My learnings on Microsoft Defender for Endpoint and Exclusions

By Christopher Brumm

Published: 2021-08-07 · Archived: 2026-04-05 15:52:22 UTC

Press enter or click to view image in full size

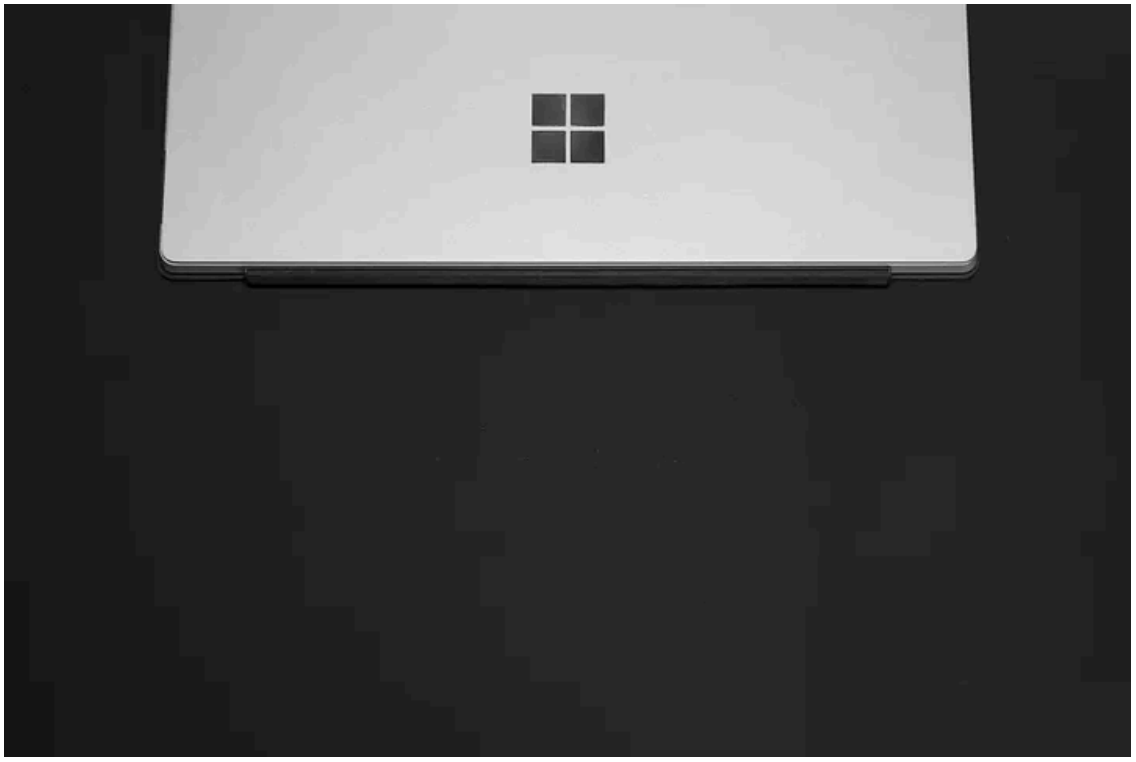

Photo by Ashkan Forouzani on Unsplash

Whenever I've had to deal with AV solutions in recent years, the topic of exclusions has always come up at some point. Usually, it was always quickly agreed that the best way to deal with exclusions is not to use them 😉. Nevertheless, I have only come across a few companies that did not have exclusions and it is always assumed to be a project risk not to migrate them when it comes to switching to a new solution.

> To be clear, my recommendation is to use every opportunity to get rid of (old) exceptions and not migrate anything that has not been proven to cause problems with the new solution.

However, it is always good to know your enemy and have a plan in case you are forced to use exclusions. So in this blog I will try to show what kinds of exclusions there are, what risks they entail and how to implement them in a practical way.

What is an Exclusion and why should I care?

Virus scanners such as Microsoft Defender AV (MDAV) have the job of detecting malware and neutralizing it. Due to the spread of Windows, this has been an ongoing issue since the 90s. While in the beginning there was a strong focus on signatures / patterns of known malware, this topic has recently become less and less important, because the viruses have been mutating for some time (like real viruses) and do so much faster than the vendors can update their signatures. Modern solutions such as Microsoft Defender for Endpoint (of which Defender AV is a part) have a wide range of detection methods in addition to signature detection and rely on machine learning and behavior monitoring methods for detection, among other things.

