can start managing your apps, services, and infrastructure with minimal extra investment. OMS is also updated periodically with new features, and can help dramatically reduce your ongoing maintenance and upgrade costs.

In addition, OMS integrates with on-premises System Center components such as System Center Operations Manager to extend your existing management investments into the cloud. System Center and OMS work together to provide a full hybrid management experience.

OMS offers the following key capabilities:

- **Insight and analytics.** This feature can collect, correlate, search, and act on logs and performance data generated by Windows operating systems and apps. It provides real-time operational insights for all your workloads and servers, on-premises and in Azure®.
- **Security and compliance.** This feature identifies, assesses, and mitigates security risks. It uses the Security and Audit solution (which collects and analyzes security events), the Antimalware solution (which provides current malware protection status), and the System Updates solution (which provides current software update status) to ensure the ongoing security of your on-premises and cloud workloads and servers.
- **Automation and control.** This feature automates administrative processes with runbooks (similar to runbooks in System Center) using Windows PowerShell®. Runbooks can access any apps, operating systems, or services that can be managed by Windows PowerShell. It also provides configuration management with Windows PowerShell Desired State Configuration(DSC), which can automatically enforce your configuration settings on-premises and in Azure.
- **Protection and recovery.** This feature can back up recovery workloads and servers. Azure Backup protects app data for on-premises and cloud-based servers. Azure Site Recovery helps provide disaster recovery by orchestrating replication, failover, and recovery of on-premises Hyper-V virtual machines.

Protect privileged identities

Privileged identities are any accounts that have elevated privileges, such as user accounts that are members of the Domain Admins, Enterprise Admins, local Administrators, or even Power Users groups. Such identities can also include accounts that have been granted privileges directly, such as performing backups, shutting down the system, or other rights listed in the User Rights Assignment node in the Local Security Policy console.

You need to protect these privileged identities from compromise by potential attackers. First, it's important to understand how identities are compromised; then you can plan to prevent attackers from gaining access to these privileged identities.

How do privileged identities get compromised?

Privileged identities often get compromised when organizations don't have guidelines to protect them. The following are examples:

- **More privileges than are necessary.** One of the most common issues is that users have more privileges than are necessary to perform their job function. For example, a user who manages DNS might be an AD administrator. Most often, this is done to avoid the need to configure different administration levels. However, if such an account is compromised, the attacker automatically has elevated privileges.
- **Signed in with elevated privileges all the time.** Another common issue is that users with elevated privileges can use it for an unlimited time. This is very common with IT pros who sign in to a desktop computer using a privileged account, stay signed in, and use the privileged account to browse the web and use email (typical IT work job functions). Unlimited duration of privileged accounts makes the account more susceptible to attack and increases the odds that the account will be compromised.
- **Social engineering research.** Most credential threats start out by researching the organization and then conducted through social engineering. For example, an attacker may perform an email phishing attack to compromise legitimate accounts (but not necessarily elevated accounts) that have access to an organization's network. The attacker then uses these valid accounts to perform additional research on your network and to identify privileged accounts that can perform administrative tasks.
- **Leverage accounts with elevated privileges.** Even with a normal, non-elevated user account in the network, attackers can gain access to accounts with elevated permissions. One of the more common methods of doing so is by using the Pass-the-Hash or Pass-the-Token attacks. For more information on the Pass-the-Hash and other credential theft techniques, see the resources on the [Pass-the-Hash \(PtH\)](#) page.

There are of course other methods that attackers can use to identify and compromise privileged identities (with new methods being created every day). It is therefore important that you establish practices for users to log on with least-privileged accounts to reduce the ability of attackers to gain access to privileged identities.

How to prevent attackers from gaining access to privileged identities

You can reduce the attack surface for privileged identities (discussed in the previous section) with each of the mitigations described in the following table.

