Technical verification report Microsoft Defender 365 for Enterprise and for Education

Version 1.0

Date 22 May 2023

Colophon

Technical verification report by	Ministry of Justice and Security, Strategic Vendor Management Microsoft, Google and AWS (SLM Rijk)
	Turfmarkt 147 2511 DP The Hague PO Box 20301 2500 EH The Hague www.rijksoverheid.nl/jenv
Contact	Henrique Barnard E h.m.barnard@minjenv.nl T 070 370 79 11
Project Name	Technical verification report Microsoft Defender 365
Authors	Privacy Company Sjoera Nas and Floor Terra, senior advisors https://www.privacycompany.eu/

Change log

Version	Date	Summary of input
0.1	4 November 2021	First Privacy Company internal rough draft
0.2	16 June 2022	Completed draft, with tech input after DSAR and audit log results
0.3	1 May 2023	Modified draft after input Microsoft and public clarifications - with track changes
0.4	22 May 2023	Interim version with more input Microsoft with track changes
0.5	22 May 2023	Clean version

CONTENTS

Figures

Tables

Summary

This report, commissioned by SLM Rijk, contains the results of a technical verification of Microsoft's data processing for security purposes through the use of Microsoft 365 Defender. This suite bundles different security services to protect devices, users, e-mail and login. The conclusions apply to both the Enterprise and the Education versions of Microsoft 365.

Nine verification questions

This report is based on nine verification questions.

1. Does Defender adequately detect critical threats and security breaches from devices and cloud apps, and with regard to malicious emails, files and URLs that are exchanged via e-mail and via OneDrive and Teams, both within the tenant and with external individuals?

2. Does Defender provide adequate warnings about suspicious logins via the Azure Active Directory?

3. How does the data anonymisation option work in Defender for Cloud Apps?

4. Does Microsoft send traffic to third parties (including through cookies, and to itself as an independent data controller) when a system administrator enables Defender? If so, are those third parties mentioned on the list of subprocessors?

5. Does Microsoft process learnings from security incidents across its Enterprise and Education user base, or are there limitations, such as prior anonymisation?

6. Does Defender create individual risk profiles and/or individual scores in the different analytic overviews and reports?

7. Does Microsoft publish adequate documentation on the personal data it collects through the tested applications, in comparison with captured network traffic and logs that are accessible for system administrators?

8. Does Microsoft give system administrators full access to all personal data it processes through the different Defender tools? Does Microsoft provide adequate explanations if it does not provide access to certain personal data?

9. Are there high risks resulting from the transfer of personal data to the USA or other third countries?

Outcomes

The main outcome of these tests is that Defender 365 does what it is supposed to do: inventory and protect devices, and detect malicious files, URLs, bulk mail and suspicious logins. Defender 365 does not recognise all possible malicious activity, for example old spyware, and gives the benefit of the doubt to 'known' sender domains. Following this logic, Defender assumes a high trust level for end users authenticated

with Microsoft's own Authenticator app. This makes Defender 365 less effective in determining the risks of 'known' and 'trusted' senders. This does not mean Defender is not capable of achieving the purposes for which it processes personal data: securing the work environment. The lack of alerts during some tests is a logical consequence of the limited test scope while Microsoft bases risk assessments on multiple criteria. These criteria could not all be replicated in the tests. For example: replicating a worldwide phishing campaign on multiple organisations was outside of the scope of the tests.

It follows from the tests that Microsoft does not automatically send personal data to external third parties, except for limited traffic to Twitter. Microsoft also qualifies itself as a third party (independent data controller) for the processing of personal and content data received through Feedback forms. It is not in line with the contractual privacy guarantees for Microsoft's Online Services that Microsoft decides itself when it is opportune to share any personal data with itself in a role as independent data controller, without any technical measure to centrally prevent this data flow.

Reply and measures Microsoft April 2023

Following a dialogue with SLM Rijk in April 2023, Microsoft clarified the purposes and scope of its data processing and transfers to the USA. Microsoft also improved its public documentation about the data processing for threat analysis. Though Microsoft previously used the words 'Insights' and 'Machine Learning', the use of these terms did not imply any personal data processing for global analytics. Microsoft clarified that it only transposes information about new threats to a list of known threats, without any customer information or personal data. Microsoft also acknowledged that device identifiers are personal data, and agreed to update its public documentation. Microsoft has committed to develop two options to block outgoing data traffic in the Defender admin console to Twitter and to centrally block the use of Feedback in Enterprise and Education tenants, or become a data processor for the Feedback data.

Remedies

In order to remedy the remaining issues with regard to the contractual privacy commitments, Microsoft and the government organisations / universities are advised to take the following remedies.

Question	Remedies Microsoft	Remedies organisations
Inventory and protect devices	No recommendations	Inform users about the data processing when they sign for receipt for their company-managed device.
Detect files, mails, apps and URLs	No recommendations	Avoid automatic deletion of suspicious mails.
		Allow employees to ask admins perform to perform a manual review on documents quarantined as malware.
		Allow the end user access to mails qualified as spam.

Detect suspicious logins	No recommendations	Instruct the admins to actively monitor for alerts on users at risk and to quickly follow up to make sure the detections are effective and the consequences of incorrect detections are minimised.
		Create a monitoring policy to restrict how admins are allowed to use the monitoring results, and inform the users about this personal data processing and the limits on its use.
		Consider using pseudonyms for employees whose identity should remain confidential
Data minimisation option Cloud Apps	Microsoft should use the term 'pseudonymisation' instead of 'anonymization'.	Use the pseudonymisation functionality to prevent unauthorised access to sensitive characteristics of end users that can be derived from their app usage.
		When using the temporary user data enrichment option to detect individual shadow IT usage, organisations must develop and communicate clear and knowable rules about the circumstances when these data can be accessed and for what specific purposes.
Traffic to third parties	Honour the two commitments to develop a setting to block the functionality of the **Twitter** feed and to offer a group policy or setting to centrally block the use of **Feedback** in Enterprise and Education tenants or become a data processor.	Use the central opt-out functionality for Twitter and Feedback.
Data reuse across userbase	No recommendations	No recommendations
Profiling	No recommendations	Provide a concise, intelligible and easily accessible internal explanation to employees about the data processing in Defender.

Docu-mentation	Microsoft must provide more information about the observed Telemetry Data from the admin portal, unless Microsoft is able to ensure that the browser telemetry data do not contain any identifying (pseudonymous) personal data.	No recommendations
Data Subject Access	Speed up the process of providing access.	No recommendations
	Stop collecting identifying data via the browser telemetry data or provide access to all personal data processed by Defender in a concise, transparent, intelligible and easily accessible form.	
	If Microsoft wants to rely on an exception to the access rights from Art. 23 (1) under i of the GDPR jo. Art. 41(1) sub i of the UAVG, Microsoft should explain the necessity in detail.	
Transfer risks	Complete the EU Data Boundary as soon as possible for all personal data, keep SLM Rijk up to date about the progress.	Accept the risks of transfer of pseudonymised server generate server logs and telemetry data to the USA until the end of 2023.
		Warn admins not to upload personal data in attachments with support tickets (transfer to the USA possible until the end of 2024).
		Consider using pseudonyms for specific employees working with secret data, and for system administrators.

1. Introduction

This report describes the results of a technical verification started in May 2022 of Microsoft's compliance with the privacy amendment negotiated in June 2019 by SLM Rijk (Strategic Vendor Management for Microsoft, Google and AWS) for specific online services that are part of Microsoft 365. In 2020, SURF, the Dutch umbrella ICT organisation for institutions of higher education, negotiated a similar privacy amendment. SLM Rijk commissioned Privacy Company to check Microsoft's compliance with the privacy guarantees for the online services for both Enterprise and Education version.

This report is a technical verification report. In reply to the findings of this report, Microsoft has provided clarifications and committed to develop data minimisation tools. As this technical verification report did not identify serious shortcomings, the results of this report will not be presented in a full DPIA report.

1.1 Previous DPIAs on Microsoft Office 365

The cluster reports build on the (repeat) Data protection impact assessment (DPIA) on Office 365 for the Web and mobile Office apps published in July 2020.[1] The DPIA concentrated on the seven most frequently used Office applications, namely: Word, PowerPoint, Excel, Outlook, Teams, SharePoint and OneDrive, in combination with the Azure AD.

The DPIA concluded there were 6 high data protection risks for end-users of the software related to the processing of data about their use of the software and services (so-called Diagnostic Data). Between April and June 2020, SLM Microsoft Rijk and Microsoft agreed on measures to mitigate the six high data protection risks, as described in Section 17.3 of the DPIA. Some of the mitigating measures could be implemented fast, other mitigating measures required more time.

Privacy Company conducted a first check on the status quo of the implementation measures in August 2020, and found Microsoft complied with the 4 mitigating measures it had agreed to implement per 1 August 2020. In a second check, performed in April 2021, Privacy Company verified Microsoft's compliance with the remaining improvement measures, also with regard to some low risks found in the separate DPIA on Microsoft Intune.[2]

1.2 Scope of this report: Microsoft 365 Defender

This report focusses on the data processing resulting from the use of Microsoft 365 Defender, hereinafter referred to as 'Defender'. This suite bundles different security services to protect devices, users, e-mail and login. Four services are in scope:

1. Microsoft Defender for Endpoint (MDE)
2. Microsoft Defender for Office 365
3. Microsoft Defender for Identity
4. Microsoft Defender for Cloud Apps3

---

[1] Data protection impact assessment Office 365 for the Web and mobile Office apps, 30 June 2020, URL: https://www.rijksoverheid.nl/binaries/rijksoverheid/documenten/rapporten/2020/06/30/data-protection-impact-assessment-office-365-for-the-web-and-mobile-office-apps/DPIA+Office+for+the+Web+and+mobile+Office+apps+30+June+2020.pdf

[2] Data protection impact assessment Intune, 30 June 2020, URL: https://www.rijksoverheid.nl/binaries/rijksoverheid/documenten/rapporten/2020/06/30/data-protection-impact-assessment-intune/DPIA+Intune+30+June+2020.pdf

[3] Defender includes two other services that Privacy Company has not investigated, Defender for Cloud and Microsoft Sentinel. See: Microsoft Defender for Endpoint, 9 April 2023, URL: https://learn.microsoft.com/en-gb/microsoft-365/security/defender-endpoint/microsoft-defender-endpoint?view=o365-worldwide.

Figure 1:Illustration Microsoft tools in Microsoft 365 Defender[4]

Some Defender services had other names until 2022, namely Exchange Online Protection (EOP) and Office 365 Advanced Threat Protection (ATP). These features are now known as Defender for Office 365. The Defender for Cloud Apps service was previously called Microsoft Cloud Apps Security (MCAS). Only the name of Defender for Endpoint was unchanged.

The 15 most relevant Defender services for this DPIA are listed below in Table 1. Many of the services for Office and Identity can also be used to enrich reports on Endpoint and Cloud Apps; there is no longer a hard distinction between the four Defender services. The four services marked with an asterisk are not separately discussed in this report, as it involves the detection of individual malicious files or activities by system administrators. Hence the customer is in control of the scope of the data processing.

Table 1: Overview of relevant data processing services in Defender

Defender	Service	Description and functionality
Endpoint	Device Inventory	*Lists devices on the network, and whether they are protected with MDE*
	Threat and Vulnerability Management	*Dashboard showing potentially malicious files and device behaviours*
Office and Identity	Users at risk	*List of potentially compromised users, both by behaviour and logins from suspicious locations.*
	Incidents & alerts	*Automatic alerts about cyber-attacks on the tenant or devices (in MDE)*
	Threat analytics	*New unknown viruses and malware in devices (MDE) and mailboxes*
	Explorer	*Suspicious emails, e.g. attachments with malware or content with phishing campaigns. Also includes tab "Top targeted users," the end users who most often receive such emails.*
	Review	*Overview of actions for system administrators, for example a* manual *review if a suspicious mail or file has been quarantined, or what to do with users blocked by*

[4] Screenshot from Microsoft instruction video about Defender for Endpoint, URL: https://www.microsoft.com/en-us/videoplayer/embed/RE4wDob?postJsllMsg=true.

		Microsoft for sending too many messages classified as bulk mail.
	Policies & rules	F*our sets of settings in Defender to help administrators manage threats: (i) Threat policies, (ii) Alert policies, (iii) Manage advanced alerts, and (iv) Activity alerts.*
	Audit	*Access to audit logs that Microsoft makes available to administrators.*
	Hunting*	*Specific security-related searches targeting users, devices (in MDE), specific alerts or configurations, e.g. look up user logins from specific countries, detect devices with outdated/vulnerable hardware, find malicious email attachments or malicious email senders.*
	Actions & submissions*	A*ctions such as quarantining a malicious file, isolating a device, soft deletion of an email, initiating an antivirus scan or blocking URLs*
	Exchange message trace*	*Search the inboxes and logs of all end users in Exchange Online, for example, see all recent outgoing mail from a specific user after an incident. Or, check all incoming mail from a sender who sent malware once.*
Cloud Apps	Cloud App Discovery dashboard and Cloud app catalogue	*Cloud App Discovery dashboard and Cloud app catalogue Overview of used apps with security score and incidents.*
	Activity Log*	*Audit log of actions with apps such as login, file discard, search, attempted access by unauthorized app.*
	Files	*Overview of files flagged as malware, with information about the owner and the app in which the file was detected.*

Out of scope

This report does not examine the processing of personal data through telemetry data from applications installed on end user devices, since this report focusses on the use of Microsoft's security cloud services. However, Section 5.2.2 does address the Telemetry Data collected by Microsoft from the browser of admins when they visit the Defender portal.

1.3 Test methodology

This report is based on two separate inspections. In November 2020 Privacy Company assessed the data protection risks of some functionality of Cloud Apps and Defender for Endpoint for a government organisation, and processed the findings in a non-public report. In 2022 Privacy Company tested the data processing via the Defender for Office 365 tools for SLM Rijk, at the time known as Exchange Online Protection and Advanced Threat Protection, including Identity Protection through the Azure AD.

In 2020, the data processing was tested by performing scripted scenarios on workstations equipped with Windows 10 Enterprise by the government organisation that commissioned these first two reports. Two sets of test scenarios were performed on 5 November 2020. During the test scenarios a normal working day was simulated,

with opening and closing of documents and sending and receiving emails, etc. Selected malware was included in several ways to trigger responses by Defender for Endpoint. This lead to a large number of reported incidents. The generated data were sent to Microsoft's Defender cloud servers in Azure. When MDE discovered a possible security incident an investigation package was sent to the central service and was made available for inspection to the system administrators. The testing of Cloud Apps Security was limited to identity management and cloud apps discovery. The test scenarios included logging-in from an unusual country. Two data subject requests were subsequently filed with Microsoft to obtain any information it may have, that was not included in the available data in the audit logs and on the dashboard. Microsoft responded in December 2020 that it had not found any further information based on the identifiers provided by Privacy Company, only the given identifiers.

The inspections in 2022 were performed in a browser, as Defender for Office 365 is a cloud-based service. Privacy Company created a VM with a Dutch government's Enterprise license in a test tenant. The VM was equipped with Windows 10 Pro version 21H1, build no. 19043.1620 and the default browser Edge (version 100.0.1185.36). The browser versions of OneDrive, Teams and Outlook services were used to store and exchange files.

To answer the fourth question (about traffic to third parties), a dedicated test script was used with the Chrome browser 99.0.4844.84 and MiTM proxy version 6.0.2 to test for the existence of third-party traffic in the Admin console.

Tests were performed between 31 March 2022 and 1 April 2022. Microsoft's public documentation and information on the restricted access Admin Console was last reviewed in April and May 2023.

To minimise noise in the captured data traffic, the recommended privacy minimisation settings were followed in Office 365 (Additional Optional Connected Experiences disabled).

1.4 About Microsoft 365 Defender

This report assesses the four major security tools included in Defender. These tools are briefly described in four paragraphs below. Essentially, Defender sends the information it collects about events to the Azure cloud servers dedicated to, and within the tenant of the customer. Microsoft also transfers some sample threat data to itself for advanced analysis, to be able to add newly discovered threats to the recognition patterns.

Microsoft describes the functionalities of Defender as follows:

> "*The **Microsoft 365 Defender portal** (https://security.microsoft.com) combines protection, detection, investigation, and response to email, collaboration, identity, device, and app threats, in a central place.*"[5]

[5] Microsoft, Microsoft 365 Defender, 7 April 2022, URL: https://docs.microsoft.com/en-us/microsoft-365/security/defender/microsoft-365-defender?view=o365-worldwide

Figure 2: Microsoft graphic Core Defender services[6]

Figure 3: Microsoft illustration scale of Defender[7]

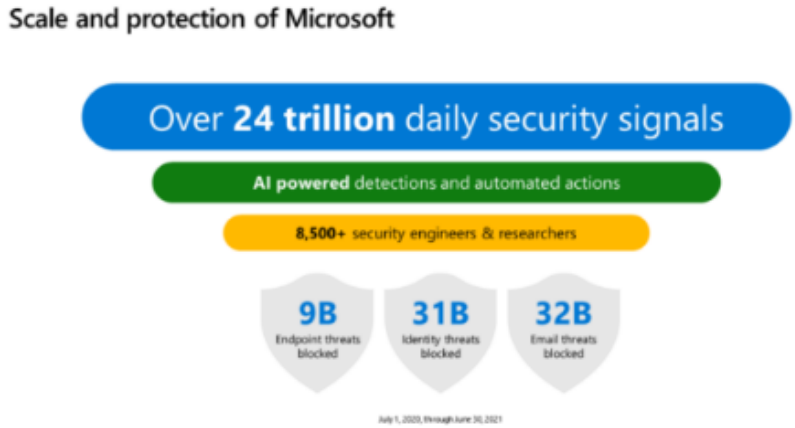

The Defender portal provides access to 27 main sources of information. Microsoft is regularly adding new functionalities and tabs.

The current sources of information[8] are:

1. Incidents & alerts
2. Hunting
3. Actions & submissions
4. Threat analytics
5. Secure score (out of scope)
6. Learning hub (out of scope)
7. Trials (out of scope)
8. Device inventory
9. Vulnerability management (out of scope)

---

[6] Screenshot taken from instruction video Microsoft about Defender for Endpoint, URL: https://www.microsoft.com/en-us/videoplayer/embed/RE4wDob?postJsllMsg=true

[7] Microsoft, Machine learning and AI Innovation at Microsoft Security Research, URL: https://www.microsoft.com/en-us/research/group/m365-defender-research/ . More details about the technical use of AI can be found in the blog, Improving AI-based defenses to disrupt human-operated ransomware, 21 June 2022, URL: https://www.microsoft.com/en-us/security/blog/2022/06/21/improving-ai-based-defenses-to-disrupt-human-operated-ransomware/

[8] As checked on 1 May 2023 on security.microsoft.com.

10. Partners and APIs (out of scope)
11. Evaluation & tutorials (out of scope)
12. Configuration management (out of scope)
13. Investigations
14. Explorer
15. Review
16. Campaigns (not tested)
17. Threat tracker
18. Exchange message trace
19. Attack simulation training (out of scope)
20. Policies & rules
21. App governance
22. Reports
23. Audit
24. Health
25. Permissions and roles
26. Settings
27. More resources

Figure 4: Collated screenshots of Microsoft 365 Defender Homepage

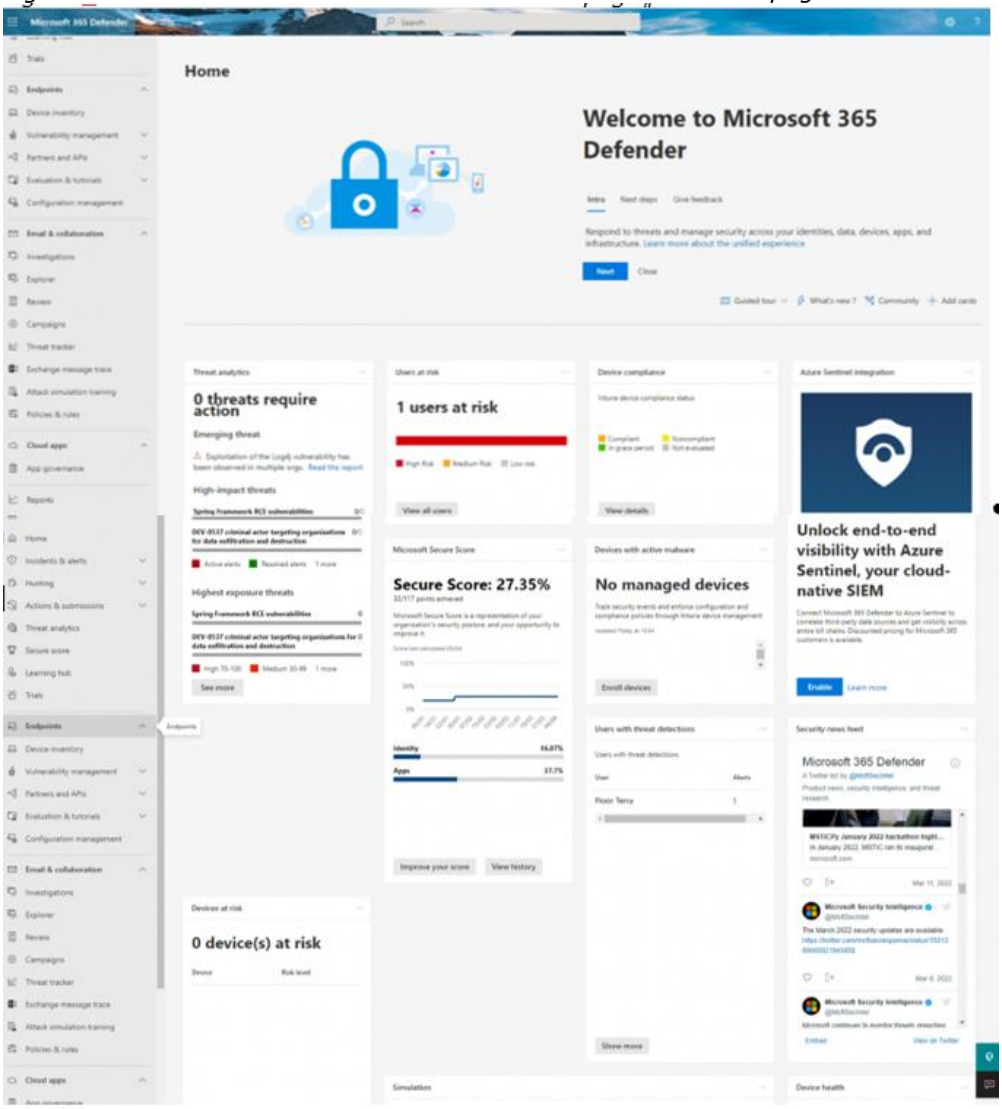

1.4.1 *Defender for Identity*

Defender for Identity uses information from the on-premises Active Directory to identify, detect, and investigate advanced threats, compromised identities, and malicious insider actions.[9]

Defender for Identity offers two types of identity risk detection or calculation: Real-time and Offline.[10]

1.4.2 *Defender for Office 365*

Defender for Office 365 protects against threats posed by email messages, links (URLs) and collaboration tools.

Defender for Office 365 includes *Exchange Online Protection (EOP)*, Microsoft's core email security tool. EOP checks all mails against blocklists for origin (spammers) and malware, and redirects all hyperlinks included in the body of e-mails to Microsoft, to be able to filter or trash malicious contents.