Malware detection by MDAV can be performed through various mechanisms. In addition to the various scheduled or on-demand scans (Quick, Full, Custom), real-time protection is also active. Real-Time Protection reviews files when they are opened and closed, and whenever a user navigates to a folder.

Like any other AV solution, there is of course the possibility to create (classic) exclusions in MDAV for files, folders, processes, and process-opened files. These are stored in the registry on the endpoint and the exclusions from a GPO can be easily displayed via Powershell.

An exclusion prevents the corresponding files or processes from being detected as malware by Defender during the scan and by Real-Time Protection, and countermeasures (such as a quarantine) from being initiated.

Since this behavior is rather bad in the first place, the question arises: Why do exclusions exist at all?

Unfortunately, many manufacturers have a list of exclusions that are "*necessary*" or "*recommended*" for the software to work. Since every admin can remember a situation where the AV agent was a bit overambitious, this is a very difficult discussion. Here are a few examples of required exclusions: SCCM, VEEAM, Exchange, Kaseya, MS SQL — Microsoft has even a list of the exclusion lists 😂

Okay there seems to be reasons — What specifically is the problem with Exclusions?

Every exclusion weakens our defense and every weakness that an attacker knows, can guess or read offers him a good opportunity. For example, if I know that the software distribution paths are excluded on all systems, this is a good place to put the tools for the next steps of my attack and run them from there.

> Our task is to challenge the necessity of any exclusion on any system and thus reduce it.

Which types of exclusions are available in Microsoft Defender for Endpoint?

Before we start looking at the different types of exclusions, it is important to understand that MDAV is part of Microsoft Defender for Endpoint (MDE) but can also be used on its own. In terms of exclusions, this means that there are methods that are limited to MDAV and methods that cover the entire suite.

> I strongly recommend using MDE, as this massively increases the protection and enables various of the features described here. A virus scanner alone — even if it is very good — is unfortunately no longer sufficient today.

Exclusions in Defender Antivirus

In MDAV there are the following types of exclusions:

- exclusions based on file name, extension and folder location
- exclusions for files opened by processes

These exclusions can be managed in several ways. Besides the tools described later, such as Intune, there are the following local options:

- Creation of a local policy
- Using the Windows Security GUI
- Powershell with the CMDlet Add-MpPreference

The easiest way to display all exclusions is the CMDlet Get-MpPreference (but for reading the exclusions you need to be local admin).