Table 2. Methods for Preventing Attackers from Gaining Access to Privileged Identities

| Attack vectors | How to mitigate |
|---|--|
| More privileges than are necessary | Implement Just Enough Administration (JEA) for all IT pros who administer Windows Server and the apps and services (such as Exchange Server or Exchange Online) running on Windows Server by using Windows PowerShell. For more information on JEA, see Just Enough Administration later in this guide. |
| Signed in with elevated privileges all the time | Implement Just in Time Administration (JIT) for all users who require elevated privileges so that the elevated privileges can only be used for a limited amount of time. Many organizations use the Local Administrator Password Solution (LAPS) as a simple yet powerful JIT administration mechanism for their server and client systems. |
| Compromised identity and Pass-The-Hash attacks | Implement Microsoft Advanced Threat Analytics (ATA) to help detect compromised identities in on-premises workloads and servers. ATA is an on-premises solution that you can use to manage physical and virtualized workloads. For more information, see Advanced Threat Analytics later in this guide. |
| Pass-The-Hash attacks | <ul style="list-style-type: none">• Implement Credential Guard to help protect credentials and credential derivatives from attacks such as Pass-the-Hash or Pass-the-Token. Credential Guard is a new feature in Windows Server 2016. For more information, see Credential Guard later in this guide.• Implement Remote Credential Guard to help protect credentials and credential derivatives from attacks such as Pass-the-Hash or Pass-the-Token that can be performed on servers that host Remote Desktop connections. Remote Credential Guard is a new feature in Windows Server 2016. For more information, see Remote Credential Guard later in this guide. |

Just Enough Administration

Just Enough Administration (JEA) is a security technology that helps restrict IT administrative rights using Windows PowerShell remoting. JEA is included in Windows Management Framework 5.0 and later versions. JEA uses the built-in capabilities of the Windows PowerShell scripting environment and implements role-based access control (RBAC). Anything that you can manage with Windows PowerShell, you can manage more securely with JEA. JEA provides a framework of reducing administrative access with more specificity than traditional access control models.

You can configure JEA as a Windows PowerShell session endpoint on any computer to manage that computer or remote computers. You can configure a PowerShell Session Configuration file (which

specifies who can connect to the JEA endpoint) and one or more Role Capability files (which specifies which actions a specific user can perform).

With JEA, you connect by using your regular, non-elevated user credentials. After JEA authorizes you, JEA runs the Windows PowerShell commands you specify by using an elevated virtual account on the targeted computer. This approach means that you never actually sign in by using elevated credentials.

JEA works using Windows PowerShell and is available on any Windows operating system that supports Windows Management Framework 5.0 and later versions (such as Windows Server 2008 R2 and later versions of Windows Server).

For more information on JEA, see the following resources:

- [Just Enough Administration](#)
- [Just Enough Administration: Windows PowerShell security controls help protect enterprise data](#)
- [Just Enough and Just in Time Administration in Windows Server 2016](#)
- [JEA GitHub repository](#)

Just in Time Administration

Just in Time (JIT) Administration is another security best practice that allows you to only use elevated identities when you are performing an IT administration task. For example, when you need to manage a DNS server you request the elevated privilege; if approved, the approval will be for a specific duration (for example, one hour) so that you can perform the IT administration task.

An example of a JIT workflow is as follows:

1. An IT administrator uses the JIT system to submit a request to get administrative access.
2. The request goes through a workflow such as two-factor authentication or manager approval and audit logs are created.
3. Upon approval, the IT administrator will either get temporary local administrative credentials to the system they need to manage or their administrative account will be put in the right group that is allowed access.
4. The IT administrator uses his administrative account (or temporary credentials) to log on to the remote system (such as Remote Desktop or Windows PowerShell) or uses a remote management tool with Run As credentials.
5. After the designated time, the privilege gets revoked.

Microsoft provides JIT administration in the following:

- [Local Administrator Password Solution \(LAPS\)](#). LAPS is available as a free download to help manage local administrator password on Windows operating systems in your organization.
- [Microsoft Identity Manager 2016](#). Microsoft Identity Manager is an on-premises identity and access management system that provides JIT administration.