Microsoft explains:

> "*The message passes through content filtering (anti-spam and anti-spoofing) where harmful messages are identified as spam, high confidence spam, phishing, high confidence phishing, or bulk (anti-spam policies) or spoofing (spoof settings in anti-phishing policies.*"[11]

Figure 5: Microsoft illustration of EOP filtering[12]

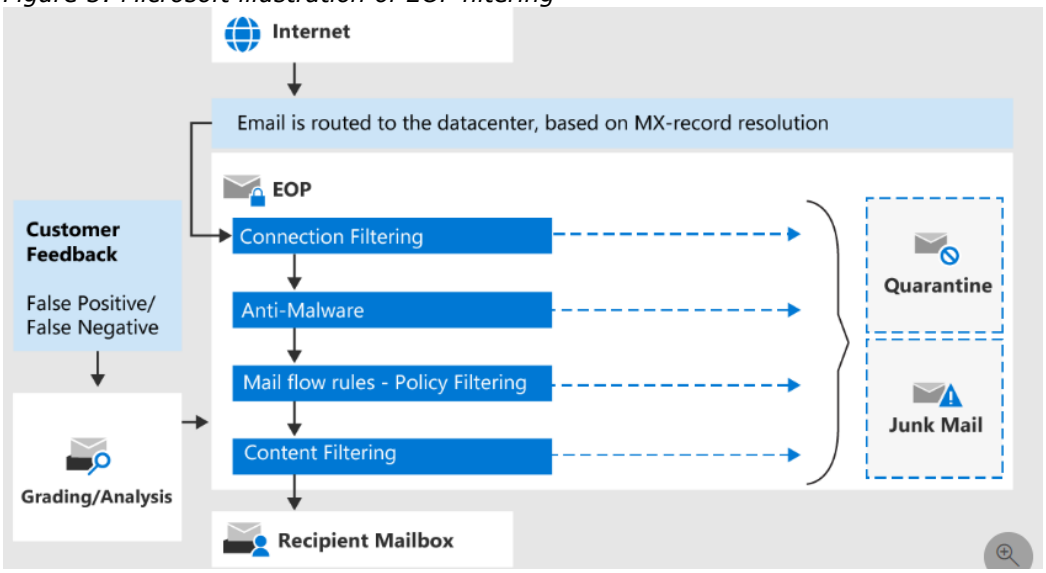

[9] Microsoft 365 Defender protection , URL: https://learn.microsoft.com/en-us/microsoft-365/security/defender/microsoft-365-defender?view=o365-worldwide

[10] Microsoft, Risk types and detection, 16 February 2023, URL: https://docs.microsoft.com/en-gb/azure/active-directory/identity-protection/concept-identity-protection-risks#risk-types-and-detection.

[11] Microsoft, Exchange Online Protection overview, 24 February 2023, URL: https://docs.microsoft.com/en-us/microsoft-365/security/office-365-security/exchange-online-protection-overview?view=o365-worldwide

[12] Idem.

Microsoft explains that Defender for Office 365 includes automated investigation and response capabilities.

> "*Microsoft Defender for Office 365 includes powerful automated investigation and response (AIR) capabilities that can save your security operations team time and effort. As alerts are triggered, it's up to your security operations team to review, prioritize, and respond to those alerts. Keeping up with the volume of incoming alerts can be overwhelming. Automating some of those tasks can help.*"[13]

Microsoft also explains it does not take automated remediation actions:

> *"In Microsoft Defender for Office 365, no remediation actions are taken automatically. Remediation actions are taken only upon approval by your organization's security team. AIR capabilities save your security operations team time by identifying remediation actions and providing the details needed to make an informed decision."*[14]

1.4.3 Defender for Endpoint

Defender for Endpoint monitors end-user devices running on multiple platforms (macOS, iOS, Linux, Android and Windows). Privacy Company has only done (limited) research on how Defender for Endpoint works, and only on Windows 10 devices. To protect the information on the device, MDE collects information about the processes active on the device, and scans incoming files and emails, trying to detect threats like viruses, hacking attempts and phishing. When MDE detects a potential threat, it will collect information about the event as well as of the state of the device at the time of the incident, including connection status, recent security events, domain names connected, and active and installed applications. To find potential threats, MDE scans emails and downloaded documents.

Microsoft explains how it sends feedback from detections both locally to the antivirus software on Windows 10 devices, and adds the newly discovered threat to the list of known threats:

"*After incriminating an entity, Microsoft Defender for Endpoint stops the attack via feedback-loop blocking, which uses Microsoft Defender Antivirus to block the threat on endpoints in the organization. Defender for Endpoint then uses the threat intelligence gathered during the ransomware attack to protect other organizations.*"[15]

---

[13] Microsoft, Automated investigation and response (AIR) in Microsoft Defender for Office 365, 28 February 2023, URL: https://learn.microsoft.com/en-us/microsoft-365/security/office-365-security/air-about?view=o365-worldwide.

[14] Idem.

[15] Microsoft, Improving AI-based defences to disrupt human-operated ransomware, 21 June 2022, URL: https://www.microsoft.com/en-us/security/blog/2022/06/21/improving-ai-based-defenses-to-disrupt-human-operated-ransomware/

Figure 6: Illustration Microsoft Defender for Endpoint[16]

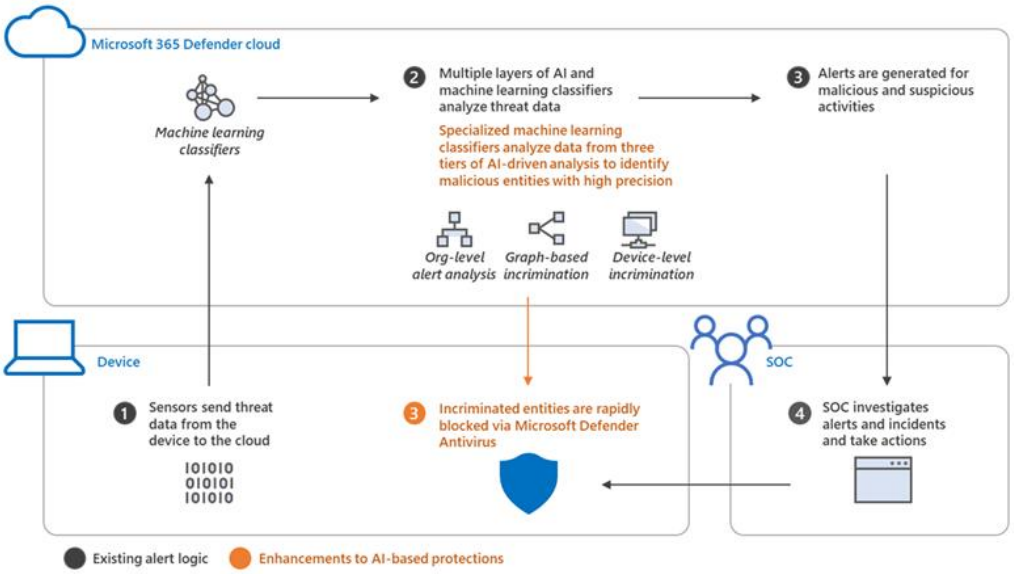

To use the service, Microsoft Intune is required as device management software.[17] SLM Rijk has published a separate DPIA on Intune.[18] The data processing by Intune is outside of the scope of this report.

1.4.4 *Defender for Cloud Apps*

Defender for Cloud Apps helps with security when working with cloud services. Defender for Cloud Apps acts as a gatekeeper, a *Cloud Access Security Broker*, to discover and provide visibility into Shadow IT and app use. It also monitors user activities for anomalous behaviours, controls access to the organisation's resources, can classify and prevent sensitive information breaches, protects against malicious actors, and helps to assess the compliance of cloud services.[19]

[16] Idem.
[17] Microsoft Defender portal, URL: https://security.microsoft.com .
[18] SLM Rijk, Data protection impact assessment Intune, 30 June 2020, URL: https://www.rijksoverheid.nl/documenten/rapporten/2020/06/30/data-protection-impact-assessment-intune.
[19] Microsoft Defender for Cloud Apps overview, 5 February 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/what-is-defender-for-cloud-apps.

Figure 7: Microsoft illustration of Cloud Apps Architecture[20]

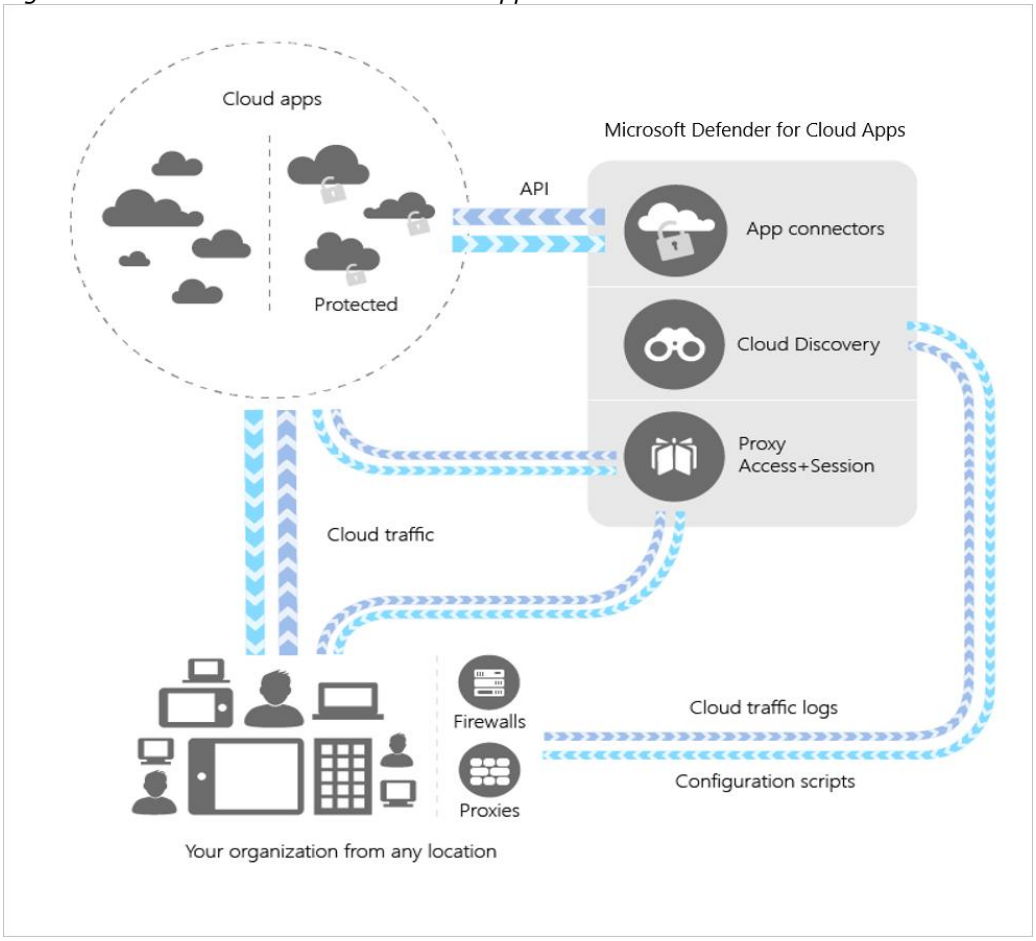

Depending on the configuration, Defender for Cloud Apps can route all traffic between the user and the selected cloud services through its system to scan all online traffic for potential malware (reverse proxy), collecting information from external cloud services through API-connectors and logs from the workstation.

Privacy Company has not tested this functionality but notes that such functionality would involve the processing of a wide array of personal data of the (external) data subjects whose data the government organisation is processing, as well as personal data relating to employees.

Microsoft explains: "*Admins can search for a specific user's metadata or user's activity. Selecting an entity opens the Users page. The Users page provides you with comprehensive details about the entity that are pulled from connected cloud applications. It also provides the user's activity history and security alerts related to the user.*"[21]

1.5 Different resources in Defender 365 portal

As Microsoft has merged the four different security tools in Defender, the results from the four tools are often presented in combined dashboards or reports.

---

[20] Idem, under Architecture, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/what-is-defender-for-cloud-apps#architecture.

[21] Microsoft, Data security and privacy practices for Defender for Cloud Apps, 24 April 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/cas-compliance-trust.

Below, the 11 most relevant modules from the 16 modules mentioned in Table 1 above are described to assess the risks of the data processing, with screenshots. This is followed by a summary list with a definition of the remaining accessible services for admins.

1. Device Inventory
2. Threat and Vulnerability Management
3. Users at risk as shown in the Home portal (analysed separately in Section 3.2 of this report)
4. Incidents & alerts
5. Threat analytics
6. Explorer
7. Review
8. Policies & rules
9. Audit
10. Cloud App Discovery and Cloud app catalogue
11. Files

1.5.1 *Device Inventory*

As shown in Figure 8 below, Defender shows overviews of connected devices and apps used by end users.

Figure 8: Screenshot Microsoft Device Inventory

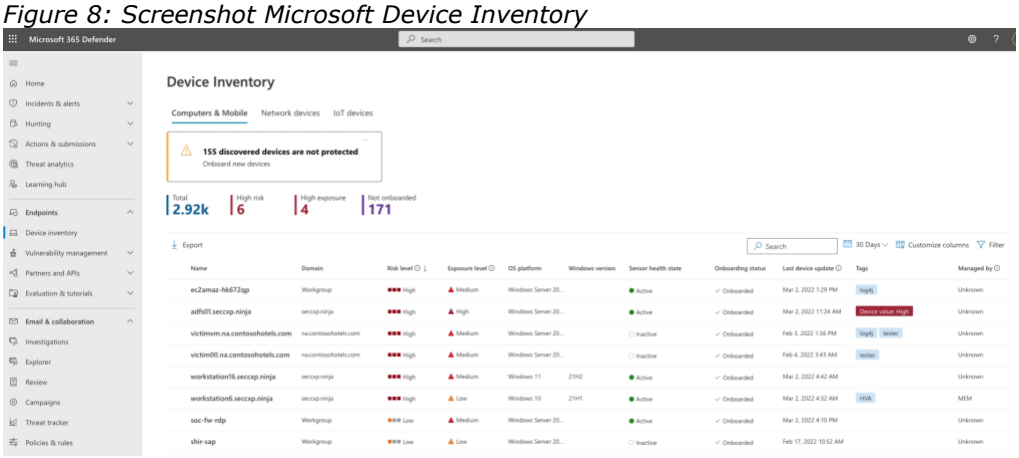

1.5.2 *Threat and Vulnerability Management*

To detect vulnerable devices, the Microsoft Defender for Endpoint sensor automatically collects vulnerability and security data from the connected devices and publishes the results in the portal. It shows devices that were in use up to 30 days before.

Figure 9: Screenshot Microsoft Threat & Vulnerability Management dashboard[22]

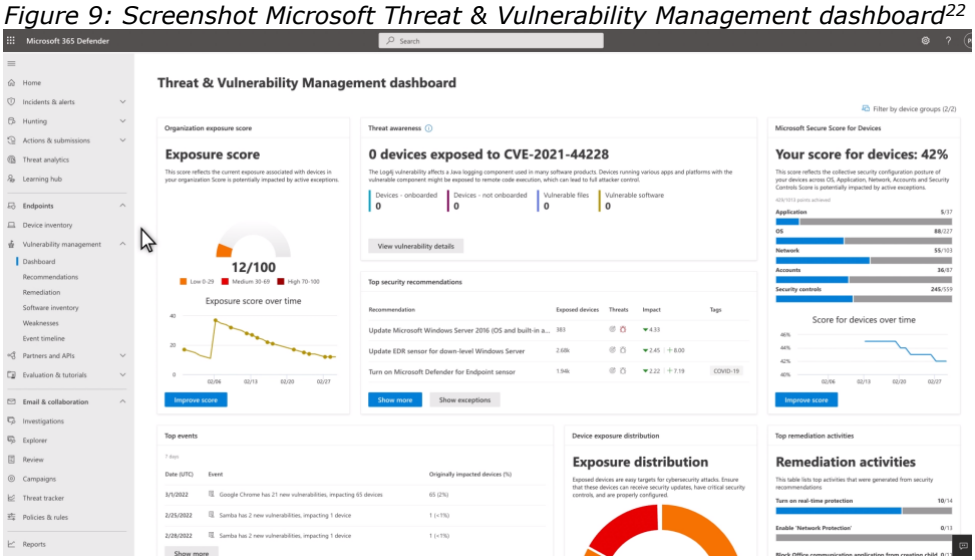

1.5.3 *Incidents & alerts*

With *Incidents & alerts* Microsoft shows automatic warnings to an organisation about cyber-attack events detected by any tool in the Microsoft 365 Defender. The incidents can be prioritized by severity, and admins can identify affected end users by device and per mailbox. They can also investigate domains and URLs associated with a Microsoft Defender for Endpoint alert.

Figure 10: Incidents & alerts

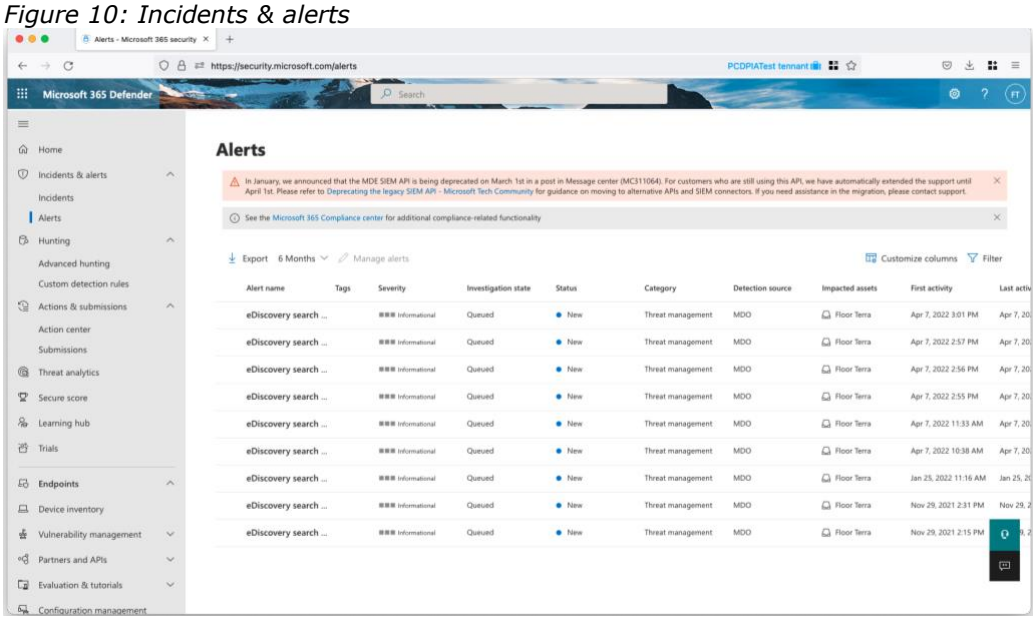

[22] Screenshot from Microsoft instruction video Microsoft Defender for Endpoint, 7 February 2023, URL: https://www.microsoft.com/en-us/videoplayer/embed/RE4wDob?postJsllMsg=true.

Figure 11: Screenshot Microsoft alert about suspicious domain[23]

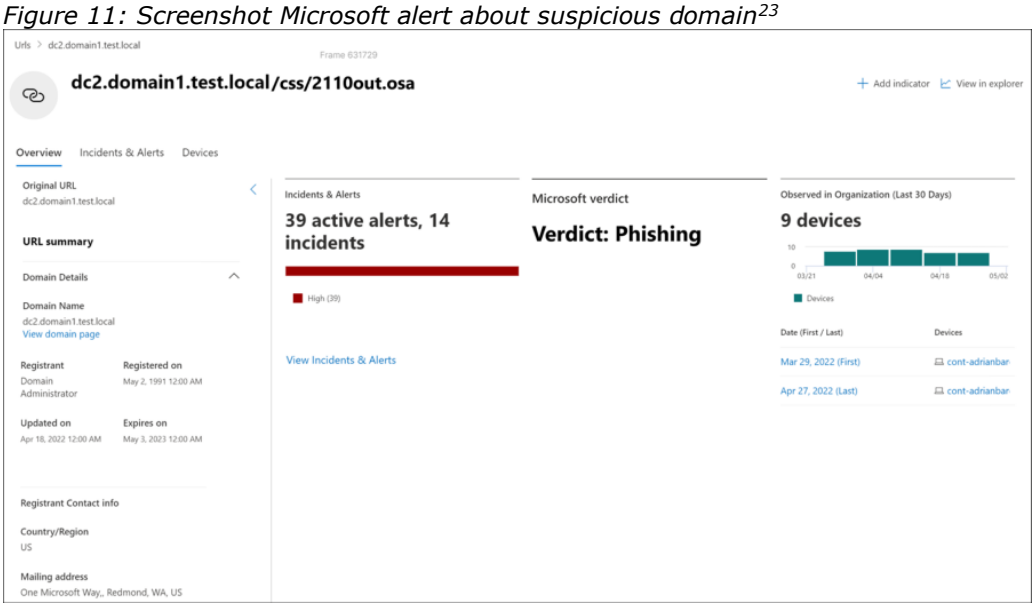

1.5.4 *Threat analytics*

Threat analytics is based on a set of reports from Microsoft security researchers. Microsoft describes that these reports cover the most relevant global threats, including:

- *"Active threat actors and their campaigns*
- *Popular and new attack techniques*
- *Critical vulnerabilities*
- *Common attack surfaces*
- *Prevalent malware."*[24]

The tool shows which of these global threats impact assets (devices or mailboxes). The admin can see devices with alerts, and devices with active and resolved alerts over time.[25]

Microsoft explains impacted mailboxes are "*mailboxes that have received email messages that have triggered Microsoft Defender for Office 365 alerts. While most messages that trigger alerts are typically blocked, user- or org-level policies can override filters.*"[26]

---

[23] Screenshot Microsoft, Investigate domains and URLs associated with a Microsoft Defender for Endpoint alert, 27 October 2022, URL: https://learn.microsoft.com/en-us/microsoft-365/security/defender-endpoint/investigate-domain?view=o365-worldwide.

[24] Microsoft, Defender for Endpoint, Track and respond to emerging threats through threat analytics, 7 February 2023, URL: https://docs.microsoft.com/en-us/microsoft-365/security/defender-endpoint/threat-analytics?view=o365-worldwide.

[25] Idem.

[26] Microsoft, Threat analytics in Microsoft 365 Defender, 8 March 2023, URL: https://docs.microsoft.com/en-us/microsoft-365/security/defender/threat-analytics?view=o365-worldwide

In the test tenant, as shown in Figure 12 below, this did not lead to any alerts, because the tests were performed with known, old malicious files and this tool only shows alerts for the current most important global threats.

Figure 12: (Global) threat analytics

1.5.5 *Explorer*

The tool *Explorer* discloses Microsoft's detection of suspicious emails, for example, attachments with malware, or content with phishing campaigns. One of the tabs discloses the 'Top targeted users', in other words, the end users that most frequently receive such mails. It is also possible to view the 'Top URLs', with an indication what Microsoft has done with the URL. Each URL shown in that list provides a hyperlink to the specific end users that have visited those URLs.
Microsoft is able to collect this information about visited URLs because Exchange Online Protection has an in-build URL replacement that redirects all clicks to Microsoft itself, in order to detect and possible junk or quarantine malicious files. See Figure 5 in Section 1.4.2 of this report.