Press enter or click to view image in full size

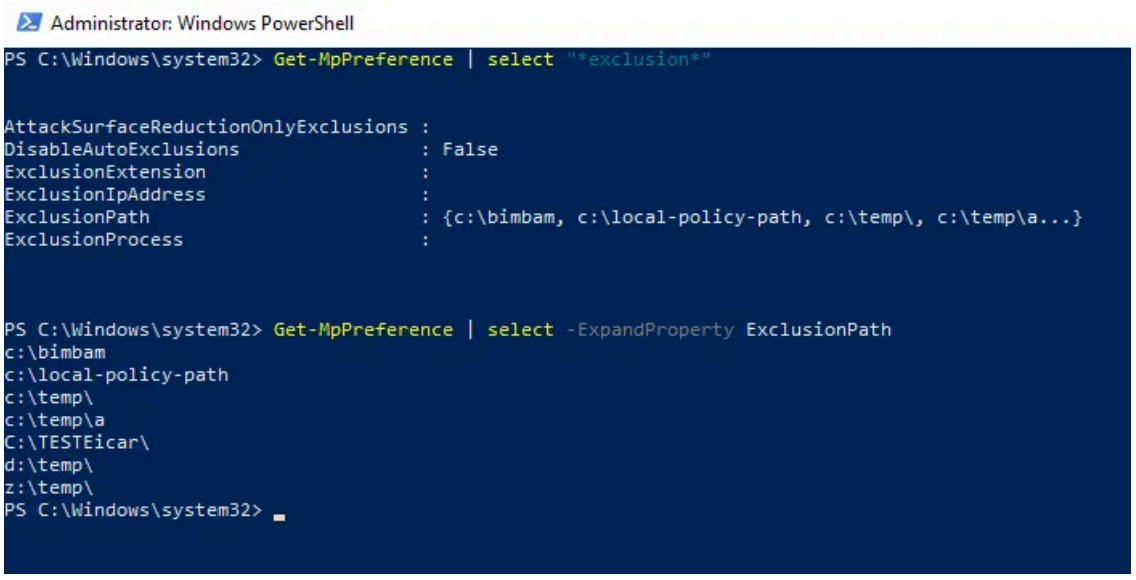

Since these exclusions are also stored in the registry on the endpoint they also can be displayed by reading the corresponding keys via Powershell:

> Interestingly, the key in the local hive can no longer be read on a Windows 11 system. I assume this is a hardening measure by Microsoft.

If you're using **process exclusions** these points are remarkable in my opinion:

- When you add a process to the process exclusion list, Microsoft Defender Antivirus won't scan files opened by that process, no matter where the files are located. The process itself, however, will be scanned unless it has also been added to the file exclusion list.
- The exclusions only apply to always-on real-time protection and monitoring. They don't apply to scheduled or on-demand scans.

Auto Exclusions in Defender Antivirus

In addition to the exceptions configured by the admin, Auto Exclusions still come into play for (2016/2019) servers ***depending on the role of the server*** if not disabled. For a domain controller, for example, exceptions are active for the NTDS database, the transaction log files, the NTDS working folder and support files.

Some notes on this (useful feature):

- These exclusions are not displayed in the above lists
- They apply only to Real-Time Protection — not to scheduled or on-demand scans
- The auto exclusion feature works only for the default installation location of the server roles
- more info here:
  Configure Microsoft Defender Antivirus exclusions on Windows Server | Microsoft Docs

Exclusions in other parts of MDE

Besides AV there are several other components that can prevent the execution of files and functions in files. They all have in common that whitelisting by a custom indicator is possible.

- endpoint detection and response (EDR)
- attack surface reduction (ASR) rules — see this great blog about ASR.
- controlled folder access

Custom indicators

Microsoft Defender for Endpoint provides centralized management of Indicators of Compromise (IoCs) in the Custom Indicators section. IoCs are actually intended to detect known malicious patterns and have them blocked, for example, by security products such as MDE.

In addition to the Alert and Alert+Block actions, the Custom Indicators section also includes the Allow action, which can be used for whitelisting. The whitelisting of files is not done by a path or filename but by hashes. Besides files it is also possible to create entries for IPs & URLs and certificates.

The file hashes can be created in MD5, SHA-1 or SHA-256. Although each of these algorithms is significantly more secure than a file or folder name, the SHA-256 hash should be used because a collision is significantly less likely due to the length of the hash (32 bits). If you use the GUI, you will also get a warning when entering MD5 hashes:

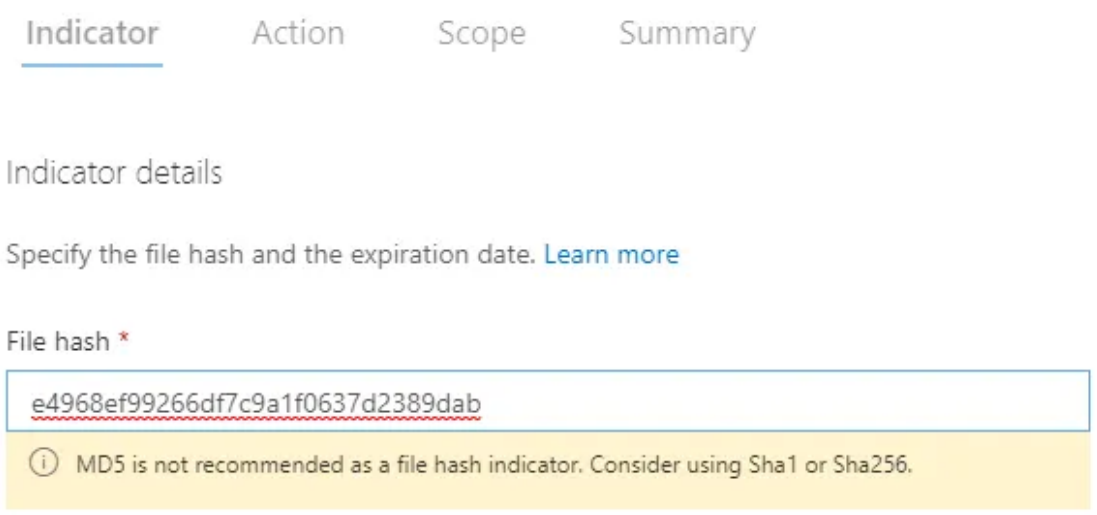

So far, I have not been able to read out the custom indicators that are effective on a system — this then reduces the probability of an intentional collision to almost 0.