Advanced Threat Analytics

[Microsoft Advanced Threat Analytics](#) (ATA) is an on-premises product that helps detect identity compromise in an organization. ATA has the ability to capture and parse network traffic for

authentication, authorization, and information gathering protocols (such as Kerberos, DNS, RPC, NTLM, and other protocols). ATA uses this data to build a behavioral profile about users and other entities on a network so that it can detect anomalies and known attack patterns. The following table lists the attack types detected by ATA.

Table 3. Attack Types Detected by ATA

| Attack type | Description |
|---------------------------|---|
| Malicious attacks | <p>These attacks are detected by looking for attacks from a known list of attack types, including:</p> <ul style="list-style-type: none"> • Pass-the-Ticket (PtT) • Pass-the-Hash (PtH) • Overpass-the-Hash • Forged PAC (MS14-068) • Golden Ticket • Malicious replications • Reconnaissance • Brute force • Remote execution <p>For a complete list of malicious attacks that can be detected and their description, see What Suspicious Activities Can ATA detect?.</p> |
| Abnormal behavior | <p>These attacks are detected by using behavioral analysis and use machine learning to identify questionable activities, including:</p> <ul style="list-style-type: none"> • Anomalous logins • Unknown threats • Password sharing • Lateral movement |
| Security issues and risks | <p>These attacks are detected by looking at current network and system configuration, including:</p> <ul style="list-style-type: none"> • Broken trust • Weak protocols • Known protocol vulnerabilities |

You can use ATA to help detect attackers attempting to compromise privileged identities.

For more information on deploying ATA, see the Plan and Design and Deploy topics in [Advanced Threat Analytics documentation](#).

Credential Guard

[Credential Guard](#) uses virtualization-based security to encrypt secrets (such as NTLM password hashes, Kerberos Ticket Granting Tickets, and credentials stored by apps) so that only privileged system processes can access them.

Credential Guard uses:

- Virtualization-based security (required)
- Secure boot (required)
- TPM 2.0 either discrete or firmware (preferred - provides binding to hardware)

The virtualization-based security requires:

- 64-bit CPU
- CPU virtualization extensions plus extended page tables
- Windows hypervisor

You can use Credential Guard to help protect privileged identities by protecting the credentials and credential derivatives on Windows Server 2016. For more information on Credential Guard requirements, see [Protect derived domain credentials with Credential Guard](#).

Remote Credential Guard

Remote Credential Guard is applicable to Remote Desktop connections. Remote Credential Guard helps protect credentials over a Remote Desktop connection by keeping the credentials on the device hosting the RDP connection and redirecting Kerberos requests back to the device that establishes the connection. If the server (or client) hosting the Remote Desktop connection is compromised, the credentials are not exposed because the credentials and credential derivatives are never sent to the device hosting the Remote Desktop connection.

To use Remote Credential Guard, the Remote Desktop client and server must meet the following requirements:

- Must be joined to an Active Directory domain and be in the same domain or a domain with a trust relationship.
- Must use Kerberos authentication.
- Must be running at least Windows 10 version 1607 or Windows Server 2016.
- The Remote Desktop classic Windows app is required. The Remote Desktop Universal Windows Platform app doesn't support Remote Credential Guard.

You can enable Remote Credential Guard by using a registry setting on the Remote Desktop server and Group Policy or a Remote Desktop Connection parameter on the Remote Desktop client. For more information on enabling Remote Credential Guard, see [Protect Remote Desktop credentials with Remote Credential Guard](#).

As with Credential Guard, you can use Remote Credential Guard to help protect privileged identities on Windows Server 2016.

Harden Windows Server

Windows Server 2016 includes built-in security mechanisms and powerful new security tools that can be configured to further lock down the server.

Control Flow Guard

New to Windows Server 2016, Control Flow Guard is built into Windows to help protect the operating system and applications from a class of memory corruption-based attacks. Visual Studio® 2017 supports Control Flow Guard as a compiler option. If you have line-of-business applications that run on Windows

Server, you can compile them with Control Flow Guard enabled to help make your applications less susceptible to memory corruption attacks.