Figure 13: Explorer (mail security incidents)

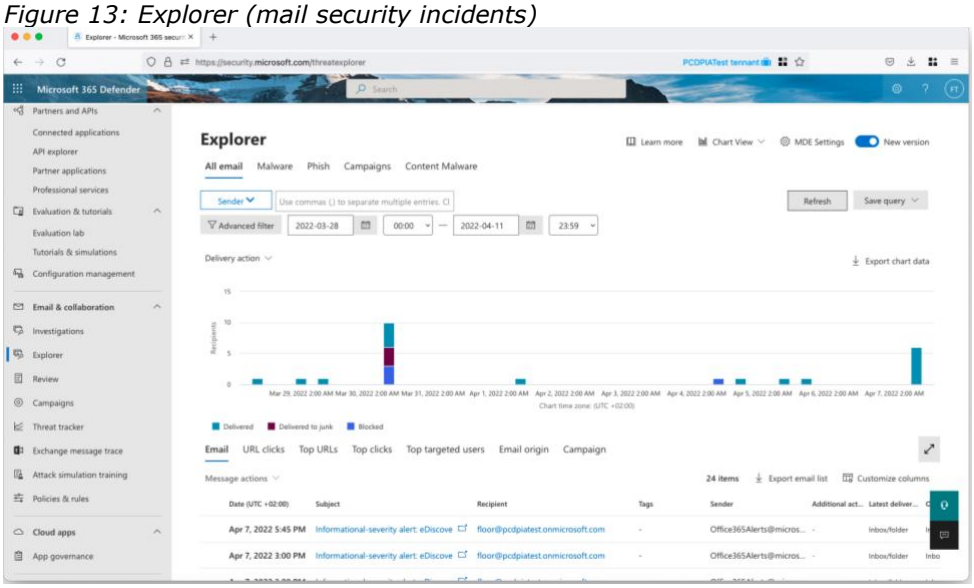

Figure 14 below illustrates details of the Top URLs in the Explorer dashboard.

Figure 14: Detail of Top URLs

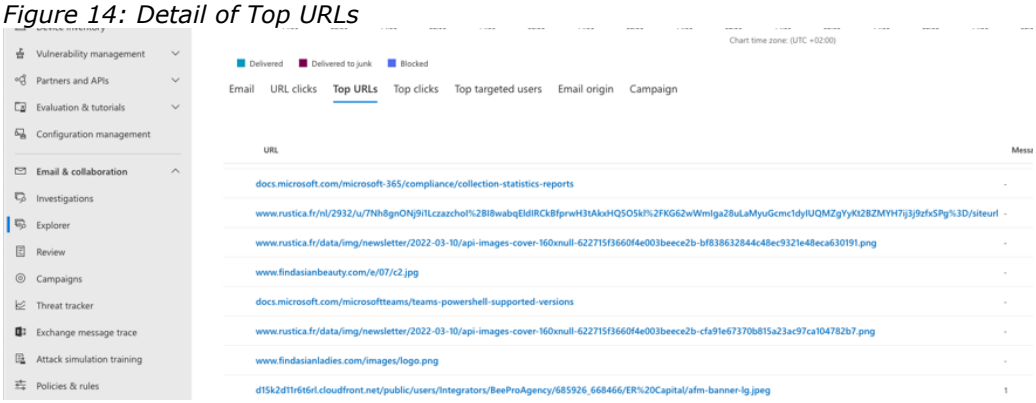

1.5.6 *Review*

As illustrated in Figure 15 below, with the dashboard *Review* system administrators get an overview of actions to take. If Microsoft has quarantined a suspicious mail or file, it asks the administrator for a 'review', that is, what the administrator wants to do with the mail or file. Another review task is to decide what to do with users that have been blocked by Microsoft for sending too many messages classified as bulk mail. See Figure 16 below.

Figure 15: Review dashboard

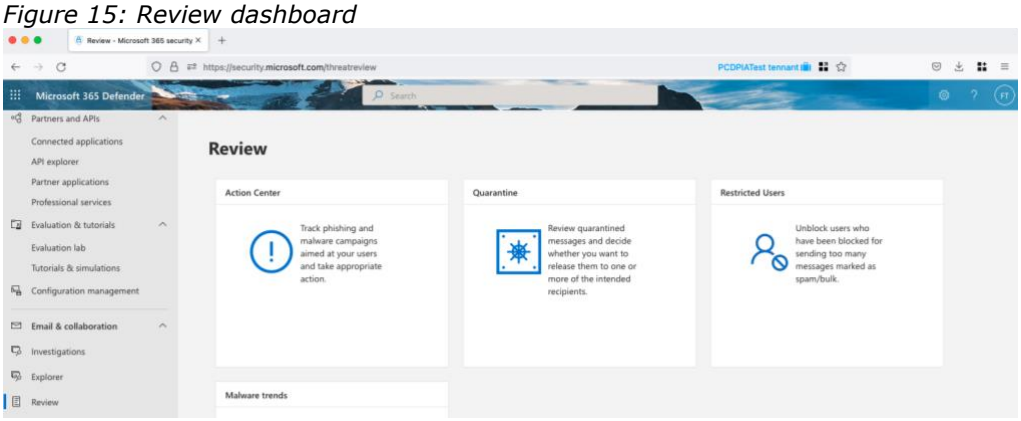

Figure 16: Request for review on email

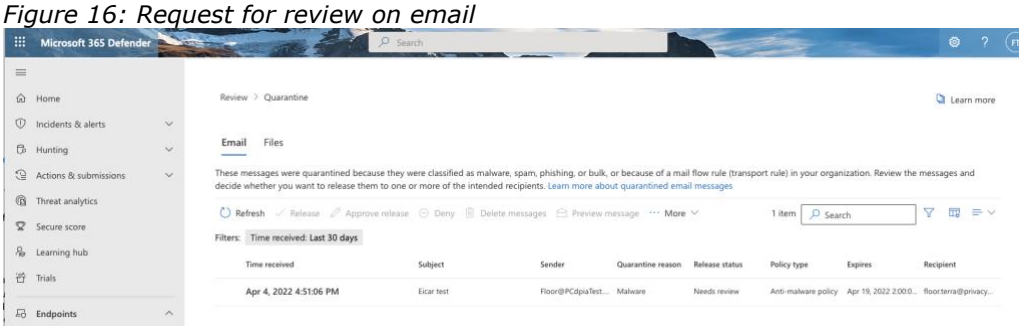

Figure 17: Request for review on files

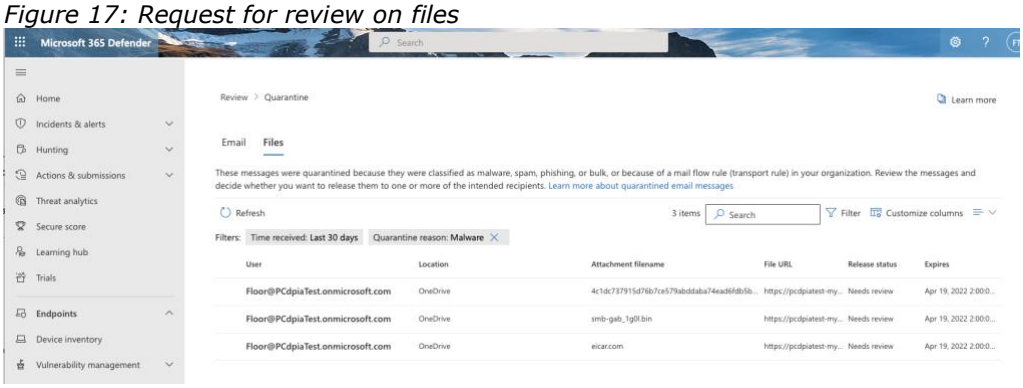

1.5.7 *Policies & rules*

With *Policies & rules* Microsoft offers four sets of settings in Defender to help admins manage threats: (i) Threat policies, (ii) Alert policy, (iii) Manage advanced alerts and (iv) Activity alerts.

With (i) Threat policies the admin decides how strict for example DKIM controls are set, or what the threshold for the number of received mails to qualify such mail as bulk mail. The admin also decides when the end user is notified if mail is qualified as spam and if mail is quarantined. By default, the option to notify senders is disabled. If enabled, the organisation risks to alert malware senders that their behaviour is detected.

With the (ii) Alert Policy the admin decides how high the threat must be before an alert is sent, in many different threat categories.

The other two options are out of scope of this report.

Figure 18: Policies & rules

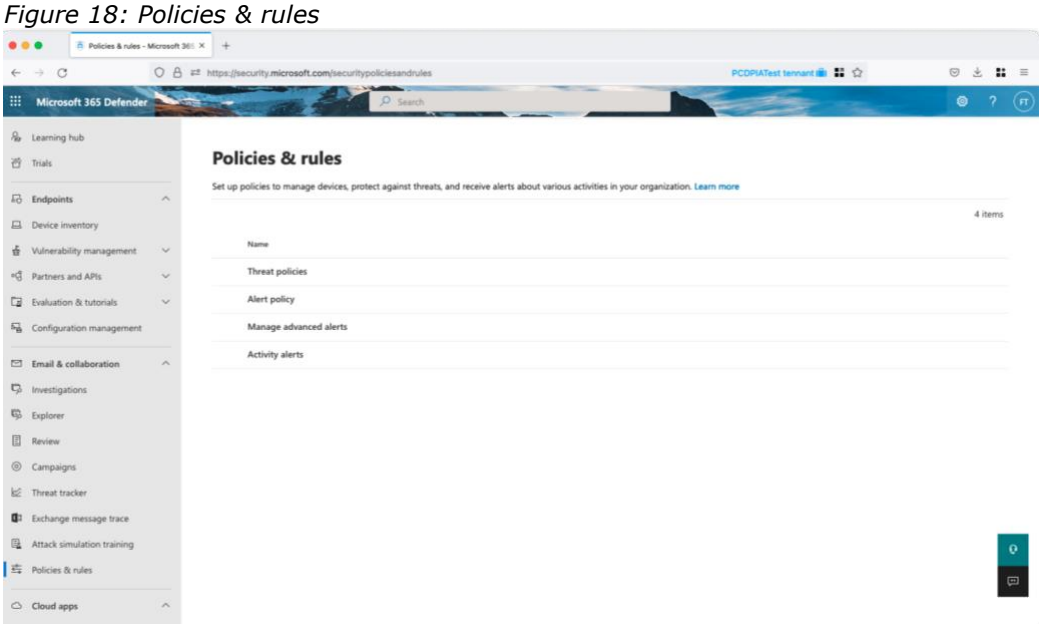

1.5.8 *Audit*

Audit provides access to the service generated server logs made available by Microsoft as audit logs. See Figure 19 below.

Figure 19: Audit logs[27]

Export

Date	IP Address	User	Activity	Item	Detail
Apr 1, 2022 3:41 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	Accessed mailbox items		Mail Items Accessed
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	Accessed mailbox items		Mail Items Accessed
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	User logged in	00000003-0000-0000-c000-00000000...	
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	User logged in	00000003-0000-0000-c000-00000000...	
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	Accessed mailbox items		Mail Items Accessed
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	Accessed mailbox items		Mail Items Accessed
Apr 1, 2022 3:40 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	User logged in	00000002-0000-0ff1-ce00-00000000...	
Apr 1, 2022 3:39 PM	45.137.101.242	Floor@PCdpiaTest.onmicrosoft.com	UserLoginFailed	00000002-0000-0ff1-ce00-00000000...	
Apr 1, 2022 3:35 PM		NOT-FOUND	SearchMtpStatus		
Apr 1, 2022 1:35 PM		NOT-FOUND	SearchMtpStatus		
Apr 1, 2022 11:35 AM		NOT-FOUND	SearchMtpStatus		
Apr 1, 2022 9:35 AM		NOT-FOUND	SearchDataInsightsSubscription		
Apr 1, 2022 9:35 AM		NOT-FOUND	SearchTIKustoClusterInformation		
Apr 1, 2022 9:35 AM		NOT-FOUND	SearchMtpStatus		
Apr 1, 2022 7:35 AM		NOT-FOUND	SearchMtpStatus		
Apr 1, 2022 5:35 AM		NOT-FOUND	SearchMtpStatus		

1.5.9 *Cloud App Discovery dashboard and Cloud app catalogue*

As shown in Figure 20 below, the Cloud Discovery dashboard has different tabs, including information about discovered apps that may be shadow ICT, and behaviour per user and per device.

Figure 20: Screenshot Microsoft Cloud Apps Discovery[28]

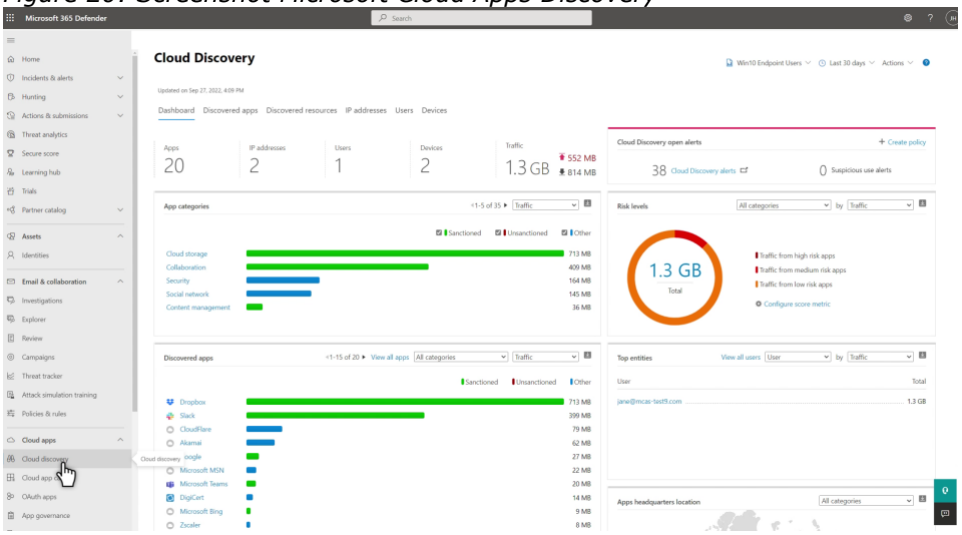

[27] Screenshot made in the government test tenant on 1 April 2022.

[28] Screenshot from Microsoft 365 Defender instruction video, URL: https://www.microsoft.com/en-gb/videoplayer/embed/RE59yVU?postJsllMsg=true&autoCaptions=en-gb

1.5.10 *Activity Log*

The Activity log allows admins to search for specific activities with policies, for example a search for activities performed on certain files. Microsoft explains: "*you can use the Activity log to find users in your organization who are using operating systems or browsers that are out of date, as follows: After you connect an app to Defender for Cloud Apps in the Activity log page, use the advanced filter and select User agent tag. Then select Outdated browser or Outdated operating system.*"[29]

1.5.11 *Files*

With Files, organisations can scan all files for modifications, for example all the files stored in OneDrive and Salesforce. The modification can be to content, metadata, or sharing permissions. Organisations can also use the Files functionality to investigate what kind of data are saved in cloud apps.[30]

1.5.12 *Description of other resources in Defender 365*

Organizations may choose to install ***Antivirus (Next Generation Protection)*** on Windows devices. The service uses information from multiple sensors in the operating system as well as machine learning, big data analytics and threat research to protect devices. This is out of scope of this report.

The resource ***Hunting*** enables admins to perform specific security related queries based on the security logging collected by the four different tools in Defender. This tool uses Kusto Query Language, that enables quick hunting over tons of records. The queries can be aimed at users, devices, specific alerts or configurations, to help organisations find and investigate security incidents.

Examples of use cases are:

- Look-up user logins from specific countries. This can be useful to investigate if hackers from a specific country such as China have attempted to log in after one account is compromised by a hacker from that country;
- Find all devices with outdated vulnerable hardware that needs to be upgraded;
- Find malicious email attachments;
- Find malicious senders of mail.

The resource ***Actions & submissions*** enables organisations to take actions such as quarantining a malicious file, isolating a device, soft deleting an email, starting an antivirus scan or blocking URLs.

Secure score provides essential security organisations to organisations based on existing Microsoft security tooling. This is similar to the Compliance score in the Microsoft 365 compliance center, with a slightly different scope. This is out of scope of this report.

---

[29] Microsoft Defender for Cloud Apps, Activities, 24 April 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/activity-filters.

[30] File filters in Microsoft Defender for Cloud Apps, 5 March 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/file-filters.

Reports offers different reports and dashboards of data that are also disclosed via other tabs, such as the homepage (with the view of *Users at risk*) and email & collaboration reports (a category that for example contains the Explorer results).

The ***Learning hub*** offers links to articles, videos, blogs and interactive training.

The resource **(Threat)** ***Investigations*** enables organisations to create scripts to automatically log investigations and start responding to the logged security incidents.

The module ***Campaigns*** provides warnings about phishing and malware attacks, as identified and categorised on the customer tenant. The test tenant was not subjected to such campaigns during the (brief) testing period, and therefore, this module was not tested.

Threat tracker provides information about cybersecurity issues via different widgets. For example: there is a wave of Russian attacks on Ukrainian sources, show if specific employees are at risk.[31] This module could not be tested in the test tenant, due to the limited testing time and limited test employees.

Exchange message trace provides a hyperlink to the Exchange Online Environment. An admin can use this tool to search for the inboxes and logs of all end users. When an incident has occurred related to a specific end user, the admin may for example want to inspect all recent outgoing mails of that user. Or, if the organisation receives malware from a specific sender, the admin may want to check all incoming mail from that sender.

Trials offers access to 90 day free trials of different Defender tools, but this is out of scope of this targeted verification report.

The module ***Health*** contains two items: Service health and Message Center. *Service Health* provides information about interruptions of Microsoft services and incident advisories. *Message Center* contains a wide variety of update messages from Microsoft, for example related to Office, Windows or Defender feature updates.

Permissions & roles enables admins to view the permissions within the organisation for access to Defender.

Settings shows some privacy relevant settings generally determined in other tools, such as the geolocation of data and the chosen data retention periods. Settings offers some choices such as enabling feature preview or adding specific mail addresses for alerts.

More resources offers hyperlinks to other Microsoft security tools, such as Azure AD Identity Protection and Azure Information Protection. The resources **App governance**, ***Device inventory***, ***Vulnerability management***, ***Partners and APIs***, ***Evaluation & tutorials***, ***Device configuration management*** and ***Attack***

[31] Microsoft, What are threat trackers?, 6 April 2022, URL: https://docs.microsoft.com/en-gb/microsoft-365/security/office-365-security/threat-trackers?view=o365-worldwide#what-are-threat-trackers.

simulation training are out of scope of this report, as they relate to either device management, Cloud App security or to trainings.

2. Verification questions

This report is based on the following nine verification questions:

1. Does Defender adequately detect critical threats and security breaches from devices and cloud apps, and with regard to malicious emails, files and URLs that are exchanged via e-mail and via OneDrive and Teams, both within the tenant and with external individuals?

2. Does Defender provide adequate warnings about suspicious logins via the Azure Active Directory?

3. How does the data anonymization option work in Defender for Cloud Apps?

4. Does Microsoft send traffic to third parties (including through cookies, and to itself as an independent data controller) when a system administrator enables Defender? If so, are those third parties mentioned on the list of subprocessors?

5. Does Microsoft process learnings from security incidents across its Enterprise and Education user base, or are there limitations, such as prior anonymisation?

6. Does Defender create individual risk profiles and/or individual scores in the different analytic overviews and reports?

7. Does Microsoft publish adequate documentation on the personal data it collects through the tested applications, in comparison with captured network traffic and logs that are accessible for system administrators?

8. Does Microsoft give system administrators full access to all personal data it processes through the different Defender tools? Does Microsoft provide adequate explanations if it does not provide access to certain personal data?

9. Are there high risks resulting from the transfer of personal data to the USA or other third countries?

The questions are answered by a short description of the relevant facts, such as available settings, public information or improvement commitments from Microsoft.

Per question, the section with the facts is followed by a description of the relevant technical findings, an assessment of the legal consequences of the technical findings and suggestions for mitigating measures.

3. Detection of malicious files, apps and URLs

This section answers the first verification question:

Does Defender adequately detect critical threats and security breaches from devices and cloud apps, and with regard to malicious emails, files and URLs that are

exchanged via e-mail and via OneDrive and Teams, both within the tenant and with external individuals??

3.1 Facts

As data processor for Office and/or Microsoft 365 Enterprise (for Education), Microsoft must assist the data controllers (the Dutch government organisations and universities) with their security posture. Based on Art. 28 of the GDPR (processor obligations) this includes (in abbreviated descriptions):

- taking all measures required pursuant to Article 32 [security measures] (Art. 28 (3) sub c);
- assisting the controller in ensuring compliance with the obligations pursuant to Articles 32 to 36 [compliance with security and DPIA obligations] (Art. 28 (3) sub f);

In order for data controllers to demonstrate that Microsoft complies with its processor obligations, the government and university admins must be able to verify the effectivity of the security tools, to establish that the data processing is necessary to comply with GDPR and specific government security obligations.

3.2 Technical findings

As described in the Section 1.3 of the Introduction, Privacy Company performed tests on two of the tools in November 2020 (Defender for Endpoint and Cloud Apps Security), and on the two other tools in April 2022 (Defender for Identity and Defender for Office 365).

For all tools, a limited test was performed with known malware to reduce the risk of triggering an actual malware incident. Both in 2020 and 2022 Privacy Company performed several tests with the EICAR[32] anti-malware test file.

In 2020 Privacy Company tested Defender for Endpoint (on the Windows device) with other examples of tools that can be used for data exfiltration: Mimikatz software to steal login credentials from a Windows PC (in 2020), FileZilla and Tor. Privacy Company also tested the data processing via Cloud Apps Security by downloading a diabetes app.

In 2022 Privacy Company downloaded *Bonzi Buddy* (a famously aggressive adware program disguised as freeware desktop virtual assistant, launched in 2000 and terminated in 2004[33]) in the OneDrive, as well as a sample of the WannaCry-malware. The WannaCry sample was not shared outside of the test tenant to reduce the risk of an actual incident.

Privacy Company performed the tests as follows:

- Use of the laptop to (attempt to) download en send the EICAR test virus via the installed Outlook app.

[32] https://www.eicar.org/?page_id=3950
[33] https://en.wikipedia.org/wiki/BonziBuddy

- Attempt to download Mimikatz on the Windows 10 device
- Install the MedM Diabetes app from the Windows Store on the Windows 10 device.
- Use of Outlook for the Web and Teams for the Web on a Windows 10 device; sending and receiving files within the tenant and with external users.
- Use of OneDrive for the Web on a Windows 10 device; uploading and downloading files and sharing the files.

The result in 2020 was that Defender for Endpoint detected, and prevented the downloading of EICAR and Mimikatz. The Diabetes app was downloaded to check what data Cloud Apps Security collected about the name/nature of the app in its logfiles. The outcome was that Defender for Endpoint does not structurally log all network traffic and installed apps on a device, only in case of security incidents.

When MDE detects a possible security incident, it creates an "*investigation package*", that contains information about the state of the device at the time of the incident. This investigation package is sent to the Defender cloud service in Azure and available for inspection to the system administrators. Microsoft can also use this information for its own advanced analysis.

The investigation package contains a lot of information about the end user device, including all installed applications, current network connections (including local and internet connections and DNS requests) and running processes (what software is active on the device). The logs in the package all contain references to device and user identifiers. Most of these logs contain very technical information that is not very revealing of the user, but some of these may be quite sensitive. The DNS cache in particular may contain information about the political, religious, or arguably sexual interests of the data subject derived from web surfing or app usage on the device.

The results in 2022 were that Defender for Office 365 and Defender for Identity showed several alerts, as shown in Figures 21

Figure 22 below. However, no limitations were observed on the sharing of *Bonzi Buddy* (the ad/spyware program).