The creation of a Custom Indicator can be done with the *Active Remediation Actions permission* via the GUI or via API. By using the API, this process can also be embedded well into a process with e.g. tickets, documentation and releases.

What are the best tools to manage exclusions?

The selection of the right tool(s) is strongly dependent on the circumstances and has a strategic component. Means: If I have today a tool for the administration of the configurations of my clients I will try on the one hand to administer also the Exclusions with it. On the other hand, I cannot manage all exclusions with all tools (equally well) and this can be the reason for a (perhaps already overdue) change of strategy.

- In the **MDE portal**, only custom indicators, i.e. hashes, can be excluded — and only there. Since the Custom Indicators are the (almost always) preferred whitelisting variant, this portal is set in any case.
- **Intune** can set all exclusions except hashes / custom indicators. Unfortunately, Intune does not seem to be able to combine lists of exclusions, which can be a challenge in large heterogeneous environments.
- **GPOs** can also be used to configure all exclusions except hashes / custom indicators — GPOs can even combine multiple lists here. For clients, however, it should be kept in mind that the distribution of computer policies via VPN is not ideal — but most of those who do this today probably already know that ;-)
- With **SCCM**, the AV exclusions can be managed well. With ASR Exclusions, however, there is currently still the restriction that no wildcards are supported — which severely limits usability.
- Last but not least, all **tools and scripts** that are able to manipulate registry values can also be used to potentially manage everything except custom indicators.

What is the best way to deal with exclusions?

As you can see, there are several ways to whitelist and block files and this creates some challenges, especially in larger environments. One strategy for dealing with file exclusions might look like this:

#1: Avoid exclusions

That means we don't want exclusions. We do not want to migrate new exclusions and we do not want to migrate existing ones. Every exclusion must be well justified.

#2: Use Custom Indicators (hashes) whenever possible!

As described above, CIs have several advantages:

- They are much safer, as collisions are very unlikely.
- The creation can be easily integrated into processes via runbooks.
- The administration is done in the Security Center, not in Config Management.

#3: Use the right (classic) exclusion type and avoid common mistakes

If hashes are not possible, proceed in the following order for Classic Exclusions:

- Processes
- (complete) Pathes
- Extensions

> Using this exclusions is something you want to avoid because they can be abused. It is not that hard to guess wich folders are excluded and it is eays to check.

Before you create an exclusion you should look at these two sources:

- Recommendations for defining exclusions
- Common mistakes to avoid when defining exclusions

In addition to lists of locations that must not be excluded, other typical errors such as the use of filenames without folders in the exclusions or the correct use of environment variables are also covered there.

#4: Make yourself familiar with the scoping of exclusions

I think Custom Indicators don't need to be scoped, as the risk is very manageable and the effort is disproportionate. (The way to do this would be Device Groups in MDE).

Get Christopher Brumm's stories in your inbox

Join Medium for free to get updates from this writer.

Remember me for faster sign in

With Classic Exclusions it is advisable to differentiate a little and find a good compromise. This looks for servers usually clearly different than for clients, since clients are clearly more homogeneous. While servers often have a handling per server type in the GPO anyway, clients will usually try to have a very uniform policy.

Depending on the structure, however, it can still make sense to maintain multiple lists and combine them as needed to reduce administration efforts. The practical example would be if only a certain user group uses a certain software and you then use a combination of two lists for these clients (GPOs and SCCM can do this). It is worth to plan a little bit to find a good compromise.

#5: Monitor and Review your Exlusions

In any environment exclusions accumulate over time and it very important to deal with these two issues:

1. How can I prevent and check if there have been (illegitimate) changes to the Exclusions?
2. How and when do I check my inventory and how can I reduce it?

