Deep Security as a Service

Best Practice Guide

About This Guide

Deep Security provides a single platform for server security to protect physical, virtual, and cloud servers as well as hypervisors and virtual desktops. Tightly integrated modules easily expand to offer in-depth defenses, including anti-malware, web reputation, intrusion prevention, firewall, integrity monitoring, and log inspection. It is an agent-based options that can all be managed through a single console across physical, virtual, and cloud server deployments.

This guide is intended to help users get the best productivity out of the product. It contains a collection of best practices that are based on knowledge gathered from previous enterprise deployments, lab validations, and lessons learned in the field.

Examples and considerations in this document serve only as a guide and not a representation of strict design requirements. These guidelines do not apply in every environment but can help guide you through configuring Deep Security for optimum performance.

Trend Micro Incorporated reserves the right to change this document and products without notice. Before installing and using the software, please review the Readme file and the latest version of the applicable user documentation.

This Best Practice Guide contains:

- ➔ Deployment considerations and recommendations.
- ➔ Guidance in sizing server and storage resources for Deep Security implementation.
- ➔ Recommended configuration to maximize system performance and reduce administrative overhead.
  Best practice tips for VDI, private and public cloud environments.

Acknowledgments

This guide was made by the following individuals who volunteered their time and expertise to this project:

Aldrin Ceriola, Jason Dablow, Erwin Dusojan, Mohamed Inshaff, Jill Maceda, Marion Mora, Winfred Lin, Reuel Morales, Raphael Bottino, Ebenizer Padu, Igor Valoto, Kyle Klassen, Fernando Cardoso, Ryoma Kobayashi, and Mary Kay Stone.

We would also like to thank the following people for their significant support and contribution during development and review:

Shiela Aballa, Rodel Villarez, Ziv Huang, Marty Tsai, Cellina Lin, Chris Lai, Paul Liang, Zion Li

Document version: 3.0

Last updated: September 24, 2019

Table of Contents

1 Environment

Deep Security consists of several components working together to provide protection. In Deep Security as a Service, we provide most of the infrastructure for you. The information provided in this section will help you determine the compatibility and recommended software for your Deep Security Agents.

1.1 Operating Systems

Refer to the Supported Features by Platform in the Deep Security Help Center for the latest information.

2 Sizing Considerations

Sizing recommendations depend on the type of environment and various other factors such as network, hardware, software and applications. See Sizing in the Deep Security Help Center.

3 Installation and Deployment

Deep Security is composed of several components that need to communicate with each other. If you're deploying in a highly segmented network environment, knowledge about the various ports it uses will be useful for preventing unintended functionality disruptions. Make sure that all required ports are open and not reserved for other purposes.

Refer to the Port numbers, URLs, and IP addresses article in the Deep Security Help Center for a list of ports required in Deep Security.

3.1 Deep Security Components

3.1.1 Deep Security Agent/Relay

A. Deployment Considerations

1. Each computer should be able to resolve the fully qualified domain name of the Deep Security Manager for a successful deployment.

2. The clock on a Deep Security Agent (DSA) or Deep Security Relay (DSR) machine must be synchronized with Deep Security Manager within 24 hours. It is recommended to sync the time with NTP server.

3. If the client machine where Deep Security Agent or Deep Security Relay will be installed has a previous OfficeScan client, the drivers (tmactmon, tmevtmgr, and tmcomm) must be fully uninstalled prior to installation. After uninstallation finishes, rebooting OfficeScan is required.

   Deep Security Agent and OfficeScan client use the same name for drivers, however, Deep Security Agent cannot use OfficeScan client's drivers and OfficeScan cannot use Deep Security Agent's.

4. The "Enable Relay" button has been removed. Instead, go to "Relay Management".

   > **NOTE** 🗎 Disabling the relay feature on a Windows 10 agent can sometimes take more than ten minutes to complete.

   The new Relay Management page does not allow users to add or modify relay group descriptions.

5. Check the fully qualified domain name (FQDN) of the machine before and after the Deep Security Agent installation. A brief network interruption occurs during the agent installation process. Sometimes, it can affect Dynamic Host Configuration Protocol (DHCP) auto-registration. It is recommended to verify the computer's FQDN (ping -a <ip or server name>) before and after the installation. Should an issue with auto-registration arise, use ipconfig /registerdns or reboot the computer.

6. Deep Security Agent installation will disable iptables (Linux) or Windows Firewall (Windows) by default to avoid conflicts. In situations where the Deep Security Agent firewall feature is NOT used, refer to the steps below to prevent the installer from disabling iptables or from making any changes to the native Windows Firewall.

   a. For Windows, refer to the article Windows Firewall settings changed after installing Deep Security Agent (DSA) to modify the Deep Security Agent MSI package and prevent it from changing the Windows Firewall.

   b. For Linux, create or touch an empty file with the following path:

```
/etc/use_dsa_with_iptables
```

If the file is present, then the Deep Security Agent scripts will not disable iptables.

```
# touch /etc/use_dsa_with_iptables

# service iptables restart

# service ip6tables restart
```

7. Deep Security Relay's communication direction must be set to agent-initiated; otherwise, the rule update might fail to apply.

8. When you install Deep Security Agent, only send the installer (msi, rpm) file to the destination machine. The plug-ins can be deployed to Deep Security Agent based on policy.

B Agent Deployment Scripts

Deep Security Manager's deployment script generator can be used to generate scripts that run on computers where the agent will be installed. The script can be modified to optionally perform subsequent tasks like activation and policy assignment.

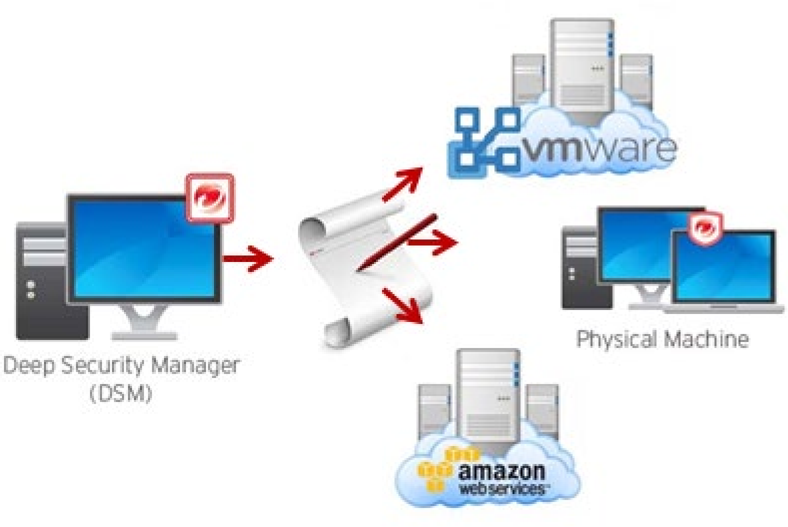

Figure 1: Deployment Script

Consider using deployment scripts in these scenarios:

- Environments where there is a need to deploy and activate multiple agents.
- Automate the activation process and deployment of policies.
- Activate and deploy to clients in environments where the server cannot communicate or discover clients directly, but clients can reach the server without problem.
- In Amazon Elastic Compute Cloud (Amazon EC2) and Azure environments, it can be bundled with an endpoint and used while instances are being auto-scaled.

Other Notes:

1. Deployment scripts only support basic function and cannot fulfill all needs for all environments. Adjust the scripts to fit your specific needs.

   Some environments might experience a delay in starting the ds_agent service. If the dsa_control activation signal is sent before the ds_agent service is started, this might prevent the activation from working successfully. Extend the sleep time in the scripts to prevent this.

   For example, in Amazon Web Services (AWS) testing, concurrent launching of 100 instances had better results when the sleep time was set to more than 60 seconds. This depends highly on AWS' system loading, disk I/O, CPU loading, network bandwidth, and database configuration.

2. All new instances must be able to access the URLs specified in the generated deployment script.

3. The agent-initiated activation feature must be configured correctly in Deep Security Manager for scripts to do activation tasks.

   The agent-initiated activation option must be enabled on the Administration > System Settings > Agents tab.

3.2 Testing Deep Security

Validate and test Deep Security features and functionality after deployment. Refer to the Testing the Deep Security modules article for guidelines on testing each module of Deep Security.

4 Configuration

Deep Security is a modular solution that can be adapted to different environments, so there is no right or wrong way to configure the product. Below are some common settings, exclusions, and other helpful configurations which appear in most Deep Security deployments. Double-check with your company's policies before adapting these recommendations.

4.1 UI Configurations

4.1.1 Dashboard

We recommend that at least the following widgets are included and placed on the area best seen on the dashboard page:

a. Alert Status – Keeps you informed of any critical items that might need immediate attention such as security updates and protection on computers going offline.

b. Computer Status – Gives you a good overview of agents' status.

c. My Account Status – Shows information about the user currently logged in.

d. Security Update Status – Shows information about out-of-date vs. up-to-date agents.

Create multiple dashboards and group them by usage (that is General, Anti-Malware, Updates and others) for easier management of large scale environments. Administrators can easily switch between them from the tabbed view. Each dashboard has a different time and computer filter, allowing multiple views into the system.

4.1.2 Alerts

By default, most alerts are enabled. In large environments, it can be beneficial to remove some alerts so only the ones that require action are triggered. Alerts should be configured to give the most relevant information, so the proper action(s) can be taken. From the alerts page, users can select Configure Alerts to enable or disable alerts.

4.1.3 Policies

Policies replicate security settings to servers and desktops that share similar security requirements. We recommend that machines with similar settings, software installed, application, or function be grouped strategically when assigning policies.

Note that the default policies built in Deep Security are meant to be examples and should not be used without prior configuration.

A. Policies vs. Computer Level Rule and Configuration Assignment

The best practice is to assign most rules through Policies for ease of management.

The advantages of using Policies are as follows:

- The user can change or test the policy settings before assigning it to the machines.

- It allows a quick removal of rules and configuration by simply taking out a machine from the policy or assigning it an entirely new one.

- It duplicates the policy and uses it as a baseline setting for future policies to be created.

When to use Computer Level rule assignment:

- Leveraging automatic assignment

- There are many varying computers (that is, each machine uses different applications, different OS updates, and so on, so they are virtually impossible to group)

NOTE 🗎 When using a combination of policy and computer level assignments, keep in mind that when you un-assign a policy from a computer, rules might still apply. This occurs if the rules were assigned independent of the policy.

B. Policy Groupings

Below are some recommended machine groupings to effectively take advantage of policies:

- By Operating System (for example, Windows 2008 Servers, Windows XP Machines, and Linux)

- By Server Function (for example, Mail Servers, Web Servers, User Laptops, and Point of Sale Systems)

- By Application installed/version (for example, OfficeScan Servers, Oracle 10 Database Servers, MS SQL 2005 Servers)

Properly grouping the machines is essential to effectively managing recommendation scans.

When a recommendation scan is performed on an individual member of a policy, the recommendations for that particular agent (Deep Security Agent) will be seen on the policy as well.

Accepting or applying the recommendations at the policy level will apply the rules to all members of the policy. The advantage of this method is the ease of maintenance. However, the disadvantage is that unnecessary rules might be assigned to certain members. For this reason, it’s recommended to group the machines accordingly, if users don't want to see the vulnerability being triggered for machines that should not be affected.