For more information about Control Flow Guard, see [Control Flow Guard](#).

Windows Defender

Windows Defender has been included in Windows operating systems since Windows 8. It helps protect Windows devices against viruses, malware, spyware, and other threats. Windows Defender has been optimized for running on Windows Server and is enabled by default in Windows Server 2016. The advanced security features it provides include:

- Virus protection and removal
- Malware protection and removal
- Spyware protection and removal
- Boot-time protection
- Real-time protection
- Cloud-based protection
- Network inspection and protection
- Free automatic updates to antimalware definitions and Windows Defender itself

You can configure Windows Defender by using Group Policy, Windows PowerShell, Windows Management Instrumentation (WMI), or interactively through the Windows Defender user interface. The Windows Server security baselines also include Microsoft recommended settings for Windows Defender.

For more information about Windows Defender, see [Windows Defender Overview for Windows Server](#).

Device Guard

Device Guard provides the ability to specify which binaries are authorized to run on your server, including user mode and kernel mode binaries (enhancing the currently available AppLocker functionality).

You can create a code integrity policy to define the set of specific kernel mode and user mode binaries that can run on your system so that malicious code is blocked from running.

Device Guard helps protect against the following threats:

- Exposure to new malware for which no malware signature is yet known
- Exposure to unsigned code as most malware is unsigned
- Malware that gains access to the kernel and then captures sensitive information or damages the system

Device Guard code integrity policies can be run in *audit mode* or *enforcement mode*. In audit mode, Device Guard will trigger an audit log event whenever a non-authorized binary is running but will not block the binary from running. The Device Guard logs are available in the following event log:

Logs\Microsoft\Windows\CodeIntegrity\Operational

Audit mode allows you to identify apps that you do want to allow in your organization. You can create a code integrity policy file based on the captured audit information in the event log.

In enforcement mode, Device Guard will actually block any binary that should be denied. You would configure Device Guard for enforcement mode after you have run a selected group of devices in audit mode to identify the apps that you want to allow.

The high-level process for deploying Device Guard is as follows:

1. Create an initial version of the Device Guard code integrity policies on a clean server.
2. Deploy Device Guard on a selected group of devices (that provide a good cross section of your devices) in audit mode.
3. Collect audit mode events for a period of time to get an accurate representation of the apps in used within your organization.
4. Create an updated version of the Device Guard code integrity policies based on the event log entries collected from the select group of devices.
5. Configure Device Guard to run in enforcement mode on a select group of devices.
6. After a period of time, broaden Device Guard deployment to larger and large groups of devices (updating the code integrity policies as necessary)

For more information on Deploying Device Guard, see [Device Guard deployment guide](#).

You can also use AppLocker with or without Device Guard. AppLocker complements Device Guard, as Device Guard does not lock down Windows Store applications. You can use AppLocker for locking down which Windows Store Apps you would like to authorize to run on your server..

For Windows Server operating systems prior to Windows Server 2016, you can use AppLocker to provide similar functionality to Device Guard. On these versions of Windows Server, you can deploy AppLocker to harden Windows Server. For more information on AppLocker, see [AppLocker](#).

Secure Boot

Secure Boot is a PC industry standard that helps ensure your device boots only software that is trusted by the device manufacturer. Secure Boot helps protect devices from rootkits and other low-level malware attacks by blocking unauthorized (non-signed) software.

When a device starts, the device firmware checks the signature for each piece of boot software to ensure they are trusted. If all boot software signatures can be confirmed, the firmware starts the operating system.

You need to ensure that Secure Boot is enabled in the device's firmware.

For more information on Secure Boot, see [Secure Boot](#).

Operations Management Suite

You can use OMS to help detect threats as well as identify devices that are not current on software updates and antimalware definitions.

Improve threat detection

Threat detection is an essential part of Windows Server security. The faster you can detect threats, the easier it is to respond before an attacker reaches full control. Microsoft provides threat detection for servers using Windows Defender Advanced Threat Protection ([ATP](#)). In addition, Windows Server includes security audit events that can be consumed by Security Information Event Management systems (SIEM) for threat detection. OMS security is one such system that provides threat intelligence and analysis.

Windows Server has always provided essential security information by recording this information in the Windows event logs. Windows Server 2016 represents the latest in this effort by providing improved auditing information that helps focus on suspicious activities that should not occur on servers, such as the following:

- Plugging a USB device into a server
- Resetting a password that is not expected
- Locking out a user account
- Remotely accessing the Security Account Manager (SAM) database

Windows Server provides these and other events that help you identify suspicious activity. For a complete list of the new events, see the **Security auditing** section in [What's new in Windows 10, versions 1507 and 1511](#).

You can use OMS to collect these new events and help improve insights into threats that may be occurring within your organization. For more information on OMS features, see [Operations Management Suite](#) earlier in this guide.

Harden Hyper-V environments

Most organizations rely heavily on virtualized datacenters. You need to ensure that these virtualized datacenters are as well-protected as your physical infrastructure.

Why harden a virtualization environment?

You need to harden Windows Server running in a virtual machine just as you would as if it were running in a physical server. In fact, because virtual environments have multiple virtual machines sharing the same physical host, it is imperative that both the physical host and the virtual machines running on the host are protected. If an attacker could compromise a host, the attacker could affect multiple virtual machines and have a greater impact on the workloads and services.

How to harden Hyper-V environments

You can use many of the same methods that you use to harden Windows Server in a physical environment in a virtual environment. However, traditionally, virtual machines are limited in their support of the necessary virtualized security hardware, such as Trusted Platform Module (TPM).

The following are Windows Server 2016 features that you can use to help harden Hyper-V environments:

- Guarded fabric and shielded virtual machines
- Virtual machine Trusted Platform Module (TPM)

- Software Defined Networking – Micro-segmentation firewall

Guarded fabric and shielded virtual machines

Guarded fabric and shielded virtual machines in Hyper-V help you strengthen the security for virtual machines running in Hyper-V environments. Because virtual machine state and memory is stored in a file, it is susceptible to attacks via the storage system, the network, or while it is backed up. This vulnerability exists for every virtualization platform today, including Hyper-V and other virtualization platforms. An attacker can directly modify the virtual machine file or copy the virtual machine files to another environment where they can inspect the virtual machine files or run that virtual machine in their environment. Shielded virtual machines can prevent such attacks.

A guarded fabric consists of the following:

- The Host Guardian Service (HGS) (typically running on a three-node cluster)
- One or more guarded hosts
- Shielded virtual machines (note that a guarded fabric can also run non-shielded virtual machines)

The following figure illustrates how the Host Guardian Service uses attestation to ensure that only known, valid hosts can start the shielded virtual machines by using the Attestation Service. The Key Protection Service determines whether to release a key that the guarded hosts need to start the shielded virtual machine.

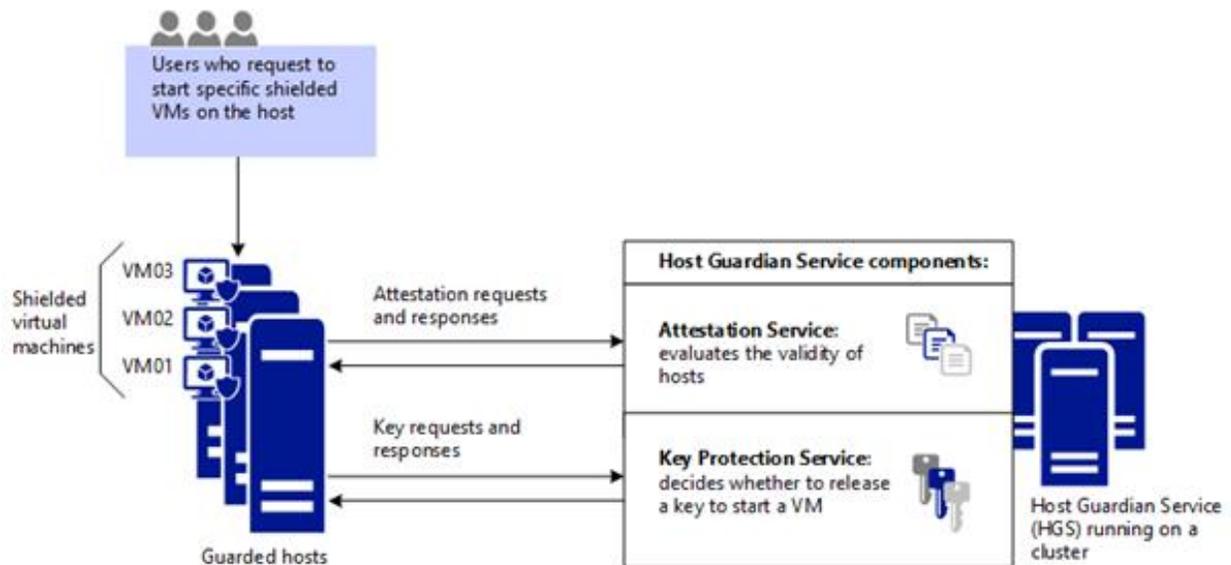


Figure 2. Host Guardian Service attestation for shielded virtual machines

The high-level steps for deploying guarded fabric and shielded virtual machines are as follows:

1. Deploy new Windows Server 2016 Hyper-V hosts, or upgrade any existing Hyper-V hosts to Windows Server 2016.
2. Deploy the Host Guardian Service on a three-node cluster.
3. Configure the attestation mode for the Host Guardian Service.
4. Configure Hyper-V hosts to attest against the Host Guardian Service.
5. Create a shielded virtual machine template.

6. Create a shielding data file, which contains sensitive information about the shielded virtual machine (such as Administrator password or RDP certificates).
7. Create your shielded virtual machines.

For more information on guarded fabric and shielded virtual machines, see [Guarded fabric and shielded VMs](#).

Virtual machine Trusted Platform Module

Windows Server 2016 supports TPM for virtual machines, which allows you to support advanced security technologies such as BitLocker® Drive Encryption in virtual machines. You can enable TPM support on any Generation 2 Hyper-V virtual machine by using Hyper-V Manager or the Enable-VMTPM Windows PowerShell cmdlet.

You can protect virtual TPM (vTPM) by using the local guardian or the Host Guarding Service. If you use the local guardian, the keys are stored locally. So, if the local Hyper-V host is compromised, the virtual machines running on that host can be compromised.

If you use the Host Guardian Service, the keys are stored in the Host Guardian Service. So while the Host Guardian Service requires more infrastructure, it also provides more protection.

Software Defined Networking Micro-segmentation

Windows Server 2016 includes [Software Defined Networking \(SDN\)](#), which provides a method to centrally configure and manage virtual network devices (such as Hyper-V Virtual Switch, Hyper-V Network Virtualization, and RAS Gateway). SDN allows you to dynamically manage your datacenter virtual network and how it interacts with your physical network.

[Datacenter Firewall](#) is one of the technologies in SDN. Datacenter Firewall provides the ability to centrally manage firewall policies that help protect virtual machines from undesired traffic originating from the Internet and your intranet networks. Datacenter Firewall is a network layer, 5-tuple (protocol, source and destination port numbers, source and destination IP addresses), stateful, multitenant firewall.

You configure Datacenter Firewall by creating [Datacenter Firewall Access Control Lists \(ACLs\)](#). Each ACL is very similar to the Windows Firewall inbound or outbound rules. You can create as many ACLs as are needed for each virtual machine. Each virtual machine may have its own unique set of ACLs.

As shown in the following figure, the Distributed Firewall Manager manages the firewall policies for each virtual machine. If you move a virtual machine from one Hyper-V host to another Hyper-V host, the Distributed Firewall Manager ensures the firewall policies are still applied to the virtual machine.

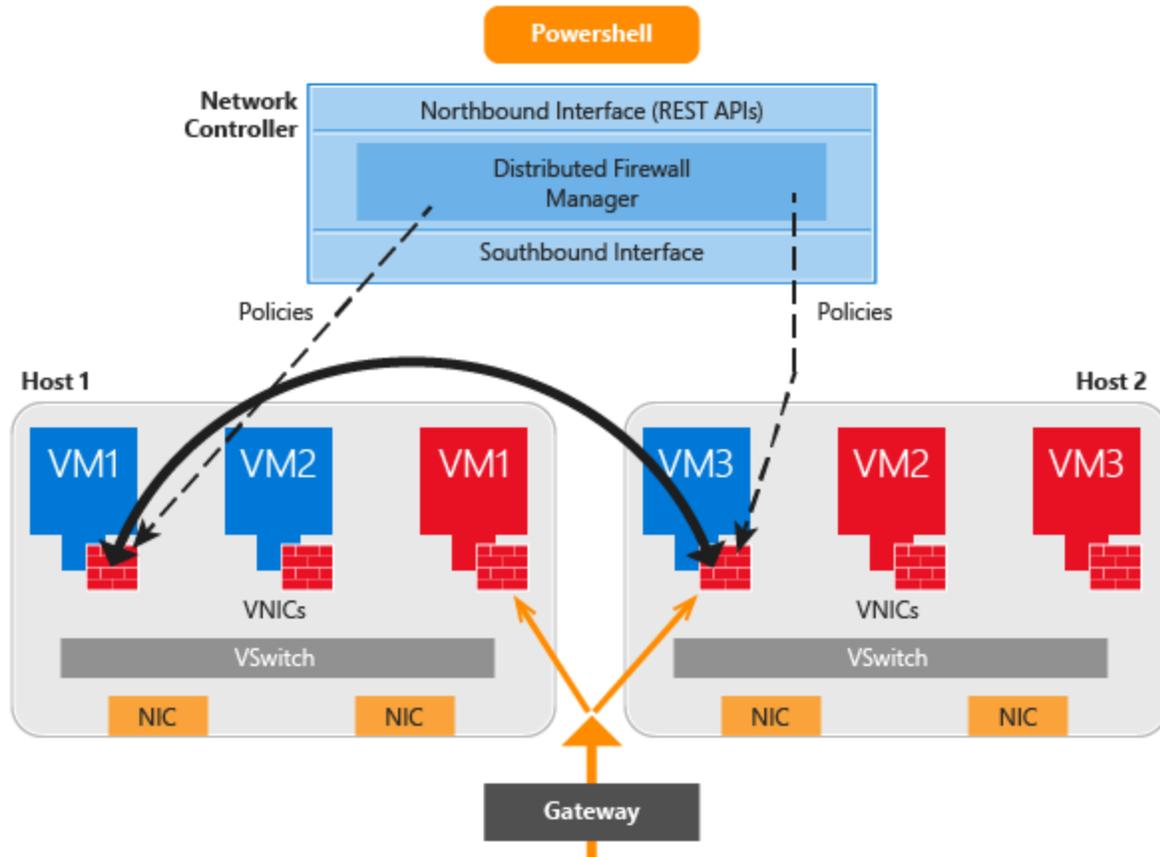


Figure 3. Datacenter Firewall

For more information about SDN and Datacenter Firewall, see:

- [Software Defined Networking](#)
- [Datacenter Firewall Overview](#)
- [Access Control Lists \(ACLs\) to Manage Datacenter Network Traffic Flow](#)
- [Configure Datacenter Firewall Access Control Lists](#)

Appendix

Windows Server 2016 Security related videos:

- [Windows Server Security overview video](#)
- [Shielded VMs video](#)
- [Just Enough Administration and Just in Time Administration video](#)
- [Windows Server Credential Guard and Remote Credential Guard video](#)
- [Windows Server Device Guard video](#)
- [Advanced Threat Analytics video](#)
- [Operations Management Suite security video](#)