For question number 1 have a look in the Monitoring section below. For Question 2 you will need a process.

How do the different mechanisms interact?

What happens if I use more than one configuration method?

By default, local changes will be merged with the lists by Group Policy, Configuration Manager, or Intune. The Group Policy lists take precedence when there are conflicts. You can configure how locally and globally defined exclusions lists are merged to allow local changes to override managed deployment settings. (Source)

I think disabling merging in general is a very useful thing. The corresponding setting can be enabled via GPO or via the Defender CSP with a custom policy in Intune. The RegKey that is set in this way now ensures that locally added exclusions are overwritten.

What happens at a conflict?

The MS documentation answers this question so:

Cert and File IoC policy handling conflict will follow the below order:

*If the file is not allowed by Windows Defender Application Control and AppLocker enforce mode policy/policies, then **Block***
*Else if the file is allowed by the Microsoft Defender Antivirus exclusion, then **Allow***
*Else if the file is blocked or warned by a block or warn file IoC, then **Block/Warn***
*Else if the file is allowed by an allow file IoC policy, then **Allow***
*Else if the file is blocked by ASR rules, CFA, AV, SmartScreen, then **Block***
*Else **Allow** (passes Windows Defender Application Control & AppLocker policy, no IoC rules apply to it)*

To simplify this, my friend Fabian has created this wonderful meme! 😍

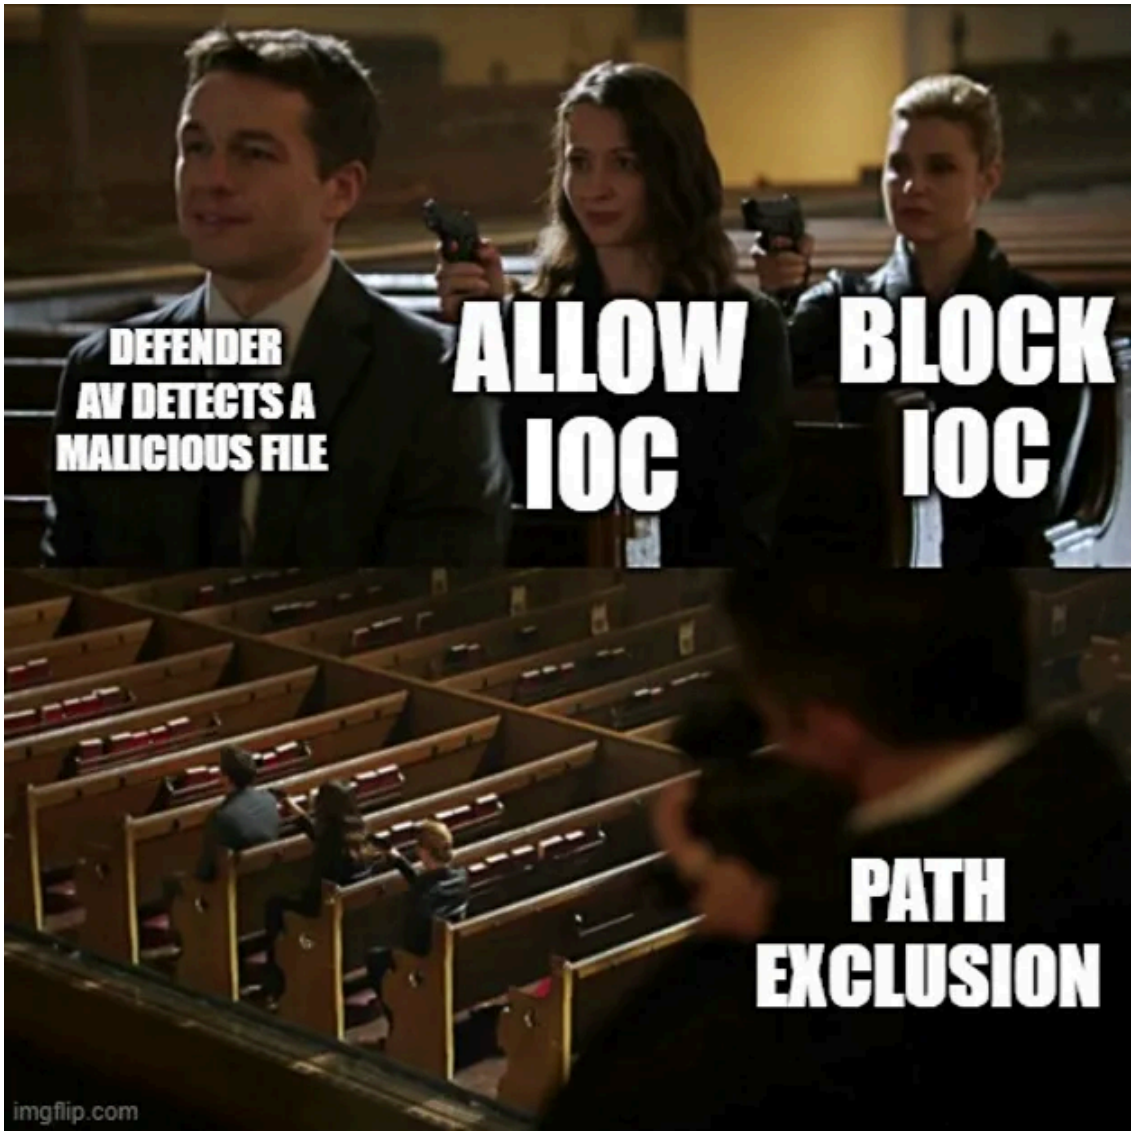

How can I prevent and check if there have been (illegitimate) changes to the Exclusions?

General Manipulation Prevention

Before we look at the AV exclusion use case, I recommend everyone check out this blog to learn what can be manipulated in general in Defender AV and how to best deal with it. It is reasonable to try to prevent manipulation of the configuration and although Tamper Protection has no influence on AV Exclusions, its use is very useful as it can prevent many basic manipulations of the Defender.

Finally, Configure Local overrides for Microsoft Defender AV settings can be used to prevent local exclusions from being generated by simply disabling the merging of the lists. This means that only the exclusions from the GPOs apply.

Detect and prevent local exclusions