Figure 21: Warning on share limitation of compromised file in OneDrive

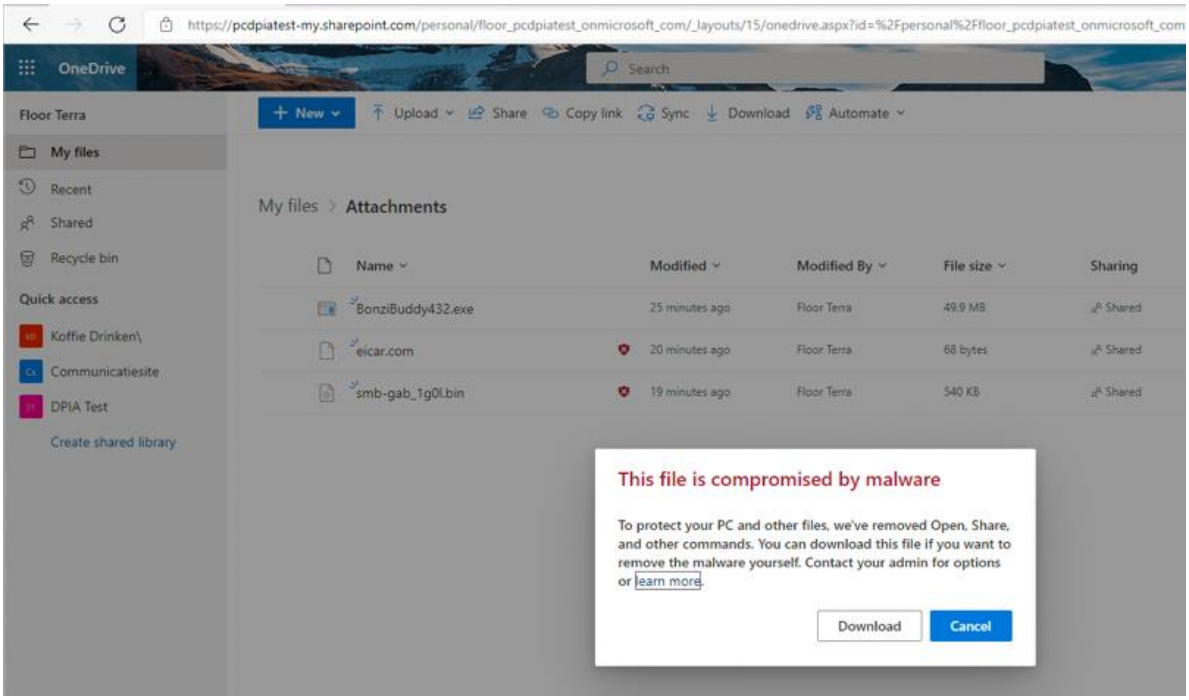

Figure 22: Defender detected an e-mail with a malware-attachment

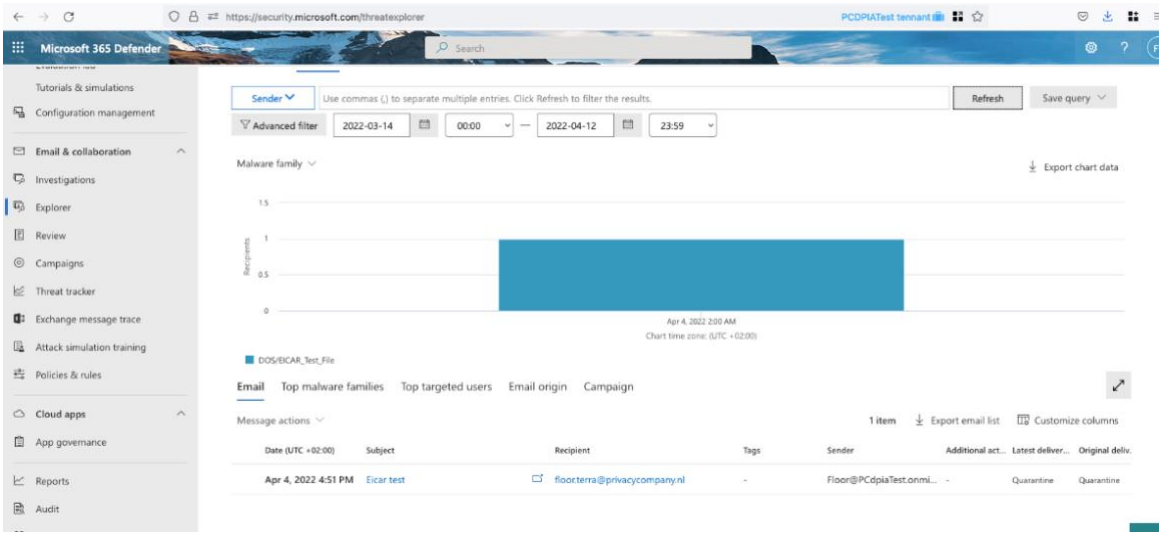

This resource enables admins to take a deep dive in the nature and origin of files marked as suspicious. As shown in Figures 23, 24 and 25 below, Microsoft shows the most important event metadata, such as affected user, type of action taken, type of incident and timestamp, with underlying information about all the details.

Figure 23: Defender alert details of EICAR detection

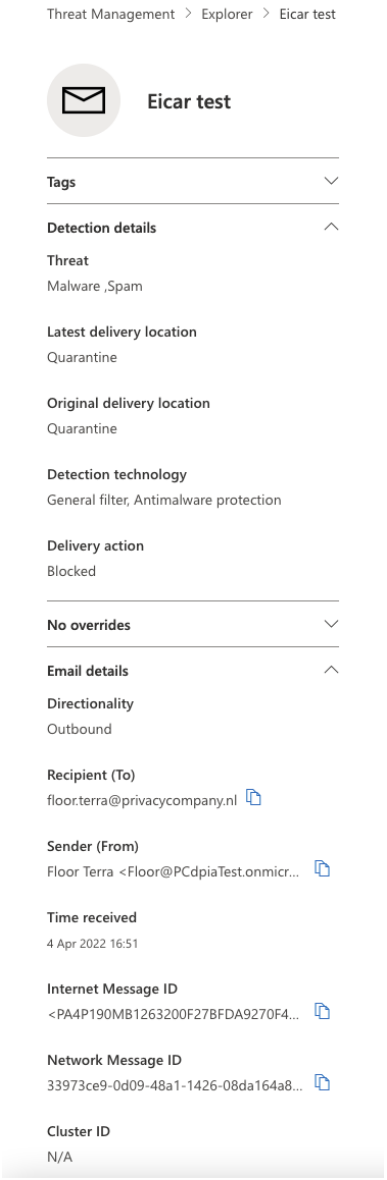

Figure 24: Alert timeline

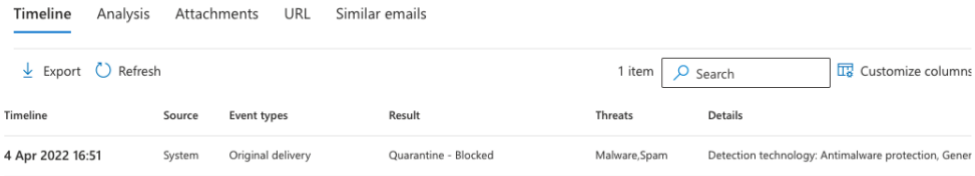

Figure 25: Analysis

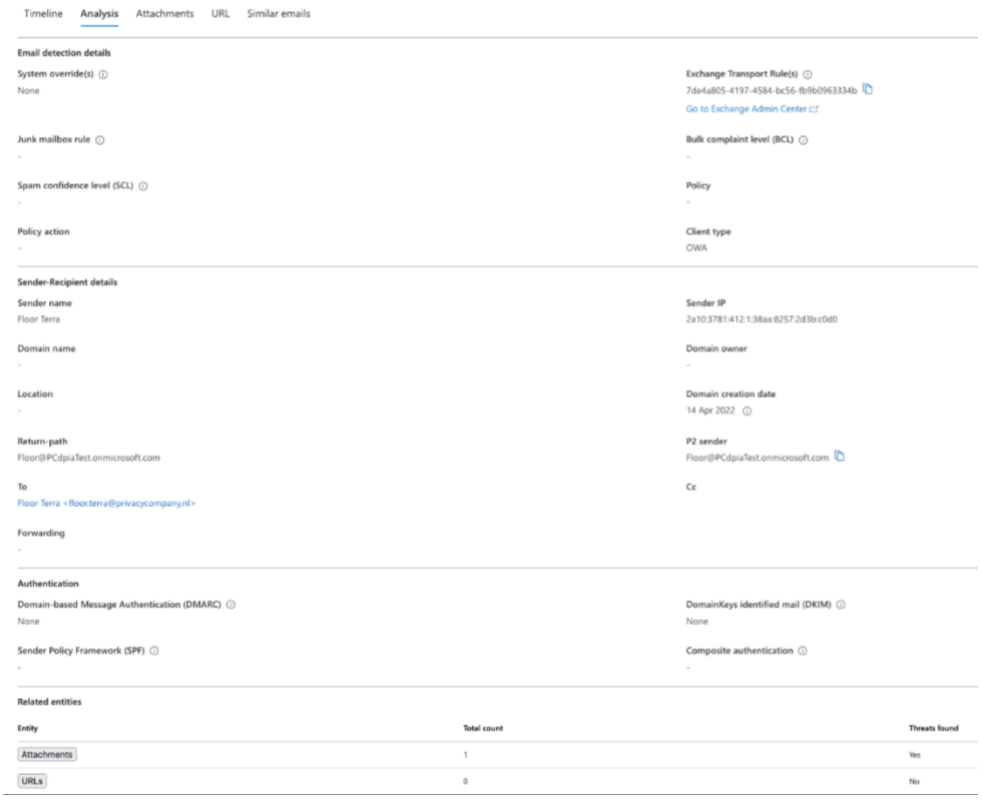

3.3 Assessment

Organisations need an adequate protection layer to protect the personal data they process on the devices, in Office and in the Active Directory. Absent adequate protection there would be imminent risks to the rights and freedoms of employees, as well as of clients and contacts of the government organisation and to the continuity and integrity of the operations of the government organisation itself.

The main outcome of these tests is that Defender 365 does what it is supposed to do: detect malicious files, apps and URLs. However, this verification report is not a security assessment. The effectiveness against relevant threats needs to be assessed by each individual organisation implementing Defender 365.

The spyware *Bonzi Buddy* was included in the test to assess detection of a class of software that is not necessarily classified as malware. It is highly unlikely that an organisation would want their employees to install *Bonzi Buddy* on their work-computers. Since this spyware was not detected, Defender 365 would not be the recommended solution against this category of risks.

For spam containing malicious URL's protection is similarly present but not absolute. E-mails containing malicious links are blocked, however some slipped through without detection. This is probably due to the fact that Microsoft takes the account properties of the sender into account, and not only the contents. In the test set-up malicious mails and spam were forwarded from existing Microsoft e-mail accounts

(live.com). This possibly lead to a lower risk score. Some of these bulk mails were blocked, but not all. It appears that Defender 365 gives the benefit of the doubt to 'known' sender domains. This makes Defender 365 less effective in determining the risks of 'known' senders.

This does not mean Defender is not capable of achieving the purposes for which it processes personal data: securing the work environment. The lack of alerts during some tests is a logical consequence of the limited test scope while Microsoft bases risk assessments on multiple criteria. These criteria could not all be replicated in the tests. For example: replicating a worldwide phishing campaign on multiple organisations was outside of the scope of the tests.

The tests show that end users receive clear information when files and e-mails are blocked. By default, they can access mails qualified as spam in a separate spam box. Malicious files that are put in quarantine can be reviewed by the admin. The admin can perform a manual check of the file before the file is permanently deleted, or automatically reject malicious files. By default quarantined files are kept in quarantine for 15 or 30 days, depending on the reason for the quarantine. They are deleted after that period. Microsoft explains that admins can change the retention period for messages quarantined by anti-spam policies (including spam, high confidence spam, phishing, high confidence phishing or bulk) and for messages quarantined by anti-phishing policies (including spoof intelligence in EOP; user impersonation, domain impersonation, or mailbox intelligence in Defender for Office 365).[34]

The sender of the suspicious files and mails on the other hand does not receive feedback from Microsoft that mail was blocked or quarantined. This default setting makes perfect sense, because if Microsoft were to send feedback, malicious senders could learn from the feedback to try to bypass the filters.

3.4 Remedies

The first assessment does not lead to recommendations to Microsoft. However, this first analysis does lead to three recommendations to the system administrators responsible for the settings in Defender.

1. Avoid automatic deletion of suspicious mails

2. Allow employees to ask admins perform to perform a manual review on documents quarantined as malware

3. Allow the end user access to mails qualified as spam

Automated deletion of spam mails is a type of automated decision making that can significantly affect the data subjects sending (legitimate) e-mails. For example, if they apply for a job but their application is deleted, or if there is a legal requirement to send a document before a fixed date, but the mail allegedly is never received.

[34] Microsoft, Quarantined email messages in EOP and Defender for Office 365, 18 March 2022, URL: https://docs.microsoft.com/en-us/microsoft-365/security/office-365-security/quarantine-email-messages?view=o365-worldwide.

4. Detection of suspicious logins in the Azure AD

This section answers the second verification question:

Does Defender provide adequate warnings about suspicious logins via the Azure Active Directory?

4.1 Facts

Defender includes detection of the following categories of risky login types.

- Activity from anonymous IP address
- Anomalous token
- Anonymous IP address
- Atypical travel
- Azure AD threat intelligence
- Impossible travel
- Malicious IP address
- Malware linked IP address
- Mass access to sensitive files
- Nation state IP
- New country
- Password spray
- Possible attempt to access Primary Refresh Token (PRT)
- Suspicious browser
- Suspicious inbox forwarding
- Suspicious inbox manipulation rules
- Suspicious sending patterns
- Token issuer anomaly
- Unfamiliar sign-in properties

Microsoft explains in a whitepaper about Azure Active Directory Data Security Considerations: "*Azure AD Identity Protection uses real-time user login data along* ***with multiple signals from company and industry sources to feed to its***

***machine learning systems to detect anomalous logins**. Personal data is scrubbed from this real-time login data before it is passed into the machine learning system, along with the remaining login data used to identify users and logins that are potentially risky."*[35]

4.2 Technical findings

In 2020 the MCAS functionality was tested to warn about suspicious logins, based on Azure AD Identity Protection. For the purpose of the test, logins were made from unusual foreign countries by signing in from a collection of Tor-exit nodes. MCAS correctly identified the infrequent country and generated an alert (in this case: Morocco).

Figure 26: Example of alert about infrequent country (2020)

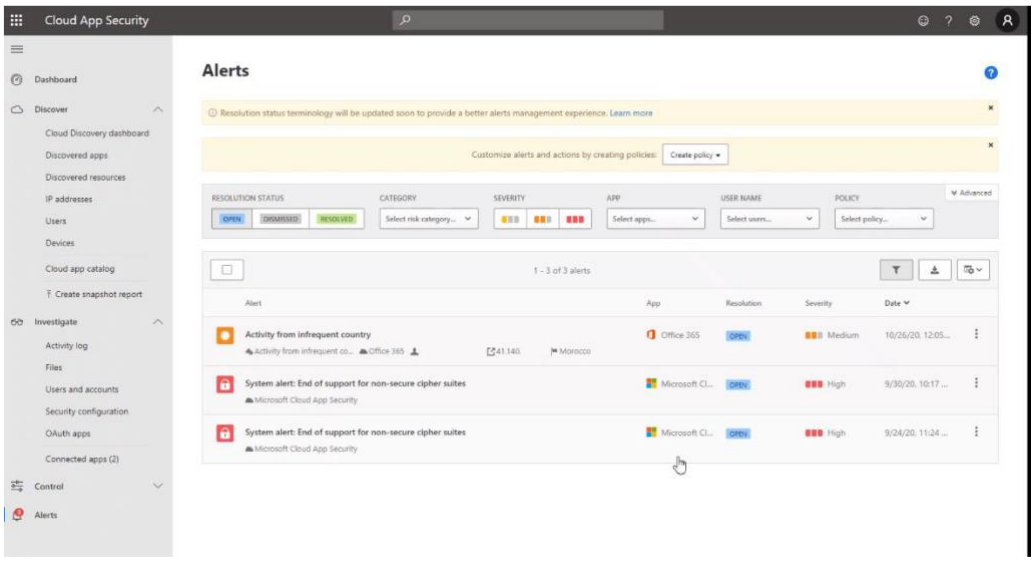

In 2022 Privacy Company performed tests simulating 'suspicious' logins by signing in from a collection of Tor-exit nodes using an account secured by the Microsoft Authenticator app. None of the logins was marked as suspicious.

Re-testing without the Microsoft Authenticator app, but with SMS as second factor did yield an alert of one user at risk. In order to simulate a login attempt on a user account from multiple locations that cannot physically be reached in that short a time span, the test account was used to log in from Tor-exit nodes from Russia, Germany, Brazil and from a consumer IP in the Netherlands. Defender detected this behaviour and labelled this attack as a 'password spray' attack.

[35] Microsoft whitepaper, Azure Active Directory Data Security Considerations, 1 July 2020, URL: https://azure.microsoft.com/en-us/resources/azure-active-directory-data-security-considerations/.

Figure 27: Detection of user at risk from Defender 365 Home dashboard

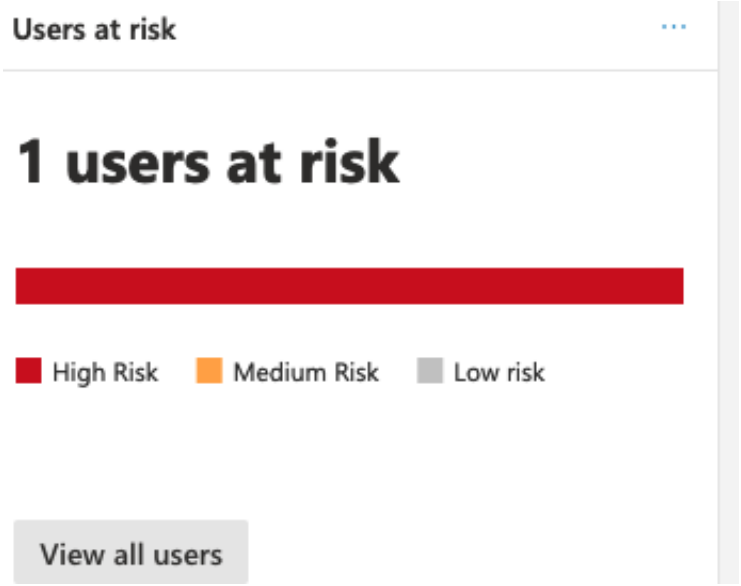

Clicking through this alert takes the administrator to the Azure admin portal showing the detected risks.

Figure 28: Overview of users at risk in the Azure admin portal

Selecting the risk detection details reveals that Microsoft labelled the logins as a 'Password spray', that is, repeated attempts to log in to the user account. This classification is likely made because password sprays are often done through a large collection of different hosts rather than through a single host.

Figure 29: Risk detection details

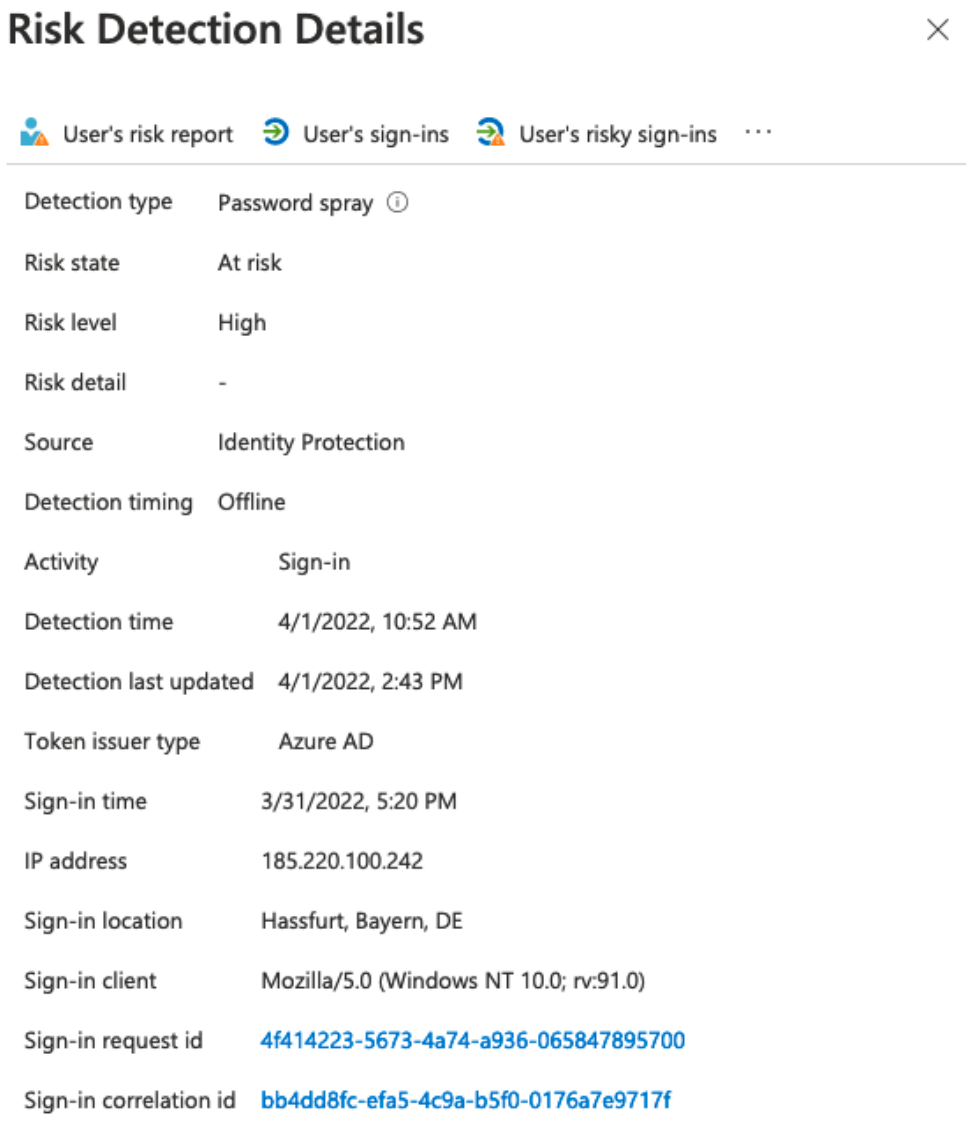

Defender 365 allows the admin to export the list of detected risky logins as a CSV or JSON file with the following output (transposed by Privacy Company for readability).

Table 2: Example of exported risky logins

Request ID	e1029815492d27508f7f8f7f6eee5aaf be568ad289d4647ebdbb4ff2616c8111
Correlation ID	bb4dd8fc-efa5-4c9a-b5f0-0176a7e9717f
Detection type	Password spray
Risk state	At risk
Risk level	High
Risk detail	-
Source	IdentityProtection
Detection timing	Offline
Activity	Sign-in

Activity time	44651,7226
Detection time	44652,4528
Detection last updated	44652,6133
User object ID	b2814ab7-0cdb-4b7c-98fe-26deea388bb3
User	Floor Terra
UPN	Floor@PCdpiaTest.onmicrosoft.com
IP address	185.220.100.242
Location	Hassfurt - Bayern - DE

4.3 Assessment

Defender does provide alerts about suspicious logins, but did not generate alerts about logins from multiple countries (through Tor exit nodes) from an end user authenticated with the Authenticator app. Apparently, Microsoft raises the trust level of a login attempt when the user is authenticated through Microsoft's own app.