NOTE 🗎 Deep Security 10 supports multiple levels of policy inheritance. A newly-created policy can be configured to inherit all or some of its settings from a parent policy. It lets you create a tree structure of security policies. For example, you can create a parent policy called "Windows Server" and two child policies, "Windows Server 2008" and "Windows Server 2003", inherited from their parent policy. Each child policy can have child policies of their own for different editions of Windows Server.

Sample Policy grouping with policy inheritance:

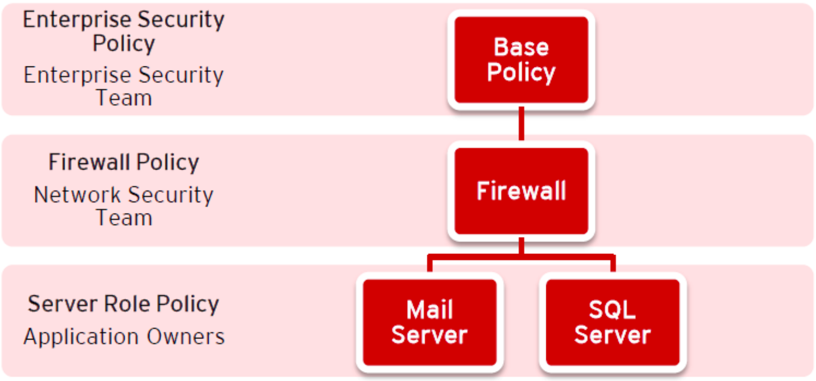

Figure 2: Policy inheritance

C. Policy Names

As a best practice, use a naming convention for policies to more easily manage multiple policies in an environment.

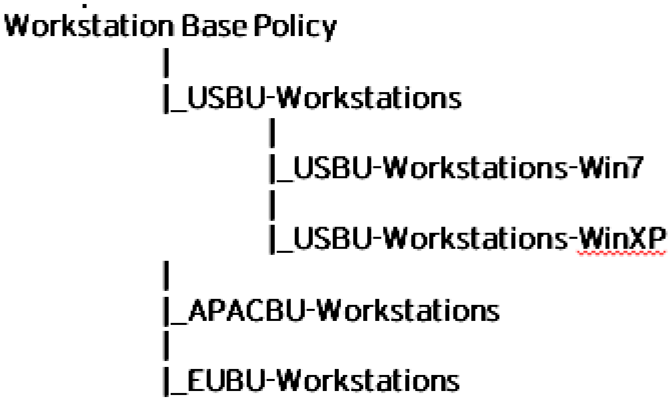

Figure 3: Sample of Naming Convention

4.1.4 Smart Folders

When using the Smart Folder function, be sure to identify Computer Name and Display Name correctly.

4.2 Module Configurations

4.2.1 Anti-Malware

A. Configuration

Go to Policies › Common Objects › Other › Malware Scan Configuration › Scan Settings.

Recommended Real-Time Scan Configuration	
Files to Scan	All Files
Directories to Scan	All directories
Actions	
Active Action	Disabled
Custom Actions:	Enabled
For Virus	Clean
For Trojans	Delete
For Packer	Quarantine
For Spyware	Quarantine
For Other Threats	Quarantine
Possible Malware upon Detection	ActiveAction
Options	
Enable Spyware / Grayware Scan	Enabled
Scan Compressed Files	Enabled
Maximum size of individual extracted files	Customized Size
Maximum Levels	2
Maximum number of files to extract	10
Scan Embedded Microsoft Office Objects	Enabled
Scan for Exploit Code in Microsoft Office Objects	Enabled
OLE Layers to Scan	3
Enable IntelliTrap*	Disabled
Enable Network Directory Scan	Enabled**
Scan Files When	Read/Write
Alert when ...	Enabled

Table 1: Real-Time Scan Configuration

* IntelliTrap helps block real-time compressed executable files and pairs them with other malware characteristics. Since IntelliTrap identifies such files as security risks and might incorrectly block safe files, you can disable IntelliTrap if users regularly exchange real-time compressed executable files. IntelliTrap only works in Real-Time mode.

**Network scanning should be disabled to maintain maximum performance during Real-Time Scan.

However, these network resources must be protected by a local AV scanner. Leave enabled if there is no other file scanner for these network shares.

Recommended Scheduled Scan Configuration	
Files to Scan	All Files
Directories to Scan	All directories
Actions	
Active Action	Disabled
Custom Actions:	Enabled
For Virus	Clean
For Trojans	Delete
For Packer	Quarantine
For Spyware	Quarantine
For Cookie	Delete
For Other Threats	Quarantine
Possible Malware – Upon Detection	Quarantine
Options	
Enable Spyware / Grayware Scan	Enabled
Scan Compressed Files	Enabled
Maximum size of individual extracted files	Customized Size
Maximum Levels	3
Maximum number of files to extract	10
Scan Embedded Microsoft Office Objects	Enabled
Scan for Exploit Code in Microsoft Office Objects	Enabled
OLE Layers to Scan	3
CPU Usage	Medium
Alert when…	Enabled

Table 2: Scheduled Scan Configuration

Recommended Manual Scan Configuration	
Files to Scan	All Files
Directories to Scan	All directories
Actions	
Active Action	Disabled
Custom Actions:	Enabled

For Virus	Clean
For Trojans	Delete
For Packer	Quarantine
For Spyware	Quarantine
For Cookie	Delete
For Other Threats	Quarantine
Possible Malware – Upon Detection	Quarantine
Options	
Enable Spyware / Grayware Scan	Enabled
Scan Compressed Files	Enabled
Maximum size of individual extracted files	Customized Size
Maximum Levels	2
Maximum number of files to extract	10
Scan Embedded Microsoft Office Objects	Enabled
Scan for Exploit Code in Microsoft Office Objects	Enabled
OLE Layers to Scan	3
CPU Usage	High
Alert when…	Enabled

Table 3: Manual Scan Configuration

When deciding which actions to take when malware is detected, note that there is a corresponding secondary action that will be triggered if the initial action fails to execute.

Primary Action (configured on the console)	Secondary Action (hardcoded)
Quarantine	Pass
Clean	Quarantine
Delete	Clean
Deny	Quarantine

B. Scan Schedule Setting

In addition to scan configurations, you can also set up a Real-Time Scan schedule. This can be useful if there is a specific timeframe in which you would like to turn off real-time scanning to improve performance.

Sample Scenario:

File Server is scheduled to have a backup of all files every day at 2:00am - 4:00am.

This server will most likely have high activity during this time, and whitelisting the 2:00am -4:00am timeslot from Real-Time Scan activity would significantly help improve performance for both the backup task and server resource.

> **NOTE** Perform a full manual scan on a server before running the actual backup task. We recommend that weekly scheduled scans are performed on all protected machines.

C. Multi-Threaded Processing

Real-Time Scan uses multi-threaded scans by default. However, for on-demand and scheduled scans, this option needs to be configured, depending on the environment.

Go to Policy/Computer > Anti-Malware > Advanced > Resource Allocation for Malware Scans.

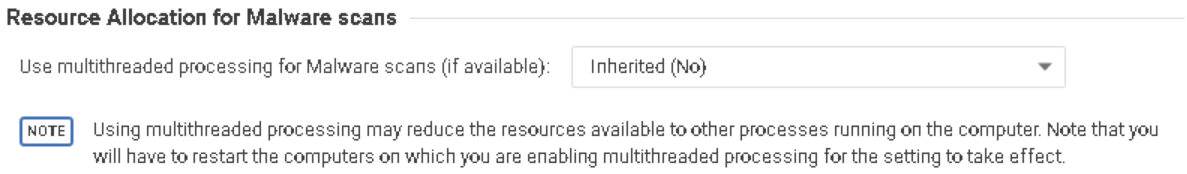

Resource Allocation for Malware scans

Use multithreaded processing for Malware scans (if available): Inherited (No)

NOTE Using multithreaded processing may reduce the resources available to other processes running on the computer. Note that you will have to restart the computers on which you are enabling multithreaded processing for the setting to take effect.

Figure 4: Resource Allocation for Malware Scans

Enable the option for physical machines using the physical Deep Security Agent to improve the performance. Note that restarting the machine is required for any change to take effect.

These are the scenarios where this setting should NOT be enabled:

- Agentless environments.
- If multi-threading is not an option, since the machine resource is limited (common for CPU-bound tasks).
- When a resource should be held by a single operator only at a time (common for IO-bound tasks).

D. Quick Scan vs. Full Scan

The Quick Scan feature improves the agent-based (Windows only) scanning time. It enables scanning for only critical files that are most likely to be infected. This allows more frequent quick scans to be scheduled with lower impact, and allows full scans to be performed on a less frequent basis (such as weekly).

Full Scan:

- Runs a full system scan on all processes and files.
- Uses the configuration set under manual scan (scans the files based on directories, extensions, files configured to be included in the scan).
- Runs at scheduled times by creating a scheduled task or manual scan (on-demand).
- Runs on all platforms supporting anti-malware.
- Takes longer to complete.

Quick Scan:

- Provides a fast, high-level scan of critical system areas for currently active threats.
- Looks for currently active malware, but will not perform deep file scans to look for dormant or stored infected files.
- Is significantly faster than a Full Scan on larger drives.
- Is only available for Windows Agent-based systems.
- Has no configurable settings and will not use any scan configuration (will not check settings like Directories to Scan or Files to Scan).
- Is only available on-demand. Quick Scans cannot be scheduled as part of a task.

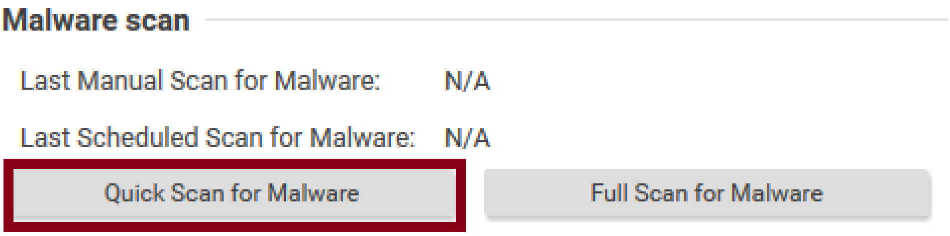

Malware scan

Last Manual Scan for Malware: N/A

Last Scheduled Scan for Malware: N/A

Figure 5: Quick Scan for Malware

E. Scan Exclusions

The Directory Lists, File Extension Lists, and File Lists can be set in the Common Objects section of Policies tab.