The behavior described above for combining exclusion lists is very useful for scoping. However, this behavior can also be used by an attacker. By default it is possible to create own local exclusions with local administration rights and in my experience there are also some accounts in many companies that have e.g. the rights to edit a GPO that affects the clients. Reason enough to have a look how we can be informed about changes.

The first approach is to look at the MDE client via a custom detection. Alex Verboon has already published something good about this.

Thus, according to my tests, the following scenarios can be identified:

- Creation of a local policy
- Using the Windows Security GUI
- Powershell with the CMDlet Add-MpPreference

However, besides the methods provided for this purpose, I can also try to put my exclusions **directly into the registry as keys**:

Theoretically this can be done in several ways with different permissions (Administrator or SYSTEM) in the hives *HKLM\Software\Microsoft\Windows Defender\…* and *HKLM\Software\Policies\Microsoft\Windows Defender\…*

In my tests I was not able to change keys in the non-policy path either as admin or as SYSTEM although SYSTEM is owner of the hive. The reason for this I suspect is that Defender AV prevents changes to this part of the registry by a kernel-mode driver.

In the hive under Policies it is possible to create further exclusions as a local admin which (after some time or a gpupdate) will be applied. Unfortunately, no logs are generated for any changes made this way or by GPO.

> This works (after a reboot) even on a non domain-joined Windows 10 device if you create the hive under policy and add the keys. This is a really serious problem from my point of view and should be fixed by MS. If you are interested in how exactly this vulnerability can be used check out this (awesome) blog by Fabian Bader!

Detect and prevent exclusions from configuration systems

Since there are many ways to configure as described above, this is not a complete list but I limit myself to the tools that are accessible to me.

Add Custom Indicators in MDE

As described above you need the *Active Remediation Actions permission* to manipulate the CIs. Each CI is directly recognizable when and by whom it was created.

Unfortunately, there are no events for the creation of CIs in the audit log yet, but since there are already filters for them, I am optimistic that it will soon be possible to specifically search for or alert on them.