When the logins from the Tor exit nodes were authenticated with SMS as second factor, Defender signalled the user was at risk.

Based on these limited tests, the service functions as expected.

In order to adequately protect the tenant against malicious logins, admins need to finetune Defender and actively monitor its alerts. Unlike Defender's ability to block files containing malware user risk detections do not necessarily lead to blocked behaviour. For certain risky behaviour types, like logging in from certain countries, admins need to take additional measures. When a user is at risk the admin needs to investigate the alert. Defender can be configured to pro-actively alert admins or specific people depending on the type and severity of the detected risks.

When user risk detections are used in conjunction with additional measures like multi-factor authentication and active monitoring by admins to follow up on alerts, this feature of Defender can be an effective security measure.

4.4 Remedies

This assessment does not lead to a recommendation for Microsoft, but organisations can take three measures to reduce possible data protection risks.

1. Instruct the admins to actively monitor for alerts on users at risk and to quickly follow up to make sure the detections are effective and the consequences of incorrect detections are minimised.
2. Create a monitoring policy to restrict how admins are allowed to use the monitoring results, and inform the users about this personal data processing and the limits on its use.
3. Consider using pseudonyms for employees whose identity should remain confidential.

5. Data minimisation options

This section answers the third question:

How does the data anonymization option work in Defender for Cloud Apps??

5.1 Facts

Microsoft offers an pseudonymisation option for user names in the dashboards from Defender for Cloud Apps. If used, username information is replaced with encrypted usernames. Microsoft calls this "*data anonymization*", but also points out the admins can resolve the real username if necessary for a specific security investigation. Such lookups/conversions are logged in the Governance log. This log allows organisations to check if admins have correctly reidentified, in line with the internal privacy policy.[36]

Microsoft describes three ways to use this option: for new snapshot reports or reports from new data sources (see Figure 30 below) or by setting the default to anonymisation under Cloud Discovery.

Figure 30: Screenshots Microsoft "anonymization" option new reports and new sources in Cloud Discovery[37]

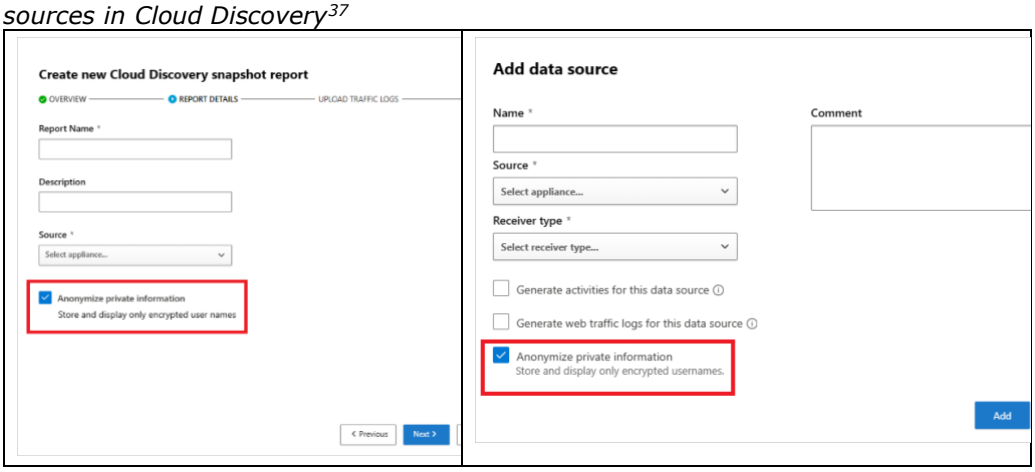

[36] Microsoft, Cloud Discovery data anonymization, 24 April 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/cloud-discovery-anonymizer, and Data security and privacy practices for Defender for Cloud Apps, 24 April 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/cas-compliance-trust

[37] Microsoft, Cloud Discovery data anonymization

Figure 31: Microsoft screenshot "anonymization" default in Cloud Discovery[38]

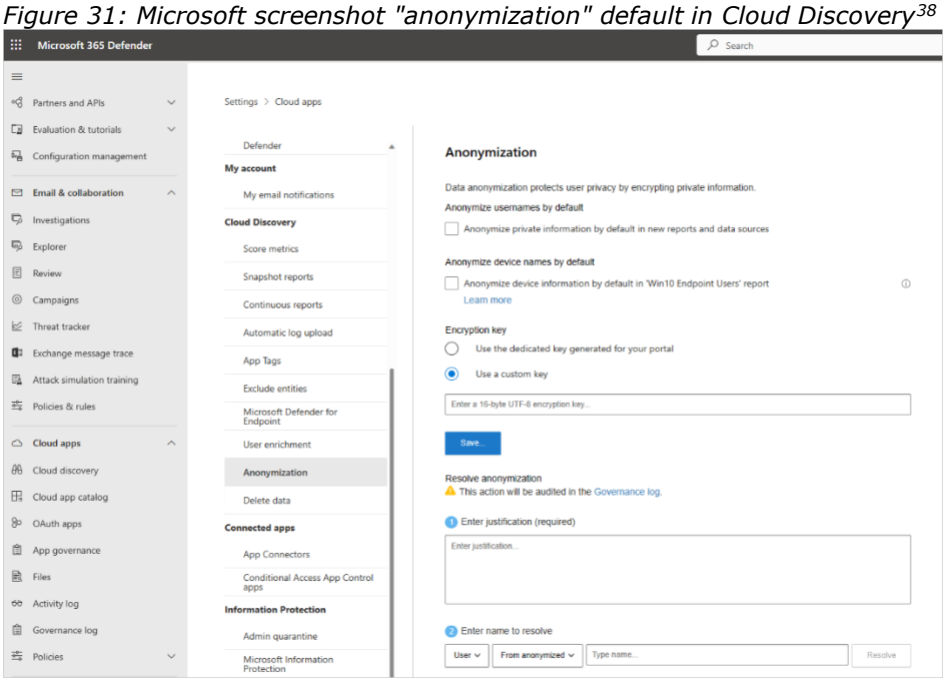

By default Microsoft does not include the identifiable user names in its cloud discovery logs. Admins can choose to include the Azure Active Directory username data in the Cloud Discovery data. This would have the opposite effect from the anonymisation option, and include the names in the logs processed by Microsoft.

Microsoft explains enabling this feature can be useful for the detection of shadow IT:

- "*You can investigate Shadow IT usage by Azure Active Directory user. The user will be shown with its UPN.*
- *You can correlate the Discovered cloud app use with the API collected activities.*
- *You can then create custom logs based on Azure AD user groups. For example, a Shadow IT report for a specific Marketing department."*

5.2 Assessment

Privacy Company has not tested these two anonymisation and identification options, but noted in 2020 that this process of hiding the user names falls under the Art. 4 (5) GDPR definition of pseudonymisation. Because the process is reversible (by the admins and by Microsoft) use of this option does not result in anonymous data. The resulting data set are still personal data, because the definition of personal data includes all data that can be directly or indirectly related to a natural person.

With regard to the directly identifying names in the 'enriched' Cloud Discovery logs, Microsoft explains in its Azure AD Whitepaper from 2020 that it does not use directly identifying data in logs for its own purposes. "*Usage data is metadata generated by the Azure AD service that indicates how the service is being used. This metadata is used to generate administrator and user facing reports and is also used by the Azure AD engineering team to evaluate system usage and identify opportunities to improve the service. This data is generally written to log files, but in some cases, is*

[38] Idem.

collected directly by our service monitoring and reporting systems. personal data is stripped out of Microsoft's usage data prior to the data leaving the originating environment."[39]

Microsoft uses the term anonymization too loosely, not in compliance with the GDPR. However, use of this pseudonymisation tool can prevent high data protection risks. Through the use of apps, end users may reveal health data, religious, political or sexual characteristics. Hiding the user names prevents admins from immediate identification of these users. Microsoft has also created audit logs that will register any attempt to reidentify the data. This is a good technical measure to ensure such reidentification is only done when justified.

If organisations want to identify specific users to detect shadow IT usage, they must ensure that access to, and use of these potentially sensitive data is strictly necessary for legitimate purposes. As this information can potentially be used in evaluation reports to negatively assess individual employees, organisations must meet a high threshold for transparency and purpose limitation.

5.3 Remedies

- Microsoft should use the term 'pseudonymisation' instead of 'anonymization'.

Organisations can take two mitigating measures:

1. Use the pseudonymisation functionality to prevent unauthorised access to sensitive characteristics of end users, derived from their app usage.
2. Prior to using the user data enrichment option, organisations must have clear and knowable rules about the circumstances when these data can be accessed and for what specific purposes. Organisations must likely involve their Workers Council when updating their internal privacy policy with these rules.

6. Traffic to third parties

This section answers the fourth verification question:

Does Microsoft send traffic to third parties (including through cookies, and to itself as an independent data controller) when a system administrator enables Defender? If so, are those third parties mentioned on the list of subprocessors?

6.1 Facts

As data processor, Microsoft may only engage authorised subprocessors to process the personal data from Dutch government organisations and universities (art 28 (3) sub d, which refers to the obligations in Art. 28 (2) and Art. 28 (4) of the GDPR.

Microsoft publishes a limitative list of subprocessors for the Online Services in its overview of Data Protection Resources, last changed on 23 November 2021.[40] The

[39] Microsoft whitepaper, Azure Active Directory Data Security Considerations, 1 July 2020.
[40] Microsoft Online Services Subprocessor List, Last updated 23 November 2021. https://servicetrust.microsoft.com/ViewPage/TrustDocumentsV3?command=Download&downloadType=Document&downloadId=926b2cf5-6b6e-43ca-9bc3-f73e961aad5f&tab=7f51cb60-3d6c-11e9-b2af-7bb9f5d2d913&docTab=7f51cb60-3d6c-11e9-b2af-7bb9f5d2d913_Subprocessor_List

most relevant change compared to the July 2020 version is the removal of the US American company *UserVoice* (specialised in digital customer research) as a subprocessor to process customer questions about the Microsoft services.[41] In 2021 Microsoft started to use its own (1st party) user discussion forums.[42]

6.2 Technical findings

Privacy Company intercepted and recorded outgoing traffic from the test VM while the Admin Console was accessed with a Chrome browser. This resulted in HTTP(S) requests to the domains listed in Table 3 below.

Requests to domains highlighted in green are the result of Chrome browser functionality and not the result of using any Defender services and are therefore out of scope. The domains highlighted in yellow are explained in more detail below the table. This includes traffic to 2 Microsoft telemetry domains, traffic to 4 Microsoft analytics domains and traffic to 6 Twitter (sub)domains.

The intercepted traffic is analysed below in four categories. Necessary functional traffic to Microsoft is not separately analysed. There is a difference between telemetry traffic to known telemetry network domains from Microsoft, and traffic related to website analytics, even though such analytics are collected from the browser in a similar way as the telemetry data. Next to these two categories, traffic to Twitter and traffic to other third parties is discussed below.

Table 3: Traffic from the Defender Admin Console

Domain	No. of requests
accounts.google.com	1
clients1.google.com	1
clients2.google.com	1
sb-ssl.google.com	1
www.google.com	20
clientservices.googleapis.com	1
content-autofill.googleapis.com	14
optimizationguide-pa.googleapis.com	2
safebrowsing.googleapis.com	3
update.googleapis.com	19
www.googleapis.com	4
clients2.googleusercontent.com	7
lh5.googleusercontent.com	2
encrypted-tbn0.gstatic.com	6
www.gstatic.com	1

[41] Microsoft's use of UserVoice was described in the June 2020 DPIA on Office for the Web and mobile Office apps for SLM Rijk, p. 48-49, URL: https://slmmicrosoftrijk.nl/?smd_process_download=1&download_id=3485.
[42] Microsoft, UserVoice pages, URL: https://support.microsoft.com/en-us/office/uservoice-pages-430e1a78-e016-472a-a10f-dc2a3df3450a.

edgedl.me.gvt1.com	19
login.live.com	1
waconafd.officeapps.live.com	4
waconatm.officeapps.live.com	2
admin.microsoft.com	37
browser.pipe.aria.microsoft.com	14
browser.events.data.microsoft.com	4
graph.microsoft.com	3
security.microsoft.com	218
login.microsoftonline.com	21
arc.msn.com	3
prod.msocdn.com	3
ecs.office.com	1
741fda1a90fc5da8166f072f2a86ddd8.fp.measure.office.com	2
8f475c157813105fc1b2e28e76fbb96a.fp.measure.office.com	2
config.fp.measure.office.com	1
upload.fp.measure.office.com	2
portal.office.com	22
webshell.suite.office.com	4
ow1.res.office365.com	2
r4.res.office365.com	1
static2.sharepointonline.com	3
abs.twimg.com	3
pbs.twimg.com	24
cdn.syndication.twimg.com	6
ton.twimg.com	2
platform.twitter.com	11
syndication.twitter.com	8
dc.services.visualstudio.com	52
spoprod-a.akamaihd.net	14
scc.azureedge.net	146
aadcdn.msauth.net	20
amcdn.msftauth.net	1
res.cdn.office.net	24
res-1.cdn.office.net	13
admincontrolsdemoapps.blob.core.windows.net	1

6.2.1 *Traffic to Microsoft analytical (sub)domains*

Requests to the subdomains at *measure.office.com* appear to collect analytical information about the use of the Compliance Centre. These requests do not contain cookies, but the requests do contain the pseudonymous identifier of the tenant, as highlighted in yellow in Figure 32 below.

Figure 32: Example of a request to measure.office.com

https://upload.fp.measure.office.com/r.gif?MonitorID=O365se&rid=c0fec6277f1366fec6f5dafbab30748f&w3c=true&prot=https:&v=20190214&tag=[{"TenantId":"bd9a989d-e990-4e6e-9566-5a8b29c3b6ff","AppId":"ComplianceCenter"}]&DATA=[{"RequestID":"8f475c157813105fc1b2e28e76fbb96a","Object":"trans.gif","Conn":"cold","Result":663,"T":128,"Rip":"2a10:3781:412::","Ep":"MDW","Fe":"cafe"},{"RequestID":"8f475c157813105fc1b2e28e76fbb96a","Object":"trans.gif","Conn":"warm","Result":215,"T":128,"Rip":"2a10:3781:412::","Ep":"MDW","Fe":"cafe"},{"RequestID":"waconafd.officeapps.live.com","Object":"trans.gif","Conn":"cold","Result":140,"T":128,"Ep":"PIE1","Fe":"WordLB1"},{"RequestID":"waconafd.officeapps.live.com","Object":"trans.gif","Conn":"warm","Result":130,"T":128,"Ep":"PNL1","Fe":"WordLB1"},{"RequestID":"waconatm.officeapps.live.com","Object":"trans.gif","Conn":"cold","Result":84,"T":128,"Ep":"PNL1","Fe":"WordLB1"},{"RequestID":"waconatm.officeapps.live.com","Object":"trans.gif","Conn":"warm","Result":18,"T":128,"Ep":"PNL1","Fe":"WordLB1"}]

As shown in Table 4 below, the full request headers do not contain any cookies.

Table 4: Full request headers to measure.office.com (no cookies)

sec-ch-ua	" Not A;Brand";v="99", "Chromium";v="99", "Google Chrome";v="99"
sec-ch-ua-mobile	?0
user-agent	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/99.0.4844.84 Safari/537.36
sec-ch-ua-platform	"macOS"
accept	*/*
origin	https://security.microsoft.com
sec-fetch-site	cross-site
sec-fetch-mode	cors
sec-fetch-dest	empty
referer	https://security.microsoft.com/
accept-encoding	gzip, deflate, br
accept-language	en-GB,en-US;q=0.9,en;q=0.8

6.2.2 *Telemetry data sent to Microsoft*

Privacy Company observed Telemetry Data being transmitted to two different domains while using the admin portal: browser.pipe.aria.microsoft.com and browser.events.aria.microsoft.com.

Table 5: Telemetry events transmitted to Microsoft

Event names	Number
Office.Taos.Shell.Impression.NavBarFull	2
Office.Taos.Shell.Monitoring	6
Office.Taos.Shell.Performance	2
Office.Taos.Shell.ServerRequest	4
awt_stats	3

generic_qos	2
impression	2
monitoring	6
performance	2
prefetch_request	2
searchbox_performance	2
serverrequest	4
session	4

The contents of two of these events are shown in more detail in Figure 33 and 34 below. The events show a hashed identity of the admin and of the organisation tenant, plus the statement that the user is not anonymous (*User.IsAnonymous": false*). This qualification is commonly used as a flag to indicate the user is logged in. Because the test account was logged in, the hashed identities are pseudonymous personal data. Microsoft is technically able to identify the admin.

Figure 33: Example event Office.Taos.Shell.Impression.NavBarFull


```
{
 "name": "Office.Taos.Shell.Impression.NavBarFull",
 "time": "2022-03-29T13:31:26.445Z",
 "ver": "4.0",
 "iKey": "o:b0c82c6598ad49f3848b1d3dc0d8dd25",
 "ext": {
  "sdk": {
   "seq": 1,
   "epoch": "1648560686441",
   "ver": "1DS-Web-JS-3.1.10"
  },
  "metadata": {
   "f": {
    "Event.Sequence": {
     "t": 4
    },
    "Event.Time": {
     "t": 9
    },
    "Data.Impression_ItemCount": {
     "t": 6
    }
   }
  }
 },
 "data": {
  "baseType": "custom",
  "baseData": {
   "properties": {
    "version": "PostChannel=3.1.10"
   }
  },
  "App.Name": "OfficeTaosShell",
  "App.Platform": "Web",
  "App.Version": "20220325.1",
```

```
    "Session.Id": "053e6039-1195-4657-bfb8-7d4a0c4f69e9",
    "Release.AudienceGroup": "Production",
    "User.PrimaryIdentityHash": "1003200138B0D6A3",
    "User.PrimaryIdentitySpace": "OrgIdPUID",
    "User.TenantId": "bd9a989d-e990-4e6e-9566-5a8b29c3b6ff",
    "User.TenantGroup": "Commercial",
    "User.IsAnonymous": false,
    "Context_Env": "NEUprod",
    "Context_Site": "ProtectionCenter",
    "Context_SiteSubId": "ShellComplianceCenter",
    "Context_WorkloadAppId": "ComplianceCenter",
    "Context_PID": "1648560686440_0.24019875064246543",
    "Context_UID": "b2814ab7-0cdb-4b7c-98fe-26deea388bb3",
    "Context_Language": "en-GB",
    "Context_Flights":
"15GA,SE404567MyDayOfficeComEnabled,SE404571FetchPhotoUsingGraph,SE404575EnableTe
nantThemeV2,SE404583OfficeHomeVisioStart,SE404593ExchangeDataOpxEnabled,SE404601E
xchangeLoginHint,SE404603OTelTelemetry,SE404609VivaInsightsRollout",
    "Context_FlightRings": "WorldWide",
    "Context_Segment": "Dynamics365,Admin,OfficeProPlus,Clp,Teams",
    "Event.Sequence": 1,
    "Event.Name": "Office.Taos.Shell.Impression.NavBarFull",
    "Event.Source": "OTelJS",
    "Event.Time": "2022-03-29T13:31:26.445Z",
    "Data.OTelJS.Version": "4.5.3",
    "Data.Event_Type": "Impression",
    "Data.Impression_Name": "NavBarFull",
    "Data.Impression_Category": "NavBar",
    "Data.Impression_Context":
"UserTheme:Mountain;ReactVersion:16.10.1;SearchUXEnabled:1;",
    "Data.Impression_ItemCount": 0,
    "Event.Id": "96ccda6f-3ec9-4d58-b318-65760320c14b.1"
  }
}
```


Figure 34: Example event Office.Taos.Shell.Monitoring


```
{
  "name": "Office.Taos.Shell.Monitoring",
  "time": "2022-03-29T13:31:26.447Z",
  "ver": "4.0",
  "iKey": "o:b0c82c6598ad49f3848b1d3dc0d8dd25",
  "ext": {
    "sdk": {
      "seq": 2,
      "epoch": "1648560686441",
      "ver": "1DS-Web-JS-3.1.10"
    },
    "metadata": {
      "f": {
        "Event.Sequence": {
          "t": 4
```

```
      },
      "Event.Time": {
       "t": 9
      },
      "Data.Monitoring_AdHoc0": {
       "t": 6
      }
     }
    }
   },
  "data": {
   "baseType": "custom",
   "baseData": {
    "properties": {
     "version": "PostChannel=3.1.10"
    }
   },
   "App.Name": "OfficeTaosShell",
   "App.Platform": "Web",
   "App.Version": "20220325.1",
   "Session.Id": "053e6039-1195-4657-bfb8-7d4a0c4f69e9",
   "Release.AudienceGroup": "Production",
   "User.PrimaryIdentityHash": "1003200138B0D6A3",
   "User.PrimaryIdentitySpace": "OrgIdPUID",
   "User.TenantId": "bd9a989d-e990-4e6e-9566-5a8b29c3b6ff",
   "User.TenantGroup": "Commercial",
   "User.IsAnonymous": false,
   "Context_Env": "NEUprod",
   "Context_Site": "ProtectionCenter",
   "Context_SiteSubId": "ShellComplianceCenter",
   "Context_WorkloadAppId": "ComplianceCenter",
   "Context_PID": "1648560686440_0.24019875064246543",
   "Context_UID": "b2814ab7-0cdb-4b7c-98fe-26deea388bb3",
   "Context_Language": "en-GB",
   "Context_Flights":
"15GA,SE404567MyDayOfficeComEnabled,SE404571FetchPhotoUsingGraph,SE404575EnableTe
nantThemeV2,SE404583OfficeHomeVisioStart,SE404593ExchangeDataOpxEnabled,SE404601E
xchangeLoginHint,SE404603OTelTelemetry,SE404609VivaInsightsRollout",
   "Context_FlightRings": "WorldWide",
   "Context_Segment": "Dynamics365,Admin,OfficeProPlus,Clp,Teams",
   "Event.Sequence": 2,
   "Event.Name": "Office.Taos.Shell.Monitoring",
   "Event.Source": "OTelJS",
   "Event.Time": "2022-03-29T13:31:26.447Z",
   "Data.OTelJS.Version": "4.5.3",
   "Data.Event_Type": "TagID",
   "Data.Monitoring_Name": "StorageApi_RequestError",
   "Data.Monitoring_Severity": "Error",
   "Data.Monitoring_AdHoc0": -2007,
   "Event.Id": "96ccda6f-3ec9-4d58-b318-65760320c14b.2"
  }
}
```

6.2.3 *Traffic to Twitter*

Figure 35 below shows that the Defender Admin Console includes content from Microsoft's Twitter security list. This inclusion leads to the transmission of data to Twitter, as highlighted in yellow in Table 2 above. Transmitted data do not include any cookies or unique identifiers, but do include the IP address the admin is connected with.

During the tests Privacy Company was not logged in to Twitter, but even if an admin is logged in to Twitter, Twitter does not collect unique identifiers through cookies or web requests.

Figure 35: Microsoft Twitter newsfeed embedded in the Admin Console

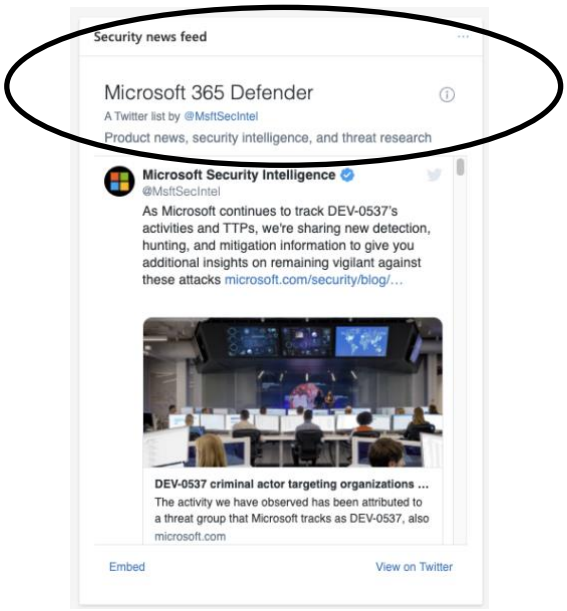

6.2.4 *Traffic to other third parties*

Privacy Company also checked if the use of Defender caused other types of third-party traffic, including to Microsoft in a role as data controller. In previous reports, a data protection risk was identified related to the use of the Feedback functionality, as the processing of these personal data is not covered by the strict purpose limitation in the privacy amendment with the Dutch government. As shown in Figure 36 and Figure 37 below, Microsoft no longer sends traffic to the third-party services of UserVoice when an admin uses the feedback form.

Figure 36: Headers traffic from Feedback form sent to office.com

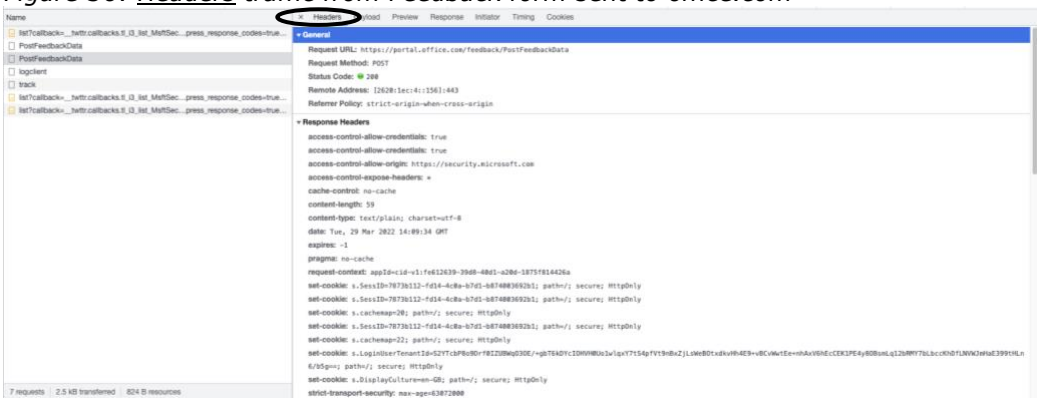

Figure 37: Payload traffic from Feedback form sent to office.com

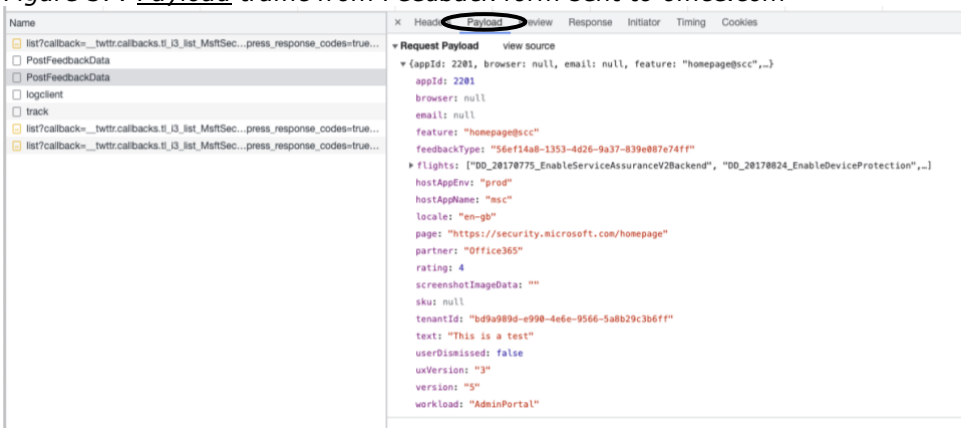

However, at the bottom of the Feedback form, Microsoft provides a hyperlink to the applicable privacy policy. This link leads to Microsoft's general (consumer) privacy statement, at https://privacy.microsoft.com/en-us/privacystatement .

Figure 38: Reference to Privacy Policy in Microsoft Feedback form

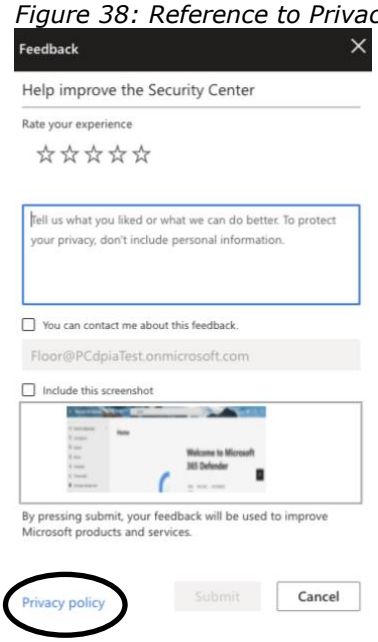

6.3 Assessment

It follows from the tests that Microsoft does not automatically send personal data to external third parties. Though Microsoft includes a Twitter feed in the Defender homepage, this does not result in the transfer of personal data to Twitter, except for the IP address, plus the context of the visit to the Defender homepage.

In reply to this verification report, Microsoft acknowledged that the traffic to Twitter was optional. Microsoft has committed to develop an option for admins to centrally block this traffic. Meanwhile, according to Microsoft, users can individually block Twitter content.

Figure 39: Illustration provided by Microsoft for individual removal of Twitter feed

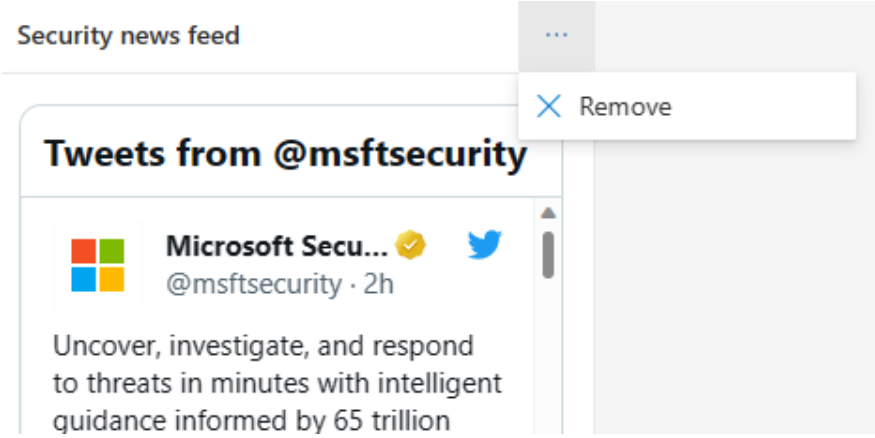

Microsoft qualifies itself as an independent data controller (and hence, also as a third party) by providing the hyperlink to its (consumer) privacy statement when an admin provides **Feedback**. This is not in line with Microsoft's commitment quoted in the 2022 public DPIA on Teams, OneDrive and SharePoint to stop sending data to third parties once a system administrator had disabled the Controller Connected Experiences.[43] However, the tests in the spring of 2021 showed that Microsoft sent data to itself as data controller (in SharePoint to search engine Bing). Microsoft replied to these findings with a commitment to disable traffic to Bing from SharePoint by mid-2022.

Even though Microsoft has replaced the third party UserVoice with its own customer feedback tools, Microsoft apparently still acts as a data controller for Feedback data. This means Microsoft contractually permits itself to process the personal data resulting from the use of the Feedback option for all seventeen purposes from its (consumer) privacy policy. This includes the display of personalised advertising. See

[43] SLM Rijk DPIA on Microsoft Teams, OneDrive, Sharepoint and Azure AD (June 2021, published 16 February 2022, p. 14-15, URL : https://slmmicrosoftrijk.nl/wp-content/uploads/2022/02/Public-DPIA-Teams-OneDrive-SharePoint-and-Azure-AD-16-Feb-2022.pdf

the previous DPIA on Office 365 for the Web and mobile Office apps, as published 30 June 2020, for a full listing of these data controller purposes.[44]

The data processing in and about the use of Feedback is not limited to the three authorised processor purposes. Microsoft has made it possible, through the programming of the Feedback form, that it receives personal data from Enterprise license end-users in a role as data controller. In doing so, Microsoft has at least partially determined the purposes for this data processing. Therefore the organisations and Microsoft can be qualified as joint controllers for the data processing via the Feedback form. Alternatively, if this disclosure does not qualify as joint processing, the Dutch government organisation that allows this data processing is disclosing personal data to a third party. That is a form of 'further processing' for which the organisation is not likely to pass the compatibility test of art. 6(4) of the GDPR.

Microsoft previously argued that it is up to end users (employees) to provide such Feedback, and therefore, that the processing does not pose a high risk for data subjects. This reasoning was not correct. There is no technical blocking option available for Feedback. Hence, organisations can only develop organisational measures such as asking end users not to use the Feedback tool. This is not likely to prevent all data processing by Microsoft, especially since Microsoft does not explicitly warn end users that use of the Feedback form may lead to breaches of work confidential information and/or personal data breaches. Microsoft's mere hyperlink to its (consumer) privacy statement does not provide sufficient warning to admins that they may not use this service.

In reply to this verification report, Microsoft explained that it "*considers an organization to be the owner of its feedback. Feedback data submitted by an organization's users is viewable and controllable by an organization's administrators.*"[45]

Microsoft refers to a new information page about the nature of Feedback data processing, the contents, the purposes and how admins can exercise control.[46]

Microsoft writes it collects and processes the Feedback data itself, and only uses it to improve Microsoft products. Microsoft explains: "*We get user feedback in the form of questions, problems, compliments, and suggestions. We make sure this feedback makes it back to the appropriate teams, who use feedback to identify, prioritize and make improvements to Microsoft products. Feedback is essential for our product teams to understand our user's experiences, and directly influences the priority of fixes and improvements.*"

Microsoft acknowledges the processing of Feedback data is completely optional. In reply to this report, Microsoft has committed to make a tenant-level configuration available to centrally disable access to from end users to the Feedback form.

---

[44] DPIA on Microsoft Office 365 for the Web and mobile Office apps, published 30 June 2020, URL: https://www.rijksoverheid.nl/documenten/rapporten/2020/06/30/data-protection-impact-assessment-office-365-for-the-web-and-mobile-office-apps.

[45] E-mail Microsoft to SLM Rijk, 14 April 2023

[46] Microsoft, Learn about Microsoft feedback for your organization, 17 February 2023, URL: https://learn.microsoft.com/en-us/microsoft-365/admin/misc/feedback-user-control?view=o365-worldwide.

In sum, the technical findings show that Microsoft does not share personal data with third parties, except with itself as controller via the Feedback functionality. However, in reply to this verification report, Microsoft has committed to allow its customers to centrally disable Twitter and Feedback.

6.4 Remedies

In order to remedy these deviations, Microsoft is advised to take the following two remedies:

1. Develop a setting to block the functionality of the Twitter feed
2. Develop a group policy or setting to centrally block the use of Feedback in Enterprise and Education tenants or become a data processor.

Additionally Microsoft must be more transparent about the Telemetry Data, but this point will be addressed separately in Section 9 of this report.

- Organisations must use the new technical opt-out functionalities designed by Microsoft to block traffic to Twitter and to Feedback.

7. Learnings from security incidents

This section answers the fifth verification question:

Does Microsoft process learnings from security incidents across its Enterprise and Education user base, or are there limitations, such as prior anonymisation?

7.1 Facts

Until 2023, Microsoft used the term 'Insights' in its documentation about Cloud App Security and Defender for Endpoint, and wrote that such information could be shared with other customers.[47]

Microsoft explained:

> *"Customer data is isolated from other customers and is not shared. However, insights on the data resulting from Microsoft processing, and which don't contain any customer-specific data, might be shared with other customers. Each customer can only access data collected from its own organization and generic data that Microsoft provides."48*

In order to create the *Insights*, Microsoft explained it removed the identifiers from the data set and qualified the resulting data as 'anonymous' data.

In 2020 Microsoft did not provide any access to the Insights in reply to the data subject access request and explained:

[47] Microsoft, URL: https://github.com/MicrosoftDocs/microsoft-365-docs/blob/public/microsoft-365/security/defender-endpoint/data-storage-privacy.mds.
[48] Idem, Answer to the question 'Is data shared with other customers?'

> *"[T]he resulting data does not retain individual-level identifiers and thus is not personal data."*

This statement did not clarify Microsoft's factual processing, as the mere removal of identifiers does not necessarily mean that the remaining data no longer are personal data. There are other ways of reidentifying data, such as when the data can be compared to data from other sources that do contain identifiers, or when the data contains sufficient details to make it possible to zoom in on an individual.

In dialogue with SLM Rijk in April 2023 Microsoft explained that the term Insights was unfortunate. Microsoft has replaced this with the term Threat Intelligence, to clarify that Microsoft only transfers the information about a new threat to its knowledge base of known threats. Microsoft confirmed in writing that it does not process any directly or indirectly identifying personal data for threat intelligence.

Microsoft now explains:

> *"Customer data is isolated from other customers and is not shared. However, threat intelligence on the data resulting from Microsoft processing, and which don't contain any customer-specific data, might be shared with other customers. Each customer can only access data collected from its own organization and generic data that Microsoft provides."*[49]

7.2 Technical findings

Privacy Company did not inspect the way Microsoft generates threat intelligence. As mentioned above, and described below, in Section 9, Microsoft did not provide any data about the Insights in reply to the data subject access requests

7.3 Assessment

Microsoft's explanation and improved public explanations have removed a potential data protection risk of further processing of pseudonymised personal data about security threats at individual customers and users outside of the customer tenant.

The processing of some pseudonymised personal data at an aggregated level (never on a per tenant basis) for the purpose of providing a secure service is not problematic. First of all, based on Articles 28 (3) sub c and 32 of the GDPR Microsoft necessarily has to process some personal data to comply with its legal obligations as a processor to protect the personal data entrusted to it by its customers with appropriate technical and organisational measures to ensure a level of security appropriate to the (global) risks. This for example includes developing and delivering security updates for all customers, based on an inventory of global threats.

Contractually the Dutch government and universities have agreed that Microsoft may 'further' process limited personal data as independent data controller for a limited list of legitimate business purposes, when proportionate.

One of these purposes is combatting fraud, cybercrime and cyber-attacks that may affect any Microsoft product or service. Microsoft has explained to SLM Rijk that it

---

[49] Microsoft Defender for Endpoint, Is data shared with other customers?, 8 February 2023, URL: https://learn.microsoft.com/en-us/microsoft-365/security/defender-endpoint/data-storage-privacy?view=o365-worldwide#is-data-shared-with-other-customers

"*does not and has not processed any security related personal data for the agreed Legitimate Business Operations purpose*."[50]

According to the privacy amendment negotiated by SLM Rijk with Microsoft, as a processor Microsoft may only process personal data to the extent necessary, and otherwise rely on pseudonymised aggregated data. The obligation to provide a secure service legitimises that Microsoft 'learns' from security incidents across its customer base, based on pseudonymised data, to develop mitigating measures such as security updates and to inform all customers about detected threats and recommended measures.

Additionally, the Dutch government and universities' privacy amendment allows Microsoft to 'further' process some personal data to combat fraud, cybercrime, or cyberattacks in a role as data controller, for Microsoft's own legitimate business purposes. In this 'controller' role, Microsoft is contractually prohibited from re-identifying pseudonymised or de-identified personal data.

In reply to this verification report Microsoft explained the use of the word 'Insights' was unfortunate. The processing only involves transferring information learned from new threats to the list with ´known threats´ Insights do not include any personal data from customers, but identify for example malware, or a malicious URL used for phishing. Microsoft has updated its information on 3 April 2023. The word "Insights" was removed and replaced with the term "threat intelligence". Microsoft explicitly documents that threat intelligence does not include any personal data from customers.[51]

Microsoft has also explained it does not and has not processed any security-related personal data for the agreed Legitimate Business Operations of combatting fraud, cybercrime or cyberattacks.[52]

An additional (security) risk is that Microsoft does not aggregate the information about the threat on a sufficiently high level to prevent leakage of critical cyber security information. For example, Microsoft recently published a threat advisory about Russian attacks on organisations in Ukraine spanning government, military, NGOs, judiciary and law enforcement and not-for-profits.[53] If Microsoft would publish a similar advisory about an attack on Dutch government institutions, this information could compromise investigations by the security services.

In reply to this possible risk, Microsoft explained "*All data transferred to the United States is maintained in protected data stores with controls to ensure that it is only processed for permitted purposes*."[54]

---

[50] E-mail Microsoft to SLM Rijk, 31 March 2023.
[51] E-mail Microsoft to SLM Rijk, 14 April 2023.
[52] Idem.
[53] Microsoft ACTINIUM targets Ukrainian organizations, 4 February 2022, URL: https://www.microsoft.com/security/blog/2022/02/04/actinium-targets-ukrainian-organizations/
[54] E-mail Microsoft to SLM Rijk, 14 April 2023.

7.4 Remedies

Following Microsoft's explanations there are no recommendations for Microsoft or for organisations.

8. Risk profiles

This section answers the sixth verification question:

Does Defender create individual risk profiles and/or individual scores in the different analytic overviews and reports?

8.1 Facts

As summarised in the Introduction Defender offers a wide variety of risk detection tools. Most of these tools are focussed on the occurrence and nature of the threats, not on individual users. There are two specific resources that show individual risk scores (i) Users at risk and (ii) Top targeted users. Besides, admins can always zoom in to the level of individual users, mailboxes or devices and create their own profiling logic.

Defender registers behaviour within Office and the Azure AD, and classifies some behaviour as risky, such as a suspicious login, clicking on a phishing link or receiving malware. The resources with 'profiles' thus consist of either passive or active risky behaviour.

As described in Section 4.1 (Detection of suspicious logins in the Azure AD) Microsoft monitors historical user login activity to warn about suspicious logins.

Microsofts overview of 'Users at risk' is based on 19 threat detection types, as shown in two screenshots in Figure 40 below.

Figure 40: Microsoft detection types

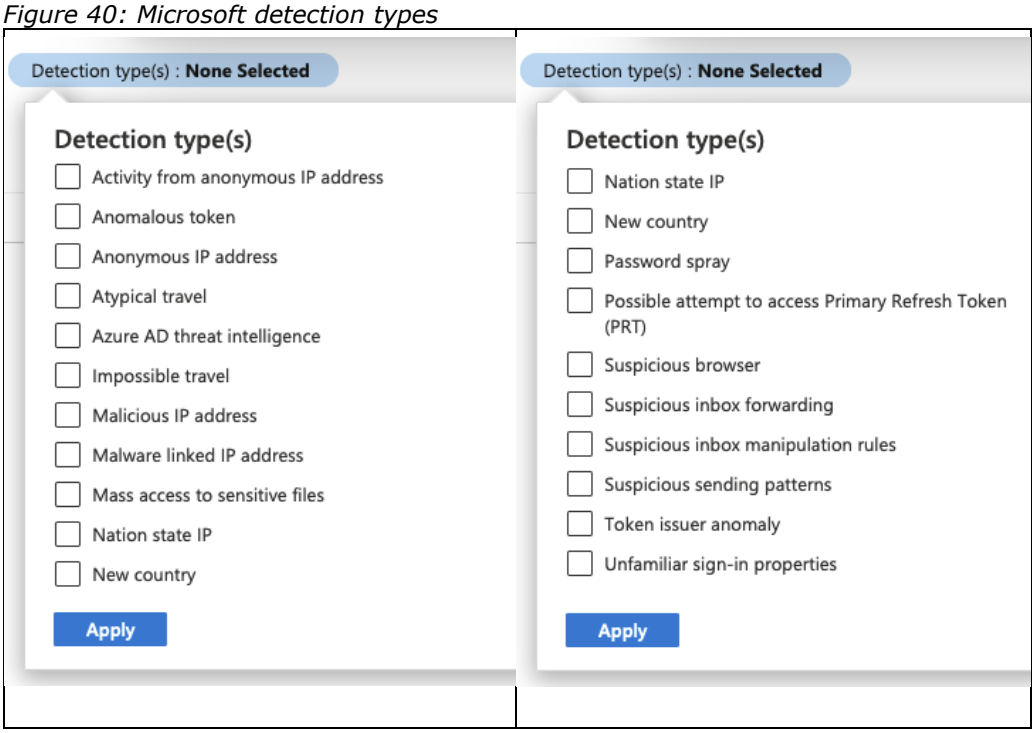

These different threat types describe generic assumptions about threats, such as an impossibly short travel time between two geolocations. The tests did not provide any indication that Defender identifies threats based on deviations of individual end user profiles. In other words, the risk profiles reflect 'known' bad behaviour patterns, as applied to individual users, but not individual behavioural profiles.

There are two apparent exceptions to this rule: (i) *Atypical travel* and (ii) *Unfamiliar sign-in properties*.

Machine learning is applied to detect 'atypical travel'. Microsoft explains that this risk is based on a combination of geographically distant locations and past user behaviour.

Microsoft writes:

> "*This risk detection type identifies two sign-ins originating from geographically distant locations, where at least one of the locations may also be atypical for the user, given past behaviour. Among several other factors, this machine learning algorithm takes into account the time between the two sign-ins and the time it would have taken for the user to travel from the first location to the second, indicating that a different user is using the same credentials.*
>
> *The algorithm ignores obvious "false positives" contributing to the impossible travel conditions, such as VPNs and locations regularly used by other users in the organization. The system has an initial learning period of the earliest of 14 days or 10 logins, during which it learns a new user's sign-in behaviour.*"[55]

The detection of unfamiliar login properties is based on a profile of the login history of a specific user and behaviour of other employees in that organisation. For example: a user logging in from a country he has never logged in from before. Microsoft may not qualify such a login as a risk if logins from that country are common for that organisation.

The second resource with user risk profiles is a tab in Explorer, as described in Section 1.5.3 of this report. This shows a graph of the 'Top targeted users', the end users that most frequently receive suspicious mails (with malware or phishing links). See Figure 13 above.

Microsoft writes in the Azure Active Directory Data Security Considerations:

> *"Azure AD Identity Protection uses real-time user login data* ***along with multiple signals from company and industry sources to feed to its machine learning systems to detect anomalous logins.*** *Personal data is scrubbed from this real-time login data before it is passed into the machine*

[55] Microsoft, Premium sign-in risk detections, 20 April 2022, URL: https://docs.microsoft.com/en-gb/azure/active-directory/identity-protection/concept-identity-protection-risks#premium-sign-in-risk-detections

learning system, along with the remaining login data used to identify users and logins that are potentially risky."[56]

The emphasised words suggest that Microsoft uses information from external sources. However, in reply to this report, Microsoft has confirmed in writing that it does not.

"Microsoft Defender for Endpoint does not use external risk profiling sources to prevent malicious login attempts. The service only uses/combines the personal data obtained by several Defender services used within one tenant, and only for the purpose of providing these services to the Customer."[57]

Admins can use the information about this labelled 'risky' behaviour to determine the threshold for decisions on individual users. As described in Section 1.5.4 of this report, Review asks admins to decide what to do with users that have been blocked by Microsoft for sending too many messages classified as bulk mail. See Figure 16 above.

For example: if a specific user receives malware, Microsoft automatically puts the mail in a quarantine folder. Admins can decide if they allow end users to view the quarantined mails, or not.

8.2 Technical findings

Privacy Company did not inspect the data processing performed by Microsoft inside the Defender services. As shown in Section 9, about the results of the data subject access requests filed for this verification report, Microsoft did not provide any data about the algorithms used to assign a risk score.

8.3 Assessment

Defender needs to process data from Microsofts raw logs such as locations of users' logins, device information and IP addresses to detect and mitigate risks and threats. Defender needs to have a flexible data processing scope to detect new and emergent risks. Defender also enables the admins to proactively hunt for risks specific to the organisation. The admins therefore also have a flexible data processing scope.

It is reassuring that Microsoft provides a description of its detection algorithms and the specific risk profiles that are applied. The descriptions are aimed at admins, not at end users. To ensure that the employees and students understand the nature of the data processing, organisations must bridge the information gap and explain in more accessible language what Microsoft does and does not do with user behaviour data, and how the organisation uses the risk profiles and alerts.

8.4 Remedies

With regard to the risk profiles organisations are advised to take the following remedy:

---

[56] Microsoft, Azure Active Directory Data Security Considerations, 1 July 2020, p. 11, URL: https://azure.microsoft.com/en-us/resources/azure-active-directory-data-securityconsiderations/.

[57] E-mail Microsoft to SLM Rijk, 18 April 2023.

- Provide a concise, intelligible and easily accessible internal explanation to employees about the data processing in Defender

9. Quality of public information

This section answers the seventh verification question:

Does Microsoft publish adequate documentation on the personal data it collects through the tested applications, in comparison with captured network traffic and logs that are accessible for system administrators??

9.1 Facts

As data processor for Microsoft 365 Enterprise (and Education), Microsoft must assist the data controllers (the Dutch government organisations and universities) with their obligation to adequately inform end users about the data processing, based on Art. 14 of the GDPR.

Microsoft does publish documentation about the nature of the data processing through Defender and an introduction to the concept of 'risk'. Until 2023 Microsoft provided limited information about the retention periods and did not publish information about the Telemetry Data it collects from browsers from admins when they use the admin portal.

Microsoft only provided brief information about the retention periods in a FAQ for admins.[58]

- *"The risky users view shows a user's risk standing based on all past sign-ins.*
- *The risky sign-ins view shows at-risk signs in the last 30 days.*
- *The risk detections view shows risk detections made in the last 90 days."*

As shown in Figure 41 below, the default retention period for the sign-ins is 180 days (based on Defender for Endpoint settings).

[58] Microsoft, How far can I go back in time to understand what's going on with my user?, URL: https://docs.microsoft.com/en-gb/azure/active-directory/identity-protection/troubleshooting-identity-protection-faq#how-far-can-i-go-back-in-time-to-understand-whats-going-on-with-my-user.

Figure 41: Default retention period

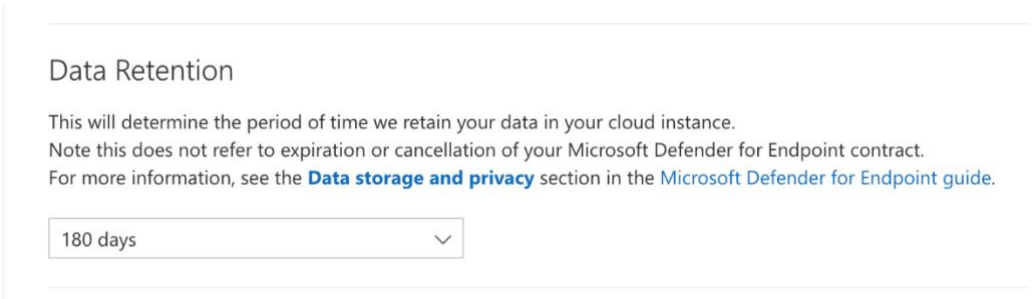

In 2023 Microsoft has improved transparency about the specific personal data it collects through the four investigated tools, and the data retention periods.

9.1.1 *Retention periods in Defender for Office 365*

By default, Microsoft retains the Defender for Office 365 data for a maximum of 30 days. Depending on the type of Plan the customer has, data can be retained up to 180 days for the Action Center.[59]

Figure 42: Microsoft table of retention periods in Defender for Office 365[60]

Defender for Office 365 Plan 1

Feature	Retention period
Alert metadata details (Microsoft Defender for Office alerts)	90 days.
Entity metadata details (Email)	30 days.
Activity alert details (audit logs)	7 days.
Email entity page	30 days.
Quarantine	30 days (configurable; 30 days is the maximum).
Reports	90 days for aggregated data. 30 days for detailed information.
Submissions	30 days.
Real-Time detections	30 days.

9.1.2 *Retention periods in Defender for Endpoint*

Microsoft writes that data from Microsoft Defender for Endpoint are "*retained for 180 days, visible across the portal. However, in the advanced hunting investigation experience, it is accessible via a query for a period of 30 days.*"[61]

Microsoft describes that it stores the Defender for Endpoint data "*in a customer dedicated and segregated tenant specific to the service for administration, tracking, and reporting purposes.*"[62]

---

[59] Microsoft, Data retention information for Microsoft Defender for Office 365, 13 March 2023, URL: https://learn.microsoft.com/en-us/microsoft-365/security/office-365-security/mdo-data-retention?view=o365-worldwide.

[60] Idem.

[61] Microsoft, Microsoft Defender for Endpoint data storage and privacy, 8 February 2023, URL: https://learn.microsoft.com/en-us/microsoft-365/security/defender-endpoint/data-storage-privacy?view=o365-worldwide

[62] Ibid.

The data include: "*file data (such as file names, sizes, and hashes), process data (running processes, hashes), registry data, network connection data (host Ips and ports), and device details (such as device identifiers, names, and the operating system version).*"[63]

9.1.3 Retention periods in Defender for Cloud Apps

Microsoft discerns between 4 types of logs.

> "*Activity log: 180 days*
>
> *Discovery data: 90 days*
>
> *Alerts: 180 days*
>
> *Governance log: 120 days*"[64]

9.1.4 Retention periods Defender for Identity

Microsoft does not publish separate retention periods for Defender for Identity. as these periods are related to the retention periods for the Azure AD determined by the organisations. Microsoft only mentions the retention period of 90 days for the audit log.

> *"Defender for Identity implements the audit of personal data changes, including the deleting and exporting of personal data records. Audit trail retention time is 90 days. Auditing in Defender for Identity is a back-end feature and not accessible to customers.*"[65]

Additionally, Microsoft does not publish information about the two types of Telemetry Data described in Section 6.2 of this report: the traffic to the analytical subdomains and the traffic to the two known telemetry domains. The contents of this traffic are limited because the data only relate to the use of the admin portal.

9.2 Technical findings

The intercepted Telemetry Data are described in Sections 6.2.1 and 6.2.2. As described above, Microsoft does not publish documentation about this category of personal data. Microsoft also did not provide access to these data in response to a Data Subject Access Request. This will be specified in Section 10 of this report.

9.3 Assessment

Microsoft generally collects the personal data it processes in Defender in an indirect manner (observed behaviour). This means Microsoft must enable its customers to comply with the indirect information obligation from Art. 14(1) sub d of the GDPR As data processor, Microsoft is obliged to comply with art 28(3) sub e of the GDPR, and

---

[63] Ibid.

[64] Microsoft, Data security and privacy practices for Defender for Cloud Apps, 24 April 2023, URL: https://learn.microsoft.com/en-us/defender-cloud-apps/cas-compliance-trust#data-retention.

[65] Microsoft Defender for Identity, Microsoft Defender for Identity data security and privacy, 5 February 2023, URL: https://learn.microsoft.com/en-us/defender-for-identity/privacy-compliance/.

assist its Enterprise customers with the exercise of data subjects rights such as the right to information.

Microsoft does not provide the required assistance with regard to the retention periods and the Telemetry Data.

9.4 Remedies

In order to remedy the lack of transparency, Microsoft is advised to take the following remedy:

- Microsoft must provide more information about the observed Telemetry Data from the admin portal, unless Microsoft is able to ensure that the browser telemetry data do not contain any identifying (pseudonymous) personal data.

10. Data Subject Access

This section answers the (eight) verification question:

Does Microsoft give system administrators full access to all personal data it processes through the different Defender tools? Does Microsoft provide adequate explanations if it does not provide access to certain personal data?

10.1 Facts

Microsoft offers two different tools for data subject access requests:

1. A DSAR tool to retrieve Content Data from a specific user.[66] This tool was renamed in September 2021 to User data search.[67]
2. A DSAR tool to retrieve Diagnostic Data from a specific user.[68]

This first tool (shown in Figure 43 and Figure 44) only provides access to Content Data in SharePoint, Exchange and Teams. That information does not specifically relate to any security processing by Microsoft and is therefore not useful for this report. Additionally, the export does not contain any files stored in OneDrive, while malicious files such as the Eicar test-file were stored in OneDrive, and cannot be retrieved via this portal.

[66] Microsoft, Data Subject Requests, Part 1: Responding to DSRs for Customer Data, Using the Content Search eDiscovery tool to respond to DSRs, URL: https://docs.microsoft.com/en-us/compliance/regulatory/gdpr-dsr-Office365#part-1-responding-to-dsrs-for-customer-data

[67] Microsoft, User Data Search, URL: https://compliance.microsoft.com/userdatasearch

[68] Microsoft, Data Subject Requests, Part 3: Responding to DSRs for system-generated Logs, URL: https://docs.microsoft.com/en-us/compliance/regulatory/gdpr-dsr-Office365#part-3-responding-to-dsrs-for-system-generated-logs. This tool can only be used by the tenant admin.

Figure 43 Export of Content Data from SharePoint, Exchange and Teams

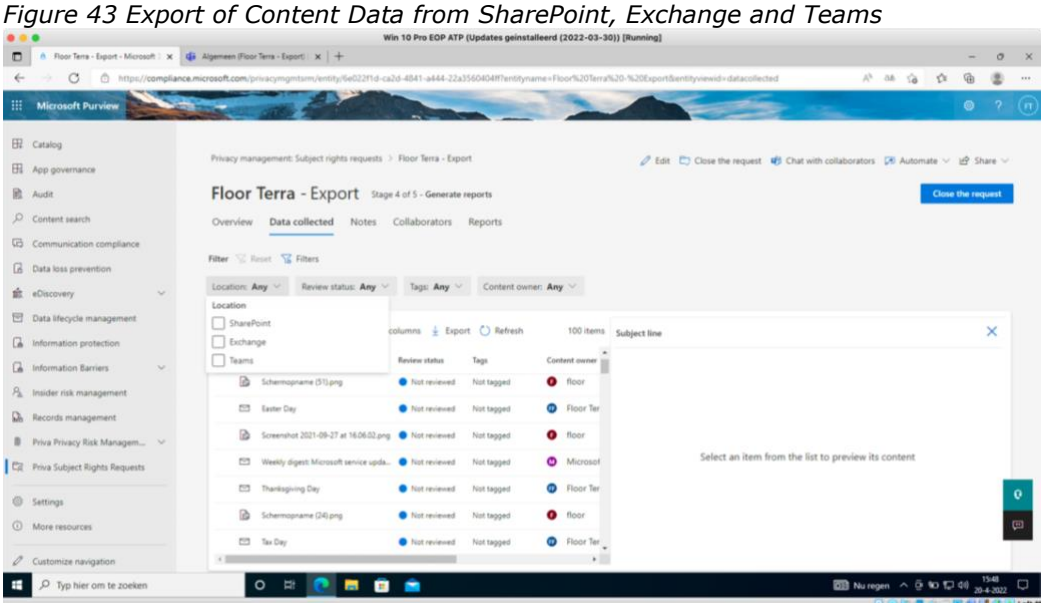

Figure 44: Content export tool does not find Eicar test file in OneDrive

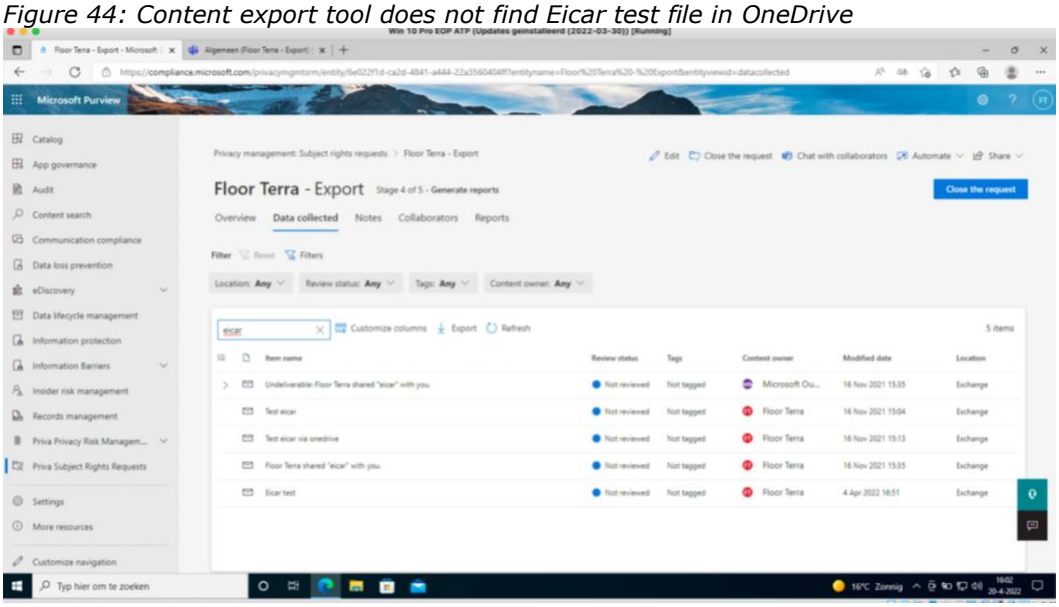

The second tool is relevant for this report. As shown in Figure 45 below, Microsoft explains that the admin needs to create a separate Azure blob. The use of Azure storage is not included in the (most expensive) E5 Enterprise license.

Figure 45: Microsoft explanation of the DSAR tool for diagnostic data

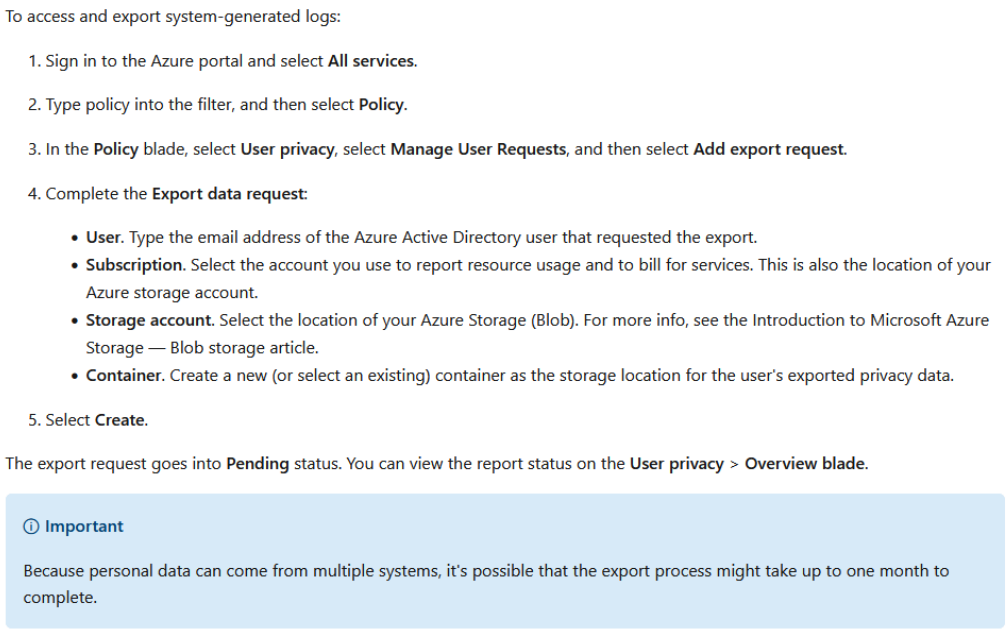

To access and export system-generated logs:

1. Sign in to the Azure portal and select **All services**.
2. Type policy into the filter, and then select **Policy**.
3. In the **Policy** blade, select **User privacy**, select **Manage User Requests**, and then select **Add export request**.
4. Complete the **Export data request**:
   - **User**. Type the email address of the Azure Active Directory user that requested the export.
   - **Subscription**. Select the account you use to report resource usage and to bill for services. This is also the location of your Azure storage account.
   - **Storage account**. Select the location of your Azure Storage (Blob). For more info, see the Introduction to Microsoft Azure Storage — Blob storage article.
   - **Container**. Create a new (or select an existing) container as the storage location for the user's exported privacy data.
5. Select **Create**.

The export request goes into **Pending** status. You can view the report status on the **User privacy** > **Overview blade**.

ⓘ **Important**

Because personal data can come from multiple systems, it's possible that the export process might take up to one month to complete.

Microsoft explains that an export usually takes 1 or 2 days, but may take up to 20 days. Microsoft warns that exported data "*do not include data that may compromise the security or stability of the service.*"[69] This warning corresponds with Microsoft's earlier statement that *service-related event level information regularly includes confidential security or other proprietary information about the operation of our services, where publication of such data could put our services, and thus our customers and Microsoft, at risk.*"[70]

Microsoft publicly commits to provide access to the *Required Service Data* in response to a Data Subject Access Request.

Figure 46: Microsoft commitment to provide Data Subject Access[71]

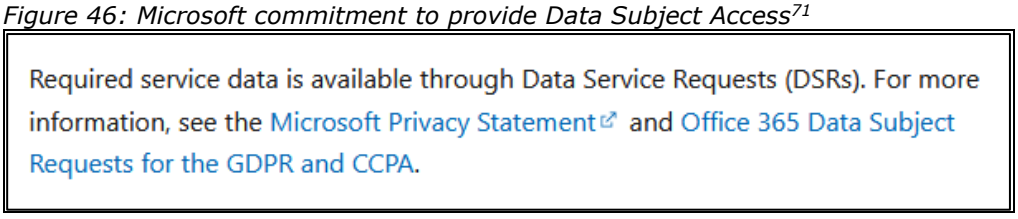

Required service data is available through Data Service Requests (DSRs). For more information, see the Microsoft Privacy Statement and Office 365 Data Subject Requests for the GDPR and CCPA.

10.2 Technical findings

In 2020 Privacy Company filed data subject access requests about Cloud Apps Security (now; Defender for Cloud Apps) and about Defender for Endpoint. Microsoft responded that only the given identifier was found in their systems, but no other personal data was discovered.

---

[69] Microsoft, What data do the export results return?, URL: https://docs.microsoft.com/en-us/compliance/regulatory/gdpr-dsr-Office365#part-3-responding-to-dsrs-for-system-generated-logs

[70] Microsoft reply to Teams, OneDrive and SharePoint DPIA for SLM Rijk, 27 June 2021.

[71] Microsoft, Required service data for Office, 30 September 2021, URL: https://docs.microsoft.com/en-us/deployoffice/privacy/required-service-data

This answer was in contrast with the intercepted network data sent by Defender for Endpoint and used by Cloud Apps Security, as well as known data sources in connection with Azure AD.

In correspondence following these 2020 data access requests, Microsoft stated that it did not have any personal data, since it deidentifies, pseudonymises and aggregates the data in compliance with ISO/IEC 19944 standards. Microsoft also argued the Device ID was not personal data, as it could only be linked to a device, not to an individual.

The DSAR exercised in 2022 for Office 365 and Defender for Identity resulted in an export to an Azure Blob Container and contains a large amount of files. Each of these files may contain many different events (actions) for different applications, but with a different structure, and often with incomprehensible random strings of characters as file names.

It took Microsoft one month to complete the DSAR request. The request was filed on 20 April 2022, and the request was completed on 20 May 2022.

Privacy Company performed a search for the observed telemetry events in the DSAR results and could not identify any of the observed events.

Microsoft does not provide any explanation about the exported results, nor a tool to structure or query the contents per app, per action or per time period. There is no explanation about types of information that has been withheld for security or confidentiality reasons.

The DSAR tool doesn't allow for filtering based on time and does contain events from outside the scripted test-scenario's, for example events on the use of Teams prior to this technical verification report, in the same tenant ID.

Table 6: DSAR Export results

Filename	**Size (bytes)**
RequestInfo.json	641
c6df84882b434278bf2a0f6517a52c68/3134bf67a6a9405fb3344e5f287f7e87/1010/Microsoft365_Outlook_OWA_ProductAndServiceUsage.json	575
2dbe5f8b339840ac9bced223f8817843/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
ca137276a7ac45458db97477cd0100db/3134bf67a6a9405fb3344e5f287f7e87/1565/ProductAndServiceUsage_OfficeApps.json	2115802
e611d051f7af45b6af02964e88f4d7aa/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
08f660c9975847a4a19af2c90512c5cf/7b35fc36d9b54e468b6a78d846a9edb5/1005/ProductAndServiceUsage.json	180586
a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Microsoft365_AdminCenter_Online_Account_UserLoginInfo.json	178
a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Microsoft365_AdminCenter_Online_Account_RememberPreferences.json	296

a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Microsoft365_AdminCenter_Online_Account_WhatsNewPreferences.json	334
a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Microsoft365_AdminCenter_Online_Account_MagicCarpetUserPreferences.json	118
a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Microsoft365_AdminCenter_Online_Account_FeatureExplorerViewModePreferences.json	89
a4370d95fb614032be19072360bb9552/6a3113dcd76843338eba1569035bd46f/1005/Microsoft365_AdminCenter_Online_CustomerContent_TID_BD9A989D-E990-4E6E-9566-5A8B29C3B6FF.json	255
95e26a9d512c4f02a3fd813765995c89/e7cd9bc9a1d14c56977ad40227589995/1005/Feedback And Ratings.json	4
00d56b4a81e04279ac3a9b7f036e1d5a/76cadbc9b7774b54be2194ca0b631c36/0/ProductAndServiceUsage.json	4
400eccd00d6449e7a1799a361d158bd8/0614e334db47470498116ce4fa328e36/225/ProductAndServiceUsage.json	248
6970f241ab054064844c106e1d290cb2/3134bf67a6a9405fb3344e5f287f7e87/1010/Microsoft365_Exchange_OutlookAddinStore_ProductAndServiceUsageData_1.json	494
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12711242.json	15162
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3517510.json	43519
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3868816.json	43004
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10042482.json	126442
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3759676.json	138696
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9569148.json	8591
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/72166.json	776562
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11714203.json	150207
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3382983.json	1166084
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/15569482.json	311
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/253221.json	8764
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3575199.json	26835
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3655982.json	190974
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10520228.json	1489
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/8725734.json	348122
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/192269.json	5107364
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/253220.json	784
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4584684.json	3862
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13584527.json	87417

a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9835187.json	194207
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13148361.json	33276
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3590448.json	62403
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/192242.json	48679
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/2385068.json	290307
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12834694.json	781
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11714213.json	46586
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10653902.json	471991
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11666262.json	129657
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3655985.json	721
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12124541.json	12372
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/283666.json	47736
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3575202.json	216748
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/14610187.json	20338
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/2748875.json	86793
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3922370.json	1155
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/111276.json	2301
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3590445.json	8764
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3389599.json	14436
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13247049.json	471
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9923815.json	750
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10830189.json	218595
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13419687.json	1637
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/286128.json	48534
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13148329.json	321746
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4569249.json	6183
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13267215.json	743

a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3274089.json	202248
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/192276.json	125376
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3575200.json	207706
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/6328826.json	743
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12883190.json	4488
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/2162603.json	6691
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10487448.json	750
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3518376.json	855
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9569208.json	393
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3590450.json	784
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3753300.json	2867031
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4839768.json	2543
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/647652.json	1711
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13259472.json	6048778
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4569511.json	24246
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11714206.json	12573578
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/8375687.json	9156
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13252635.json	521
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10748310.json	11811914
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/2385483.json	78312
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4662187.json	128666
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4201140.json	249490
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12670722.json	493495
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10584507.json	102182
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11557923.json	6165
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/12015245.json	14844
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/4581962.json	5409
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13356765.json	8293
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9837044.json	418091

a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3389420.json	1128
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3656010.json	2791567
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10525443.json	6038
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3656011.json	2805454
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11714216.json	445650
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/3742824.json	1881
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/11714236.json	8786
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13607420.json	869
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13219196.json	786
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/9562245.json	917
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/13247057.json	521
a42a5d8f0c984421955d01a4e24b5336/c814ff83c1324f809ff3d7fa05a821f7/10265806.json	17104
135e4fd3170e4b44beb0b270c60f289e/7b35fc36d9b54e468b6a78d846a9edb5/1005/ProductAndServiceUsage.json	180586
af1378eb1f7a4fcdb7c3b2cb71053bb5/efde18ca75354e008cf4fa343e3e7b31/commandId.txt	165
bfd06a9ff6e243ffaf0b13f385768313/7b35fc36d9b54e468b6a78d846a9edb5/1005/ProductAndServiceUsage.json	180586
ea3537d4ba704f589610a4a7a8a8e8a8/7b35fc36d9b54e468b6a78d846a9edb5/1005/ProductAndServiceUsage.json	180586
ba149b6472004e638175b98ca076318c/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
abab1cc2bcac4dbd9d4018a387fe7637/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
d364c307aafe4289b78b28125a20e3aa/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
37a0e5bcebe0456d907041c22341ed81/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165
a0c8f27fe183466b820da95d8be178dd/3134bf67a6a9405fb3344e5f287f7e87/commandId.txt	165

Microsoft provides two types of DSAR files: the first category contains information about the individual use of services. The second category is 'other', with many different types of information.

First category

The first category is called 'ProductAndServiceUsage'. These are not the Telemetry Data sent from the end user browser or device, but the Diagnostic Data directly generated by Microsoft itself, on its cloud servers. None of these events contain any information relating to the use of Defender. Therefore, these data are out of scope of this report. The only type of events in this category are log-ins of the admin to the Azure AD, as shown in Figure 47 below. Directly identifiable data in this event have

partially been pseudonymised or replaced with generic terms such as 'REDACTED' or 'POSSIBLE_ARM_OBJECT_ID', but the tenant domain and userDisplayName are still transferred in the clear, as highlighted in yellow.

Figure 47: Example of event related to use of Azure AD


```
{
 "time": "",
 "correlationId": "",
 "properties": {
  "PreciseTimeStamp": "2022-04-12 14:02:28.1861400",
  "clientTime": "2022-04-12 14:02:18.1940000",
  "clientTimeZone": "-120",
  "source": "BladeOpened",
  "extension": "Microsoft_AAD_IAM",
  "name": "RiskDetectionDetailsBlade",
  "assetType": "",
  "action": "Click",
  "actionModifier": "mark",
  "data":
"{parameters:{correlationId:_LTR_POSSIBLE_ARM_OBJECT_ID_,requestId:_LTR_P
OSSIBLE_ARM_OBJECT_ID_,id:e1029815492d27508f7f8f7f6eee5aafbe568ad289d46
47ebdbb4ff2616c8111,location:{state:Bayern,countryOrRegion:DE,country:DE,city:
Hassfurt},source:IdentityProtection,userPrincipalName:_REDACTED_EMAIL_@PCdpi
aTest.onmicrosoft.com,userType:0,additionalInfo:[{value:Mozilla/5.0 (Windows NT
10.0; rv:91.0),field:userAgent}],lastUpdatedDateTime:2022-04-
01T12:43:07.4399802Z,activity:0,crossTenantAccessType:0,currUserIdFilterValue:_
REDACTED_CURRENT_USER_OBJECTID_,detectionTimingType:3,activityDateTime:2
022-03-31T15:20:33.6600585Z,resourceTenantId:null,detectedDateTime:2022-04-
01T08:52:01.8730000Z,userDisplayName:Floor
Terra,tokenIssuerType:0,riskEventType:passwordSpray,userId:_REDACTED_CURRE
NT_USER_OBJECTID_,homeTenantId:_LTR_POSSIBLE_ARM_OBJECT_ID_,riskDetail:
0,ipAddress:185.220.100.242,riskState:1,riskLevel:2,riskType:0,asUtc:false}}"
 }
},
```


Second category

It is not possible to summarise the contents of the other types of events provided in reply to the DSAR request with a few common characteristics.

The Telemetry Data are encoded in several different formats, but the structure of each event is always similar. Each outgoing event in the data traffic has a header of standardised fields, and some contents that are unique for that type of event. The header contains the event name, date and timestamp and some other fields. That makes it possible to analyse the entire flow of Telemetry Data by event type.
The DSAR files in this second category however do not contain any such headers or otherwise identifiable common structure to distinguish per event type. Therefore, it is not possible to structurally compare the observed outgoing Telemetry Data with the results in this second category.

Privacy Company performed a query (grep) in the DSAR results for all outgoing Telemetry Data in the intercepted network traffic, and did not find any match. It is plausible that Microsoft only provides access to server-side generated diagnostic data via the DSAR tool, and not to any of the received telemetry data from the end user device. It follows that the DSAR is incomplete, as Microsoft does not provide any access to the Telemetry Data that are clearly being sent to Microsoft in the outgoing

data traffic. Microsoft has offered several reasons for this omission to the Dutch government, explained in the assessment below.

Many types of events contain pseudonymised identifiers for the user and for the tenant, such as "UserId": "b2814ab7-0cdb-4b7c-98fe-26deea388bb3", as shown in Figure 48 below. None of the events contain the names of tested malware, or malicious mails and websites.

Figure 48: Example of telemetry event with pseudonymised identifiers


```
{
 "time": "",
 "correlationId": "",
 "properties": {
  "CountryCode": "NL",
  "Lcid": "EN-GB",
  "PreferredLanguage": "en-GB",
  "SnapshotDate": "2022-04-01T00:00:00.0000000",
  "TenantId": "bd9a989d-e990-4e6e-9566-5a8b29c3b6ff",
  "UsageLocation": "NL",
  "UserID": "b2814ab7-0cdb-4b7c-98fe-26deea388bb3",
  "WindowsLiveNetId": "1003200138B0D6A3",
  "WindowsLiveNetIdLong": "11538011191551484 51"
 }
},
```


Answering the DSAR took considerably longer than earlier requests filed by Privacy Company with Microsoft: previously, in the same minimalistic test environment, Microsoft was able to complete the output within a couple of days. As shown in Figure 49 below, Microsoft tells requesters it usually is able to comply within a few days, but warns it may take longer, without specifying the reasons.

Figure 49: Microsoft information about time to answer DSARs

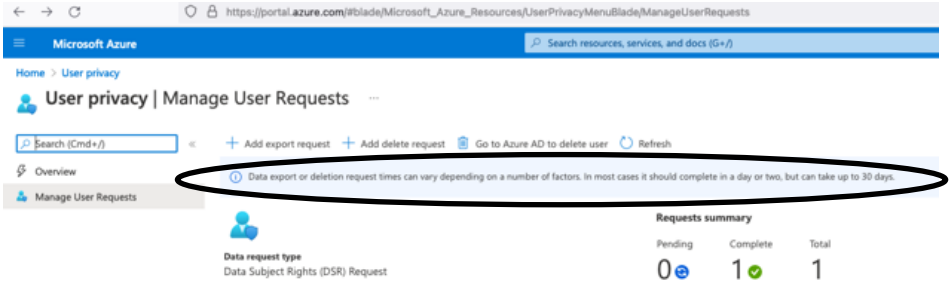

10.3 Assessment

Microsoft's assertion in 2020 that the Device ID was not personal data, was incorrect. Devices protected by organisations with an Enterprise or Education license are generally only used by a single user. In the test tenant Privacy Company could guarantee this single-user usage. This means the Device ID can be related to an individual and all data connected to the Device ID have to be qualified as personal data.

In reply to this report, in April 2023 Microsoft has confirmed this reasoning, and will soon publish an overview of identifiers collected in Defender logs.

As explained in the Technical Findings, Microsoft's DSAR results for the Diagnostic Data are a hodgepodge of files with varying structures. This does not comply with the Art. 12 GDPR requirements to provide the requested information in an *intelligible and easily accessible form*.

Microsoft clearly has processes in place to replace directly identifying personal data such as email addresses and confidential data such as document names with generic replacements, or specific pseudonymous identifiers. These processes are not yet flawless: as shown in Figure 47 the events may still contain some direct identifiers. This is not in line with Microsoft's assurance that the Diagnostic Data should not contain any identifiable user data, different from the functional data that are necessary to provide the contracted (cloud) services, and apart from a specific exception for OneDrive.[72]

Events in the second category of 'other' data only appear to contain pseudonymised identifiers. None of the DSAR files contain names of e-mails or documents with malware or URLs of malicious websites,

At first sight, it appears Microsoft does not provide access to any of the intercepted Telemetry Data, only to the server-side generated data. This conclusion is based on the fact that none of the outgoing telemetry events intercepted from the outgoing data traffic appear in the DSAR results.

It is possible that Microsoft applied a data minimisation procedure to the incoming telemetry events, and only gave back minimised contents (without the header) of the stored Telemetry Data. However, even if such a procedure were applied, the results do not comply with the requirements for Data Subject Access as defined in Article 12 (form and comprehensibility) and 15 (right to access) of the GDPR.

Microsoft explains that it routinely deletes the personal data that are not Customer Data at most within 180 days.[73] This includes all types of Diagnostic Data (telemetry and server-generated service logs). It is possible that some telemetry data were already deleted during the 30 days Microsoft took to provide the requested access.

In view of the very limited activities deployed in the test tenant, and the fact that only 3 users (1 admin and 2 end users) were configured, the period of 30 days to answer the request seems overly long. Moreover, this time period puts Microsoft's customers, the government organisations that are the data controllers, in an impossible position to answer data subject access requests in time. At most, the controller will have one or two hours left to provide the data subject with Microsoft's answers to stay within the primary GDPR deadline of 30 days. Of course the employer can communicate that an extension of this deadline is necessary, but there does not appear to be a legitimate reason for such an extension.

---

[72] Microsoft letters of 8 and 27 June 2021, in reply to the second technical verification report on the core Office 365 Suite as mobile apps and as Office for the Web, and in reply to the findings of the first Cluster DPIA on Teams, OneDrive, SharePoint and the Azure AD on all platforms.

[73] Microsoft, Data retention, deletion, and destruction in Microsoft 365, 17 November 2021, URL: https://docs.microsoft.com/en-us/compliance/assurance/assurance-data-retention-deletion-and-destruction-overview

The lack of access to the Telemetry Data runs afoul of Microsoft's public commitment to provide access to the Required Service Data, as quoted in Figure 46 above. Such access is even more necessary for data subjects to exercise their rights now that Microsoft does not want to publish event-level information, nor provide access through a Data Viewer Tool, because this would "*include confidential security or other proprietary information about the operation of our services, where publication of such data could put our services, and thus our customers and Microsoft, at risk*."

Apparently, with this argument, Microsoft claims that it can rely on an exception to the GDPR transparency rights in Art. 23 (1) sub i of the GDPR, as translated for controllers in the Netherlands in Art. 41(1) sub i of the UAVG. This provision allows data controllers not to provide information about, or access to personal data as far as this is necessary and proportionate to safeguard the rights and freedoms of others. In order to successfully claim that the right or freedom of another, such as Microsoft itself, prevails over the rights of the requesting data subject, Microsoft must evidence what important interests necessitate the exception and why these outweigh the rights of the requesting data subject.[74] However, Microsoft does not provide any public information about this categorical refusal.

Generally speaking, nor controllers nor data processors can successfully claim company confidentiality to refrain from disclosing what categories of personal data they process. As the EDPB notes in its Guidelines on the restrictions under Art. 23 GDPR[75] restrictions on the data subjects rights can only be exceptional, not structural, and a strict necessity and proportionality test is required:

> *"The case law of the CJEU applies a strict necessity test for any limitations on the exercise of the rights to personal data protection and respect for private life with regard to the processing of personal data: 'derogations and limitations in relation to the protection of personal data (..) must apply only insofar as is strictly necessary'.*[76] *The ECtHR applies a test of strict necessity depending on the context and all circumstances at hand, such as with regard to secret surveillance measures."*[77]

In sum, in its role as data processor, Microsoft is obliged to assist its Enterprise customers that are the data controllers with any requests for access ex Art. 15 GDPR. Microsoft has created two dedicated tools for admins to help them answer such requests. One tool to retrieve Content Data, and a second tool to retrieve Diagnostic Data relating to a specific user. However, the results of this second tool do not comply with Microsoft's processor commitments in multiple ways. First of all, the format and structuring of the data are not uniform or intelligible, while (second) the contents structurally omit many categories personal data, are stripped of context and cannot

[74] See for example the ruling from the administrative court Utrecht from June 2020, par. 19, URL: https://uitspraken.rechtspraak.nl/inziendocument?id=ECLI:NL:RBMNE:2020:2222

[75] EDPB, Guidelines 10/2020 on restrictions under Article 23 GDPR, Version 1.0, Adopted on 15 December 2020, URL: https://edpb.europa.eu/sites/default/files/consultation/edpb_guidelines_202010_article23_en.pdf

[76] The EDPB refers to CJEU, judgment of 16 December 2008, case C-73/07, Tietosuojavaltuutettu v. Satakunnan Markkinapörssi Oy and Satamedia Oy, ECLI:EU:C:2008:727, paragraph 56.

[77] The EDPB refers to ECtHR, Szabo and Vissy v. Hungary, 12 January 2016, paragraph 73.

be compared to any telemetry events and third, Microsoft does not provide a specific explanation why some data are not provided due to a legal exception, such as infringement on the rights of Microsoft.

10.4 Remedies

In order to remedy these deviations, Microsoft is advised to take the following 3 remedies.

1. Speed up the process of providing access to the Diagnostic Data (taking less time to answer)
2. Stop collecting identifying data via the browser telemetry data or provide access to all personal data processed by Defender in a concise, transparent, intelligible and easily accessible form
3. If Microsoft wants to rely on an exception to the access rights from Art. 23 (1) under i of the GDPR jo. Art. 41(1) sub i of the UAVG, Microsoft should explain in detail why restrictions of the right to access some events would be necessary to safeguard the protection of the data subject or the rights and freedoms of others, including Microsoft itself.

11. Transfer risks

This section answers the final (ninth) verification question:

Are there high risks resulting from the transfer of personal data to the USA or other third countries?

11.1 Facts

Microsoft is moving data processing to the EU, with its EU Data Boundary program.[78] Currently, EU customers can already choose to have most Content Data processed, and service generated server logs generated in the EU. At the end of 2023, all Telemetry Data should also exclusively be processed in the EU. By the end of 2024, EU Enterprise and Education customers can choose to have all support tickets dealt with by exclusively EU based support employees.

In reply to questions from SLM Rijk, Microsoft explained that there are some exceptions to the EU Data boundary: some Content Data with pseudonymised identifiers from Defender for Endpoint, and a generic exception of all pseudonymised logs.

> *"All Customer Data for Defender services is stored at rest in the EU except for a small amount of customer content for Microsoft Defender for Endpoint that is transferred to the United States for the limited purpose of improving the Defender services' ability to protect from threats. The content transferred is data such as command lines, filenames, file paths, and URLs. Prior to transfer, the content is processed to pseudonymize potential identifiers.*
>
> *In addition, pseudonymized personal data is also transferred to the United States to be used for improving product protection. All data transferred to the*

[78] : Microsoft, Continuing Data Transfers that apply to all EU Data Boundary services, 27 April 2023, URL: https://learn.microsoft.com/en-us/privacy/eudb/eu-data-boundary-transfers-for-all-services.

> *United States is maintained in protected data stores with controls to ensure that it is only processed for permitted purposes.*"[79]

Even though Microsoft already generates the service generated server logs in the EU, Microsoft currently systematically transfers aggregated logs to the USA.

> *"Microsoft online services create system-generated logs as part of the regular operation of the service. Currently all system-generated logs are aggregated globally in the United States. These system-generated logs may contain pseudonymized personal data. Examples of system-generated logs that may contain pseudonymized personal data include:*
>
> - *Product and service usage data such as user activity logs*
> - *Data specifically generated by interaction of users with other systems*

Microsoft has explained it will stop processing these aggregated logs in the USA upon completion of the EU Data Boundary.

Microsoft writes:

> "*This EU Data Boundary documentation reflects the current state of the EU Data Boundary as of the date of publication. As noted in many cases in this article, we are continuing to deploy more services, service capabilities, and associated data within the EU Data Boundary and will update this documentation accordingly and note the last updated date.*"

Pseudonymization, as defined in Article 4(5) of the GDPR, is the processing of personal data so that it can no longer be attributed to a specific data subject without using additional information. In other words, it takes personally identifiable information within a data record and replaces it with one or more artificial identifiers, or pseudonyms, thus protecting the data subject's identity.

Microsoft has explained to SLM Rijk that it requires all personal data in system-generated logs to be pseudonymized. Microsoft uses various techniques to pseudonymize personal data in system-generated logs, including encryption, masking, tokenization and data blurring. Regardless of the specific method of pseudonymization, this protects user privacy by enabling authorized Microsoft personnel to use system logs containing only personal data that has undergone the security step of pseudonymization.

SLM Rijk has asked Microsoft to clarify what it meant with pseudonymisation of Content Data, as Diagnostic Data may also contain Content Data and personal data in file and folder names. Microsoft explained that it considers all data collected by Defender as Customer Content Data, even though the data Defender collects typically look like a form of system data.[80]

---

[79] Idem.

[80] E-mail Microsoft to SLM Rijk, 14 April 2023.

11.2 Technical findings

There are no technical findings.

11.3 Assessment

There is a (theoretical) risk that Microsoft is compelled to hand over personal data obtained from customers for this security purpose to law enforcement authorities or security agencies and secret services. This transfer risk has been thoroughly analysed in the Data Transfer Impact Assessment (DTIA) accompanying the 2022 public DPIA on Teams, OneDrive and SharePoint. The DTIA concludes that the chance is very slim that Microsoft is compelled to disclose personal data from EU public sector customers. Though Microsoft cannot disclose if it has **received** any specific legal demands subject to a secrecy obligation, it can explain that it has never **disclosed** personal data. Microsoft publicly explains:

"*Microsoft does **not provide**, and has never **provided**, EU public sector customer's personal data to any government."* This 'never' explicitly includes secret services, even though Microsoft is prohibited from disclosing if it has received orders under secrecy obligation.[81]

As a result of the EU Data Boundary, Content Data from EU customers are already exclusively processed in the EU, as well as the system generated server logs. By the end of 2023, Microsoft also expects to process all telemetry data and the aggregated data from the server logs exclusively in the EU.

If US authorities want access to these data, they can invoke the US Cloud Act. This act offers meaningful options for Microsoft to resist disclosure. If the names of specific employees or system administrators are too sensitive to transfer, the organisation can pseudonymise those names with identity federation.

By the end of 2024, all support tickets from EU Enterprise and Education customers will also exclusively be dealt with by agents based in the EU. If customers want to minimise the risk of unauthorised access to the contents of these support tickets in third countries, they can instruct staff never to include personal data in support tickets (except for the name of the person filing the ticket).

Additionally, Microsoft pseudonymises the personal data it processes for security purposes, provides legal guarantees of contesting each order, has a proven track record of such resistance, pledges to pay the customer damages if such compelled disclosure occurs and publishes detailed transparency reports twice per year.

Finally, it is likely that the European Commission will adopt a new (third) adequacy decision about the data protection rules in the USA in the summer of 2023.

On 25 March 2022, President Joe Biden and European Commission President Ursula von der Leyen signed an agreement 'in principle' to work out legal measures to ensure adequate protection of the data from the commercial sector in the USA. [82] On 7 October 2022, Biden signed a new Executive Order implementing this agreement with new binding safeguards for the data collection by US intelligence agencies, and

---

[81] Microsoft, Compliance with EU transfer requirements for personal data in the Microsoft cloud, November 2021, URL: https://go.microsoft.com/fwlink/p/?LinkID=2184913.

[82] European Commission press release, European Commission and United States Joint Statement on Trans-Atlantic Data Privacy Framework, 25 March 2022, URL: https://ec.europa.eu/commission/presscorner/detail/en/IP_22_2087

introducing a new redress procedure.[83] Following this EOP, the European Commission prepared a new draft adequacy decision.[84]

The Commission has asked the EDPB for an Opinion. The EDPB has issued its Opinion in February 2023. The EDPB notes the substantial improvements offered by EOP, but expresses concerns, asks for clarifications from the Commission and calls on the Commission to monitor the implementation in future joint reviews.[85] The LIBE committee of the European Parliament has taken a critical stance.[86] After the non-binding vote in the plenary EP, the ministers of the Member States are invited to agree (the Council). A possible new adequacy decision is not expected before July 2023.

If the EC succeeds in adopting a new adequacy decision, Microsoft will not have to sign up or certify for adequacy. The adequacy decision will apply to all transfers to the USA, also when based on SCCs and BCR.[87] Max Schrems immediately announced that he would likely challenge the arrangement once again at the European Court of Justice: "*noyb expects to be able to get any new agreement that does not meet the requirements of EU law back to the CJEU within a matter of months e.g. via civil litigation and preliminary injunctions. The CJEU may even take preliminary action if a deal is clearly violating previous judgements.*"[88]

In sum, based on the extremely small chance that the risks of unauthorised access to the personal data materialise at Microsoft and the upcoming new adequacy decision for the USA, the risk of undue access to these pseudonymised security data can be qualified as a low data protection risk.

11.4 Remedies

Microsoft is advised to clarify the documentation about the EU Data Boundary. It must be clear what data will be transferred to the USA and possibly other third countries after completion of the EU Data Boundary.

---

[83] Executive Order of the President, Enhancing Safeguards for United States Signals Intelligence Activities, URL: https://www.whitehouse.gov/briefing-room/presidentialactions/2022/10/07/executive-order-on-enhancing-safeguards-for-united-states-signalsintelligence-activities/ .

[84] Press release European Commission, Commercial sector: launch of the adoption procedure for a draft adequacy decision on the EU-U.S. Data Privacy Framework, 12 December 2022, https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/eu-us-data-transfers_en.

[85] EDPB, Opinion 5/2023 on the European Commission Draft Implementing Decision on the adequate protection of personal data under the EU-US Data Privacy Framework, 28 February 2023, URL: https://edpb.europa.eu/system/files/2023-02/edpb_opinion52023_eu-us_dpf_en.pdf.

[86] LIBE list of proposed amendments on draft report, 9 March 2023, URL: https://www.europarl.europa.eu/doceo/document/LIBE-AM-745289_EN.pdf.

[87] Stated by Commissioner Didier Reynders in reply to questions from the European Parliament, 24 June 2022, URL: https://www.europarl.europa.eu/doceo/document/E-9-2022-001307-ASW_EN.html.

[88] Noyb, "Privacy Shield 2.0"? - First Reaction by Max Schrems, 25 March 2022, URL: https://noyb.eu/en/privacy-shield-20-first-reaction-max-schrems

Organisations are advised to consider the use of pseudonymous accounts for specific employees working with for example secret information, and use pseudonyms for the system administrators.