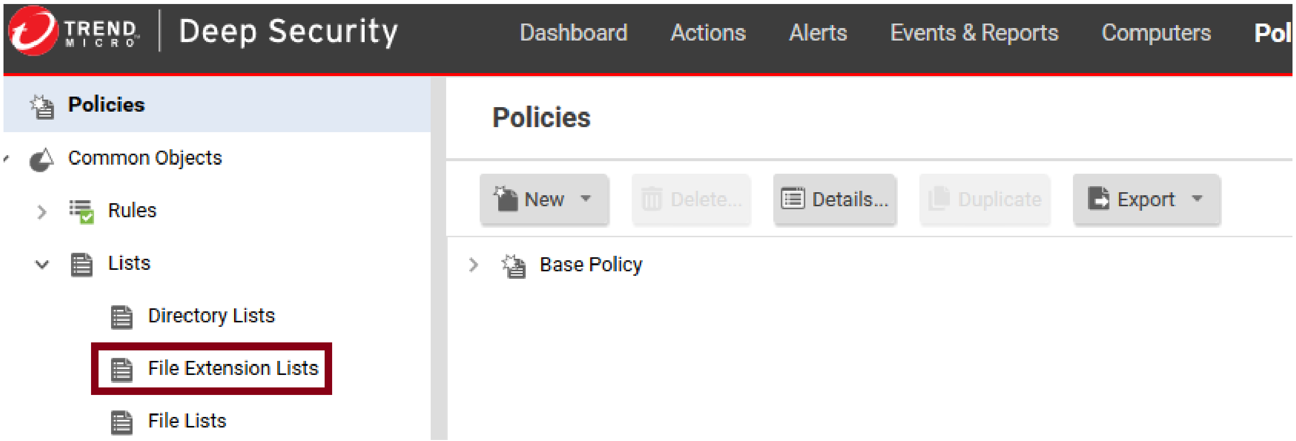

Figure 6: Directory, File Extension, and File Lists

These lists are then referenced on the Exclusions tab in the Malware Scan Configurations.

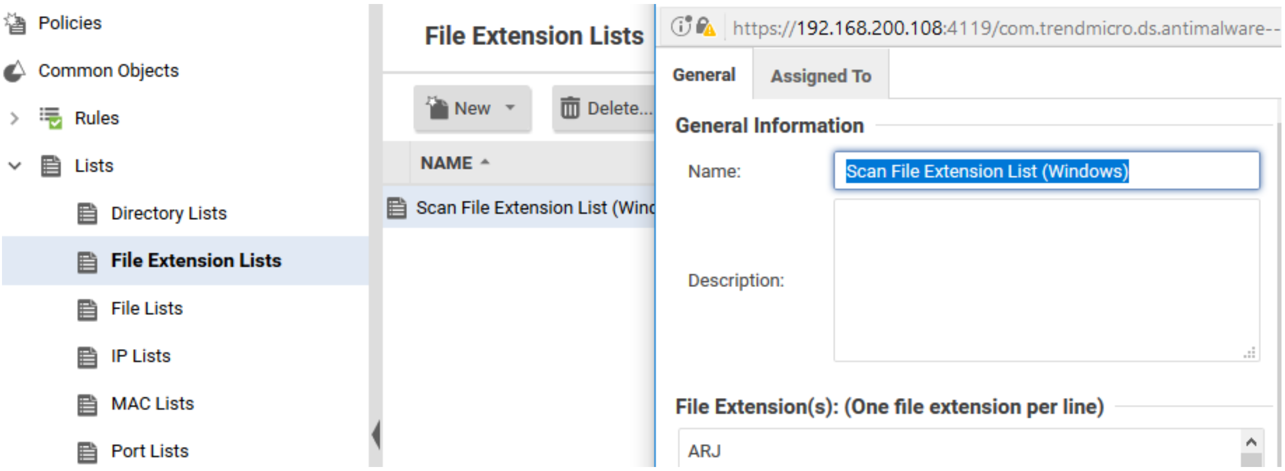

Name: Scan File Extension List (Windows)

Description:

File Extension(s): (One file extension per line)

ARJ

Figure 7: Exclusions tab of Malware Scan Configuration

Use this list as a starting point and refine it based on your environment and paths.

General Exclusions and Excluding Windows Update or Automatic Update Files

Files:

pagefile.sys

NTUser.pol

registry.pol

${Windir}\Software Distribution\Datastore\DataStore.edb

${Windir}\Software Distribution\Datastore\Logs\Edb*.log

${Windir}\Software Distribution\Datastore\Logs\Res1.log

${Windir}\Software Distribution\Datastore\Logs\Res2.log

${Windir}\Software Distribution\Datastore\Logs\Edb.chk

${Windir}\Software Distribution\Datastore\Logs\tmp.edb

${Windir}\Software Distribution\Datastore\Logs\hiberfil.sys

${Windir}\Software Distribution\Datastore\Logs\pagefile.sys

${Windir}\Software Distribution\Datastore\Logs\Edbres00001.jrs

${Windir}\Software Distribution\Datastore\Logs\Edbres00002.jrs

${Windir}\Security*.edb

${Windir}\Security*.sdb

${Windir}\Security*.log

${Windir}\Security*.chk

Directories:

${allusersprofile}\

${Windir}\system32\GroupPolicy\

${Windir}\Cluster\

Extension Exclusions:

*.pst

Microsoft Windows Server Domain Controllers

Files:

TEMP.edb

EDB.chk

Directories:

${Windir}\SYSVOL\

${Windir}\NTDS\

${Windir}\ntfrs\

${Windir}\system32\dhcp\

${Windir}\system32\dns\

Microsoft SQL Server

Large databases should not be scanned because it might hinder performance. Since Microsoft SQL Server databases are dynamic, exclude the directory and backup folders from the scan list. If it's necessary to scan database files, a scheduled task can be created to scan them during off-peak hours.

Directories:

${ProgramFiles}\Microsoft SQL Server\MSSQL\Data\

${Windir}\WINNT\Cluster\ # if using SQL Clustering

Q:\ # if using SQL Clustering

File Servers

Access to files over shared drives can degrade performance. To scan some file types, only a fraction of content is required. Other file types require a full scan or even a decompression.

Trend Micro recommends that file servers are excluded from scanning and perform scanning on the local file server itself. With exclusions in place, there is no need to scan the file as it is accessed, which increases performance.

It is also recommended to use agent protection for file servers for better performance.

> **NOTE** 🗎 If there are any custom applications not mentioned here, please contact the software vendor to get their recommended scan exclusions. You can also refer to the Recommended scan exclusion list for Trend Micro Endpoint products.

F. Quarantine Settings:
The default quarantined file settings are the recommended settings. To access the settings, go to Deep Security Manager > Policies > Common Objects > Other > Malware Scan Configuration > Advanced

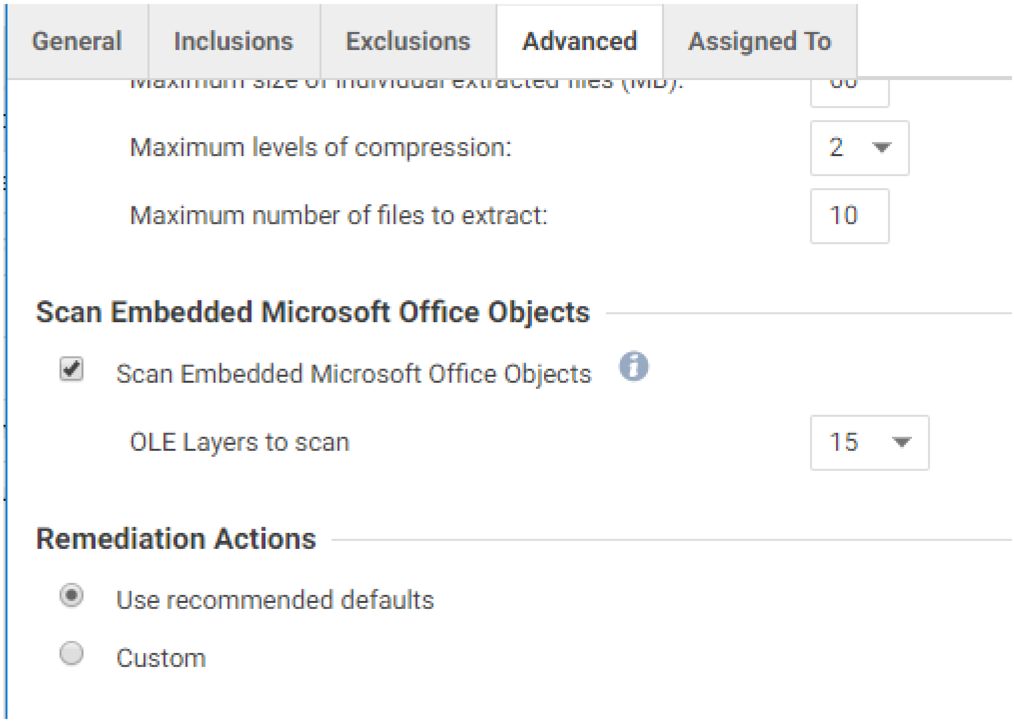

Maximum size of individual extracted files (MB): 60

Maximum levels of compression: 2

Maximum number of files to extract: 10

Scan Embedded Microsoft Office Objects

☑ Scan Embedded Microsoft Office Objects

OLE Layers to scan 15

Remediation Actions

● Use recommended defaults

○ Custom

Figure 8: Advanced Anti-Malware Settings

G. Maximum performance configuration for anti-malware:

To maximize the performance of the anti-malware feature, the following actions are recommended:

1. Use Scan files during "Read" for file scanning.
2. Add UNC path in the exclusion list. At the same time, check that the real-time scan is enabled so that all VMs are being protected.
3. Set up the proper exclusion list to exclude the folder, file, or extensions.
4. Set the scan limitation to prevent scanning a file larger than the specified size.

Read vs Write:

Read:
The system scans the virus when reading any files. This means you might download a test virus to your disk and until someone wants to run it, the system can catch the test virus when any file events are reading.

Write:
Write is important, but it affects the performance. Some FTP clients and browsers download a file by splitting the body to several pieces.

For example: Download a 100 MB file.
The browser can download 1 MB each time (write 1 MB and close the file, write another 1 MB and close the file and so on until the entire file has been downloaded). This means the file has to be scanned 100 times. The worst case scenario is if the malware hides itself in the last bytes. If we do not scan on write and the file is malware, it could be safe because it has been put on the disk without execution. We can scan it when it starts to be launched (read) and prevent its execution if a malware is found.

Write might detect malware in time, but it greatly affects the performance.

H. Use Security Enhancement Feature:

Server Platform

Server platforms use "Default Real-Time Scan Configuration" which turns security enhancement off by default. If you would like to enable security enhancement on the server platform, here are some suggestions:

1. Test on your staging environment first, before applying to your production environment.

2. You could add your critical applications to *Behavior Monitoring Protection Exceptions*. This can avoid false positives and impact your business. On the other hand, you will lose protection if your critical applications have been compromised.

NOTE: If you have added your critical applications to Exclusions > Process Image File List before, you don't have to add it to *Behavior Monitoring Protection Exceptions*. AM won't monitor any activities your application has if you've added it to *Process Image File List* exclusion.

3. Once any false positives occur, disable Security Enhancement first, including behavior monitoring, endpoint correlation and process memory scan in real-time AM configuration.

Availability

Security Enhancement protection relies on Trend Micro backend services. If you lose network connection of these backend services, Security Enhancement might not be able to protect you from advanced threats, such as ransomware attacks. Confirm that Trend Micro backend services are reachable from your environment and the proxy configuration is correct.

Security Enhancement relies on monitoring system activities, including file events generated by any process. Security Enhancement detects malicious behavior by tracing these system events. If you change the setting of AM config Inclusions > Scan Settings > Directories to scan from All directories to a specific Directory List, then only file events coming from this Directory List will be monitored. Monitoring capability outside the Directory List will be lost, and so detection capability and protection will be lost as well. For example, if you configure Directory List to "C:\MyFolder", a ransomware that has encrypted your files located outside C:\MyFolder won't be detected.

False Alarm mitigation

If your legitimate program is detected as a malicious program by Security Enhancement, add it to Behavior Monitoring Protection Exceptions. If that doesn't work, try disabling Endpoint Correlation, because this feature does not support exclusion. Unless you disable Endpoint Correlation, you won't be able to add your program to any exclusion list, including Behavior Monitoring Protection Exceptions.

Clean malicious program manually

Behavior monitoring, including ransomware protection, cannot quarantine malicious programs. It can only terminate that malicious process, without changing its program files. If a malicious program installs a run key or adds itself to the system schedule task, then it might be launched again after the system reboots or the task scheduled. Behavior monitoring will continue to terminate it periodically.

Once the system admin confirms the program is malicious, they must delete that malicious program manually to avoid further damage.

4.2.2 Web Reputation

The default security level "Medium" is suitable for most users. However, if you want further security, you can adjust it to the "High" level.

Web Reputation queries will go to the Smart Protection Server (if enabled) or to our cloud WRS servers. It's recommended to set up a local Smart Protection Server in house to limit the amount of required internet queries, which can lead to performance degradation.

If you are using products from Websense, be aware that there are potential incompatibilities between Deep Security Web Reputation and Websense's URL filtration. We recommend disabling the Web Reputation if the protected computer is behind a Websense edge appliance.

If you have specific web pages to allow or block, configure them in the Exceptions tab. By default, Web Reputation is enabled to port 80 and 8080. If you have an HTTP proxy server using other ports, configure it in the Advanced tab.

1. Create a new Port List from Shared > Port Lists including you proxy port (e.g. 3128).

2. Choose the created Port List at Web Reputation > Advanced > Ports.

Other setting recommendations:

- The Block pages that have not been tested by Trend Micro option should be unchecked. Otherwise, it could cause false positives.

- Include internal company URLs in the Allowed list under Exceptions. Wildcards are supported.

- Ensure that your company's firewall/proxy allows traffic going to https://ds10.icrc.trendmicro.com when using the global Smart Protection Server.

4.2.3 Firewall

Firewall configuration and administration must be performed carefully. There are no single set of rules that can fit all environments. This guide aims to give users best practice tips and recommendations that can be used as references and guidelines when building your own rules.

A. Inline vs. Tap Mode

- Use Inline Mode (Deep Security Manager > Policies > Settings > Network Engine > Network Engine Mode). When operating Inline, the live packet stream passes through the network engine. Stateful tables are maintained, firewall rules are applied, and traffic normalization is carried out. As a result, Intrusion Prevention rules can be applied to payload content.

- Use Inline Mode with rules set to Detect, when there is a need to test the configuration and rules before deploying them into the production environment. This way, the real world process of analyzing the traffic takes place without having to perform any action, such as blocking or denying of packets.

Running Deep Security in Tap Mode is NOT recommended. It is not the best practice to perform tests or evaluate Deep Security. Traffic patterns in this mode do not represent how the network will behave should the administrator decide to switch to Inline Mode.

B. Firewall Rule Actions

Know the difference between the firewall rule actions before creating your rules. Each rule can take one of the following actions:

- **Deny** – Explicitly blocks traffic that matches the rule.

- **Force Allow** – If a packet matches a Force Allow rule, it is passed but still filtered by Intrusion Prevention. No events are logged. This action type must be used for UDP and ICMP traffic.

- **Bypass** – Allows traffic to bypass both Firewall and Intrusion Prevention analysis. It should be created in pairs (for both incoming and outgoing traffic). Use this setting for media-intensive protocols only.

- **Log only** – If a packet matches a Log Only rule, it is passed and an event is logged. No other action will be taken.

- **Allow** – If a packet matches an Allow rule, it is passed and any other traffic not covered by a rule will be implicitly denied. Use this with caution.

C. Restrictive vs. Permissive Firewall

Typically, firewall policies are based on one of two design strategies. Either they permit any service unless it is expressly denied, or they deny all services unless expressly allowed. Decide what type of firewall you would like to implement to reduce administrative overhead in terms of creating and maintaining the rules.

Permissive Mode (Reactive)

- Permits all traffic by default and only blocks traffic it believes to be malicious based on signatures or other information.

- Easy to implement, however, it provides minimal security and requires complex rules.

- Rarely used, except in cases where you are not using the firewall but want to leverage it to block a port.

- Deny rules are used to explicitly block traffic.

Restrictive Mode (Proactive)

- The recommended best practice from a security perspective.

- Stops all traffic by default, and only allows traffic explicitly permitted.

- If the primary goal of your planned firewall is to block unauthorized access, the emphasis needs to be on restricting, rather than enabling, connectivity.
- Easier to maintain and more secured.
- Allow rules are used only to permit certain traffic across the firewall and deny everything else.

NOTE 🗎 Allow rules explicitly allow traffic that matches it to pass. In addition, it implicitly denies everything else that is not defined. Be careful when creating allow rules without defining the related rules correctly. Doing so can cause it to block all traffic apart from what the Allow rule is created for.

D. Stateful Inspection

Stateful configurations should be used when the firewall is ON.

The Stateful filtering engine inspects and validates each packet on an individual basis, which involves analyzing the packet within the context of traffic history, correctness of the packet's header values, and protocol state transitions. This enables protection against attacks such as denial of service, provided that a default configuration with Stateful TCP/ICMP/UDP is enabled and only solicited replies are allowed.

If the UDP Stateful option is enabled, Force Allow **must** be used when running UDP servers (like DHCP).

If there is no DNS or WINS server configured for the Deep Security Agents, a Force Allow, Incoming UDP Ports 137 rule might be required for NetBIOS.

Stateful logging should be disabled unless required for ICMP/UDP protocols.

E. Interface Isolation

Interface Isolation allows you to force a computer to use only one interface at a time. This prevents attackers from bridging across two interfaces. It is commonly used to protect users with wireless laptops.

Configure this via Policy > Firewall > Interface Isolation.

- Enter string patterns that will match the names of the interfaces on a computer in order of priority.
- Limit the number of active interfaces to one at any given time.
- It is not recommended to enable this at the global level. Enable it through the policy instead.

NOTE 🗎 Interface patterns accept wildcards such as asterisk (*) as well as regex expressions.

F. Other Recommendations

- Bypass Rules

  Bypass rules operate like force allow but skips the rest of the packet processing pipeline, so intrusion prevention is also skipped. Use this action for traffic that you prefer to allow across both the firewall and intrusion prevention.

  We recommend creating a pair of rules for each type of traffic. For example, create a rule bypassing the incoming traffic (request), and another to bypass outbound traffic (response).

- Rule Priority

  Rule priority determines the order in which filters are applied so high priority rules get applied before low priority rules. When actions share the same priority, the order of precedence for rules are **Bypass, Force Allow,** and then **Deny**. However, a deny action with a higher priority will take precedence over a bypass action with a lower priority.

  Note that **Allow** rules can only have a priority of **0**. Keep this in mind when using **Allow** rules to implicitly deny traffic (any traffic not matching the **Allow** rules are denied). This means when a **Deny** rule is added on the list, it will take precedence over all the existing **Allow** rules in place. Use Force Allow for traffic that should always be allowed (such as ARP).

  To simplify the administration of firewall rules, consider reserving certain priority levels to specific actions. For example, apply a default Priority 3 to rules that use **bypass**, Priority 2 for **Force Allow** rules and Priority 1 for **deny** rules. This reduces the potential for rule conflicts.

- ARP Traffic

  Always allow ARP. If a computer relies on dynamic ARP, include an appropriate rule to allow ARP. ARP forms the basis of the TCP/IP stack. ARP facilities provide translation from IP addresses to Ethernet addresses, which are essential for sending packets to other systems on the local LAN segment. Without this conversion, there can be no other form of peer-to-peer IP communication.

  Deep Security Manager should not instruct a Deep Security Agent to drop ARP packets, unless it's actually desired (configuration uses static ARP tables). To ensure this, follow these guidelines:

  - Enable the Trend Micro-provided ARP force allow rule.

  - Do not prevent broadcast ARP packets.

- Out Of Allowed Policy

  Out of Allowed Policy (Open Port) events can help quickly identify misconfigurations in rules. Generating these events for TCP, UDP, and ICMP advanced settings can assist with building and adjusting your policy.

  To configure this, go to Policy > Firewall > Advanced > Generate Firewall Events for packets that are Out of Allowed Policy.

- Use Port, IP, and MAC lists

  These lists are objects that can be reused by multiple rules. Using these lists in the configuration of multiple firewall rules facilitates configuration changes since only a single common list must be updated. Modifications done on any of the lists are picked up by all the rules where they are used or assigned.

- Number of rules

  Avoid assigning more than 300 rules, because doing so can affect system performance.

- Document all firewall rule changes

  Use the Description field of the firewall rule to note why, when, and for what purpose the rule was created for. Note when and why rules are created and deleted for easier maintenance.

- Advanced Network Engine settings
  To configure this, got to Policies > Policy > Settings > Network Engine > Advanced Network Engine Settings.

Established Timeout:

This parameter defines the maximum time an idle connection can be kept. Certain applications can require an active connection for longer, so increase the value when you have such applications and there are "Out of Connection" events.

Cold Start Timeout:

Specify the amount of time to allow non-SYN packets that could belong to a connection that was established before the stateful mechanism was started. Increasing this value can avoid an "Out of Connection" event after restarting Deep Security Agent or deploying a new profile.

Maximum TCP Connections and Maximum UDP Connections:

Maximum simultaneous TCP/UDP Connections are allowed. Consider heap memory size before adjusting these parameters.

Enable Debug Mode:

In debug mode, the Deep Security Agent and Deep Security Virtual Appliance captures a certain number of packets (specified by the setting below: "Number of Packets to retain in Debug Mode"). When a rule is triggered and debug mode is on, Deep Security Agent and Deep Security Virtual Appliance will keep a record of the last number of packets that passed before the rule was triggered. It will return those packets to the Deep Security Manager as Debug Events.

Number of Packets to retain in Debug Mode:

Specify the number of packets to retain and log when the Debug Mode is on.

Log All Packet Data:

Record the packet data for events that are unassociated with specific firewall or intrusion prevention rules. These are the log packet data for events such as Dropped Retransmit or Invalid ACK.

Minimum Fragment Size:

If legitimate traffic is blocked and there are First Fragment Too Small firewall events, change its value to "0" to disable the checking.

Minimum Fragment Offset:

If legitimate traffic is blocked and there are Fragment Offset Too Small firewall events, change its value to "0" to disable the checking.

For more tips and information about the Deep Security Firewall, refer to the article Understanding the features of Deep Security firewall.

4.2.4 Intrusion Prevention

A. Modifying Rules

Intrusion Prevention (formerly called Deep Packet Inspection) rules should never be modified at the global level (Deep Security Manager > Policies > Common Objects > Rules > Intrusion Prevention Rules) because there is no way to restore them. Configuration should be done by overriding the Policy or Computer. This way, the default master copy of the rules is kept on a global level and can be used as a reference, should there be a need to revert back changes.

Incorrect rules can cause downtime. You can create a rule based on the signature only. For those advanced rules (Start/End/Patterns or XML format), please contact Trend Micro Technical Support to obtain a qualified rule.

B. Using Detect Only or Prevent Mode

- If a specific rule is causing false positives, place that rule in **Detect Only** Mode or un-assign it.
- Any rule requiring configuration should be assigned **Detect Only** Mode until the rule can be configured for that computer.
- For new deployments, we recommend setting rules to **Inline Detect** Mode for easier identification of false positives.
- Once the tests and additional configurations have been made, switch a rule to **Prevent** Mode to start blocking the packets that match the rule.

C. HTTP Protocol Decoding

The HTTP Protocol Decoding filter is the most important filter in the Web Server Common Application Type. This filter is responsible for decoding the HTTP traffic before the other rules inspect it. In addition, this filter allows control over various components of the decoding process.

This rule is required should you choose to use any of the Web Application Common or Web Server Common filters that requires it. The Deep Security Manager automatically assigns this rule when it's required by other rules. Because each web application is different, the policy that uses this filter should run in detect-only mode for a period of time, before switching to Prevent Mode to determine if any configuration changes are required. Changes are often required to the list of illegal characters.

Refer to the following articles for more details on this rule and how to tune it:

HTTP protocol decoding in Deep Security

Modifying the list of URI characters that Deep Security Agent considers illegal

Troubleshooting the "Illegal Character in URI" error in Deep Security

D. Cross-Site Scripting and Generic SQL Injection Rules

Two of the most common application-layer attacks are SQL injection and cross-site scripting (XSS). SQL injection rules and cross-site scripting intercept the majority of attacks by default. Adjust the drop score for specific resources if they are causing false positives.

Both rules are smart filters that require custom configuration for web servers. If you that have output from Web Application Vulnerability Scanners, leverage that information when applying protection. For example, if the user name field on login.asp page is vulnerable to SQL Injection, ensure that the SQL Injection rule is configured to monitor that parameter with a low threshold to drop on.

More details can be found in this article: Understanding the Generic SQL Injection Prevention rule.

E. Filtering SSL Data Streams

Deep Security Manager supports intrusion prevention analysis of SSL traffic and is able to filter SSL encrypted data streams. The Deep Security Agent does not filter SSL connections that use compression.

This can be assigned and configured on individual computers. Open the Details window of the computer you wish to configure, and go to Intrusion Prevention > Advanced > SSL Configurations > View SSL Configurations.

> **NOTE** 🗎 This feature might cause a performance impact. It is not recommended for servers with high numbers of connections per second.

If this feature is activated, it's recommended to disable the inspection of HTTP responses to avoid performance degradation. As all web attacks that we protect against are included in the HTTP request and not the HTTP response, disabling inspection on responses will improve performance.

To configure this:

a. Go to the computer or policy > Intrusion Prevention.

b. Select a rule with Web Server Common app type, right-click Application Type > Properties.

c. Go to Configuration tab and uncheck Inherited.

d. Uncheck Monitor responses from Web Server.

e. Update the changes to the computer/policy.

F. Other Recommendations

- Set the rules to only log dropped packets to save disk space.
- If rules are manually assigned, do not assign more than 300 rules as it affects system performance.
- Use Recommendation Scan to apply the necessary rules for the best protection and performance.
- Only select the Always Include Packet Data option (Rule Properties > General > Events) if you're interested in examining the source of attacks. Otherwise, leaving the packet data logging on will create much larger log sizes.
- Application types under intrusion prevention rules should be checked prior to use.

  For example, Trend Micro OfficeScan and Trend Micro OfficeScan NT Listener application types are inspecting incoming ports 8080, 4343, 26964, 24880, and 46485 by default.
- OfficeScan ports can be changed, especially the random 5-digit client port. These rules should be re-configured to match your OfficeScan settings before assigning them.
- One port cannot be assigned to more than eight application types; otherwise the rules will not work on that port.

G. Interface Tagging

By default, firewall and intrusion prevention rules are assigned to all interfaces on the computer. You can use Interface Types to assign firewall or intrusion prevention rules to a specific interface on a machine that has multiple interfaces (for example, if there are some specific rules you would like to apply to only the wireless network interface).

To configure, go to Policy > Interface Types > Network Interface Specificity.

Think about the difference in protection for different interfaces when creating policies. Consider populating the Interface Type based on the different networks available to all potential Deep Security Agent protected machines.

H. Ransomeware Detection

Refer to the article Ransomware Detection and Prevention in Deep Security, which includes recommendations on rolling out the rule.

I. TippingPoint Network Security

Many customers are benefitting from both TippingPoint network security and Deep Security host security. Intrusion prevention (IPS) rules now show the TippingPoint ID of the equivalent TippingPoint rule.

4.2.5 Integrity Monitoring

Monitoring the operating system, application files and directories is an excellent way to maintain the integrity of the data on your server. Unexpected changes to these files can be a good indicator that something suspicious has occurred and should be investigated. Rules created for Integrity Monitoring should be as specific as possible to improve performance and avoid conflicts or false positives. Do not try to create a rule that monitors the entire hard drive.

A. Using Integrity Monitoring to protect against malware

Integrity Monitoring can monitor files and registries. Malware typically infects a system by modifying certain registry keys and various system files. The default Deep Security rules allow you to monitor the integrity of a machine by observing what is most commonly changed by malware in an infected system. Here are a few sample rules that are applicable for all types of situations in Windows platform:

- Rule 1002773 – Microsoft Windows – 'Hosts' file modified
- Rule 1002776 – Microsoft Windows – 'All Users' Startup programs modified
- Rule 1002778 – Microsoft Windows – System DLL or EXE file modified

Unless new software or a security patch is installed, there is no reason why any of these files should be modified. If such an event is raised, the administrator can check what is happening on the machine to determine whether or not it's compromised.

It's also possible to create custom rules to monitor specific threats. If a user knows the behavior of a particular virus they are trying to contain in an environment, they can create a special monitoring rule that checks for certain registry keys or files created by the virus. This can determine if the spread of the virus is being contained.

Note that Integrity Monitoring detects changes made to the system, but cannot prevent or undo these changes.

B. Baselines

Baselines are automatically created when integrity monitoring rules are assigned to a computer. Retrieving baselines is necessary to recognize any abnormal behavior that might occur. Trend Micro recommends enabling Scan Computers for Integrity Check for computers.

C. Rules from a Recommendation Scan

Recommended integrity monitoring rules typically result in too many monitored entities and attributes. Decide what is critical and what should be monitored, then create custom rules or tune out of the box rules.

Pay attention to the rules that monitor frequently changed properties, like process IDs or open ports, as they can be very active and might require some adjusting.

D. Trusted-Source-Based Event Tagging

When the Integrity Monitoring feature is used, depending on the rules and settings, it might be difficult to search and determine which events are good and informational, and which events need further investigation.

The Deep Security auto-tagging feature helps to group and label multiple events to suppress security events for legitimate changes.

To configure this feature, go to Deep Security Manager > Events and Reports > Integrity Monitoring Events > Auto-Tagging > Trusted Source.

Deep Security allows administrators to automatically tag authorized changes by using internal reference servers, Certified Safe Software Service that Trend Micro hosts in the cloud, or by comparing it with other computers in a group. Certified Safe Software Service is a cloud-based database of signatures that Trend Micro has certified as known-good files. More information on how to enable Trusted-Source-Based Event Tagging can be found in the Online Help and Administrator's Guide of Deep Security.

Selecting the Trusted Source:

- Local Trusted Computer

  Use this when implementing a Golden Host model, where applications and files installed on the Golden Host are used as a basis for comparison.

  This model is most useful when:

  - There are in-house applications installed on the local trusted computer.
  - Software, service packs or patches are installed on the local trusted computer and can be used as a reference for other computers.
  - The local trusted computer is malware-free and secure.
  - The local trusted computer contains Integrity Monitoring rules that are similar to the computer that will use it as reference.

  Best Practices:

  - The security events from the trusted computers must be collected before the security events from other computers. You can use scheduled task to automatically scan trusted computers.
  - Create two scheduled integrity monitoring scans. The first scan only checks the trusted computers while the second scan checks the others.
  - To only trust events that have been generated as part of a maintenance window, leverage the Pause Collection functionality available in the Auto-Tag Rule properties. This functionality disables automatic additions of new information to the Known Good Store based on changes on the trusted source, when the collection has been paused. When paused, the events from the

associated computers related to previously trusted events will continue to be tagged. However, new information will not be added to the Known Good Store until collection is resumed.

Certified Safe Software Service

Use this when there are no local reference servers and users are free to install and upgrade software by themselves or at any given time. In this scenario, files are compared against Trend Micro's database of known-good files.

Best Practices:

- Ensure the Deep Security Manager has connection to the internet to query this cloud-based service.
- Certified Safe Software Service only supports SHA-1. If this service will be used, the Policy > Integrity Monitoring > Advanced tab > Content Hash Algorithms should be set to SHA-1.
- Among the three trusted-source-based event tagging mechanisms, Certified Safe Software Service is the safest and most secure because there is no need to maintain a reference server. Trend Micro is responsible for ensuring that the cloud service only contains known-good files.
- Certified Safe Software Service should have top priority over other auto-tag rules.

Trusted Common Baseline

Use this when a group of computers can use each other as reference. The baselines of the computers in this group will be added to the common baseline. The computers in this group should be secure and free of malware, as changes in one computer will automatically be added to the baseline. When a similar event occurs on another computer in the group, the event will automatically be tagged.

Best Practices:

- The trusted common baseline auto-tagging rule should be in place before any integrity monitoring rules are applied to the computers in the common baseline group.
- Group the computers sharing the same operating system and function (for example, Microsoft SQL servers running on Windows 2008 R2).
- Setup and maintenance of trusted common baseline is easier compared to local trusted computer, but the level of protection is lower because all computers in the group are considered trusted. Trusted common baseline should be set to the lowest priority.

E. Real-Time File Integrity Monitoring

In Deep Security 11.0 and later, the updated file monitoring engine is shared with application control and allows real time detection of file changes for both Linux and Windows. Previously, Linux integrity scans were scheduled only. This enhancement improves Deep Security's ability to meet compliance requirements.

- Beginning in Deep Security 11.0, integrity monitoring supports real-time monitoring of file changes for both Linux and Windows.

> **NOTE** 🗎 The Deep Security Agent for both 64-bit Windows and 64-bit Linux now depends on the application control plugin to trigger the real-time file system events that are sent to the integrity monitoring plugin.

- 32-bit Windows platforms will run in legacy mode and will not provide the user and process information for real-time change events. As in previous releases, real-time integrity monitoring is not available on 32-bit Linux platforms.

- Only real-time file events include information about the user/process that made the change. As in previous releases, other type of integrity monitoring events such as change to services or running process will not include this information.

F. Change details

In Deep Security 11.0 and later, the updated file monitoring engine will capture "who" made changes to a monitored file. This attribute is critically important for users to investigate and respond appropriately to change events and can help meet compliance requirements.

4.2.6 Log Inspection

Events from the Windows event log and other application specific logs are a great source of information for the status of your server and applications. Log Inspection is an automated solution to inspect these log files for suspicious events and alerts, and is a valuable feature to include for your defense in depth strategy.

This feature is especially useful for creating easier access to important events in monitored log files, without manually tracing through it.

- Log Inspection rules must be properly configured. Most recommended rules work well, but Windows Event rules should be adjusted to gather security events relevant to your requirements. Events for this feature can overwhelm the Deep Security Manager database if too many log entries are triggered and stored.

- Severity Clipping

  - Send Deep Security Agent events to syslog when they equal or exceed the following severity level: This should typically be changed when a syslog server is used. This setting determines which event triggered by those rules is sent to the syslog server (if syslog is enabled).

  - Store events at the Deep Security Agent for later retrieval by Deep Security Manager when they equal or exceed the following severity level: This setting determines which log inspection events are kept in the database and displayed on the log inspection events screen. Custom rules can be made to monitor logs that are not in the included set of rules.

4.2.7 Application Control

When Application Control function is enabled for blocking executables, the below extensions will be blocked specifically:

.class

.jar

.war

.ear

.php

.py

.pyc

.pyo

.pyz

If your environment has an extension mentioned above but is considered a safe file, add it to the whitelist in software inventory.

In Deep Security 11.0, Application Control was enhanced with a new Block by Hash feature that allows administrators to submit known bad hash values to Deep Security for Application Control blacklist enforcement.

The control recognizes a new "Global rule set" that includes a list of hash values to be blocked. This rule set takes precedence over any other rules from existing shared or local rule sets, and will be enforced by every Deep Security Agent enabled with Application Control. This feature provides a simply way for users to block unwanted or bad software from running at a global system-wide level. The design allows the workflow to be fully automated; with APIs for creating the Global rule set, adding and deleting hash values.

Application Control creates a software change event log whenever new executable files are detected on protected systems. Sometimes these changes are generated as part of the normal operation of trusted software. For example, when Windows self-initiates a component update, thousands of new executable files may be installed. Application control will now auto-authorize many of these file changes when created by well-known Windows processes and no longer create corresponding change log events. By removing the "noise" associated with expected software changes, users will have clearer visibility into changes that may require their attention.

Before you deploy this feature, make sure you already checked the support matrix of application control vs features in the Supported features by platform article in the Deep Security Help Center.

NOTE 🗎 When using Application Control, if you create a golden image, update it with required patches, create a shared ruleset, and then apply that shared ruleset to other computers. When you install those same patches on the other computer, they will be allowed to execute because they are in the shared ruleset. However, the patch updates will appear on the Software Changes page. To avoid this, Application Control must be set to Maintenance Mode when applying patches.

4.3 Administration and System Settings

4.3.1 Recommendation Scan

The recommendation engine is a framework that exists within Deep Security Manager, which allows the system to suggest and automatically assign security configurations. The goal is to make the configuration of computers easier and only assign security that is required to protect that computer.

Recommendation scans affect the performance impact of Deep Security Manager, so schedule these when no other tasks are running.

A. Run recommendation scans weekly

Recommendation scans can impact Deep Security Manager's performance, so avoid scanning with high frequency. Systems that don't change often (servers) can be scanned less frequently. Systems that lack control over when changes occur (workstations) should be scanned more frequently.

Ongoing scans for recommendations are not advised, this setting should be set to "no".

(Policy/Computer › Settings › Scanning › Recommendations › Perform ongoing scans for recommendations)

If ongoing scans are set to automatically start, administrators have no control over when it will occur. The best practice is to create a new scheduled task with type "Scan Computers for Recommendations" to take place once a week instead.

B. Run scans after a major change (application of a patch, installation of new application, etc.)

Scans should be performed after major changes to the computer to determine if any additional protection is required.

C. Run scans after applying a new Deep Security Update.

This allows you to use the recently released rules, and get the latest updates assigned or unassigned.

D. Assign recommended rules to the policy, not to the computer.

As a best practice, recommended rules should be assigned to the policy and not directly to computers.

Recommended rules can only be applied automatically to the machine where the recommendation scan was ran. Refer to the Policy section for additional details.

E. Run the scan on computers with similar functions.

In environments with similar computers, scans can be performed on a subset of computers to gather baseline recommendations for them all.

F. Automatic Assignment of Intrusion Prevention Recommendations

This option is disabled by default (**Policy/Computer > Int rusion Prevention > General > Recommendations).** It's not recommended to enable this option on the computer level. An exception to this would be when the machine is on its own and cannot be associated with other machines in a group. When this is enabled, intrusion prevention rules will automatically be enabled on the machine when the rule is found to be applicable, or a matching application is found on the machine related to the rule.

See Policy vs Computer Level for more details.

Disabling this setting gives administrators better control on assigning and un-assigning recommended rules.

4.3.2 System Settings

A. Communication Direction

This option can be set at the policy or computer level. The default **Agent-Initiated** method is recommended and used in most production deployments.

To configure this setting, go to Policy/Computer > Settings > Computer > Communication Direction.

B. Heartbeat Settings

This can be configured at the policy or computer level. Look for it **in Policy/Computer > Settings > Computer > Heartbeat.**

Heartbeat Interval

- Servers - 10 Minutes
- Desktops – 60 Minutes

The most important factor in choosing the interval setting is the acceptable amount of time between when an event triggers, and when the events are delivered to the Deep Security Manager. Choosing a high frequency can have a negative impact on the Deep Security Manager's performance.

Why do servers require a lower heartbeat (more frequent interval)?

They are typically more critical assets, and administrators might want to be notified of relevant events more frequently.

If protection is in place when roaming, why would administrators want a laptop to connect to Deep Security Manager while off network?

To have the ability to update the policy on the laptop when roaming. Also, events are stored in the Deep Security Manager with the event timestamp, not the timestamp when they were delivered to Deep Security Manager. Historical events can often be overlooked for devices that haven't performed a heartbeat in the last 24 hours.

Number of heartbeats that can be missed before an alert is raised

By default, the value is "2". If a heartbeat is missed after two attempts, the agent will be tagged as offline. We recommend increasing this value in most environments, so agents that are actually online won't be tagged as frequently.

In addition, if a heartbeat fails, events are stored locally to Deep Security Agents until the connection is restored.

If SIEM or Syslog servers are used to store events, heartbeat setting are less of a concern. Agents send events to Syslog in real-time, without batching and waiting for the next heartbeat.

C. Agent-Initiated Activations

This option is most common for environments with large distributed installations, where it's more desirable for the activation to be initiated by the agent, rather than the Deep Security Manager. This is the preferred method of communication for DSaaS.

- Very useful when a large number of computers are added to a Deep Security installation and script can be used to automate the activation process. See Deployment Scripts.

- For Agent-Initiated Activation to succeed, the **Allow Agent -Initiated Activation** option must be enabled on the **Administration > System Settings > Agents** tab.

  During the activation, the agent can determine the assigned policy and apply it. Additionally, agents can request scans or updates after they have been activated. This can be used to tightly integrate scans to other changes, such as patch management. Refer to the product Online Help or Administrator's Guide for additional details.

- Allow reactivation of cloned VMs.

  This is used in environments with VM clones (for instance, cloning new VM/instance from pre-activated VM, templates, or AWS images).

  Below are some notes to consider:

  - VM/Instance must be managed under Cloud Account.

  - VM/Instance must have unique system IDs (BIOS UUID, MAC addresses, hostname, IP).

    - Ensure the network communication in the environment has no communication issues. This helps prevent the host from going offline or getting a mismatch.

    - Cloned VM – Original VM must remain activated

    - Clone activation will not migrate any policies or settings from the original VM.

- Allow reactivation of Unknown VMs

  This allows previously activated VMs, which have been removed from their cloud environment and deleted from Deep Security Manager, to be reactivated if they are added back to the inventory of VMs.

  This is useful if the server deleted the agent by accident or if the server deactivated the agent, but the agent did not receive the deactivation request.

  Below are some notes to consider:

    - VM MUST have a valid server certificate but no activation record on current Deep Security Manager server(s).

    - Unknown activation will not migrate any policies or settings from the original VM.

D. Send Policy Changes Immediately

By default, this setting is turned on. If there are changes made to any setting within the Deep Security environment, all affected computers are immediately updated.

Change the setting by going to Policy/Computer > Settings > Computer > Automatically send policy changes to computers.

It's recommended that this option is disabled. Instead, use a scheduled task to update and send policy changes to agents manually. Manual or scheduled updates give the administrator more control to follow the existing change control process. Scheduled tasks can be set to update machines during maintenance windows, off hours or other times with low traffic.

To monitor when machines were last updated, administrators can use the "Last Successful Update" information on the **Computers** tab of Deep Security Manager.

E. Agent Self-Protection

When the agent is installed, Deep Security Agent can protect its services, installation directories and status from any modification, including shutdown from the self-protection setting.

If this setting is turned on, enable and set a password for the local override setting by going to **Policy/Computer > Settings > Agent Self Protection .**

F. Scheduled Tasks

Tasks can be configured to automate certain common tasks by a schedule. Below is a list of recommended tasks to establish:

- Check for Security Updates (Frequency: Once Daily)

- Scan Computers for Malware (Frequency: Once Weekly, or in accordance to company policy)

- Scan Computers for Recommendations (Frequency: Once Weekly)

- Send Policy (Frequency: Once Weekly, and run as needed)

When scheduling recommendation scans, the best practice is to set the task by group (per policy, or for a group of computers, no more than 1,000 machines per group) and spread it to different days (database server scans are scheduled every Monday, mail server scans are scheduled every Tuesday, and so on).

G. Log Retention

The best practice is to run the data pruning feature built into Deep Security Manager. If there is a compliance requirement to keep log sets for a longer period of time, the recommendation is to use third-party SIEM products to store the data.

Event retention is relevant to maintain a reasonably sized database. Default retention time settings are outlined in Log and event storage best practices in the Deep Security Help Center.

H. Using Tags for Events

Tagging events allows administrators to manually tag events with pre-defined or custom labels. This makes log monitoring and review more efficient.

To configure tags and auto-tag rules, go to **Policies > Common Objects > Other >Tags**

See also Trusted-Source-Based Event Tagging.

5 Disaster and Recovery

5.1 Recovering a physical machine (with Deep Security Agent) in a Disaster

Sometimes, assigning an incorrect policy or rule can completely isolate a machine from the network. To remove a faulty rule or policy, do one of the following:

1. If rules have been applied to the policy only, remove the faulty rule from the policy and trigger a Send Policy to the affected machines.

   a. Go to Policy and double-click the affected policy.

   b. Click Firewall/IP > Assign/Unassign the rule and press Save.

   c. On the affected machines, right-click Send Policy.

2. If rules have been applied directly on the machines, open the details for each affected machine and remove the faulty rule.

   a. Go to the affected machine and double-click for details.

   b. Select Firewall/IP > Assign/Unassign the rule and press Save.

   c. Go to Overview > Actions > Send Policy or right-click on the affected machine under Computers > Actions > Send Policy.

3. If you do not know which rule is at fault, remove the entire policy from the machine.

   a. Right-click the affected machine**, then go to Actions > Assign Policy > None**

   b. Right-click the affected machine**, then go to Actions > Send Policy**

4. If the rule involved is a firewall or intrusion prevention rule, you can also consider turning the firewall and intrusion prevention state to "Off". You can do this locally on the affected machine or on the Policy under the **General** tab.

5. If Deep Security Manager cannot communicate with the agents, log on locally to the machine and trigger an agent reset to completely clear all configurations on the agent and deactivate it.

   On the command prompt of the local agent, run:

   dsa_control /r

   The "Reset" action does the following:

   - Cleans up all Deep Security Agent configuration settings and Deep Security Agent memory
   - Removes relation between Deep Security Agent and Deep Security Manager
   - Removes corresponding entries from the database

   Refer to Agent Self Protection for more details.

6. Reactivate using a new policy without the recent change.

5.2 Isolating a Deep Security Issue

1. It's recommended to first isolate the module causing the issue, as opposed to deactivating or uninstalling the agent. Check the related event logs for information and clues regarding the issue.

   If no related logs are observed and multiple features are used, turn off the suspected module one by one to find the culprit.

For example, if the issue involves HTTP blocked traffic, first turn off WRS and then the firewall.

2. For issues involving WRS:

   - If traffic to a certain site is blocked, consider adding it to the "Allowed" URLs by going to the **Policy/Computer > Web Reputation > Exceptions** tab. Enter the URL in the allow list, save, and send the policy.

   - If adding the site to the allow list does not help, turn off the web reputation (**Policy/Computer > Web Reputation > General > Web Reputation State**).

   - If WRS is turned off and the issue still persists, check other enabled features.

3. For issues involving the firewall:

   - Note if a new rule or a modification on a rule has taken place. Un-assign the suspected rule and verify if the issue persists.

   - If you are not sure which rule is causing the issue, consider removing the policy assigned to the affected machine. Verify if the issue still persists.

   - If no recent change has occurred but traffic is blocked, turn the firewall off. To do this, go to **Policy/Computer > Firewall > General > Firewall State**.

   - If the firewall is disabled and the issue persists, verify that Firewall Stateful Configurations are also set to "None" **(Policy/Computer > Firewall > General > Firewall Stateful Configurations).**

   - If both settings are turned off and the issue persists, switch the Network Engine to "Tap" mode. Go to Policy/Computer > Settings > Network Engine > Network Driver Mode.

   Should the issue still persist, check the other features that are enabled.

4. For issues involving intrusion prevention:

   - Note if a new rule update has been applied or a modification on a rule has taken place. Un-assign the suspected rule or roll back the security update. Verify if the issue persists.

   - If you are not aware which rule is causing the issue, consider removing the policy assigned to the affected machine. Verify if issue still persists.

   - If no recent change or update has been applied but traffic is blocked, switch the behavior from "Prevent" to "Detect" or turn off the intrusion prevention. Both settings can be found under **Policy/Computer > Intrusion Prevention > General >Intrusion Prevention State/Behavior** .

   - If intrusion prevention is turned off and the issue still persists, switch the Network Engine to "Tap" mode. Go to **Policy/Computer > Settings > Network Engine > Network Driver Mode.**

   - If the issue still persists, check the other features that are enabled.

5. For issues involving anti-malware:

   Performance Related:

   If there are performance or access issues when the AM module is turned on, consider adding the directory or file being scanned to the exclusion list first. To do so, go to the Scan Configuration used by the Computer/Policy (**Policy/Computer > Anti-Malware > General> Select Scan type > Configuration > Edit > Exclusions).** Verify if the issue still persists.

   If adding the file or directory to the exclusion does not work, remove the policy assigned to the affected machine.

   - If the issue persists, turn off anti-malware protection. Go to **Policy/Computer > Anti-Malware > General > Anti-Malware State**.

   - If the issue continues, de-activate the agent.

   - Should the issue still persist, check the other features that are enabled.

   Detection Issues:

   - If the issue involves undetected malware, verify the anti-malware state and make sure there are no errors. Check for failed events under **Policy/Computer > Anti-Malware > Events.**

     Consult the following articles for Anti-Malware state verification:

     Verifying a successful Deep Security Virtual Appliance (DSVA) installation

     "Anti-Malware Driver Offline" status appears when logging on to the Deep Security Manager (DSM) console with vShield manager

   - Verify Smart Protection settings and ensure there are no connection failures (**Policy/Computer > Anti-Malware > Smart Protection).**

   - Should the issue persist, contact Trend Micro Technical Support.

6. For issues involving integrity monitoring:

   - Note if a new rule update has been applied or a modification on a rule has taken place. Note the additional modifications made and review the configuration changes. You can also un-assign the suspected rule or roll back the security update. Verify if the issue persists.

   - If no recent change or update has been applied but alerts continue to be generated, turn off the integrity monitoring by going to **Policy/Computer > Integrity Monitoring > General > Integrity Monitoring State** .

   - Should the issue still persist, check the other features that are enabled.

7. For issues involving log inspection:

   - Note if a new rule update has been applied or a modification on a rule has taken place. Note the additional modifications made and review the configuration changes. You can also un-assign the suspected rule or roll back the security update. Verify if the issue persists.

   - If no recent change or update has been applied but alerts continue to get generated, turn off the log inspection by going to **Policy/Computer > Log Inspection > General >Log Inspection State**.

- Should the issue still persist, check the other features that are enabled.

6 Other Deployment Scenarios

6.1 Environments using Teamed NICs

Windows NIC teaming software creates a new virtual master interface, which adopts the MAC address of the first subordinate interface. By default, the Windows Agent will bind to all virtual and physical interfaces during installation. As a result, in a teamed NIC environment, the agent will bind with the physical interfaces as well as the virtual interface created by the teaming software. The agent cannot function properly with multiple interfaces that have the same MAC address.

To function properly in a teamed-NIC environment, the agent must be bound only to the virtual interface created by the teaming software. For more information, refer to this article: Deep Security Agent and Vulnerability Protection Agent are unable to attach to Intel Teamed NIC Virtual Adapter.

1. Using the agent in a teamed-NICs environment on Windows 2003 requires SP 2 or later.

2. The Deep Security Agent's network driver is bound to the network interfaces for only the installation or upgrade period. After installation, it is impossible for the bindings to be automatically adjusted when you add or remove network interfaces to or from a teamed-NIC.

   Doing so can lead to network connectivity problems, or to the system not being properly protected. After adding or removing a network interface in a teamed environment where the agent's network driver is installed, verify that the driver is only bound to the virtual interface and not bound to any physical adapters.

3. On Solaris systems with multiple interfaces on the same subnet, the operating system may route packets through any of the interfaces. Because of this, any Firewall Stateful Configuration options or intrusion prevention rules should be applied to all interfaces equally.

6.2 Solaris Zones

Keep in mind that Solaris Zones allows multiple instances of Solaris to run in one shared kernel.

The Deep Security Agent for Solaris is only supported to run with the Global/Root Zone. Refer to the article Installing the Deep Security Agent (DSA) on Solaris in the global zone for more details:

6.3 Microsoft Cluster Servers

Cluster servers involve two separate installations of the underlying operating system with shared resources (databases, disks, IP addresses) that get swapped back and forth when the cluster performs a failover.

Deep Security can be configured to protect one node in the cluster or both. In this environment, consider the following:

- That you are installing Deep Security Agent to a local disk, and not a shared disk.

- If the cluster software uses a network heartbeat with a dedicated network interface card, no rules should be assigned to this interface. You can also create bypass rules so the heartbeats aren't inspected.

Installing or uninstalling Deep Security Agent might cause temporary disconnection of the cluster due to binding or unbinding the drivers. Choose a suitable time for this to happen.

6.4 Microsoft Hyper -V

When deploying Deep Security on a Microsoft Hyper-V environment, the Deep Security Agent should be installed within each guest operating system in each virtual machine (VM). This provides the maximum amount of context and security for each guest.

- Recommendation scans can be used to determine the applicable set of intrusion prevention, integrity monitoring and log inspection rules required per guest.

- Anti-malware, web reputation, and firewall policies can also be individually configured per guest using the Deep Security Agent deployment.

If you wish to protect the Parent Partition (also known as the Management Operating System), additional steps are required so that the network traffic is not inspected twice.

It's recommended to choose one of the following options:

- Do not use intrusion prevention in the parent partition.

- Use intrusion prevention in the parent partition, but use the firewall policy assigned to the agent in the parent partition to bypass incoming and outgoing traffic for the IPs of the VMs being hosted.

This can be done with two bypass rules - one for incoming and another for outgoing - that operate on the destination IP range of guests for incoming traffic, and the source IP range of guests for outgoing traffic.

Bypass skips the intrusion prevention rule processing, preventing a duplicate inspection of the traffic in both the parent partition and guest virtual machine.

It is also recommended to use bypass rules, like the second option above, if there is a firewall policy on the parent partition.

6.5 Citrix

Citrix XenDeskop

1. Install a deactivated Deep Security Agent on a Master image.

   Install the agents in the master image (deactivated) and then perform agent-based activation in the provisioning process. Use an Event-Based Task to assign the correct policy based on the attributes available (such as Computer Name).

   Activating an agent using the Deep Security Manager console has its limits. These limits are sometimes due to network topology or firewall constraints that could prevent manager-initiated activation jobs. In some environments (non-persistent), using different scripting techniques (like ScriptLogic, PowerShell or Batch script) are the most common ways to automatically activate and configure Deep Security Agent.

   Alternatively, you can achieve the same result when Deep Security Agent is installed into a Master image and the streaming (PVS) Citrix VDI desktops are running on top of VMware vSphere Environment.

   - If a computer with the same name already exists: Re-activate the existing computer

   - Reactivate cloned agents

   - Reactivate unknown agents

With this setting, the previously activated Deep Security Agent in master image is reactivated the first time that it contacts the Deep Security Manager from the streaming VM on which it was launched. This all occurs in the context of the first heartbeat. It does not require a second heartbeat to complete. The Deep Security Manager does not need to establish a connection to the agent for this function to work (as long as the agent is able to connect to the Deep Security Manager).

Conflict Resolution with Personal vDisk (Persistent)

For Deep Security Agent to work properly with Citrix Personal vDisk (PvD) in XenDesktop, you must add a folder and file rules in the master image to always overwrite PVD content and allow the agent to generate ds_agent.config file.

The folder rule forces the files in the master image C:\programdata\Trend Micro\Deep Security Agent\dsa_core directory to always overwrite PVD content. While this rule works, it will not allow the agent operation to generate ds_agent.config file. Without this file, the agent service will not be able to start properly, so the agent automatic reactivation will fail. Therefore, a file rule must be created as well to allow the ds_agent.config file to be created.

The file rule will be combined with the folder rules during PVD update.

files_rules.txt addition

```
[Rule-Begin]

Type=File-Catalog-Construction

Action=Catalog-Location-Guest-Modifiable

name="%ALLUSERSPROFILE%\Trend Micro\Deep Security Agent\dsa_core\ds_agent.config"

name="%PROGRAMFILES%\Trend Micro\Deep Security Agent\AgentData\dsa_core\ds_agent.config"

[Rule-End]
```

custom_folders_rules.txt addition

```
[Rule-Begin]

Type=Conflict-Resolution

Action=Rebuild-Dst

name="%ALLUSERSPROFILE%\Trend Micro\**\*"

name="%PROGRAMFILES%\Trend Micro\**\*"

[Rule-End]
```

2. Deep Security Agent and the Citrix target device driver

   On Citrix PVS 6.0 Environment, if you are installing (In-Guest) Deep Security Agent, the Citrix Target device driver might not be able to connect successfully to the provisioning server due to a possible conflict.

If you are installing Deep Security Agent on a Windows operating system that is connected to a PVS server using disk provisioning, the temporary workaround is to change the tbimdsa driver loading order during system startup from PNP_TDI to NDIS.

To do so, manually change the loading order of tbimdsa driver used by Deep Security Agent.

a. Go to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\tbimdsa

b. Add or modify String "Group" value to "NDIS"

c. Add or modify DWORD "Start" value to "0"

   By changing the Group from "PNP_TDI" to "NDIS" and Start value from "3" to "0", it allows tbimdsa driver to load after Citrix driver has loaded.

d. Reboot the machine. The PVS Target Device will be able to connect to the vDisk upon boot-up.

Refer to this article: Citrix Target Device driver cannot connect to the provisioning server when in-guest Deep Security Agent (DSA) is installed on Citrix PVS 6.0 environment (https://success.trendmicro.com/solution/1098061) for more details.

Citrix XenApp

1. Citrix XenApp's API Hooks

   Citrix's API hooks can prevent the Deep Security Agent service from starting. To resolve this, the ds_agent.exe must be added into XenApps exclusion list. See How to Disable Citrix API Hooks on a Per-application Basis in the Citrix documentation.

2. Anti-Malware Exclusion for Citrix

   Trend Micro recommends that Citrix files are excluded from scanning by Deep Security. For a more comprehensive list of recommended scan exclusion, refer to the following article: Citrix-recommended exclusions on Deep Security.

6.6 Private, Public & Hybrid Cloud Environments

Amazon Web Services (AWS)

Deep Security Manager can now be connected to Amazon Web Services to provide instance discovery and collect additional information about these instances. This can be used to automate security (for example, assigning a policy based on an Amazon Security Group).

Assign a dedicated account for Deep Security so that you will be able to refine the rights and permissions or revoke the account at any time. It's recommended that you give Deep Security an Access and Secret key with only read-only permissions.

vCloud Environment

Deep Security Manager can now connect to the vCloud director to discover the machines that need to be protected. If this is used with a public cloud, it can help with agent management. If vCloud is used within a private or community cloud where Deep Security Manager is deployed, the vCloud support can work together with the vCenter integration to provide agentless protection to vCloud.

The vCloud director (vCD) workloads are presented in Deep Security in the following hierarchy:

1. vCloud Director Instance
2. Virtual Datacenter
3. vApp
4. Virtual Machine (being the endpoint that can be protected)

This allows the administrator to select virtual machines from certain vDC/vApp's to be protected.

1. Multiple vCD instances can be presented, but make sure the following rules are applied:
   - All vCenters that vCD used for resources are already configured in the administrative side of the portal.
   - Present vCD instances at vCD System object. This will allow all workloads to be discovered in vCD.
2. The following vCloud Director settings must be configured correctly:
   - vCD public URL
   - vCD public REST API base URL (System > Administration > Public Addresses)
3. The vCloud organization accounts that will be used by Deep Security Manager to access vCloud must have the "Administrator View" right. This can be verified by checking the user's role properties in vCloud, then by going to the Rights for this Role > All Rights > General folder.

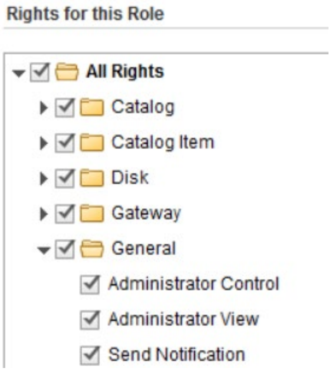

Rights for this Role

- ☑ All Rights
  - ☑ Catalog
  - ☑ Catalog Item
  - ☑ Disk
  - ☑ Gateway
  - ☑ General
    - ☑ Administrator Control
    - ☑ Administrator View
    - ☑ Send Notification

Figure 9: Rights for this Role

4. Consider the following settings when adding the vCloud Director Instance:
   - The name should be descriptive.
   - Enter the address of the vCloud Director instance as follows:

     vcloud.mycompany.com
   - There is no need to add "http" or "https" in the From field of the address.
   - There is no need to add the organization name at the end of the URL.

5. When importing the vCloud resources into Deep Security Manager, the user name must include "@orgName". For example, if the vCloud account's user name is “kevin” and the vCloud Organization you've given the account access to is “CloudOrgOne”, then the Deep Security user must enter “kevin@CloudOrgOne” as their user name.

6. When adding more than one vCloud Director instance, ensure that the corresponding Provider Virtual Datacenter resources have been added to Deep Security Manager. This includes:

    - All vCenter instances used for Provider Virtual Datacenters

    - All vShield Manager instances used for Provider Virtual Datacenters

7. Public Catalog VMs must have the vShield driver installed as part of the template configuration before adding the vApp/VM to the catalog.

8. Configure the vCenter Database to Assign Unique UUIDs to New Virtual Machines. Refer to VM BIOS UUIDs are not unique when virtual machines are deployed from vApp templates (2002506) in the VMware documentation.

    Other Useful References:
    Changing the UUID of vCenter Servers in Deep Security

6.7 IBM Rational ClearCase

If the Deep Security on a Linux has IBM Rational ClearCase installed, it can result in server freeze. Please refer to the IBM's recommendation for running the AV on the server: Support Policy for Anti-Virus and ClearCase.

6.8 Docker support

8.12.1.Supported Docker Platform

Deep Security as a Service supports Docker environment on only the Linux platform, not Windows. Deep Security Agent should be installed into Docker host. Installing Deep Security Agent to Docker container is not supported.

Docker host OS must be one of supported OS by Deep Security Agent, such as:

- Red Hat Enterprise Linux/CentOS
- SUSE Linux Enterprise Server
- Amazon Linux
- Ubuntu
- Debian

Docker optimized host OS, such as below, are NOT supported:

- Atomic Host
- Snappy Ubuntu Core
- Photon OS

However, Docker container can be based on any distributions, such as Alpine Linux, Busybox, and others.

8.12.2.Container Protection

Deep Security Agent provides intrusion prevention, web reputation, and real-time anti-malware features to Docker containers by installing Deep Security Agent to Docker host.

A. Intrusion Prevention

Create **/etc/use_dsa_with_iptables** as below before installing Deep Security Agent to your Docker so Deep Security Agent does not remove iptables configuration, which is required for Docker networking.

```
# touch /etc/use_dsa_with_iptables
```

Note that assigned IPS rules take effect to ALL Docker containers on the same host because Deep Security Agent is working at the host level.

Because iptables is enabled, you should also allow communication ports required for Deep Security Agent in iptables rule, such as 4118/tcp. Please refer to the following solution for required ports.

“Communication ports used by Deep Security”

(https://success.trendmicro.com/solution/1060007)

B. Recommendation Scan

Recommendation scan does not completely work for applications in Docker containers. You should manually choose which IPS rules to assign. Some IPS rules might be recommended even if the application is not vulnerable because Deep Security Agent can only find running processes in containers but not detect the application version.

C. Inter-Container Traffic

Deep Security’s network security features (firewall, intrusion prevention, web reputation) does NOT affect inter-container traffic on a host in either classic link model (by --link option) or networking model.

D. Host Port and Container Port

Because Deep Security Agent works at host, container service port and host service port should be the same so IPS rules can inspect traffic as expected. For example, if you run web server on container port 80/tcp, you should bind it to host port 80/TCP. Otherwise, you should modify port configuration at all application types to add host service port.

Some container orchestration platforms, such as Amazon ECS or Kubernetes, can use dynamic host port by its configuration. You should configure to use static host port same as container service port.

8.12.3.Host Protection

Deep Security Agent provides all security features to Docker host.

A. Firewall

You do not need to enable Deep Security Firewall feature in general, because Docker container only exposes ports which are explicitly configured. If you must use it, there are several points and limitations as shown below.

If you use firewall on Docker host, you should allow BOTH incoming and outgoing traffic for incoming connection to Docker container application. For example, if you are running web server in a Docker container on port 80/TCP, you must allow incoming traffic to Docker host on port 80/TCP AND outgoing traffic to Docker container on port 80/TCP.

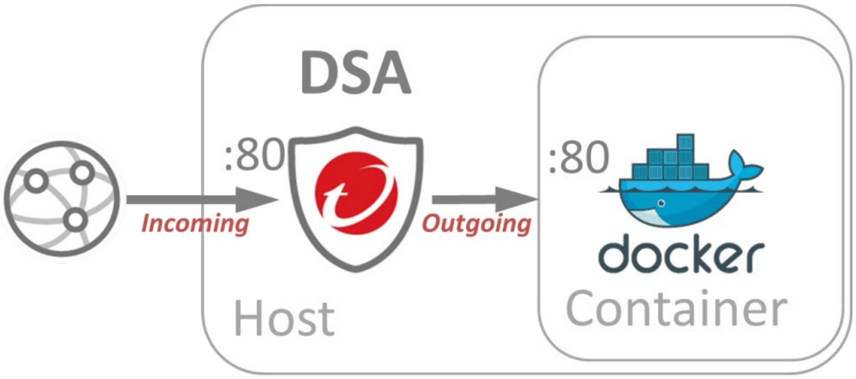

Figure 10: Firewall on Docker Host

You might also need to allow some ports at host to be used by Docker and related orchestration framework. Below are some examples.

a. Docker Remote API port 2375/TCP (HTTP) or 2376/TCP (HTTPS), Docker registry 5000/TCP

b. Amazon ECS Agent 443/TCP, 51678/TCP, 51679/TCP
http://docs.aws.amazon.com/AmazonECS/latest/APIReference/API_PortMapping.html

c. Kubernetes apiserver 8080/tcp(HTTP) or 6443/tcp(HTTPS), etcd 2379/tcp, kubelet 10250/tcp and 10255/tcp

https://kubernetes.io/docs/admin/kube-apiserver/

https://kubernetes.io/docs/admin/kubelet/

d. Docker swarm mode 2377/tcp, 7946/tcp+udp and 4789/tcp+udp
https://docs.docker.com/engine/swarm/swarm-tutorial/#/open-ports-between-the-hosts

You cannot use random (dynamic) host port mapping with Deep Security Firewall because rule cannot be adapted to dynamic ports.

8.12.4.Deployment Scripts

Amazon ECS (EC2 Container Service)

When deploying Deep Security Agent to Amazon ECS cluster instance, you can modify Deployment Script to be put into instance user-data as below.

```
 #!/usr/bin/env bash
echo ECS_CLUSTER=ClusterName >> /etc/ecs/ecs.config
touch /etc/use_dsa_with_iptables
curl https://app.deepsecurity.trendmicro.com:443/software/agent/amzn1/x86_64/ -o
/tmp/agent.rpm -s
rpm -ihv /tmp/agent.rpm
...(snip)...
```

Other reminders:

- You can put your *ClusterName* for the instance to join to ECS cluster.
- Create /etc/use_dsa_with_iptables not to disable iptables.
- Use curl instead of wget because wget is not installed by default in Amazon ECS-optimized AMI.

6.9 Automation Activation from Gold Image

See Bake the agent into your AMI or WorkSpace bundle in the Deep Security Help Center.