Press enter or click to view image in full size

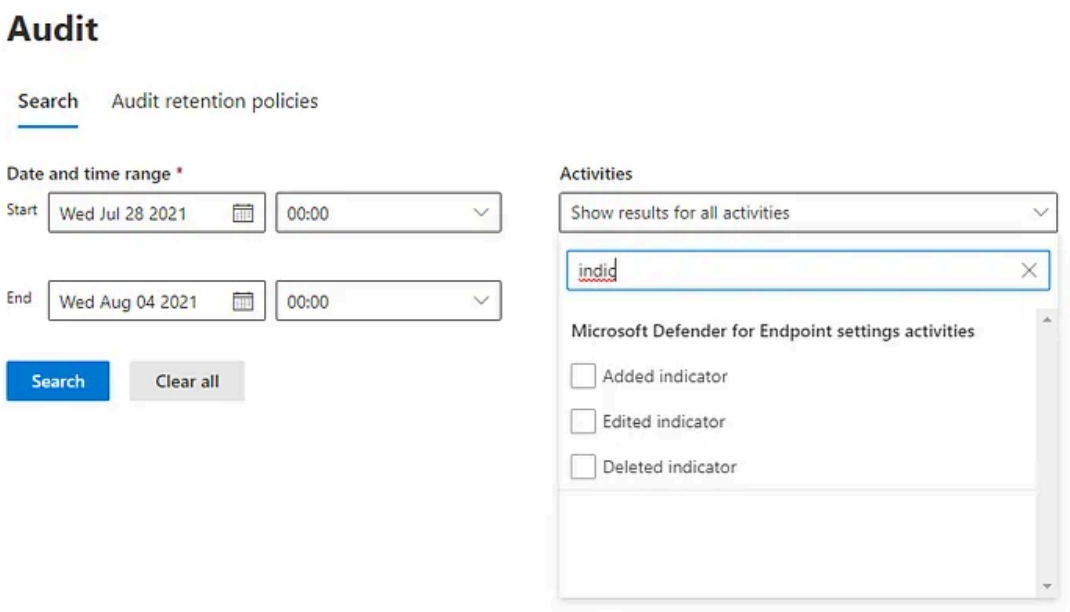

Add Exclusions in Intune

Of course, it is also possible to add further Exclusions with Intune. These can be created by (at least) the following roles:

- **In Intune:** *Endpoint Security Manager and Custom Roles*
- **In AAD:** *Global Admin, Intune Admin, Security Admin*

These changes are included in the Intune Audit Log and can be queried alerted accordingly if needed.

Add Exclusions via GPO

Finally, there is the possibility to add exclusions via GPO. By the above described behavior for merging the lists it is possible to add additional exclusions with each existing or new GPO that acts on the system.

Unfortunately, neither the modification of the GPOs nor the configured systems have an event that indicates that an exclusion has been added. You will only see that a GPO has been modified or applied. To determine this change, only a regular export and comparison would help at the moment.

To go into a little more detail, a look at Advanced Hunting:

There are some entries in the DeviceRegistry events table from the *HKLM\Software\Policies\Microsoft* policy hive that even contain changes to the *Windows Defender* keys, but nothing from the exclusions hive.

Conclusion

In this blog I have tried to summarize my learnings around the topic of Exclusions in Defender for Endpoint. We first looked at what types of exclusions there are and how they can be managed, including a strategy of which ones to use. Then we dived a bit deeper into the mechanics of exclusions and what happens when conflicts occur and what scoping options are available. The last part then dealt with the risks, hardening and monitoring.

I am afraid that this topic will continue to be relevant and hope that I have overlooked something, especially in the areas of hardening and monitoring. It is quite obvious that an attacker with admin rights can only be stopped with difficulty, but the fact that he can create exclusions unseen and thus create a space in which he can reload all the necessary tools makes the situation even more difficult.

For me, the critical part was formulating a realistic strategy for dealing with exclusions:

> #1: Avoid exclusions
>
> #2: Use Custom Indicators (hashes) whenever possible!
>
> #3: Use the right (classic) exclusion type and avoid common mistakes
>
> #4: Make yourself familiar with the scoping of exclusions
>
> #5: Monitor and Review your Exlusions

What's your strategy for dealing with exlusions? Contact me: https://twitter.com/cbrhh

Acknowledgements

Thanks to Fabian Bader and Nadine Kern for reviewing and checking all that stuff. 💖 And don't forget to read Fabians Blog !

More Sources:
