DPIA Amazon Web Services (AWS)

Data protection impact assessment on the processing of personal data with Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon RDS with a MySQL database

Version 1.2 Public

Date	23 June 2023
Status	Public

DPIA by

Ministry of Justice and Security
Strategic Vendor Management Microsoft, Google Cloud and Amazon Web Services

Turfmarkt 147
2511 DP The Hague
PO Box 20301
2500 EH The Hague
www.rijksoverheid.nl/jenv

Contact

E slmmicrosoft@minjenv.nl
T 070 370 79 11

Project Name

DPIA report on **AWS cloud services**

Authors

Privacy Company
Sjoera Nas and Floor Terra, senior advisors
https://www.privacycompany.eu/

Change log

Version	Date	Summary of input
0.1	15 February 2021	First completed draft of part A and technical appendix for AWS
0.2	1 March 2021	Input SLM processed, with track changes
0.3	2 March 2021	Clean version for AWS
0.4	22 October 2021	Comments AWS processed in track changes, new information added to Sections 7 and 8
0.5	4 November 2021	Comments SLM processed, first draft part B and D with track changes
0.6	8 November 2021	First completed clean version for SLM
0.7	22 November 2021	Version ready for comments by AWS
0.8	22 May 2023	Updated version after completion of negotiations with AWS, including results of separate DTIA
0.9	29 May 2023	Comments SLM processed with track changes
1.0	30 May 2023	Completed report for final check by AWS
1.1	20 June 2023	Track changes after input AWS
1.2	23 June 2023	Public version

Contents

Figures

Summary

Amazon Web Services Inc. (AWS) provides many different cloud services, as infrastructure, as platform and as software. This Data Protection Impact Assessment (DPIA) assesses the risks of the use of three combined services: Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon Relational Database Service (Amazon RDS, in this case with a MySQL database).

Scope of AWS DPIA

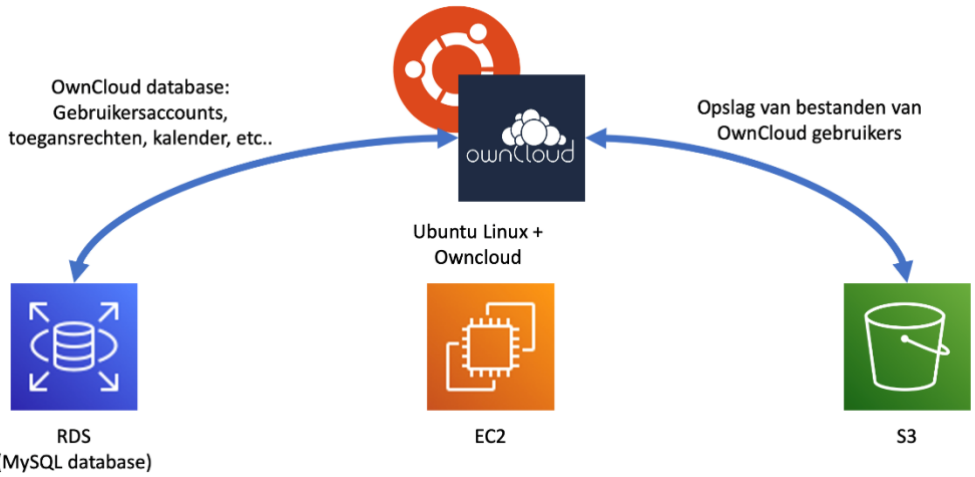

This report was commissioned by SLM Microsoft, Google Cloud and Amazon Web Services Rijk (the Strategic Vendor Management office housed at the Ministry of Justice and Security). This DPIA has to be read together with the separate Data Transfer Impact Assessment (DTIA) performed in May 2023 on the transfer of personal data through these 3 services.

Outcome: no more high data protection risks

As a result of discussions between the Dutch government and AWS, AWS has taken organisational and contractual measures to mitigate 7 previously identified high data protection risks. The outcome of this DPIA is that there are now no more known high risks if Dutch government organisations follow the recommended mitigating measures in this DPIA. Government organisations can also mitigate or accept the 9 known remaining low data protection risks. These low risks, and mitigating measures, are shown in the table at the end of this summary.

Personal data

This DPIA is based on a legal analysis of the available public documentation about the tested AWS services, answers from AWS to detailed questions from Privacy Company and a technical examination of the data processed by AWS in its log files.

To better understand the data processing in log files by Amazon, a Privacy Company employee has performed a number of scripted scenarios and has accessed the available log files from the Admin Console. In addition, a Data Subject Access Request was filed with AWS to obtain the Diagnostic data relating to the behaviour of the test admin.

This report distinguishes between 5 categories of personal data processed by AWS:

1. Content Data (customer uploaded Content Data in the VMs and storage spaces)

2. Account Data (including Contact Data)
3. Diagnostic Data (including Configuration and Security Data)
4. Support Data
5. Website Data (the restricted access Admin Console)

Purposes, roles and legal grounds

AWS contractually qualifies as data processor for the personal data in the Content Data, Account, Diagnostic, Support and restricted access Website Data.

The contract with the Dutch government includes a limitative list of 3 purposes, with identified sub purposes, for which AWS may process personal data as a processor. The 3 main purposes are:

1. Providing and maintaining the Services used by Customer and its Authorised Users, including through Customer's use of settings, administrator controls or other Service functionality (such as the AWS management console and APIs made available by AWS for the Services).
2. Securing the Services and the AWS Network, including by providing security features and services.
3. Providing Customer-requested support and perform basic troubleshooting.

The Dutch government specifically authorises AWS to further process limited personal data as an independent data controller for an exhaustive list of compatible purposes, when such processing is strictly necessary and proportionate. These purposes range from billing and calculating employee compensation to combatting fraud, and from responding to data subject access requests for personal data in AWS's controller role to improving the performance and core functionality of the services. Where possible, AWS will use pseudonymised data for these purposes.

One of these agreed compatible purposes involves compliance with legal obligations, including with possible disclosure orders from government authorities. AWS contractually guarantees it will make every reasonable effort to challenge any overbroad or inappropriate order, and redirect the government authority back to the customer. If AWS is nonetheless compelled to disclose personal data, AWS will promptly notify the customer if legally permitted. AWS commits it has not built in any *backdoors* or similar programming in the services that could be used by AWS or by third parties to obtain unauthorized access to the system.

Another of these agreed compatible purposes is abuse detection, prevention and protection. AWS ensures that it will not proactively scan the Content Data in the tested services except under very limited circumstances (e.g., scanning of a percentage of outgoing emails for email abuse in line with industry standards).

9 low risks and mitigating measures

No.	Low risks	Recommended measures government organisations	Recommended measures AWS
1.	**Disclosure or access to Content Data as a result of transfer to the USA**	Apply encryption. For S3: Encrypt files stored in S3 with keys outside of AWS's control if the files contain personal data AWS may not access. For EC2/RDS: AWS Nitro is designed to prevent AWS from accessing the Content Data inside of the VM.	Continue to organise external audits on compliance with access policies and disclosure to authorities and continue to disclose the findings via Artifakt.
			Continue to publish transparency reports about requests and disclosures.
2.	**Disclosure or access to Account, Diagnostic, Support and Website Data as a result of transfer to the USA**	Pseudonymise admin employee accounts, for example using identity federation.	Increase transparency about government access to personal data other than Content Data
		Do not host a website on an EC2 instance if the identifiability of visitors through their IP addresses is sensitive, or use a proxy.	
		Ask the AWS account manager to configure the alert for support employees so that only employees based in EU Member States and in countries for which an adequacy decision is available such as Canada or Japan may respond to tickets.	Ensure that the account managers are able to set the alert per customer.
3.	**Loss of control cookies and website telemetry restricted access Website Data**	Pseudonymise admin employee accounts, for example using identity federation.	Ensure that no analytical website telemetry data with user account names or account identifiers are being sent from the website when an admin selects the *'Customize cookies*' option in the cookie consent banner.
		Select the third option *'Customize cookies'* in the cookie banner on the restricted access websites (Admin and Support).	
4.	**Loss of control subprocessors**	Where available, opt-out of Service Improvement for Services.	Continue to organise external audits on compliance of subprocessors with the agreed data protection guarantees.
5.	**Loss of transparency Diagnostic, Support and Website Data**	Advise admins never to upload personal data in Support Requests.	Publish more detailed and up-to-date documentation, including essential information such as the processing of IP addresses of visitors to AWS hosted websites / applications.
			Update an overview of the categories of Personal Data that AWS processes as a controller for the agreed compatible purposes.
		Pseudonymise admin employee accounts, for example using identity federation.	Warn admins not to upload personal data in attachments to Support Requests, encourage and enable masking of personal data in screenshots.
6.	**No access for data subjects to some Account, Diagnostic,**	Inform employees about access to the data in the available admin log files and in the Support Centre.	Assist admins as controllers to honour data subject access rights for all personal data in Diagnostic, Website (Admin Console and Support Centre),

	Support, and Website Data		Support and Account Data and explain to admins when such access is denied on a case by case basis.
7.	**Chilling effects employee monitoring system**	Complement internal privacy policy for the processing of employee personal data with rules for what specific purposes specific personal data in the log files may be (further) processed and analysed. This includes listing the specific risks against which the logs will be checked, and which measures the organisations will take to ensure purpose limitation.	-no measures necessary-
8.	**Loss of control over personal data in inaccessible AWS security logs**	(SLM Rijk) conduct audits on compliance with purpose limitation, data minimisation and retention periods.	Continue to organise relevant audits on compliance with purpose limitation, data minimisation and retention periods.
9.	**The sending of unsolicited marketing mail to procurement officers**	Instruct procurement officials (Commercial Contacts) to opt-out from marketing communications through the AWS Email Preference Center.	-no more measures necessary, AWS will ask admins for consent for commercial newsletters and mails (no more opt-out).

Conclusion

As a result of the negotiations with the Dutch government, AWS has become a data processor for all personal data in and about the use of Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon RDS. If Dutch government organisations follow the recommended measures from this DPIA, they can use these 3 AWS services without any known high data protection risks.

If government organisations encrypt the Content Data with self-managed keys and apply the other risk mitigating measures such as the use of pseudonymous account data for the admins, the transfer risks are no longer qualified as high.

Introduction

This report, commissioned by Strategic Vendor Management Microsoft, Google Cloud and Amazon Web Services Rijk office (SLM[1]) housed at the Ministry of Justice and Security, is a data protection impact assessment (DPIA) about the use of three Amazon Web Services (hereinafter: AWS): Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon Relational Database Service (Amazon RDS, in this case, with a MySQL database).

This DPIA report is written by the Dutch privacy consultancy firm Privacy Company.[2]

Data Protection Impact Assessment (DPIA)

Under the terms of the General Data Protection Regulation (GDPR), an organisation may be obliged to carry out a data protection impact assessment (DPIA) under certain circumstances. According to Article 35, GDPR, a DPIA is required where the organisation is processing personal data on a large scale when it involves likely high risk to the rights and freedoms of the data subjects. The GDPR does not specify what constitutes a large scale. The DPIA is intended to shed light on, among other things, the specific processing activities, the inherent risk to data subjects, and to propose safeguards to mitigate these risks. The purpose of a DPIA is to ensure that any risks attached to the process in question are mapped and assessed, and that adequate safeguards have been implemented to mitigate those risks.

A DPIA used to be called PIA, *privacy impact assessment*. According to the GDPR a DPIA assesses the risks of data processing for the rights and freedoms of individuals. Data subjects have a fundamental right to protection of their personal data and some other fundamental freedoms that can be affected by the processing of personal data, such as for example freedom of expression.

The right to data protection is therefore broader than the right to privacy. Consideration 4 of the GDPR explains: "*This Regulation respects all fundamental rights and observes the freedoms and principles recognised in the Charter as enshrined in the Treaties, in particular the respect for private and family life, home and communications, the protection of personal data, freedom of thought, conscience and religion, freedom of expression and information, freedom to conduct a business, the right to an effective remedy and to a fair trial, and cultural, religious and linguistic diversity*".

This DPIA follows the structure of the DPIA Model mandatory for all Dutch government organisations.[3]

Umbrella DPIA versus individual DPIAs

Pursuant to article 35 of the GDPR, a DPIA is mandatory if an intended data processing constitutes a high risk for the data subjects whose personal data are being processed. The Dutch Data Protection Authority (Dutch DPA) has published a list of 17 types of

---

[1] SLM is the abbreviation of the Dutch words Strategisch Leveranciersmanagement Microsoft.

[2] https://www.privacycompany.eu/

[3] *Model Gegevensbeschermingseffectbeoordeling Rijksdienst (PIA)* (September 2017). For an explanation and examples (in Dutch) see: https://www.rijksoverheid.nl/documenten/rapporten/2017/09/29/model-gegevensbeschermingseffectbeoordeling-rijksdienst-pia.

processing for which a DPIA is always mandatory in the Netherlands.[4] If a processing is not included in this list, an organization must itself assess whether the data processing is likely to present a high risk.

The European national supervisory authorities (hereinafter referred to as the Data Protection Authorities or DPAs), united in the European Data Protection Board (EDPB) have also published a list of 9 criteria.[5] As a rule of thumb if a data processing meets two of these criteria a DPIA is required.

In GDPR terms SLM **is not the data controller** for the processing of personal data via the use of AWS. The data controller is the individual government organisation that decides to switch from on-premise IT to cloud services from AWS. However, as central negotiator for many cloud services, SLM has a moral responsibility to assess the data protection risks for the employees and negotiate for a framework contract that complies with the GDPR. Therefore, SLM commissions umbrella DPIAs to assist the government organisations to select a privacy-compliant deployment, and conduct their own DPIAs where necessary. Only the organisations themselves can assess the specific data protection risks, related to the technical privacy settings, nature and volume of the personal data they process and vulnerability of the data subjects.

This umbrella DPIA is meant to help the different government organisations with the DPIA they must conduct when they deploy AWS, but this document cannot replace the specific risk assessments the different government organisations must make.

Scope of this DPIA: AWS Virtual Machine EC2, S3 and RDS MySQL database

AWS provides many kinds of cloud services. This DPIA assesses the data protection risks of the use of a Virtual Machine (VM) in the AWS cloud infrastructure (EC2), with a pre-configured MySQL database structure (RDS) to store data in Simple Storage Service (S3).

This DPIA covers:

- Description of the AWS VM, storage and database services, plus the related services that are inextricably linked to the use of these services: the AWS management console and AWS authorisation management;
- Role and purposes of AWS as provider of the AWS VM, storage and database services
- Transfer of personal data, both via content and log files, to countries outside of the European Economic Area
- Exercise of data subject rights
- Retention and backup periods, in the storage forms.

Out of scope

The following topics are outside of the scope of this DPIA:

- Functional data (organisations need to send this kind of personal data to AWS to communicate with the VM and database in the cloud)

[4] Dutch DPA, list of processings for which a DPIA is required, in Dutch only, Besluit inzake lijst van verwerkingen van persoonsgegevens waarvoor een gegevensbeschermingseffectbeoordeling (DPIA) verplicht is, URL: https://autoriteitpersoonsgegevens.nl/sites/default/files/atoms/files/stcrt-2019-64418.pdf.

[5] The EDPB has adopted the WP29 Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679, WP248rev.01, 13 October 2017, URL: https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=611236

- The legal grounds for the processing of (special categories of) personal data by the individual government organisations
- Use of other Amazon Software-as-a-service (SaaS), Platform-as-a-service (PaaS) in the cloud
- Use of advanced technical support (AWS Developer, Business and Enterprise Support Services).
- Use of additional security logging and monitoring services, such as Amazon CloudWatch, Amazon Inspector, Amazon GuardDuty, Amazon Audit Manager and AWS Security Hub
- As part of IAM-tooling: Single Sign-On

Methodology

This DPIA is based on different sources of information and fact-finding strategies. Privacy Company has carefully studied all available public documentation from AWS about the tested AWS services, including information about log files.

Privacy Company has asked questions and engaged in a dialogue with representatives of AWS.

Additionally, in order to better understand the data processing in log files by Amazon, a Privacy Company employee has performed a number of scripted scenarios and has accessed the available log files from the Admin Console. In addition, a Data Subject Access Request was filed with AWS to obtain the Diagnostic data relating to the behaviour of the test admin. The results of this DSAR request are discussed in Section 3.3 of this report.

In May 2023 Privacy Company also performed a DTIA on the transfer of personal data through the 3 tested services. The DTIA was based on the excel form provided by the Swiss lawyer David Rosenthal[6], with a few changes made by Privacy Company.

Chronology

This DPIA was conducted between **September 2019 and May 2023**.

On 23 September 2019, SLM sent a list of questions to AWS. On 30 September 2019, the test admin from Privacy Company filed a DSAR for all personal data relating to his activities recorded by AWS. On 9 October 2019 AWS partially answered two of the questions, with added information on 16 October 2019 (corrected hyperlinks with accessible information). The outcomes of the DSAR are discussed in Section 3.3 of this DPIA.

After a kick-off meeting on 6 November 2019, a follow-up technical workshop was cancelled last-minute by AWS, as well as the next date the technical workshop was scheduled for, in February 2020.

On 28 November 2019 SLM sent a (expanded) list of questions. In January 2020, Privacy Company signed NDAs. AWS provided answers to the expanded list of questions on 6 August 2020. Follow-up meetings with AWS took place on 28 September, 3 November and 2 December 2020. During these meetings, Privacy Company explained its initial findings and requested technical information about the Diagnostic Data collected by AWS. By mail of 18 December 2020, AWS provided one quotable sentence about the purposes of the processing of Diagnostic Data. On 4

[6] The excel sheet has been expanded since. See:
https://www.rosenthal.ch/downloads/Rosenthal_EU-SCC-TIA.xlsx

February 2021 AWS sent its AWS Information Request Report, made available to the public on 31 January 2021. Because a long time had passed since Privacy Company had initially executed the test scenarios (in September 2019) and the production of Part A of the DPIA, limited test scenarios were repeated on 11 February 2021 and data were captured from the log service CloudTrail. On 20 October 2021 a further validation of the default settings of CloudTrail was performed.

The updated draft DPIA still concluded there were 7 high and 4 low data protection risks. For some high risks no meaningful mitigating measures could be identified that could be applied by the customers (government organisations).

In May 2022, AWS sent a new proposal to SLM. In August 2022 SLM Rijk asked Privacy Company to help assess the proposed changes, and restart the DPIA project. Since March 2023, AWS helped clarify the impact of the changes in a series of constructive meetings scheduled by SLM.

On 30 May 2023, SLM Rijk sent the final version of the draft DPIA to AWS for a review on company confidential information. AWS provided extensive input, including references to updates of public documentation. Many of AWS's requests have been honored to remove confidential information. In those cases the text has been replaced by [**Confidential**].

Input from AWS

On 1 October 2021, AWS responded with detailed comments on the factual findings of the initial DPIA. AWS's main comments were summarized in bullet points, followed by a short description how the comment was processed in the previous updated DPIA report.

However, in view of the contractual measures, the 2021 input is no longer relevant for most comments. Therefore, this section only mentions a few key changes as a result of the input from AWS.

Initially, AWS proposed to further define and delimit the purposes for which it would process certain personal data as independent data controller. AWS was willing to commit that as a data controller, it would not use personal data for profiling, advertising or marketing, or any other purposes beyond defined purposes.

After 1 March 2021, when AWS was provided with the draft factual findings, AWS updated its publicly available Data Processing Addendum, subprocessor list and SCC. AWS requested to take those changes into account.[7] Throughout the updated DPIA report, the references to the content of these documents were updated.

- AWS asked Privacy Company to use its own definitions, instead of the term Diagnostic Data.[8] AWS used the terms Customer Content, Service Attributes and Account Information.[9] Meanwhile, AWS has changed its the names of its data types, and now uses the term 'Metadata'. This includes security logs and customer chosen configuration data. AWS repeated its request to use its own terms, instead of Diagnostic Data.

The objection is noted in Section 2.3, but the distinction between Content, Account, Diagnostic, Support and Website Data is made in all DPIA reports about cloud

[7] AWS response to part A of the DPIA, 1 October 2021, Par. 4 and 13.

[8] Idem, Par. 1.

[9] AWS separate input on data types, 1 October 2021.

providers for SLM Rijk, and is maintained as a systematic clarification of the different categories of personal data. AWS's original term 'Service Attributes 'could also apply to non-personal data.

AWS objected to the qualification as joint controller with the Dutch government organisations. AWS's input with regard to its role was added to Section 5 of the report. However, in view of the contractual measures, which limit the role of AWS to a processor role for almost all personal data, this debate no longer has the same relevance for this DPIA.

- AWS explained that CloudTrail was enabled by default[10]

Privacy Company mistakenly wrote that enabling of CloudTrail required a financial investment, while it follows from AWS's input that CloudTrail logs are available for 90 days "*without the need to manually setup CloudTrail*."[11] On 20 October 2021 a further validation of the default settings of CloudTrail was performed. The correction was added to Sections 2.3.2 and Section 3.1.2 of this report.

- AWS stated that customers do not have a right to access all information processed by AWS, and that not all data mentioned in the report as 'missing' from the Data Subject Access Request (DSAR) were personal data or should be provided, as "*disclosure in the context of a DSAR could result in unlawful disclosure of personal data and be antithetical to the objective of data protection.*"[12]

AWS's objection, including the quote mentioned above, was added to Section 3.3, while the merits of this statement are assessed in Section 15 of the DPIA.

- AWS objected against some of the separate purposes mentioned in Section 4.2.1 of the report for the Content Data. AWS writes: "*Maintenance of services, detection of misuse of services, and compliance with laws are related and necessary functions of service provision.*"[13] AWS similarly objected to the dissection of purposes mentioned in its general Privacy Notice for the other categories of personal data, "*because customers do not determine means and purposes of processing by AWS as a controller*."[14] Additionally, AWS wrote that the sections on purposes "*include erroneous summary conclusion of fact and law.*"[15]

AWS's objections were added to Section 4.2. However, in view of the list of specific processor purposes agreed with the Dutch government, this section of the report was seriously abridged.

- AWS objected against the qualification of automated scanning of Content Data to detect prohibited content.

The argument and response are no longer relevant, and are deleted from the report. Contractually, automated scanning of Customer Content for the purpose of identifying potentially abusive content or activity is excluded, except under very limited

---

[10] Idem, Par. 3.

[11] AWS CloudTrail features, undated, URL: https://aws.amazon.com/cloudtrail/features/

[12] Idem, Par. 6, p. 2.

[13] Idem, Par. 7, p. 2.

[14] Idem, Par. 8, p. 2. Similarly in Par. 3.3, page 4.

[15] Idem, Par. 11 and 12, p. 3.

circumstances (e.g., scanning of a percentage of outgoing emails for email abuse in line with industry standards).

- AWS asked to take the updated EDPB guidance on measures to supplement transfer tools into account, that do allow for a risk assessment of the likelihood of access. AWS explained that the combination of SCC and additional measures taken by the customer can be used to guarantee an adequate level of data protection. AWS publishes transparency reports about the types and numbers of requests it receives for Content Data from law enforcement authorities.[16]

In 2023, Privacy Company helped SLM to perform a separate DTIA on the risks of transfer of personal data in the 3 tested services. AWS provided very helpful and detailed input on the (very extensive) spreadsheet. The DTIA takes the expanded explanation about encryption into account with regard to the Content Data. With regard to the other categories of personal data the risks affect a select group of system administrators, and they can use pseudonymous accounts with identity federation. Additionally, the results of a C5:2020 audit were taken into account, that showed no findings of noncompliance with regard to disclosure to law enforcement authorities. The DTIA also mentions the likelihood that the European Commission will adopt a new adequacy decision for the USA in the summer of 2023.

- AWS objected against the assumption that AWS staff would have access to encryption keys.[17] In 2023 AWS explained that the Nitro System for EC2 instances offers extra protection against tampering. For example protection of the server from unauthorized modification of system firmware thanks to the Nitro Security Chip.[18] AWS also provided explanation about its Key Management System (KMS). "*AWS KMS is designed so that no one, including AWS employees, can retrieve customer plaintext KMS keys from the service. AWS KMS uses hardware security modules (HSMs) that have been validated under FIPS 140-2, or are in the process of being validated, to protect the confidentiality and integrity of a customer's keys. Customers plaintext KMS keys never leave the HSMs, are never written to disk, and are only ever used in the volatile memory of the HSMs for the time needed to perform the customer's requested cryptographic operation.*"[19] AWS also explained that the US Cloud Act does not give the authority to request a service provider to decrypt encrypted data.

In reply to this input, the risk was reevaluated, and new descriptions were added. Based on the review of the design the risk of forced decryption is now assessed to be near zero. Even though the probability of decryption by AWS is likely to be extremely low, the solution does not meet the letter of the three possible guarantees provided by the EDPB: "*the keys are retained (i) solely under the control of the data exporter,*

[16] Amazon Information Request Report, 31 January 2023, URL: https://d1.awsstatic.com/Security/pdfs/Amazon_Information_Request_Report.pdf..
[17] AWS response to part A of the DPIA, 1 October 2021, Par. 19.
[18] AWS whitepaper, The Security Design of the AWS Nitro System, 18 November 2022, URL: https://docs.aws.amazon.com/pdfs/whitepapers/latest/security-design-of-aws-nitro-system/security-design-of-aws-nitro-system.pdf.
[19] AWS Key Management Service (AWS KMS), undated, URL: https://aws.amazon.com/kms/. Public report NC group on the AWS Nitro System API & security claims, 11 April 2023, URL: https://research.nccgroup.com/2023/05/03/public-report-aws-nitro-system-api-security-claims/

or (ii) by an entity trusted by the exporter in the EEA or (iii) under a jurisdiction offering an essentially equivalent level of protection to that guaranteed within the EEA [numbering added by Privacy Company]."[20]

- AWS asked to consider the application of identity federation and IAM roles as mitigating measures.

A new section 8.2 was added to the report, with a factual explanation about the existence of pseudonymisation options for the system admins. The results of the application of such measures are described in Section 17 of the DPIA.

Outline of this DPIA

Following the structure of the Dutch central government model DPIA, this DPIA report uses a structure of four main divisions, which are reflected here as "parts".

A. Description of the factual data processing
B. Assessment of the lawfulness of the data processing
C. Assessment of the risks for data subjects
D. Description of mitigation measures

Part A explains the data processing by AWS and by government organisations resulting from the use of the EC2, S3 and RDS services. Part A starts with a description of the main categories of personal data processed by AWS, and an analysis of the enrolment framework. Section 2 describes the tested services, and the contents of the different log files generated by AWS.

This starts with a description of the way the data are collected and processed and describes the categories of personal data and data subjects that may be affected by the processing, the purposes of the data processing, the different roles of the parties, the different interests related to this processing, the locations where the data are stored and the retention periods. In this Section, factual contributions and intentions from AWS are included.

Part B provides an assessment by Privacy Company, with input from SLM of the lawfulness of these data processings through AWS EC2, S3 and RDS. This analysis starts with an analysis of the extent of the applicability of the GDPR and the ePrivacy Directive, in relation to the legal qualification of the role of the AWS as provider of the cloud services. Subsequently, conformity with the key principles of data processing is assessed, including transparency, data minimisation, purpose limitation, and the legal ground for the processing, as well as the necessity and proportionality of the processing. In this section the legitimacy of any transfers of personal data to countries outside of the EEA is separately addressed, as well as how the rights of the data subjects are respected.

Part C assesses the risks to the rights and freedoms of the data subjects created by the processing activities identified in Part A of this DPIA. It identifies specific risks that may result from these processings by determining (1) the likelihood that such a risk

[20] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, Version 2.0, Adopted on 18 June 2021, URL: https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf.

may occur, and (2) the severity of the impact on the rights and freedoms of the data subjects if the risk occurs.

Finally, **Part D** contains concrete measures that can be taken by either AWS or the individual governmental organisations to mitigate the risks identified in Part C. These measures might either reduce the chance the risks occur, or the impact they might have, or both. Part D also contains an assessment of any residual risk attached to the use of EC2, in the tested combination with S3 and RDS, for risks that cannot or will not be mitigated by applying the suggested measures.

Part A. Description of the Data Processing

This first part of the DPIA provides a description of the characteristics of the personal data that can be generated and processed if a government organization uses a VM with a database in AWS's cloud. Specifically, the use of Amazon Elastic Compute Cloud (Amazon EC2) together with Amazon Simple Storage Service (Amazon S3) and Amazon RDS hosting of a MySQL database.

Part A continues with a description of the different categories of personal data that may be processed, the categories of data subjects that may be affected by the processing, the purposes of the data processing by AWS, the locations where data may be stored, processed and analysed, and the roles of the institutions and AWS as processor or as (joint) data controllers.

Finally, this Section provides an overview of the different interests related to the processing, and of the retention periods.

1. Description of tested AWS Services

This Section describes the different AWS services examined for this DPIA.

In general, when analysing the data processing by a cloud provider, three main categories of collecting data can be distinguished:

1. **Content Data** stored on the cloud provider's servers, often called 'Content', 'Customer Content' or 'Customer Data'
2. **Diagnostic Data**. All data generated or collected by the cloud provider about the use of its servers and services, including visits to its public website to look-up for example relevant privacy policies and help explanations, and the (restricted access) websites requiring log-in by an administrator, but only to the extent that these data are stored by the cloud provider and not merely transported; and
3. **Functional Data**, data that are temporarily processed by the cloud provider to execute desired functionalities. The key difference between Functional Data and Diagnostic Data as defined in this report, is that functional data are and should be transient.[21] This means that these data should be immediately deleted or anonymised upon completion of the transmission of the communication. Otherwise they qualify as Content Data or Diagnostic Data. As long as the cloud provider does not store these Functional Data, they are not Diagnostic Data.

This DPIA is focussed on the data protection risks of the processing of Diagnostic Data, but also addresses the risks of transfer of personal data in Content, Account, Support and (restricted access) Website Data to the USA, and mitigating measures such as pseudonymisation and encryption.

---

[21] Compare Article 6(1) of the EU ePrivacy Directive (2002/58/EC, as revised in 2009 by the Citizens Rights Directive) and explanation in recital 22: "*The prohibition of storage of communications and the related traffic data by persons other than the end users or without their consent is not intended to prohibit* ***any automatic, intermediate and transient storage*** *of this information in so far as this takes place* ***for the sole purpose of carrying out the transmission*** *in the electronic communications network and* ***provided that the information is not stored for any period longer than is necessary for the transmission and for traffic management purposes****, and that during the period of storage the confidentiality remains guaranteed.*"

Figure 1: Content Data, Functional Data and Diagnostic Data

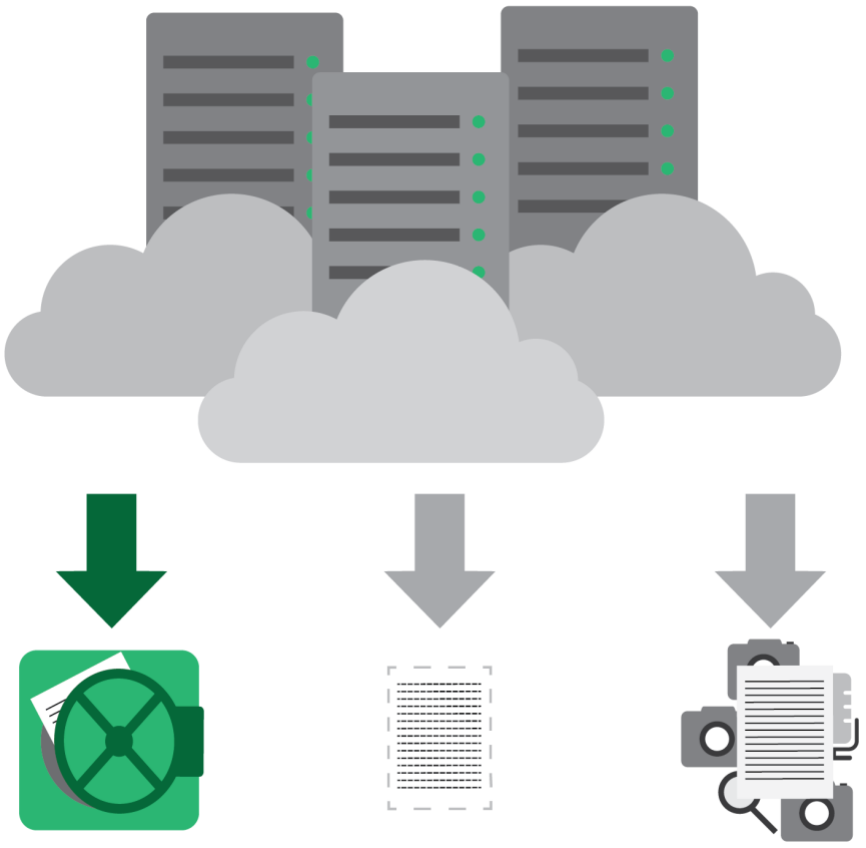

1.1 Amazon Elastic Compute Cloud

This DPIA assesses the data protection risks of the processing of personal data on and about the use of virtual machines in the cloud (VMs). AWS calls these VMs: Elastic Compute Cloud (EC2).

A VM is an emulation or virtual representation of a computer system that provides the functionality of a physical computer (also known as *bare metal server*). One physical cloud server can host many virtual machines, with different operating systems and applications. The different VMs are *sandboxed*, isolated from each other, and prevented from interfering with the physical computer. Customers can also choose a *virtual private cloud* (isolated from the AWS public cloud).

This is a *Infrastructure as a Service (IaaS)* offering, which provides scalable computing capacity using server instances in AWS' data centres. EC2 is designed to make web-scale cloud computing easier for developers.[22]

Amazon EC2 provides the following features:

- Virtual computing environments, known as instances.

[22] AWS, Amazon EC2, URL: https://aws.amazon.com/ec2/?nc2=h_ql_prod_fs_ec2

- Preconfigured templates for the instances, known as Amazon Machine Images (AMIs), that package the bits an organisation needs for the server.
- Various configurations of central processing units (CPU), memory, storage, and networking capacity for the instances, known as instance types.
- Secure login information for the instances using key pairs. AWS stores the public key, and the organisation stores the private key in a secure place.
- Storage volumes for temporary data that is deleted when an organisation stops or terminates the instance, known as instance store volumes.
- Persistent storage volumes for data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes.
- Multiple physical locations for resources, such as instances and Amazon EBS volumes, known as Regions and Availability Zones.
- A firewall that enables organisations to specify the protocols, ports, and source IP ranges that can reach the instances using security groups.
- Static Ipv4 addresses for dynamic cloud computing, known as Elastic IP addresses.
- Metadata, known as tags, that an organisation can create and assign to the organisation's Amazon EC2 resources.
- Virtual networks an organisation can create that are logically isolated from the rest of the AWS cloud, and that an organisation can optionally connect to the organisation's own network, known as virtual private clouds (VPCs).

As provider of cloud infrastructure services, AWS is the global market leader. According to estimates from Synergy Research Group, AWS's market share in the worldwide cloud infrastructure market amounted to 32 percent in the first quarter of 2023, while Microsoft Azure had 23 percent, and Google Cloud 11%.[23] Other global competitors are Oracle Cloud, IBM Cloud and Alibaba Cloud. All these providers offer cloud services in their global network of data centres, frequently with options for customers to select the data region where they want to store the Content Data *at rest*.

1.2 Amazon Simple Storage Service

Amazon Simple Storage Service (S3) is an object storage service that offers scalability, data availability, performance, and security.

Figure 2: AWS explanation about S3 access points[24]

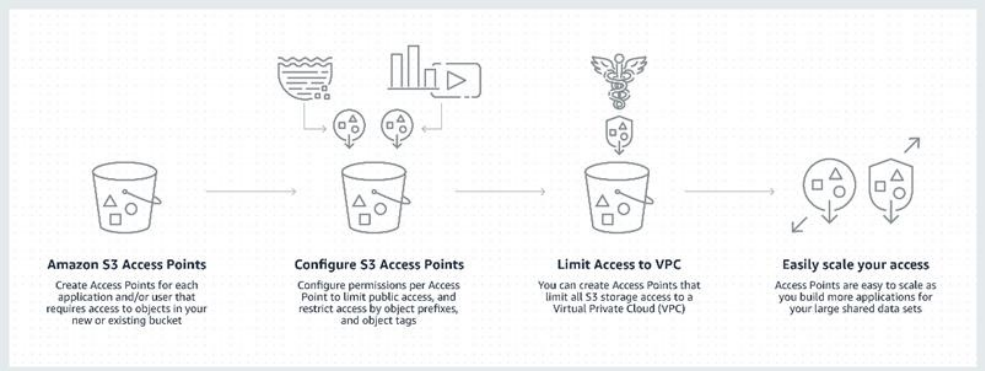

[23] Statista, Big Three Dominate the Global Cloud Market, quoting estimates from Synergy Research Group, 28 April 2023, URL: https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/ .

[24] AWS, How do S3 Access Points work? URL: https://aws.amazon.com/s3/features/access-points/

S3 provides a web services interface that can be used to store and retrieve data from anywhere on the web.

Customers of all sizes can use it to store and protect any amount of data for a range of use cases as websites, mobile applications, backup and restore, archive, enterprise applications, Internet of Things (IoT) devices and big data analytics.

To provide customers with the flexibility to determine how, when, and to whom they wish to expose the information they store in AWS, S3 APIs provide both bucket and object-level access controls, with defaults that only permit authenticated access by the bucket and/or object creator.

Figure 3: AWS illustration of the combination of EC2 with S3[25]

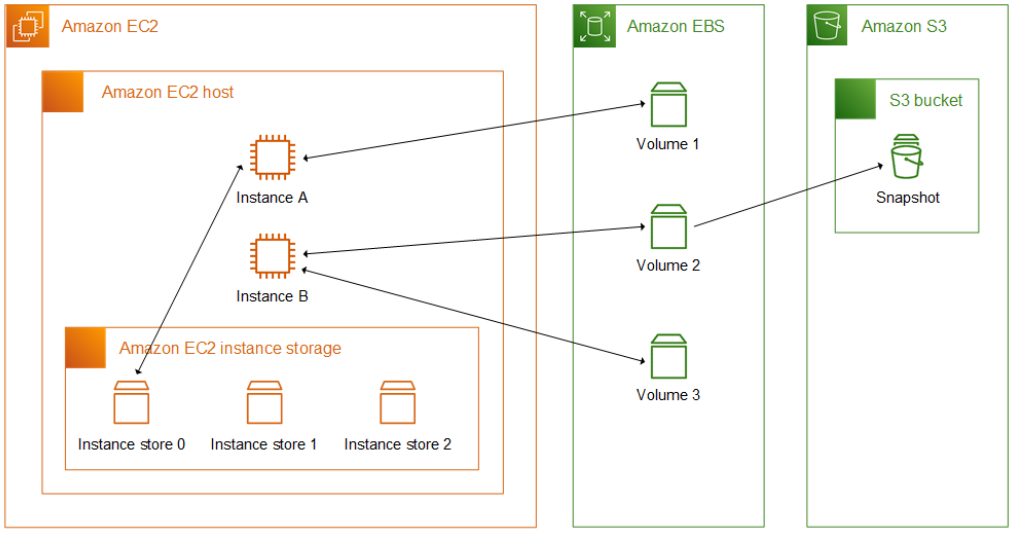

Global competitors with similar data storage services are: Microsoft Azure Storage Services, Google Cloud Storage, Oracle Object Storage, IBM Cloud Object Storage, Rackspace Cloud Files and Alibaba Object Storage Service.

1.3 Amazon RDS with a MySQL database

Instead of just storing a blob of unstructured data in AWS S3 bucket, for this DPIA a test set-up was created with a MySQL Database with test data, as offered as a database service by AWS.

Privacy Company installed a VM with the operating system Ubuntu Linux. This is a standard VM image offered by AWS. Additionally, a preconfigured MySQL RDS instance was configured. Owncloud was installed as a web application on the EC2 VM and configured to use the MySQL database as storage for structured data. The S3 bucket was used to store user files.

Owncloud is an open source, self-hosted, file storage cloud service, similar to Microsoft OneDrive and Dropbox.

[25] AWS storage, URL: https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/Storage.html

Figure 4: Test set-up for this DPIA

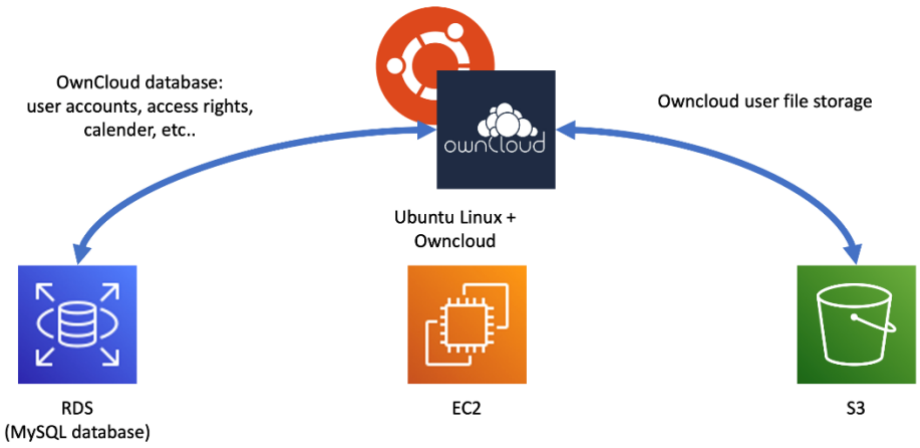

MySQL is an open source relational database. Open source databases are databases in which the underlying code can be freely viewed, downloaded, modified and re-used. Modifying the source code enables users to adjust it to match their preferences. Therefore, open source databases are more flexible and offer more freedom than traditional license models of closed source databases.

Managing an open source database takes a lot of time. AWS offers preconfigured Relational Database Services to reduce time by automating the administration tasks of databases such as a MySQL database. This makes it easier for system administrators to set up, operate, scale, patch and backup relational databases in the cloud. Amazon RDS automates, amongst other things, automatic failover, backups, software patching, encryption, access management, security and monitoring of the MySQL database and recovery. The system administrator configures the access and authorisation policy and controls monitoring/logging.

Global competitors for Amazon RDS are: Microsoft Azure SQL Database, MongoDB Atlas, Google BigQuery, IBM Db2, SAP HANA Service and Couchbase.

According to a 2020 study from market research firm Global Market Insights, the Europe cloud computing market is set to grow from its current market value of more than $25 billion to over $75 billion by 2026. Most of this market share will be earned through managed database services.[26]

2. Legal: personal data and enrolment framework

The Dutch government DPIA model requires that this Section provides a list of the kinds of personal data that will be processed, and per category of data subjects, what kind of personal data will be processed by the product or service for which the DPIA is conducted. Since this is an umbrella DPIA, this information is presented in two different Sections: a general legal description of the categories of personal data and data subjects, and, in Section 3, a description of the technical findings on the Diagnostic Data collected in log files.

[26] Euractiv, Europe Cloud Computing Market to witness steady growth of 12% during 2020-2026, URL: https://pr.euractiv.com/pr/europe-cloud-computing-market-witness-steady-growth-12-during-2020-2026-208846 The researchers explain: *"The SaaS cloud computing delivery model held the majority of Europe cloud computing market with around 65% share in 2018."*

2.1 Definition of personal data

According to article 4(1) (a) GDPR,

> "'*personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.*"

The concept of processing is defined in Article 4(2) of the GDPR:

> "*'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction.*"

The GDPR explains that pseudonymised data are personal data, to which the GDPR applies. Recital 26 explains:

> "*Pseudonymised personal data that can be linked to a natural person through the use of additional data should be regarded as data relating to an identifiable natural person. In order to determine whether a natural person is identifiable, account must be taken of all means that can reasonably be expected to be used by the controller or by another person to directly or indirectly identify the natural person, for example selection techniques. In determining whether any means can reasonably be expected to be used to identify the natural person, account shall be taken of all objective factors, such as the cost and time of identification, taking into account available technology at the time of processing and technological developments.*"

2.2 AWS terminology

In its Privacy Notice, AWS uses the term 'Personal Information', and describes three different ways it collects these data:

- *Information You Give Us: We collect any information you provide in relation to AWS Offerings. Click here[27] to see examples of information you give us*
- *Automatic Information: We automatically collect certain types of information when you interact with AWS Offerings. Click here[28] to see examples of information we collect automatically.*
- *Information from Other Sources: We might collect information about you from other sources, including service providers, partners, and publicly available sources. Click here[29] to see examples of information we collect from other sources."[30]*

[27] AWS provides an internal hyperlink to a list of examples in its Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/ .

[28] Idem.

[29] Idem.

[30] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/ .

As a result of the negotiations with the Dutch government, the Privacy Notice no longer applies to the data processing via the 3 tested services, with the exception of 2 categories of personal data:

1. Processing of visitor data of the public (commercial) AWS website.
2. Use of contact data of procurement officers for AWS's own commercial purposes (Commercial Account Data).

In the contract with the Dutch government AWS uses the GDPR definition of personal data [**Confidential**]. As required by the European Commission SCC, AWS also uses the term personal data in the Standard Contractual Clauses.

2.3 Five categories of personal data

In order to provide a precise analysis of the different types of personal data processing and their impact, this DPIA distinguishes between five categories of personal data that AWS can process as a result of the use of the tested set-up (use of EC2 and S3 in combination with a MySQL RDS). See Figure 5 below.

1. Content Data (called Customer Data by AWS)
2. Account Data (defined as Authorised User Account Data by AWS)
3. Diagnostic Data (including security logs and Configuration Data)
4. Support Data
5. Website Data (limited to the restricted access websites, i.e. the Admin and Support Console)

Figure 5: Five different categories of personal data

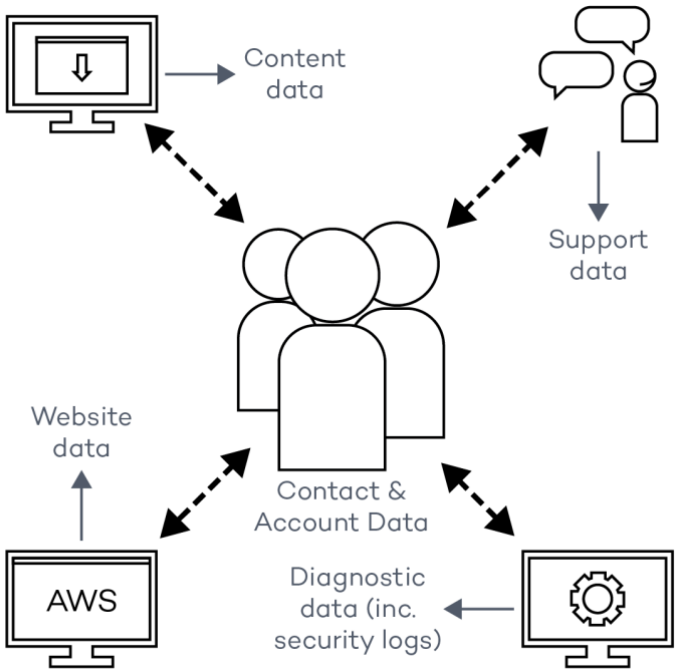

The category of Diagnostic Data includes personal data in resource identifiers, metadata tags, access controls, rules, usage policies, permissions, and similar items, to the extent configuration choices are personal data.

2.3.1 Content Data

Content data are the personal and non-personal data uploaded by customers to a (storage and or database on a) VM (Amazon EC2 instance). AWS explains (customer) data as:

> "*the personal data that is uploaded to the Services under Customer's AWS Enterprise Accounts*"

Content Data do not include the Account Data, and may also include non-personal data.

In its Data Privacy FAQ, AWS explains that such content may be any type of data, such as

> *"software (including machine images), data, text, audio, video, or images that a customer or any end user transfers to us for processing, storage, or hosting by AWS services in connection with that customer's account, and any computational results that a customer or any end user derives from the foregoing through their use of AWS services*."[31]

AWS allows its customers to select the geolocation where Customer Data will be processed within the AWS Network. AWS does not offer such a data region selection for the Diagnostic Data about the use of its services, including visits to its publicly accessible website, the Admin Console and use of its Support Services. See Section 7 of this DPIA.

When assessing the data protection risks of processing specific Content Data on a cloud server, government organisations must take the nature of the data into account, such as classified information, personal data of a sensitive nature and special categories of data.[32] The possible contents of these categories are described below, but it is up to the individual government organisations to complete this umbrella DPIA with a specification of the specific data they intend to process in the 3 tested AWS services.

Classified information
Some Dutch government employees will, depending on the capacity in which they work, frequently process Classified Information. The Dutch government defines four classes of Classified Information, ranging from confidential within the ministry to extra secret state secret.

Classified Information is not a separate category of data in the GDPR or other personal data legislation. Nonetheless, information processed by the government that is qualified as classified information, whether it qualifies as personal data or not, must legally be protected by special safeguards. The processing of this information when related to an individual, can also have a privacy impact. If the personal data of an employee, such as an account ID, or unique device identifier, can be connected to the

[31] AWS Data Privacy FAQ, How does AWS classify customer data?, URL: https://aws.amazon.com/compliance/data-privacy-faq/

[32] In reply to this DPIA AWS suggested that these examples of types of data are speculative, as organisations may also process non-classified, non-personal data on the cloud services. AWS response to part A of the DPIA, 1 October 2021, Par. 4.3. As umbrella DPIA, this report cannot determine the precise nature of the data processing. By nature, a DPIA focusses on risks related to the processing of personal data, not on other types of information that may be processed by a cloud provider.

information that this person works with Classified Information, the impact on the private life of this employee may be higher than if that person would only process 'regular' personal data. Unauthorised use of this information could for example lead to a higher risk of being targeted for social engineering, spear phishing and/or extortion.

Overview of possible personal data
The following list may serve as a source of inspiration for DPOs and Privacy Officers to get a good overview of data that may be processed by their organisation on an AWS cloud service.

- *Basic personal data (for example place of birth, street name and house number (address), postal code, city of residence, country of residence, mobile phone number, first name, last name, initials, email address, gender, date of birth), including basic personal data about family members and children;*
- *Authentication data (for example user name, password or PIN code, security question, audit trail);*
- *Contact information (for example addresses, email, phone numbers, social media identifiers; emergency contact details);*
- *Unique identification numbers and signatures (for example Social Security number, bank account number, passport and ID card number, driver's license number and vehicle registration data, IP addresses, employee number, student number, patient number, signature, unique identifier in tracking cookies or similar technology);*
- *Pseudonymous identifiers;*
- *Financial and insurance information (for example insurance number, bank account name and number, credit card name and number, invoice number, income, type of assurance, payment behavior, creditworthiness);*
- *Commercial Information (for example history of purchases, special offers, subscription information, payment history);*
- *Biometric Information (for example DNA, fingerprints and iris scans);*
- *Location data (for example, Cell ID, geo-location network data, location by start call/end of the call. Location data derived from use of wifi access points);*
- *Photos, video and audio;*
- *Internet activity (for example browsing history, search history, reading, television viewing, radio listening activities);*
- *Device identification (for example IMEI-number, SIM card number, MAC address);*
- *Profiling (for example based on observed criminal or anti-social behavior or pseudonymous profiles based on visited URLs, click streams, browsing logs, IP-addresses, domains, apps installed, or profiles based on marketing preferences);*
- *HR and recruitment data (for example declaration of employment status, recruitment information (such as curriculum vitae, employment history, education history details), job and position data, including worked hours, assessments and salary, work permit details, availability, terms of employment, tax details, payment details, insurance details and location and organizations);*
- *Education data (for example education history, current education, grades and results, highest degree achieved, learning disability);*
- *Citizenship and residency information (for example citizenship, naturalization status, marital status, nationality, immigration status, passport data, details of residency or work permit);*
- *Information processed for the performance of a task carried out in the public interest or in the exercise of an official authority;*
- *Special categories of data (for example racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning*

health, data concerning a natural person's sex life or sexual orientation, or data relating to criminal convictions or offences); or

- *Any other personal data identified in Article 4 of the GDPR.*

2.3.2 *Account Data*

Account Data (*Authorised User Account Data*) are defined as

> personal data about the use by Customer or its Authorised Users of the Services that is specifically identified with Customer's AWS Enterprise Account..

This category includes the AWS account information from existing and past customer' system administrators, both as owners of the 'root' account, as well as regular system administrators with an admin account. However, this definition excludes Commercial Contact Data, explained below.

AWS explains:

> "*Account information means information about you that you provide to us in connection with the creation or administration of your AWS account. For example, Account information includes names, usernames, phone numbers, email addresses and billing information associated with your AWS account*."[33]

AWS explains that each Enterprise customer needs at least two AWS accounts, a root user account and a system administrator AWS account. Without an AWS Account administrators cannot access the Admin Console or file Support Requests. Customers must actively provide personal data for such an account such as a name and e-mail address. The root account manages the authorisations for the different administrators.

> *"When you create an AWS account, a root user account is automatically created for your AWS account. This user account has complete access to all your AWS services and resources in your AWS account. Instead of using this account for everyday tasks, you should only use it to initially create additional roles and user accounts, and for administrative activities that require it."*[34]

For this DPIA, AWS Identity and Access Management (IAM) was tested to manage the accounts. IAM is the web service used to securely manage the access to the tested AWS resources. IAM contains a database with users (username and password), roles, related account settings and access tokens. E-mail addresses are not required for IAM, since system accounts may also be used to give other systems programmatic access to AWS.

Instead of using AWS IAM customers can also use their own or another third party identity management system. This is what AWS calls identity federation. With identity federation, organisations can systematically determine what information about admin account is provided to AWS. They can for example replace directly identifying user names and e-mail addresses by pseudonymised alternatives, such as 'admin01@governmentorganisation.nl'. If the government organisation for example uses Microsoft's Azure Active Directory, they can use the pseudonyms in the Azure

[33] AWS Customer Agreement, Section 14 Definitions.

[34] AWS Whitepaper, Navigating GDPR Compliance on AWS, December 2020, URL: https://docs.aws.amazon.com/whitepapers/latest/navigating-gdpr-compliance/navigating-gdpr-compliance.pdf

AD, and enable admins with Single Sign On to use that pseudonymous account to sign in to the AWS console. See Section 8.2 of this report for a technical explanation.

AWS also processes a second category of Account Data: of commercial sales contacts that procure the licenses for an organisation (Commercial Contact Data). The processing of these data is covered by AWS's general Privacy Notice.

AWS explains in its Privacy Notice that it may enrich such (Commercial) Contact Data with information from external (data broker) sources. "*We might collect information about you from other sources, including service providers, partners, and publicly available sources. Click here to see examples of information we collect from other sources.*"[35] The hyperlink refers to a list with the following examples of other sources:

- "*marketing, sales generation, and recruitment information, including your name, email address, physical address, phone number, and other similar contact information;*
- *subscription, purchase, support, or other information about your interactions with products and services offered by us, our affiliates (such as AWS training courses), or third parties (such as products offered through the AWS Marketplace) in relation to AWS Offerings;*
- *search results and links, including paid listings (such as Sponsored Links); and*
- *credit history information from credit bureaus.*"[36]

Figure 6: AWS promotional message sent to root admin account[37]

AWS also explains in its Privacy Notice it can use Commercial Contact Data to send promotional messages. Privacy Company received some invitations for AWS events. See Figure 6 above.

[35] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/
[36] Idem.
[37] Example of a message received from aws-marketing-email-replies@amazon.com on Monday, 22 February 2021 at 12:06.

According to the interface with communication preferences (See Figure 7 below), AWS can send 8 different kinds of newsletters, announcements and surveys to the individual admins.

Figure 7: Default settings until 2023 in the AWS Email Preference Center [38]

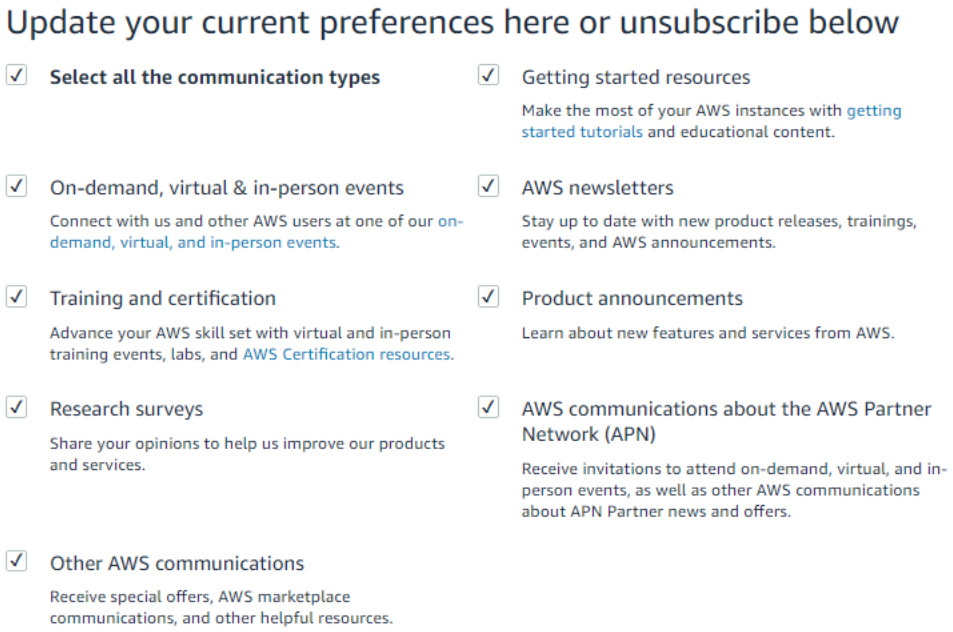

AWS offers a single button to opt-out from all these communication types. See Figure 8 below.

Figure 8: Opt-out from marketing mails AWS[39]

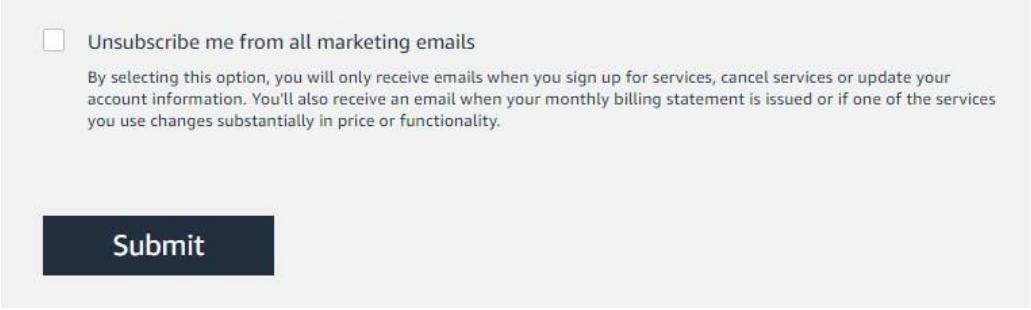

Contractually AWS may only send personalised marketing communications to the government organisations' system administrators if they have provided consent. Based on the GDPR such consent can only be given if it is based on adequate information, and has to be expressed in an unambivalent, active way (never through an opt-out).

2.3.3 Diagnostic Data

Diagnostic data are the metadata that AWS collects and stores in different system-generated cloud server logs about the behaviour of system administrators and

[38] AWS Communication Preferences, URL: https://pages.awscloud.com/communication-preferences.html
[39] Idem.

performance of the VM. For AWS the category of Metadata includes Operational logs, Security Logs and individual usage data (in Service logs for the customer), as well as Configuration Data (previously known as 'Service Attributes')

In this DPIA and the DTIA these data are called Diagnostic Data. This includes the security logs that are collected by AWS, including for example data about access attempts by third parties.

AWS collects Diagnostic Data about the individual use of its cloud services by end-users such as system administrators in two ways, namely by:

1. collecting system-generated logs on its own AWS cloud servers about the activities of the administrators, and;
2. collecting SIEM logs about activities by admins and external visitors to information hosted on its network, VMs and databases.

Additionally, AWS collects Diagnostic Data about visits to its public and restricted access websites. These data are described separately in Section 2.3.5 below.

Based on the results of the technical tests executed for this DPIA, Sections 3.1.1 to 3.1.4 of this DPIA report describe the different Diagnostic Data in more detail (where available), and analyse if these log files contain personal data.

AWS allows customers to turn on IAM logging through AWS CloudTrail.[40] By default, AWS makes these CloudTrail logs available to all customers for a period of 90 days.[41] The contents of these logs are described in Section 3.1.2 of this report.

Next to CloudTrail, administrators of AWS S3 *buckets* can also view access log files with Diagnostic Data.[42] The contents of these log files about access to the contents of S3 buckets (*read, write, update*, etc) are described in Section 3.1.1 of this report.

2.3.4 Support Data

Support data are generated when a government organisation files a Support Request with an admin AWS account through the restricted access Support Portal website. A Support Request may include an attachment, such as a screenshot of contents of a database. AWS also collects metadata about the Support Requests filed by its customers. These data always relate to an identifiable AWS customer (admin account).

This report does not analyse the data processing as a result of the use of the separate AWS Developer, Business and Enterprise Support Services.

---

[40] AWS, AWS Logging and monitoring in AWS Identity and Access Management, undated, URL: https://docs.aws.amazon.com/IAM/latest/UserGuide/security-logging-and-monitoring.html

[41] AWS CloudTrail features, undated, URL: https://aws.amazon.com/cloudtrail/features/

[42] AWS, Amazon S3 Server Access Log Format, undated, URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/LogFormat.html

Figure 9: Screenshot of creation of AWS Support Ticket[43]

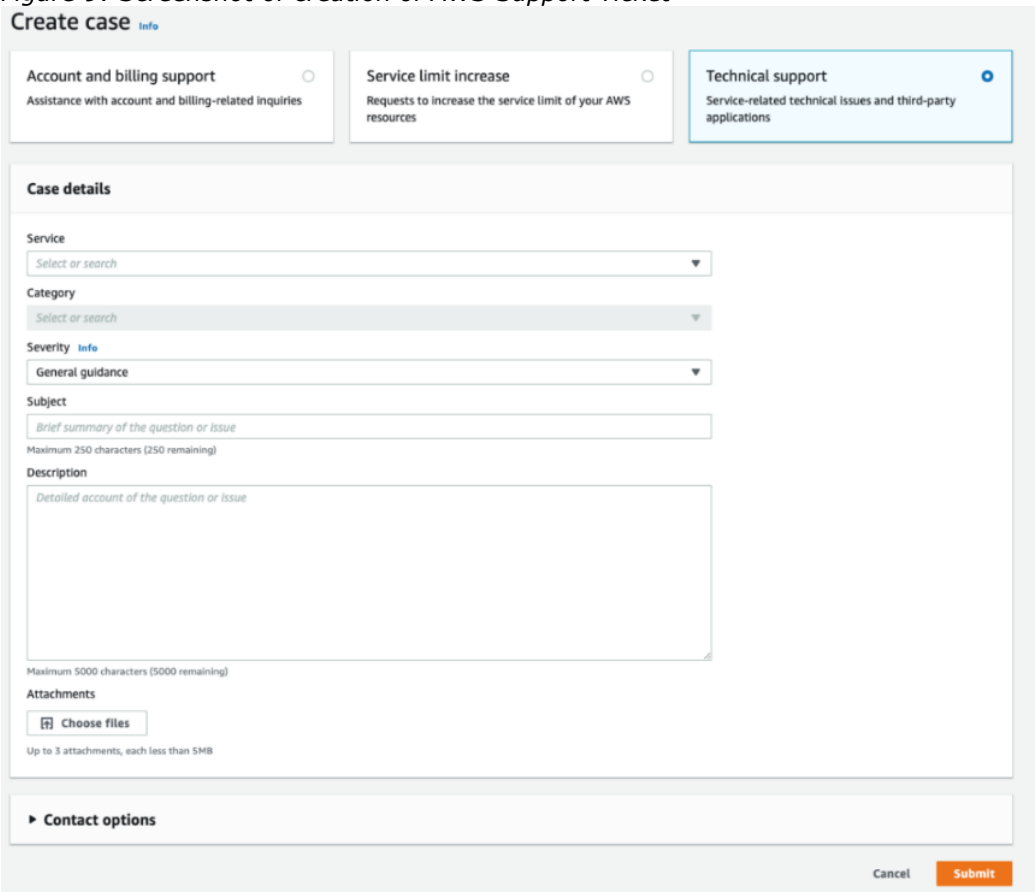

The actual contents of the Diagnostic and the Website data that may be part of Support Data will be explained in more detail in Sections 3.1 and 3.2 of this report.

2.3.5 *Website Data*

This report identifies two categories of Website Data:

1. Restricted access Website Data (Admin console and Support portal)
2. Publicly accessible AWS Website Data

Both categories of Website Data include data collected by cookies and pixels. Technically, as described above, Website Data are a form of metadata on the behaviour of system administrators, and therefore part of the broad category of Diagnostic Data. However, for analytical clarity, and because of differences in applicable privacy terms and inspection methods, this report separately analyses the data recorded through the use of the AWS-websites.

AWS publishes a Cookie Notice. According to this Notice, AWS collects the following Diagnostic Data about the interactions of each visitor to its website:

- "Network and connection information, such as the Internet protocol (IP) address used to connect your computer or other device to the Internet and information about your Internet service provider

[43] Screenshot made by Privacy Company on 12 February 2021.

- Computer and device information, such as device, application, or browser type and version, browser plug-in type and version, operating system, or time zone setting
- The location of your device or computer
- Authentication and security credential information
- Content interaction information, such as content downloads, streams, and playback details, including duration and number of simultaneous streams and downloads
- The full Uniform Resource Locators (URL) clickstream to, through, and from our site (including date and time) and AWS Offerings, content you viewed or searched for, page response times, download errors, and page interaction information (such as scrolling, clicks, and mouse-overs)." 44

AWS does not provide separate information about the Diagnostic Data it collects in its webserver access logs about the use of an AWS-account on the restricted access websites.

2.4 Enrolment framework

The enrollment framework for the use of the different AWS cloud services consists of a number of documents, shown in Figure 10 below [**Confidential**].

The publicly available enrollment framework consists of the following documents:

- Service Terms[45]
- AWS GDPR Data Processing Addendum including (the new) Standard Contractual Clauses[46]
- AWS Supplementary Addendum to the DPA[47]
- AWS Customer Agreement[48]
- AWS Acceptable Use Policy[49]
- Separate service level agreements, four of which are relevant for this DPIA[50]

---

[44] AWS Cookie Notice, last updated 3 September 2020, URL: https://aws.amazon.com/legal/cookies/.

[45] AWS Service Terms, last updated 30 May 2023, URL: https://aws.amazon.com/service-terms/ . A previous version of these terms, from July 2019, consulted for this DPIA, contained an explicit hierarchy provision: *"In the event of a conflict between the terms of these Service Terms and the terms of the AWS Customer Agreement or other agreement with us governing your use of our Services (the "Agreement"), the terms and conditions of these Service Terms apply."*

[46] AWS GDPR Data Processing Addendum, URL: https://d1.awsstatic.com/legal/aws-gdpr/AWS_GDPR_DPA.pdf , last updated 16 September 2021 (a previous version was dated 22 May 2018). *"Entire Agreement; Conflict. This DPA incorporates the Standard Contractual Clauses by reference. Except as amended by this DPA, the Agreement will remain in full force and effect. If there is a conflict between the Agreement and this DPA, the terms of this DPA will control, except that the Service Terms will control over this DPA. Nothing in this document varies or modifies the Standard Contractual Clauses."*

[47] AWS, Supplementary Addendum to AWS Data Processing Addendum, 1 page, URL: https://d1.awsstatic.com/legal/aws-dpa/supplementary-addendum-to-the-aws-dpa.pdf.

[48] AWS Customer Agreement, last updated 20 April 2023, URL: https://aws.amazon.com/agreement/.

[49] AWS Acceptable Use Policy, last updated 1 July 2021, URL: https://aws.amazon.com/aup/

[50] AWS Service Level Agreements (SLAs), URL: https://aws.amazon.com/legal/service-level-agreements/ .

- Amazon Compute Service Level Agreement[51], identical to Amazon EC2 Service Level Agreement[52]
- Amazon RDS Service Level Agreement[53]
- Amazon S3 Service Level Agreement[54]
- AWS Key Management Service Level Agreement[55]

+

- AWS Privacy Notice[56]
- AWS Site terms[57]
- AWS Cookie Notice[58]

During this DPIA, AWS was asked to specify the hierarchy of all relevant elements of the contract framework.

AWS provided the following answer:

> "*The AWS Customer Agreement applies to the use of the AWS services by customers. Applicable policies for the use of the AWS services, such as the AWS Acceptable Use Policy and AWS Service Terms, are incorporated by reference into the AWS Customer Agreement. Under the terms of the AWS Customer Agreement, if terms of the AWS Customer Agreement are inconsistent with the terms of the policies, the AWS Customer Agreement will control, except that the terms of the Service Terms will control over the AWS Customer Agreement. AWS becomes a processor when an AWS customer uses AWS services to process personal data uploaded to its AWS account. The AWS Data Processing Addendum applies to such processing.*"[59]

In a subsequent response, AWS added:

> "*The AWS GDPR DPA is part of the AWS Service Terms (see Section 1.14.1), and the AWS Service Terms are incorporated into the AWS Customer Agreement (Section 1.1) or other agreement governing the customer's use of AWS services.*"[60]

According to the public contract information, AWS is a data controller for the processing of Contact Data, Website Data, Support Data and Diagnostic Data.[61] AWS

---

[51] Amazon Compute Service Level Agreement and Amazon EC2 Service Level Agreement, last updated 25 May 2022, URL: https://aws.amazon.com/compute/sla/

[52] The Amazon EC2 SLA can also be reached through the URL: https://aws.amazon.com/ec2/sla/

[53] Amazon RDS Service Level Agreement, last updated 19 May 2022, URL: https://aws.amazon.com/rds/sla/

[54] Amazon S3 Service Level Agreement, last updated 5 May 2022, URL: https://aws.amazon.com/s3/sla/

[55] AWS Key Management Service Level Agreement, last updated 29 November 2022, URL: https://aws.amazon.com/kms/sla/

[56] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/

[57] AWS Site Terms, last updated 30 September 2022, URL: https://aws.amazon.com/terms/.

[58] AWS Cookie Notice, last updated 30 December 2022, URL: https://aws.amazon.com/legal/cookies/

[59] AWS general response of 9 October 2019 to DPIA questions SLM Cloud Rijk of 23 September 2019, p. 2 of 3.

[60] AWS response to DPIA questions, 17 July 2020, answer to Q1a.

[61] As confirmed by AWS in its reply to Part A of the DPIA, Annex with data types, 1 October 2021. "*In so far there are personal data processed by AWS in Service Attributes, the AWS*

explained: "*To the extent AWS collects personal data as a data controller, the processing and transfer of such data is subject to the AWS Privacy Notice. The AWS Privacy Notice describes how AWS collects and uses personal information in relation to AWS websites, applications, products, services, events, and experiences that reference the AWS Privacy Notice.*"[62]

In the agreement with the Dutch government AWS is a data processor for all personal data except for the public website data and the commercial contact data (of procurement officers). These contractual stipulations prevail over any public documentation.

Figure 10: Privacy hierarchy enrolment framework with the Dutch government ***[Confidential]***

Only when AWS qualifies itself as data controller, its (general) AWS Privacy Notice applies. **[Confidential]**

In sum, the enrolment framework consists of two pillars, with different applicable guarantees for the personal data processing, depending on AWS's role as a data processor or as a data controller. See Section 5 of this DPIA report for an assessment of the GDPR role(s) of AWS and the government organisations.

2.5 Possible categories of data subjects

This umbrella DPIA can only indicate categories of data subjects that may be involved in the processing but cannot assess the specific risks of the actual data processing per organisation that uses or will use the AWS cloud services.

As described in Section 2.3.1 of this report, each government organisation determines itself what personal Content Data it wants to store in a database on a cloud server, relating to what natural persons. Such databases may include lists of Customer's customers, patients, clients, employees, suppliers and/or end-users. The government organisation must keep an overview of the different categories of data subjects in its data processing inventory (Article 30 GDPR), and must list these categories of data subjects in the Standard Contractual Clauses, when such clauses apply to the transfer of personal data to the USA. In the new SCC from the European Commission, the Annex continues to contain a separate list of the category or categories of data subjects.[63]

As mentioned in Section 2.3.1 of this DPIA, SLM recommends organisations to check the following generic list of possible data subjects.:

- Employees, contractors and temporary workers (current, former, prospective) of data exporter;

Privacy Notice applies. This notice provides information, safeguards, and measures on for example the use, location, security, and transfer of service attributes. When AWS collects its customers personal data and determines the purposes and means of processing that personal data, AWS acts as a personal data controller."

[62] Idem. Similarly, in the Privacy Notice: "*This Privacy Notice describes how we collect and use your personal information in relation to AWS websites, applications, products, services, events, and experiences that reference this Privacy Notice (together, "AWS Offerings").*"

[63] European Commission Implementing Decision (EU) 2021/914 of 4 June 2021 on standard contractual clauses for the transfer of personal data to third countries pursuant to Regulation (EU) 2016/679 of the European Parliament and of the Council, URL: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32021D0914&from=EN.

- Dependents of the above;
- Data exporter's collaborators/contact persons (natural persons) or employees, contractors or temporary workers of legal entity collaborators/contact persons (current, prospective, former);
- Users (e.g., customers, clients, patients, visitors, etc.) and other data subjects that are users of data exporter's services;
- Partners, stakeholders or individuals who actively collaborate, communicate or otherwise interact with employees of the data exporter and/or use communication tools such as apps and websites provided by the data exporter;
- Stakeholders or individuals who passively interact with data exporter (e.g., because they are the subject of an investigation, research or mentioned in documents or correspondence from or to the data exporter);
- Minors; or
- Professionals with professional privilege (e.g., doctors, lawyers, notaries, religious workers, etc.).

Generally speaking, two categories of data subjects may be affected by the processing of Diagnostic Data and Contact Data by AWS: system administrators and legal/sales contact persons.

3. Results technical investigation

In order to better understand the data processing in log files by Amazon, a Privacy Company employee has performed a number of scripted scenarios and has subsequently accessed the available log files in the (web) Admin Console. In addition, a data subject access request was filed with AWS to obtain the Diagnostic data relating to the behaviour of the test admin. The results of the DSR request are discussed in Section 3.3 of this report. This Section explains in detail what Diagnostic Data were observed in the technical inspection of the data, and why some of these stored Diagnostic Data on the use of the tested AWS services are personal data as defined in article 4(1) of the General Data Protection Regulation (GDPR).

3.1 Diagnostic Data

As explained in Section 2.3.2, AWS collects Diagnostic Data about the use of its cloud services by the system administrators of the Dutch government. Sections 3.1.1 to 3.1.4 discuss how Privacy Company obtained access to Diagnostic Data in the context of this DPIA and contains an overview of the contents of these Diagnostic Data. Section 3.2 describes the contents of the Website Data, and Section 3.3 describes the outcomes of a Data Subject Access Request filed by (Privacy Company's) system administrator of the test set-up.

In its general Privacy Notice, AWS provides a list of examples of Diagnostic Data it may collect from end-users. AWS also refers to its separate Cookie Policy in this list.

"*We collect information automatically when you:*

- *visit, interact with, or use AWS Offerings (including when you use your computer or other device to interact with AWS Offerings);*
- *download content from us;*
- *open emails or click on links in emails from us; and*
- *interact or communicate with us (such as when you attend an AWS event or when you request customer support).*

Examples of the information we automatically collect include:

- *network and connection information, such as the Internet protocol (IP) address used to connect your computer or other device to the Internet and information about your Internet service provider;*
- *computer and device information, such as device, application, or browser type and version, browser plug-in type and version, operating system, or time zone setting;*
- *the location of your device or computer;*
- *authentication and security credential information;*
- *content interaction information, such as content downloads, streams, and playback details, including duration and number of simultaneous streams and downloads;*
- *AWS Offerings metrics, such as offering usage, occurrences of technical errors, diagnostic reports, your settings preferences, backup information, API calls, and other logs;*
- *the full Uniform Resource Locators (URL) clickstream to, through, and from our website (including date and time) and AWS Offerings, content you viewed or searched for, page response times, download errors, and page interaction information (such as scrolling, clicks, and mouse-overs);*
- *email addresses and phone numbers used to contact us; and*
- *identifiers and information contained in cookies (see our Cookie Notice)."*[64]

Though AWS provides documentation about the existence and contents of the logs that it makes available for administrators, for example in the whitepaper on Navigating GDPR compliance,[65] there was no public documentation about other Diagnostic Data AWS collects, such as network logs and the webserver access logs of AWS own websites.

AWS explains administrators have access to three kinds of logs with Diagnostic Data: application logs, resource logs, and AWS service logs.[66] AWS also provides admins with access to telemetry data and to monitoring tools.

3.1.1 S3 Access logs

S3 Access logs are logs of all external access to S3 storage buckets.

Figure 11 below shows the interface for system administrators to the S3 buckets. This interface give access to the available access logs and to the customer files.

[64] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/ .
[65] AWS Whitepaper, Navigating GDPR Compliance, last updated April 2022, p. 13, 'Compliance Auditing and Security Analytics', URL: https://docs.aws.amazon.com/whitepapers/latest/navigating-gdpr-compliance/navigating-gdpr-compliance.pdf#monitoring-and-logging
[66] AWS, SEC 4: How do you detect and investigate security events?, URL: https://wa.aws.amazon.com/wat.question.SEC_4.en.html

Figure 11: AWS S3 management interface

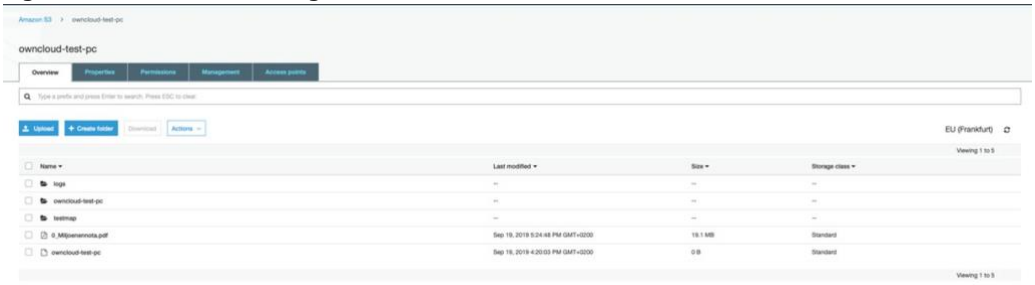

AWS describes the S3 access logging in the whitepaper Navigating GDPR compliance:

> *"When you enable logging, you can get detailed access logs for the requests that are made to your Amazon S3 bucket. An access log record contains details about the request, such as the request type, the resources specified in the request, and the time and date the request was processed. For more information about the contents of a log message, see Amazon Simple Storage Service Server Access Log Format in the Amazon Simple Storage Service Developer Guide.*
> *Server access logs are useful for many applications because they give bucket owners insight into the nature of requests made by clients that are not under their control.* ***By default, Amazon S3 does not collect service access logs, but when you enable logging****, Amazon S3 delivers access logs to your bucket on an hourly basis.*
> *This information includes:*
>
> - *Granular logging of access to Amazon S3 objects*
> - *Detailed information about flows in the network through VPC-Flow Logs*
> - *Rule-based configuration verification and actions with AWS Config Rules*
> - *Filtering and monitoring of HTTP access to applications with WAF functions in CloudFront."*[67]

These S3 access logs are not enabled by default, but AWS strongly recommends enabling this logging to get detailed information regarding the requests that are made to the S3 bucket.

> *"Server access logs are useful for many applications because they give bucket owners insight into the nature of requests made by clients that are not under their control."*[68]

The log format is similar to a standard web server access log, both in contents and structure. The S3 access logs contain the following unique user identifiers:

- Bucket owner
- Remote IP
- Identity of the requesting user or "-" if unknown
- Object path (containing both the 'filename' and the names of the folders the files is stored in, only if the filename contains personal data
- The full request URI, only if the URI contains personal data

Other data such as the timestamp, the URL referrer and operations on the object are also personal data, when they relate to activities performed by natural persons that can be identified with the set of unique identifiers mentioned above. Of course, it is also possible that a machine performs actions that are registered in these logs, such

[67] AWS Whitepaper, Navigating GDPR Compliance, p. 14.
[68] Idem.

as a request filed by an automated system without human interaction to a storage bucket. If the actions are performed by a machine, the data in these logs are not personal data.

AWS refers to detailed information about the contents of log messages in the Amazon Simple Storage Service Developer Guide.[69] Lists of other data observed in the S3 access logs are included in the technical appendix shared with AWS.

Figure 12: Example of an access log file via S3 interface[70]

```
3fc3e29a3f524e650fe467f0fa2f4462d8a5b688ca0a08d14943e9b842a01e11 owncloud-test-pc
[17/Sep/2019:13:22:18 +0000] 172.16.151.180
        3272ee65a908a7677109fedda345db8d9554ba26398b2ca10581de88777e2b
61 4B26DC1F09596910 REST.PUT.OBJECT logs/2019-09-17-13-22-18-1BFE49397877EC74
        "PUT    /owncloud-test-pc/logs/2019-09-17-13-22-18-1BFE49397877EC74
HTTP/1.1" 403 AccessDenied 243 644 7 - "-"
        "aws-internal/3 aws-sdk-java/1.11.590                    Linux/4.9.184-
0.1.ac.235.83.329.metal1.x86_64  OpenJDK_64-Bit_Server_VM/25.222-b10  java/1.8.0_222
vendor/Oracle_Corporation" –
        TJ89kDYr6ne1C5yv50gbxrFkgiqUN9zojDmK3u+YJLfoe5MAYBaVU+LmfADurC
CRzJfnSaSJc7k= SigV4 ECDHE-RSA-AES128-SHA AuthHeader s3.eu-central-1.amazonaws.com
TLSv1.2
3fc3e29a3f524e650fe467f0fa2f4462d8a5b688ca0a08d14943e9b842a01e11 owncloud-test-pc
[17/Sep/2019:13:50:23 +0000] 35.159.50.194
        3fc3e29a3f524e650fe467f0fa2f4462d8a5b688ca0a08d14943e9b842a01e11
0DC8778B6ACAEF1D    REST.HEAD.OBJECT owncloud-test-pc    "HEAD    /owncloud-test-pc
HTTP/1.1" 200 - - 0 13 –
        "-"    "aws-sdk-php2/2.7.5    Guzzle/3.8.1    curl/7.58.0    PHP/7.2.19-
0ubuntu0.18.04.2" –
        UqqwOVpev89oQN3rRh2mapILXsVeDwyMu6OcQTAovvB1tgnY3a6v7Xxse6LV
hGq3hogd7xhCrJ0= SigV4 ECDHE-RSA-AES128-GCM-SHA256 AuthHeader
        owncloud-test-pc.s3.amazonaws.com TLSv1.2
3fc3e29a3f524e650fe467f0fa2f4462d8a5b688ca0a08d14943e9b842a01e11 owncloud-test-pc
[17/Sep/2019:13:50:23 +0000] 35.159.50.194
        3fc3e29a3f524e650fe467f0fa2f4462d8a5b688ca0a08d14943e9b842a01e11
71A498415FF4C8C4  REST.GET.ACL owncloud-test-pc  "GET  /owncloud-test-pc?acl HTTP/1.1"
200 - 480 - 7 –
        "-"    "aws-sdk-php2/2.7.5    Guzzle/3.8.1    curl/7.58.0    PHP/7.2.19-
0ubuntu0.18.04.2" –
        pHyqnLVRwltq9Ssyqi8RwH8hXyuDhYcDVrSQ1x+9Ev
```

3.1.2 CloudTrail logs

Next to the S3 access logs AWS provides a logging service for all other AWS services. This native logging facility is called Amazon CloudTrail. CloudTrail logs are available for all customers for a period of 90 days "*without the need to manually setup CloudTrail*."[71]

When a customer enables CloudTrail logging, by default all API calls to AWS services are logged.

CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of a customer AWS account. AWS explains:

> "*With AWS CloudTrail, you can continuously monitor AWS account activity. A history of the AWS API calls for your account is captured, including API calls made through the AWS Management Console, the AWS SDKs, the command*

[69] Amazon S3 Server Access Log Format, URL: https://docs.aws.amazon.com/AmazonS3/latest/userguide/LogFormat.html

[70] Log file accessed by Privacy Company on 17 September 2019.

[71] AWS CloudTrail features, undated, URL: https://aws.amazon.com/cloudtrail/features/.

> *line tools, and higher-level AWS services. You can identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators enable and disable CloudTrail logging."*[72]

AWS recommends storing the CloudTrail logs in a separate S3 bucket, with restricted access and encryption of the data at rest, to prevent tampering with the logs.

> "*The permissions on the bucket should prevent deletion of the logs, and they should also be encrypted at rest using Server-Side Encryption with Amazon S3-managed encryption keys (SSE-S3) or AWS KMS–managed keys (SSE-KMS). CloudTrail log file integrity validation can be used to determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it."*

AWS does not provide detailed public documentation what data it collects in its own system generated logs, and what data it makes available for admins.

In the first test runs, executed in September 2019, Privacy Company did not access the available data in CloudTrail, but filed a Data Subject Access Request (DSAR) to see what data about the test set-up AWS possibly logged. AWS did not provide answers to this question. See Section 3.3 of this DPIA report for the outcomes.

In the second test-run, on 11 February 2021, Privacy Company did access CloudTrail. Privacy Company configured a custom trail. Custom Trails (different from the default CloudTrails) are stored in a S3 bucket with an unlimited retention period. The customer can configure S3 lifecycle rules to migrate or delete old log files, for example expiry after a set period. In October 2021 Privacy Company conducted a third test-run, to verify the default availability for a period of 90 days when no Custom Trails are configured.

Privacy Company did not create a lifecycle rule. See Figure 13 below.

Figure 13: No default lifecycle rules in test set-up[73]

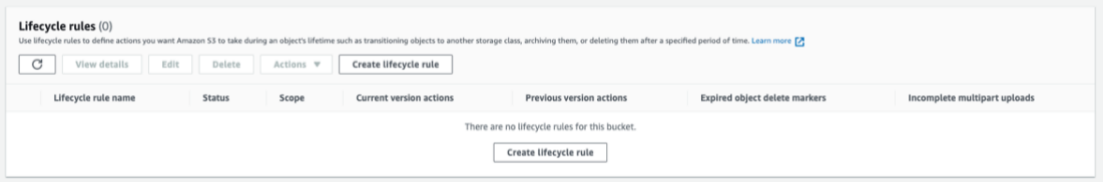

Privacy Company did not test any optional log monitoring services like CloudTrail Insights or CloudWatch. Privacy Company also did not install a CloudWatch Agent on the VM to enable additional logging that would be visible in CloudTrail.

AWS explains the difference between CloudTrail and CloudWatch:

> *"AWS CloudTrail logs can also trigger preconfigured Amazon CloudWatch events. You can use these events to notify users or systems that an event has occurred, or for remediation actions. For example, if you want to monitor activities on your Amazon EC2 instances, you can create a CloudWatch Event*

[72] AWS Whitepaper, p. 13.

[73] Screenshot made by Privacy Company in the test set-up on 12 February 2021.

rule. When a specific activity happens on the Amazon EC2 instance and the event is captured in the logs, the rule triggers an AWS Lambda function, which sends a notification email about the event to the administrator. The email includes details such as when the event happened, which user performed the action, Amazon EC2 details, and more."

AWS also collects data through AWS Identity and Access Management (IAM), a web service used to securely control access to the tested AWS resources.

Log entries for IAM logging contain the following unique user identifiers:

- Username (in this case: *root*)
- principalId (*unique identifier for the entity that made that call*. The entity is the user, in this test the root account)
- accountId (in this test, the same as the principalID)
- public IP address of the user

Other data such as the timestamp, region, user agent, the URL referrer and the URL of the page they are logging into (the captured activity) are also personal data, when they relate to activities performed by natural persons that can be identified with the set of unique identifiers mentioned above.

Figure 14: Example of a IAM log file in CloudTrail[74]


```
{
 "eventVersion": "1.08",
 "userIdentity": {
 "type": "Root",
 "principalId": "539301343337",
 "arn": "arn:aws:iam::539301343337:root",
 "accountId": "539301343337",
 "accessKeyId": ""
 },
 "eventTime": "2021-02-11T17:05:26Z",
 "eventSource": "signin.amazonaws.com",
 "eventName": "ConsoleLogin",
 "awsRegion": "us-east-1",
 "sourceIPAddress": "82.217.32.212",
 "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 11_2_0) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/88.0.4324.150 Safari/537.36",
 "requestParameters": null,
 "responseElements": {
 "ConsoleLogin": "Success"
 },
 "additionalEventData": {
 "LoginTo": "https://console.aws.amazon.com/console/home?nc2=h_ct&src=header-signin&state=hashArgs%23&isauthcode=true",
 "MobileVersion": "No",
 "MFAUsed": "No"
 },
 "eventID": "d37ad8cb-171e-4ec4-ad16-35a025473651",
 "readOnly": false,
 "eventType": "AwsConsoleSignIn",
 "managementEvent": true,
 "eventCategory": "Management",
```


[74] Log file accessed by Privacy Company on 11 February 2021.

```
 "recipientAccountId": "539301343337"
}
```


Since CloudTrail logs contain all API calls of system management activities, these logs also include logs about changes of the RDS configuration (see Figure 15 below, EventName: ModifyDBInstance).

Log entries for Amazon RDS logging contain the following unique user identifiers:

- Username (in this case: *root*)
- principalId (*unique identifier for the entity that made that call*. The entity is the user, in this test the root account)
- accountId (in this test, the same as the principalID)
- public IP adress of the user

Other data such as the timestamp, region, user agent, the URL referrer and the URL of the page they are logging into (the captured activity) are also personal data, when they relate to activities performed by natural persons that can be identified with the set of unique identifiers mentioned above.

By default, AWS does not log access by users (admins) to applications inside the RDS instances. However, customers can choose to use additional services that may show such access (outside of the scope of this DPIA).

Figure 15: Log file Amazon RDS


```
{
 "eventVersion": "1.08",
 "userIdentity": {
 "type": "Root",
 "principalId": "539301343337",
 "arn": "arn:aws:iam::539301343337:root",
 "accountId": "539301343337",
 "accessKeyId": "ASIAX3EG5OBUU5DBWEKJ",
 "sessionContext": {
  "sessionIssuer": {},
  "webIdFederationData": {},
  "attributes": {
  "mfaAuthenticated": "false",
  "creationDate": "2021-02-11T17:05:26Z"
  }
 }
 },
 "eventTime": "2021-02-11T22:46:18Z",
 "eventSource": "rds.amazonaws.com",
 "eventName": "ModifyDBInstance",
 "awsRegion": "eu-central-1",
 "sourceIPAddress": "82.217.32.212",
 "userAgent": "aws-internal/3 aws-sdk-java/1.11.848 Linux/4.9.230-
0.1.ac.223.84.332.metal1.x86_64 OpenJDK_64-Bit_Server_VM/25.202-b08 java/1.8.0_202
vendor/Oracle_Corporation",
 "requestParameters": {
 "dBInstanceIdentifier": "database-1",
 "applyImmediately": true,
 "allowMajorVersionUpgrade": false,
 "deletionProtection": false,
 "maxAllocatedStorage": 1000,
```

```
"dryRun": false
},
"responseElements": {
"dBInstanceIdentifier": "database-1",
"dBInstanceClass": "db.m5.large",
"engine": "mysql",
"dBInstanceStatus": "available",
"masterUsername": "admin",
"endpoint": {
 "address": "database-1.ckhzxphuziuy.eu-central-1.rds.amazonaws.com",
 "port": 3306,
 "hostedZoneId": "Z1RLNUO7B9Q6NB"
},
"allocatedStorage": 100,
"instanceCreateTime": "Feb 11, 2021 10:17:54 PM",
"preferredBackupWindow": "23:01-23:31",
"backupRetentionPeriod": 7,
"dBSecurityGroups": [],
"vpcSecurityGroups": [
 {
 "vpcSecurityGroupId": "sg-889486eb",
 "status": "active"
 }
],
"dBParameterGroups": [
 {
 "dBParameterGroupName": "default.mysql8.0",
 "parameterApplyStatus": "in-sync"
 }
],
"availabilityZone": "eu-central-1b",
"dBSubnetGroup": {
 "dBSubnetGroupName": "default-vpc-b230c7d8",
 "dBSubnetGroupDescription": "Created from the RDS Management Console",
 "vpcId": "vpc-b230c7d8",
 "subnetGroupStatus": "Complete",
 "subnets": [
 {
  "subnetIdentifier": "subnet-03508f69",
  "subnetAvailabilityZone": {
  "name": "eu-central-1a"
  },
  "subnetOutpost": {},
  "subnetStatus": "Active"
 },
 {
  "subnetIdentifier": "subnet-62687d1f",
  "subnetAvailabilityZone": {
  "name": "eu-central-1b"
  },
  "subnetOutpost": {},
  "subnetStatus": "Active"
 },
 {
  "subnetIdentifier": "subnet-7b415936",
```

```
  "subnetAvailabilityZone": {
  "name": "eu-central-1c"
  },
  "subnetOutpost": {},
  "subnetStatus": "Active"
 }
 ]
},
"preferredMaintenanceWindow": "sun:00:41-sun:01:11",
"pendingModifiedValues": {},
"latestRestorableTime": "Feb 11, 2021 10:40:00 PM",
"multiAZ": false,
"engineVersion": "8.0.20",
"autoMinorVersionUpgrade": true,
"readReplicaDBInstanceIdentifiers": [],
"licenseModel": "general-public-license",
"iops": 3000,
"optionGroupMemberships": [
 {
 "optionGroupName": "default:mysql-8-0",
 "status": "in-sync"
 }
],
"publiclyAccessible": false,
"storageType": "io1",
"dbInstancePort": 0,
"storageEncrypted": true,
"kmsKeyId":       "arn:aws:kms:eu-central-1:539301343337:key/94b6e856-f709-47ef-89ab-
f15e9b2034f9",
"dbiResourceId": "db-JYSIPXENDOHADI2E65OVHOEQO4",
"cACertificateIdentifier": "rds-ca-2019",
"domainMemberships": [],
"copyTagsToSnapshot": true,
"monitoringInterval": 60,
"enhancedMonitoringResourceArn":              "arn:aws:logs:eu-central-1:539301343337:log-
group:RDSOSMetrics:log-stream:db-JYSIPXENDOHADI2E65OVHOEQO4",
"monitoringRoleArn": "arn:aws:iam::539301343337:role/rds-monitoring-role",
"dBInstanceArn": "arn:aws:rds:eu-central-1:539301343337:db:database-1",
"iAMDatabaseAuthenticationEnabled": false,
"performanceInsightsEnabled": true,
"performanceInsightsKMSKeyId":   "arn:aws:kms:eu-central-1:539301343337:key/94b6e856-
f709-47ef-89ab-f15e9b2034f9",
"performanceInsightsRetentionPeriod": 7,
"deletionProtection": false,
"associatedRoles": [],
"httpEndpointEnabled": false,
"maxAllocatedStorage": 1000,
"tagList": [],
"customerOwnedIpEnabled": false,
"networkType": "IPV4"
},
"requestID": "9e9bc44b-e1c7-412a-ae45-ef475e0af5f4",
"eventID": "f62e1158-144c-4342-a213-ea865353c8b1",
"readOnly": false,
"eventType": "AwsApiCall",
```

```
"managementEvent": true,
"eventCategory": "Management",
"recipientAccountId": "539301343337"
}
```


3.1.3 *Other logs centrally generated and processed by AWS*

Privacy Company has attempted throughout the production of this DPIA to discover what (network) log files AWS collects independently from any configurations and procurement of extra services by its customers. Privacy Company requested access to specific details on AWS internal logging that AWS considers confidential.

AWS confirms the existence of network and service logs by reference to summaries of these logs in the AWS SOC Type 1 and Type 2 audit reports. The AWS SOC 2 Type I Privacy Report provides*:*

> *"AWS maintains centralized repositories that provide core log archival functionality available for internal use by AWS service teams. Leveraging S3 for high scalability, durability, and availability, it allows service teams to collect, archive, and view service logs in a central log service. (…) Processes are implemented to protect logs and audit tools from unauthorized access, modification, and deletion."*

The AWS SOC 2 Type 2 report provides:

> *"Production hosts at AWS are equipped with logging for security purposes. This service logs all human actions on hosts, including logons, failed logon attempts, and logoffs. These logs are stored and accessible by AWS security teams for root cause analysis in the event of a suspected security incident. Logs for a given host are also available to the team that owns that host. A frontend log analysis tool is available to service teams to search their logs for operational and security analysis. Processes are implemented to protect logs and audit tools from unauthorized access, modification, and deletion."*[75]

Privacy Company has not been allowed to see any of the contents of the logs created by AWS about the test set-up of this DPIA. AWS did - [**Confidential**] confirmed that it collects log files with personal data as part of its security program.

AWS also pointed to a recently completed new audit, against the German C5:2020 standard, about the period from 1 October 2021 to 30 September 2022.[76] In the two operational controls OPS-11 and OPS-12 the processing of Diagnostic Data is covered. Based on an interview with the responsible AWS security assurance manager, the auditors confirm compliance with OPS-11, that AWS solely collects and uses the Diagnostic Data for the three purposes of billing, incident management and security management.[77]

---

[75] AWS response to DPIA questions, 17 July 2020, answer to Q3g.

[76] AWS information about the C5:2020 audit, URL: https://aws.amazon.com/compliance/bsi-c5/. The contents of the standard are documented in the Criteria Catalog at: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/CloudComputing/ComplianceControlsCatalogue/2020/C5_2020.pdf.

[77] [**Confidential**]

The report also states that AWS complies with the other criteria in OPS-11:

- Exclusively anonymous metadata to deploy and enhance the cloud service so that no conclusions can be drawn about the cloud customer or user;
- No commercial use;
- Storage for a fixed period reasonably related to the purposes of the collection;
- Immediate deletion if the purposes of the collection are fulfilled and further storage is no longer necessary; and
- Provision to cloud customers according to contractual agreements.[78]

The auditors confirm that AWS also complies with the additional criterion that

> *"personal data are automatically removed from the log data before the CSP processes it as far as technically possible. The removal is done in such a way that allows the CSP to continue to use the log data for the purpose for which it was collected."*[79]

The auditors also confirm that AWS complies with OPS-12, that AWS retains the Diagnostic Data for the specified periods, and deletes the data when further retention is no longer necessary for the purpose of the collection. The auditors do not explain if they have verified this criterion by looking at the available data and comparing those with specific retention periods.

To Privacy Company, none of the categories of collected data [**Confidential**] seem excessive for security purposes. The collection rules seem in line with cloud provider industry practices. However, since the security logs may also include IP addresses of visitors to Dutch government applications or websites (if they are hosted on AWS), both AWS and/or the Dutch government organisations are required to inform data subjects about the existence, purpose and length of this data processing, regardless of their GDPR role as processor or controller (see Section 5 of this DPIA for an assessment of the factual roles). Though AWS has attempted to give Privacy Company some insight in its policy and rules governing the security purposes for the processing, Privacy Company was not allowed to access and assess the actual data collection by AWS, not from the data generated by Privacy Company itself, nor from any other kind of test set-up. Therefore this DPIA cannot confirm there are no excessive personal data in these logs.

3.1.4 System usage statistics

AWS also shows system usage statistics to admins about the basic system load, such as CPU load, disk operations and network load. These graphs do not include personal data. See Figures Figure 16 and Figure 17 below.

[78] BSI controls catalogue, p. 64.

[79] [**Confidential**]

Figure 16: AWS Console monitoring usage of EC2 services[80]

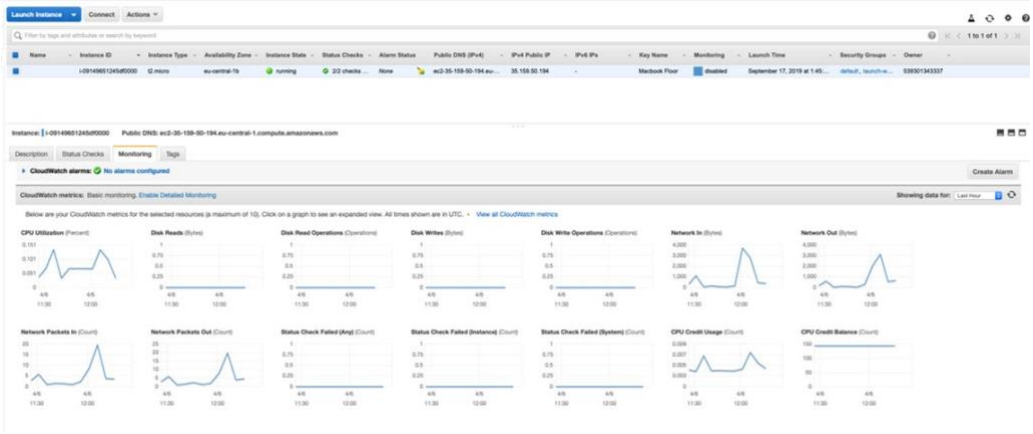

Figure 17: AWS Console monitoring usage of RDS[81]

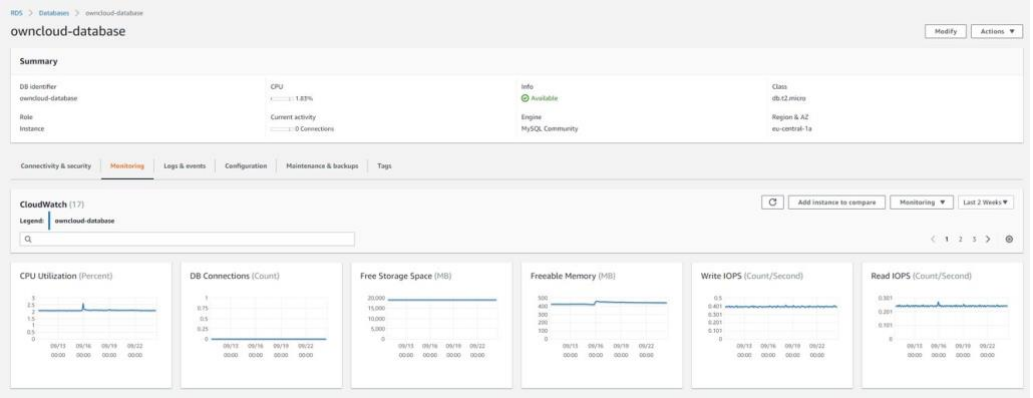

3.2 Website Data

AWS describes in its Cookie Notice that it collects information with the help of cookies, pixels and other similar technologies "*to recognize your browser or device, learn more about your interests, provide you with essential features and services Admin Console, and (…).*" These three purposes can only be achieved if the data stream contains unique user identifiers. AWS explains that the information collected through cookies includes online identifiers such as Authentication and security credential information, the IP address and the unique properties of the browser, in combination with detailed information about time, location, URL referrer and computer and device information.

- *"Network and connection information, such as the Internet protocol (IP) address used to connect your computer or other device to the Internet and information about your Internet service provider*
- *Computer and device information, such as device, application, or browser type and version, browser plug-in type and version, operating system, or time zone setting*
- *The location of your device or computer*
- *Authentication and security credential information*
- *Content interaction information, such as content downloads, streams, and playback details, including duration and number of simultaneous streams and downloads*

[80] Screenshot captured by Privacy Company on 17 September 2019.

[81] Screenshot captured by Privacy Company on 22 September 2019.

- *The full Uniform Resource Locators (URL) clickstream to, through, and from our site (including date and time) and AWS Offerings, content you viewed or searched for, page response times, download errors, and page interaction information (such as scrolling, clicks, and mouse-overs).*"[82]

Some of the categories of data mentioned in this list are probably collected through webserver access logs and not through cookies and pixels. Such logs are often configured to record the log files in a text file in a Common Log Format. Commonly, web server access logs collect and store the client IP addresses, user agent strings, date, time, server name, server IP and services running, among many others. Such logs are necessary to be able to detect security incidents, but can also be used to monitor for errors and improve the user interface. These logs show who visited the website, where the visitors came from, what pages they visited, where they went to, and in case of access to the restricted access Admin Console and Support Centre, their authentication and security credentials.

3.2.1 AWS restricted access website

Privacy Company tested the restricted access Admin Console pages for EC2, S3, RDS and the Support Centre pages within the Admin Console.

First the admin has to sign in to the Admin Console page, as root user, or as IAM user. See Figure 18 below. The text in this banner explains that the site uses essential cookies, with a reference to its Cookie Notice. The text does not mention other cookies, such as performance cookies.

The Cookie Notice does not contain an overview of what these essential cookies are. AWS writes:

> "*Our cookies allow you to take advantage of some* ***essential and useful features****. Blocking some types of cookies may impact your experience of our sites."*[83]

At the bottom of the Cookie Notice, AWS publishes a list of 26 companies that may set third party cookies.

> "*Below is a list of the third parties that may set cookies when you use AWS Offerings. You can learn more about how these third parties use information collected through cookies by reviewing the privacy policies on their sites.*"[84]

[82] AWS Cookie Notice, 'Information we collect through cookies'.

[83] Idem.

[84] Idem.

Figure 18: AWS sign-in to the Admin Console pages[85]

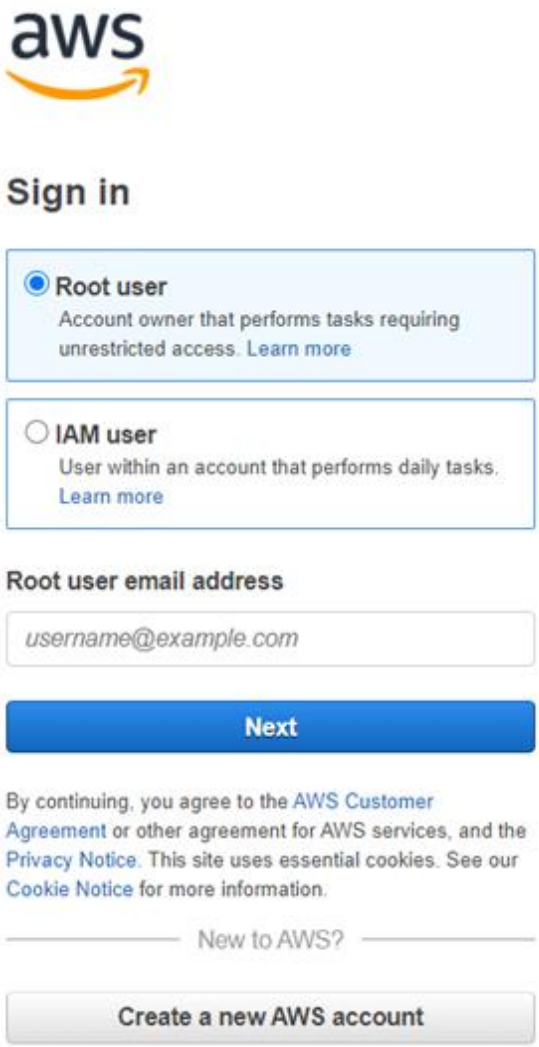

AWS does not explain for what purposes these third party cookies are set, and also fails to provide a hyperlink to the privacy and cookie explanations of these third parties. Some of the cookies clearly serve targeted advertising purposes, such as cookies from Oracle's BlueKai, Drawbridge, Google, Tapad and The Trade Desk.

This lack of information would be problematic, if any of those third party cookies were actually set visiting the AWS restricted access website (Admin Console and Support Centre). However, as documented in the technical appendix shared with AWS, during repeated testing only first party cookies were set and read, some with a very long retention period in the user's browser (until 2041).

In the dialogue with SLM Rijk, AWS explained it will never set advertising or third-party cookies on its restricted access websites. By default, AWS does however set Essential, Functional and Performance cookies.

AWS contractually commits to comply with the privacy by design and data minimisation principles from the GDPR, including minimisation of the processing of

[85] https://signin.aws.amazon.com, last checked 22 May 2023.

Website Data *to the extent strictly necessary* for the three agreed main processor purposes.

Figure 19: AWS Admin Console with cookie banner

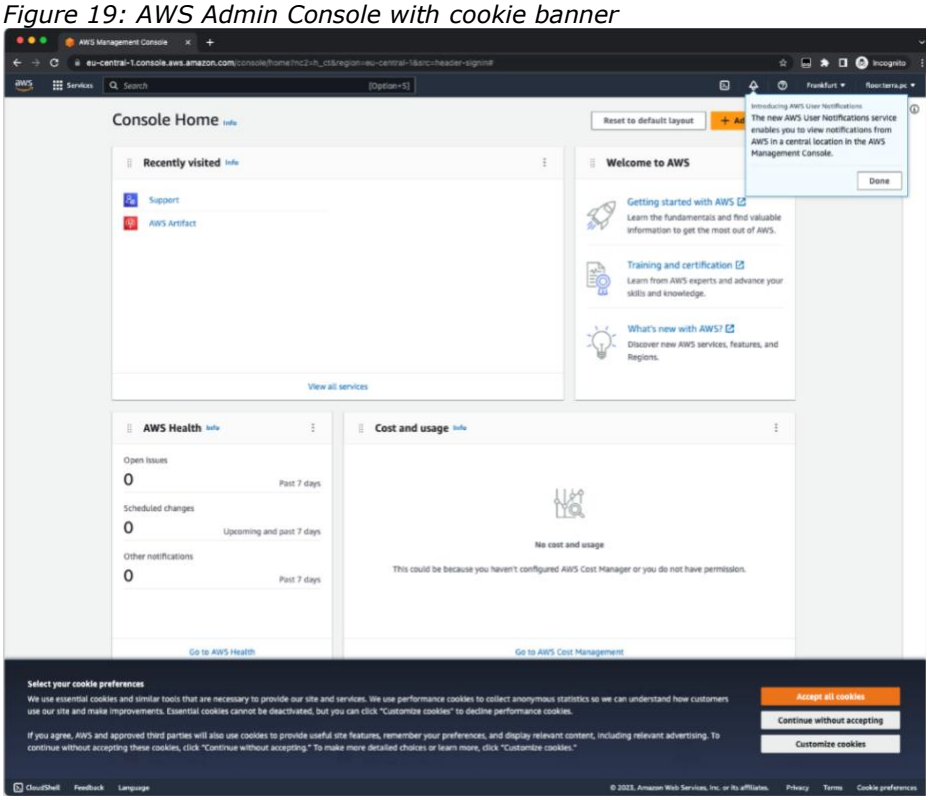

Figure 20: AWS Support Portal with cookie banner

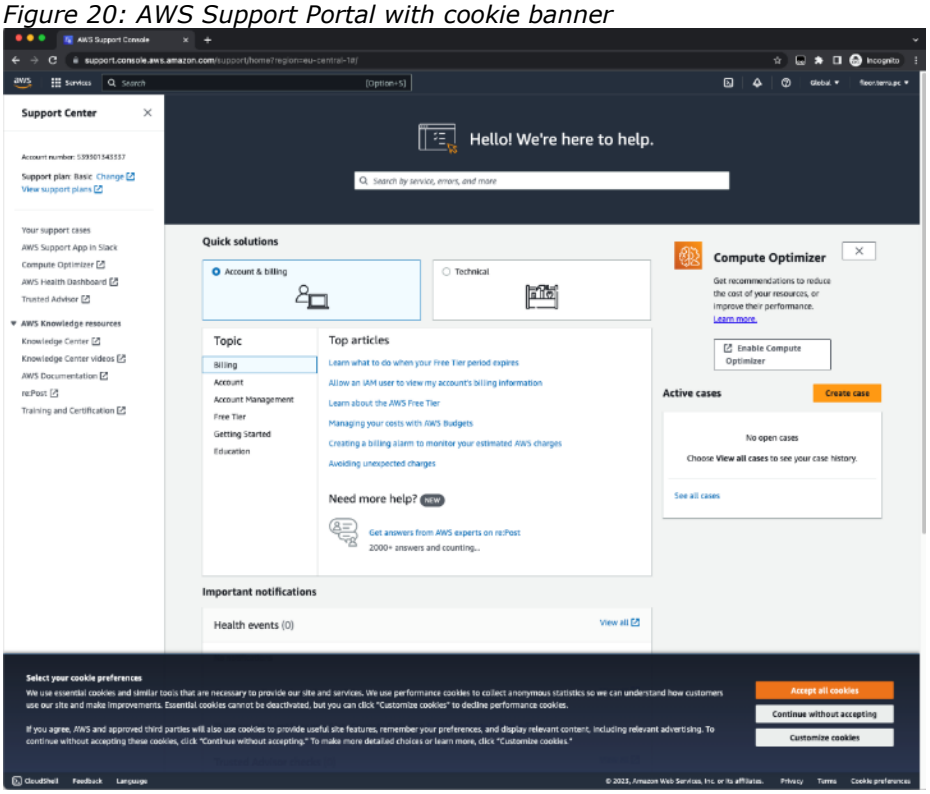

After signing in, AWS shows a cookie banner at the bottom of the Admin Console and the Support Centre page with a request to choose between 3 cookie options:

1. Accept all cookies
2. Continue without accepting
3. Customize cookies

See Figure 19 and

Figure 20 above. When an admin selects the second option (*Continue without accepting*), AWS sets 23 first-party session and permanent cookies. Twenty of these cookies have a unique value, but the cookies do seem to serve functional purposes such as session continuation, and not analytics (called performance cookies by AWS). See Figure 21 below.

Figure 21: AWS first party cookies set on restricted access Support Portal

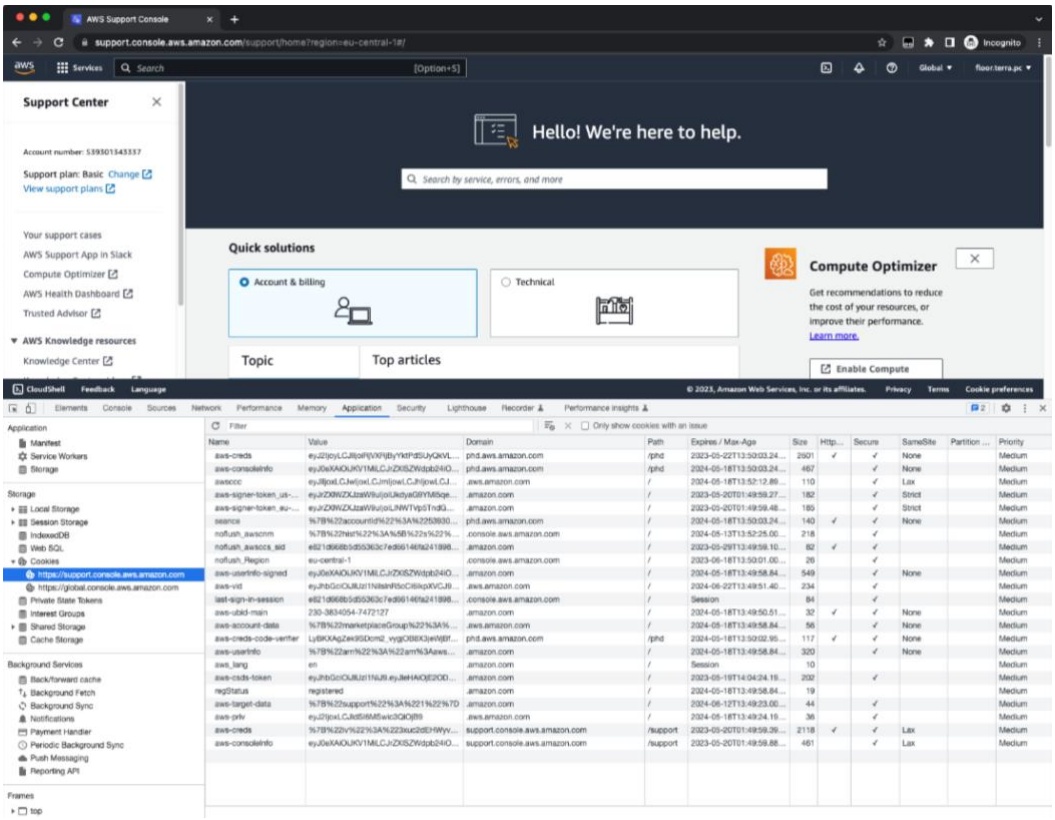

When retesting the effectivity of the cookie banner on 19 May 2023 and 20 June 2023, Privacy Company detected a secondary data stream. See Figure 22 below.

Figure 22: AWS collection of website telemetry data after selecting 'Continue without accepting'

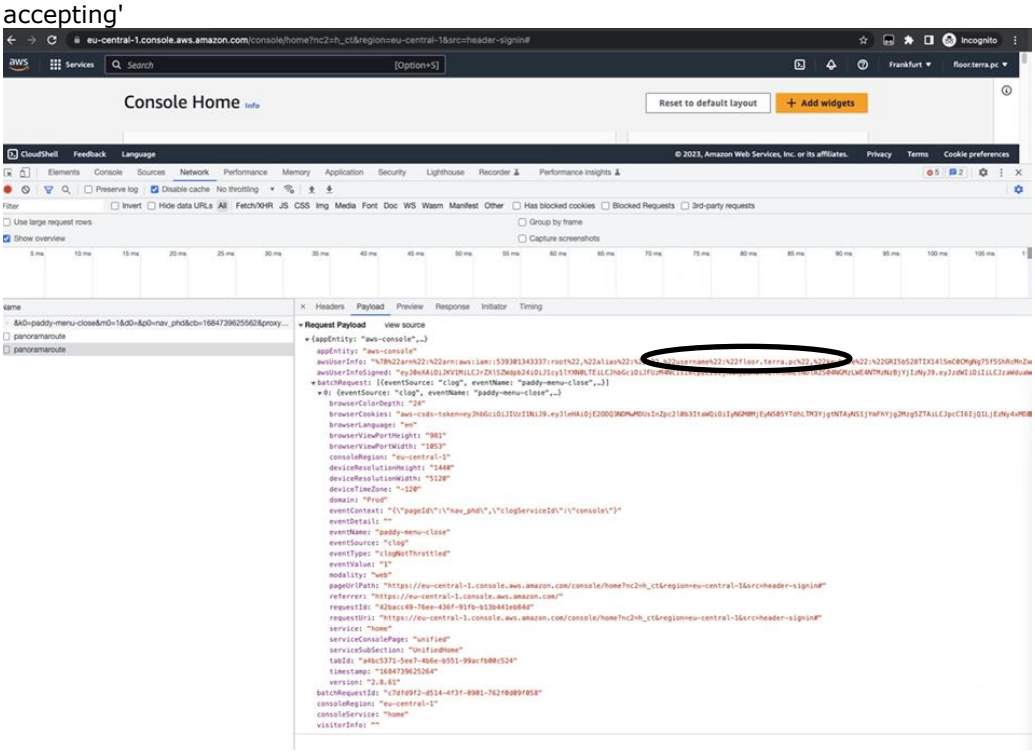

It is not clear if this stream was present before. AWS has programmed its restricted access website to generate and transmit telemetry data via the browser of visitors in events called 'telemetry' and 'panoramaroute'. Via these events, AWS apparently collects analytical data, regardless of the choice an admin made to continue without accepting cookies. These data include the full user account name, in this case: floor.terra.pc.

As shown in Figure 23 below, AWS itself collects these data, through its domain A2z.com/panoramaroute.

Figure 23: Destination of analytical website telemetry: AWS panoramaroute

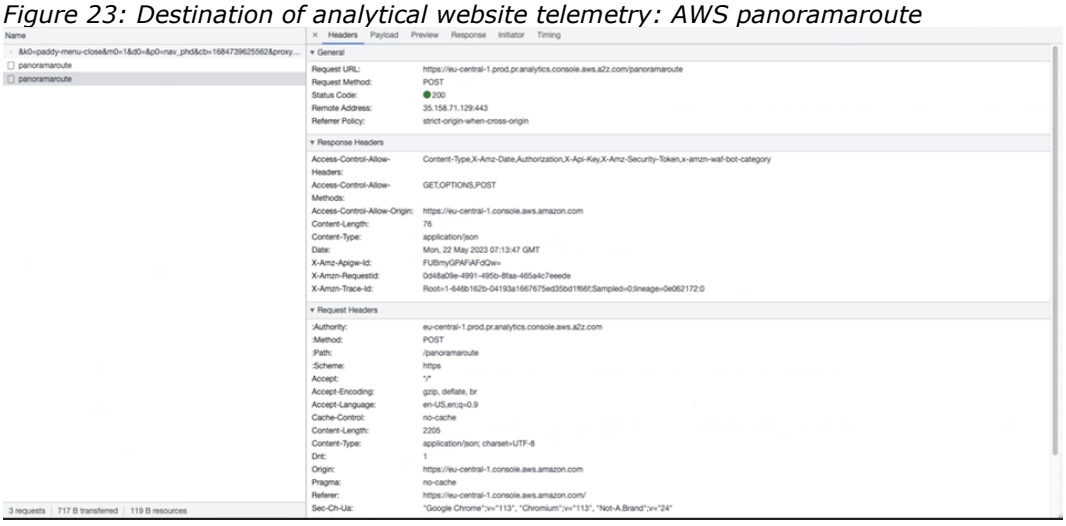

In reply to this finding, AWS pointed out that it does write in the cookie banner that admins must select 'Customize cookies' to opt out from the performance cookies, a.k.a. the analytical cookies.

When an admin selects that third option, AWS shows a screen with four cookie options. See Figure 24 below.

Figure 24: AWS screen with options to customise different cookies

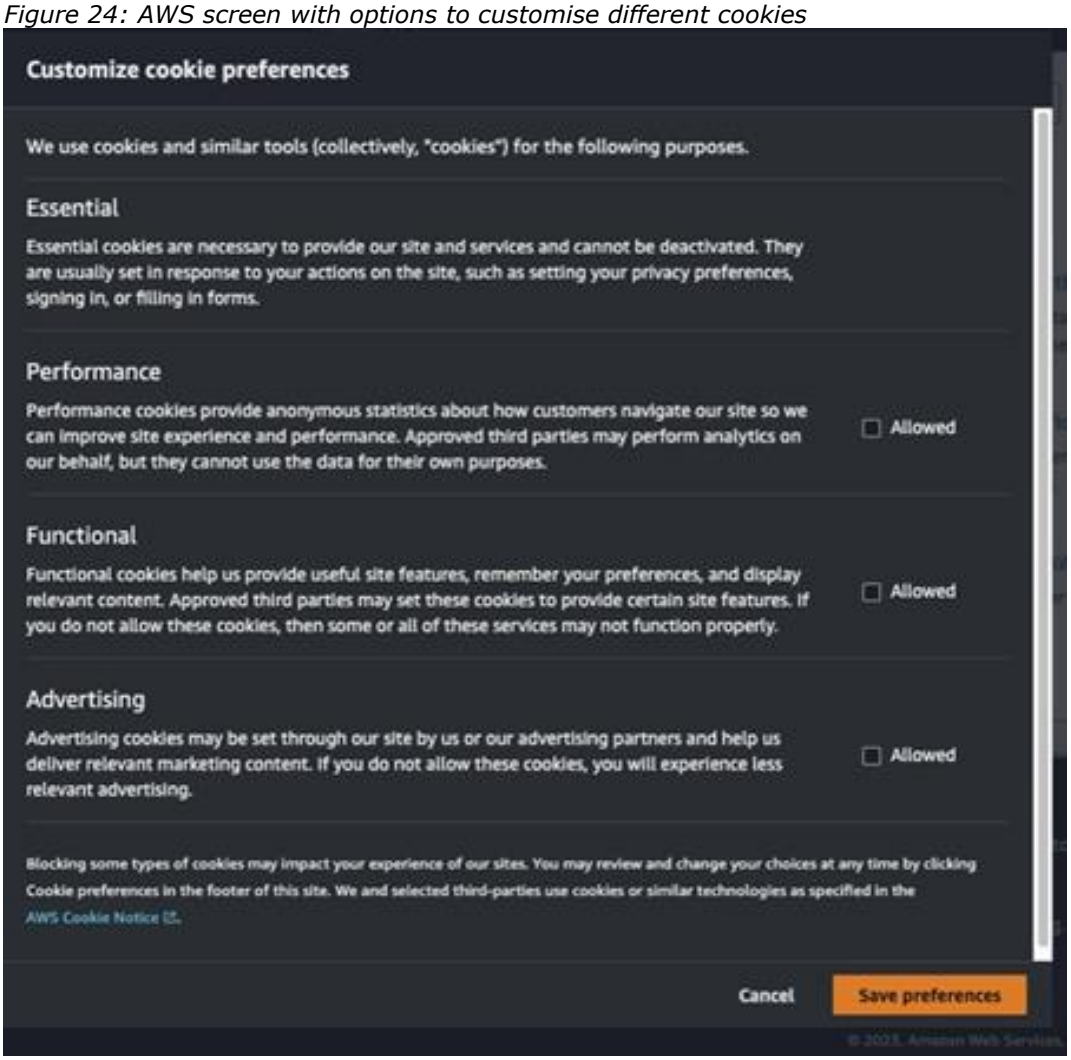

According to AWS's explanation the third option in the cookie banner 'Customize cookies' would allow admins to reject analytical cookies.

However, if an admin chooses this third 'customize cookies' option, and accepts the default setting in this pop-up screen, AWS still collects analytical personal data through website telemetry. These data include the full user account name. See Figure 25 below for the contents of the outgoing data traffic, and Figure 26 for the inclusion of the user account name. This same data collection was observed in a retest on 20 June 2023. AWS replied that it was still investigating this issue. This issue will be discussed in the ongoing dialogue between SLM Rijk and AWS.

Figure 25: Outgoing telemetry data after accepting only Essential cookies

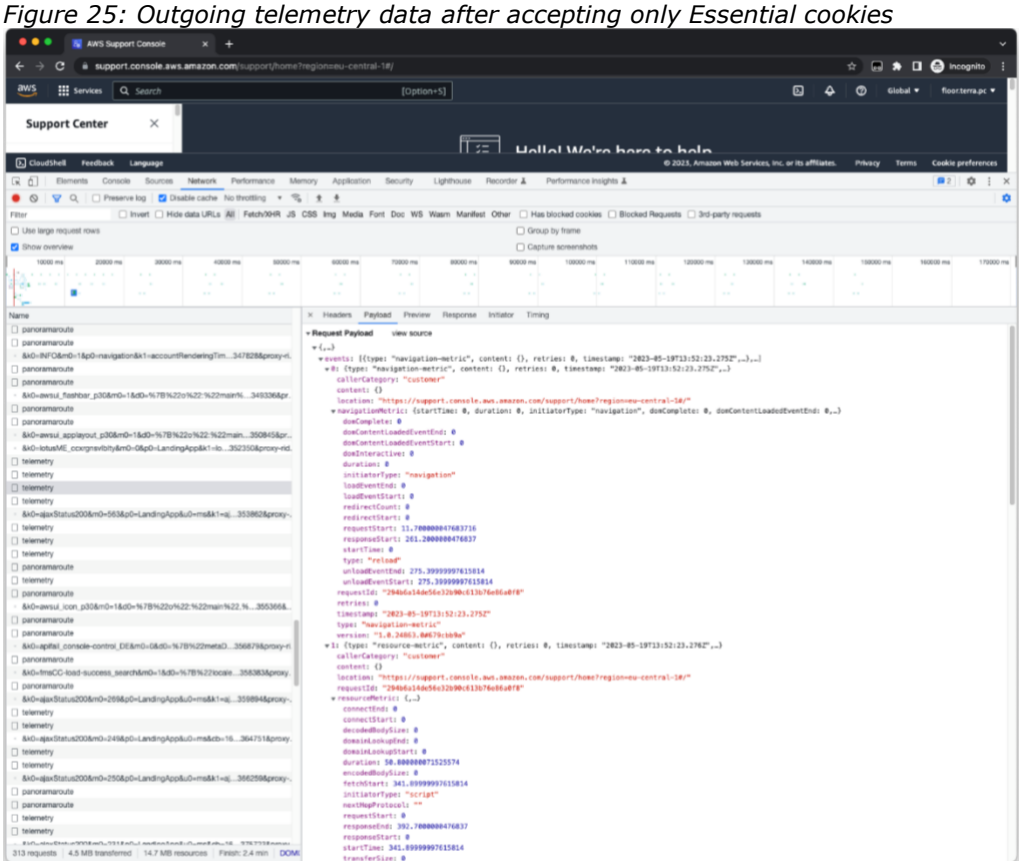

Figure 26: User account name in telemetry event panoramaroute

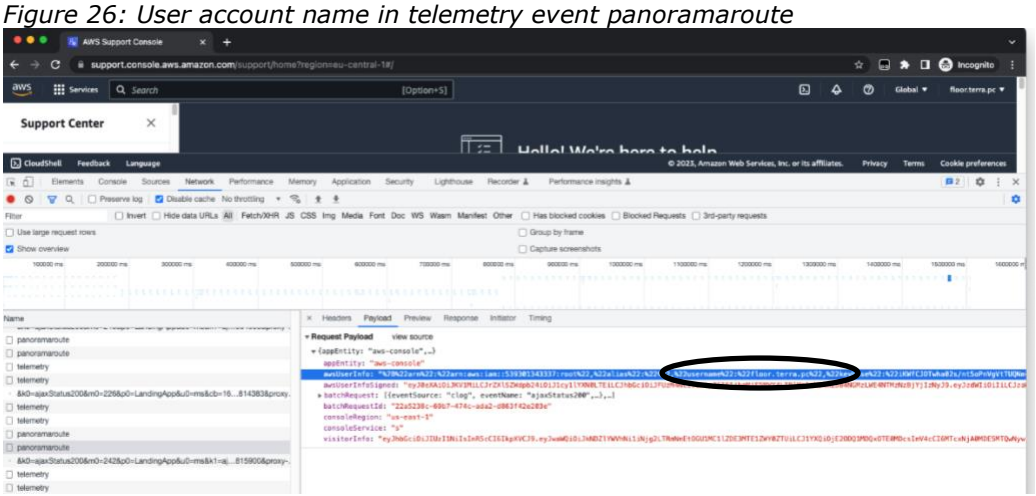

3.2.2 *Public access website*

When visiting the AWS publicly accessible website with a clean browser, it shows the same pop-up with the same request to customize cookie preferences, but a slightly different text, namely: "*We use essential cookies and similar tools that are necessary to provide our site and services. We use performance cookies to collect anonymous statistics so we can understand how customers use our site and make improvements. Essential cookies cannot be deactivated, but you can click "Customize cookies" to decline performance cookies.* ***If you agree, AWS and approved third parties will also use cookies to provide useful site features, remember your preferences, and display relevant content, including relevant advertising***. *To continue*

without accepting these cookies, click "Continue without accepting." To make more detailed choices or learn more, click "Customize cookies."

If a visitor clicks 'Customize cookies' AWS shows four options. The default setting on the public website is set to the level 'Performance' (where the default on the restricted access websites is 'Essential').

Figure 27: AWS Cookie preferences pop-up public website

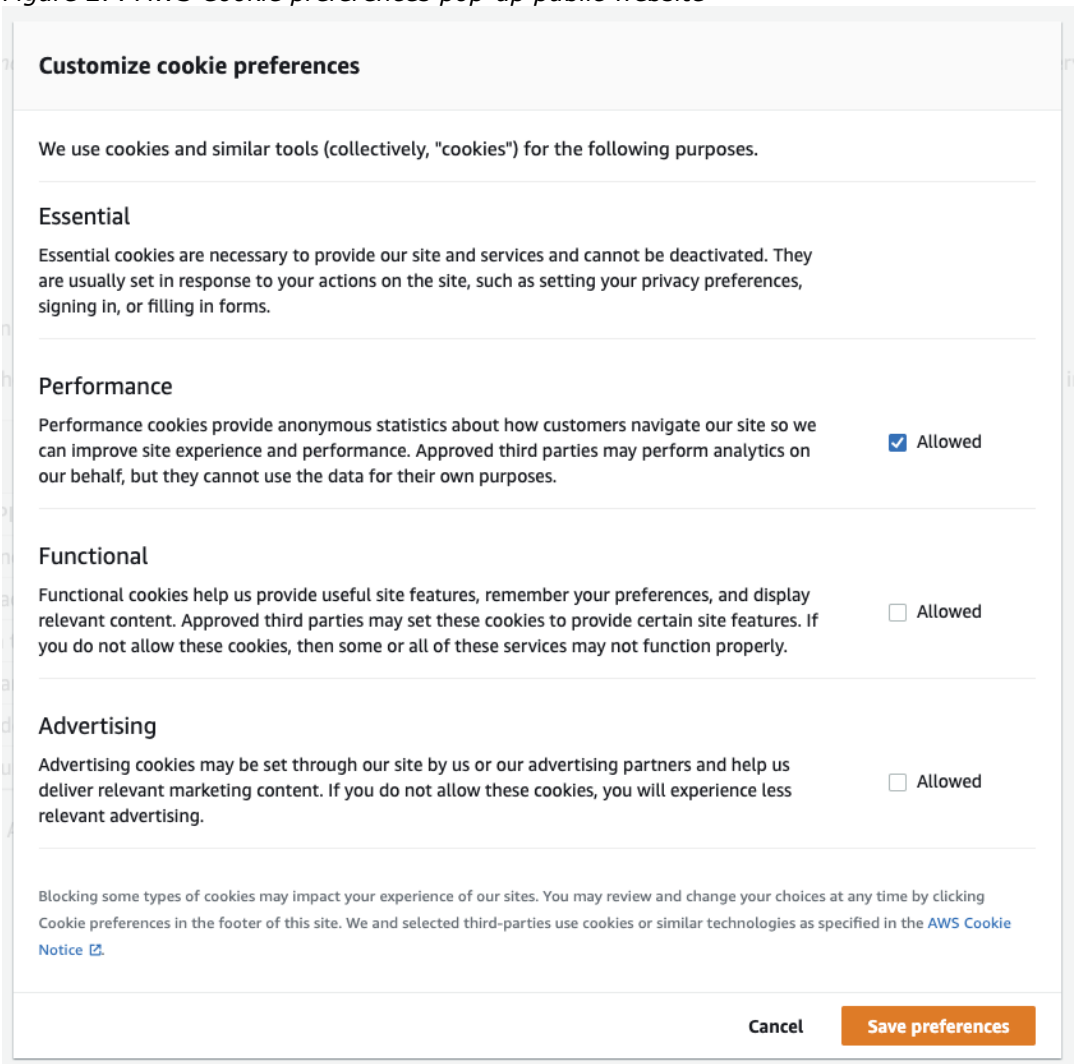

AWS explains that at this 'Performance' level, third party analytical cookies may be set and read, as long as these parties are prohibited from using the data for their own purposes. Privacy Company visited a selection of documentation pages on the AWS website without changing the default cookie settings.[86] However, no third-party cookies where observed, in spite of the long list of third parties that may set or read cookies according to AWS's Cookie Notice.

[86] AWS Documentation pages, with hyperlinks to all kinds of Guides and API References, Tutorials and Projects, SDKs and Toolkits and General Resources, URL: https://docs.aws.amazon.com/.

3.3 Results Data Subject Access Request

Privacy Company filed a Data Subject Access Request (DSAR) by e-mail of 30 September 2019 to an Amazon representative, after having performed the test scenarios. On 2 October 2019, Privacy Company received a standard reply.[87] In this mail, AWS refers to five sources that offer some access to some personal information:

1. *Access, view and edit your personal information via the My Account page of the AWS Management Console;*[88]
2. *Access payment information, charges and account activity, and security credentials via the AWS Account pages;*[89]
3. *Access and update communications preferences by visiting the AWS Communications Center;*[90]
4. *The possibility to file a Data Request Form;*[91]
5. *File a separate request for a log of the console activity for your AWS account by replying to this case. The relevant department will review your request and respond to your log request accordingly.*

In this mail AWS also refers to its Privacy Notice. This notice contains general information about the purposes of the data processing. However, as will be discussed in the next Section 4 about the purposes for the processing, the Privacy Notice does not provide a list of what categories of personal data are processed for what specific purposes. The reference to the Privacy Notice thus creates a puzzle for the requesting data subject with regard to the categories of personal data and the purposes for which they are processed. Privacy Company used option 4, and accessed the Data Request Form. This is a standard letter with a tick list to request access to specific data. The form (see Figure 28 below) can also be used to request rectification, erasure, objection, stop processing or portability of specific personal data. This form however does not allow data subjects to request access to Diagnostic Data collected by AWS that are not already accessible through the Management Console. According to AWS's standard reply, data subjects should file a separate request for these data.

On 10 October 2019 AWS explained that access requests should be made via the process described in the AWS Privacy Notice[92] and offered to give more context in a face-to-face meeting. Privacy Company used all available sources and avenues, and organised multiple technical meetings, but did not succeed in getting full access to all personal data. As concluded in Section 3.1.3 Privacy Company has no reason to assume any excessive data collection, but AWS has not provided any contractual guarantees (be it in a role as controller or as data processor) about data minimisation.

The 'missing' data from the DSAR response include:

1. use of the AWS Admin and Root Account outside of available logs in CloudTrail;
2. information recorded in the webserver access logs with information about IP address, end user, device and activities;

[87] E-mail AWS of 2 October 2019 to the test admin of Privacy Company.
[88] AWS Management Console, URL: https://console.aws.amazon.com/console/home.
[89] AWS Account pages, URL: https://aws.amazon.com/account/.
[90] AWS Communications Preference Center, URL: https://pages.awscloud.com/communication-preferences.html.
[91] AWS Data Request Form, URL: https://s3-us-west-2.amazonaws.com/aws-support-documents/Forms/AWSDataRequestForm.pdf.
[92] AWS Privacy Notice, URL: https://aws.amazon.com/privacy/.

3. information about filed Support Requests;
4. information collected by AWS in its network logs, and;
5. the necessary extra information elaborated in Article 15 (1) GDPR, in particular the provisions c (recipients), d (retention period), (f) the right to lodge a complaint with the DPA, (g) information about external sources and (h) automated decision-making, including profiling, in particular with regard to credit-scoring.

Figure 28: AWS Data Request Form

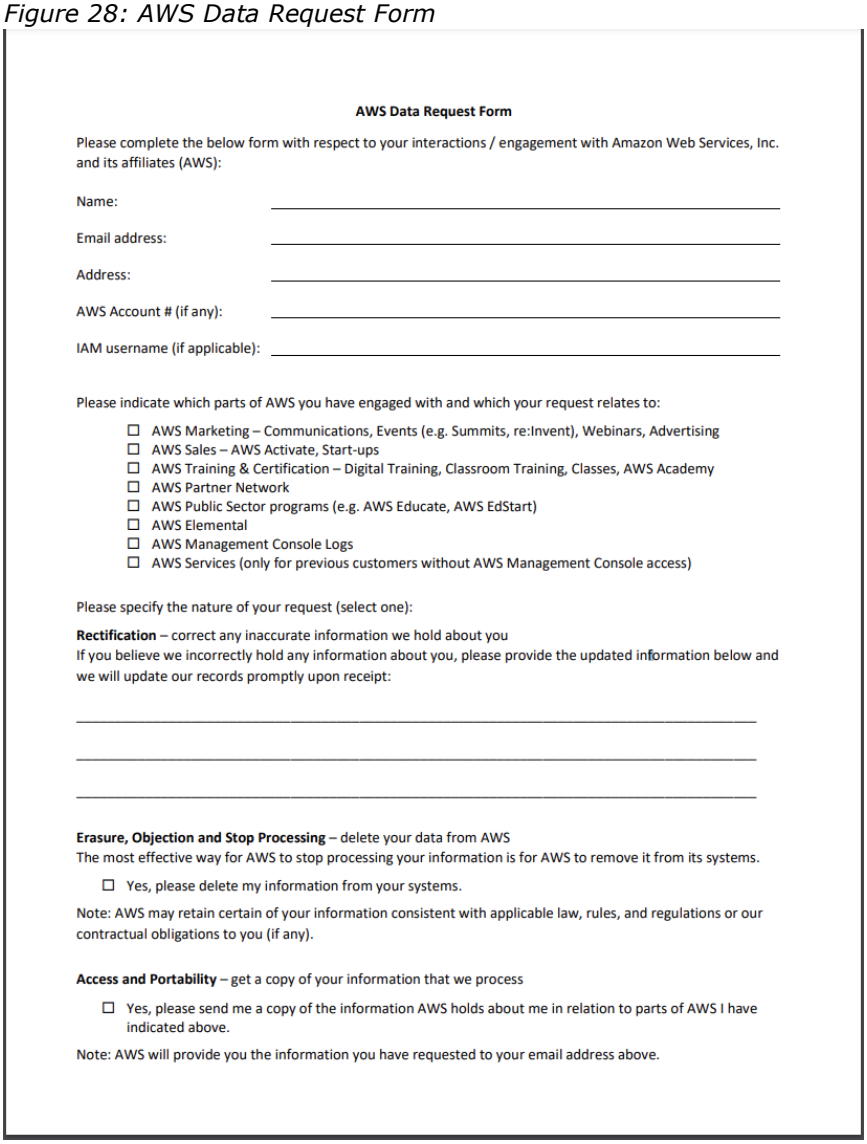

AWS Data Request Form

Please complete the below form with respect to your interactions / engagement with Amazon Web Services, Inc. and its affiliates (AWS):

Name: _____

Email address: _____

Address: _____

AWS Account # (if any): _____

IAM username (if applicable): _____

Please indicate which parts of AWS you have engaged with and which your request relates to:

- ☐ AWS Marketing – Communications, Events (e.g. Summits, re:Invent), Webinars, Advertising
- ☐ AWS Sales – AWS Activate, Start-ups
- ☐ AWS Training & Certification – Digital Training, Classroom Training, Classes, AWS Academy
- ☐ AWS Partner Network
- ☐ AWS Public Sector programs (e.g. AWS Educate, AWS EdStart)
- ☐ AWS Elemental
- ☐ AWS Management Console Logs
- ☐ AWS Services (only for previous customers without AWS Management Console access)

Please specify the nature of your request (select one):

Rectification – correct any inaccurate information we hold about you
If you believe we incorrectly hold any information about you, please provide the updated information below and we will update our records promptly upon receipt:

_____

_____

_____

Erasure, Objection and Stop Processing – delete your data from AWS
The most effective way for AWS to stop processing your information is for AWS to remove it from its systems.

- ☐ Yes, please delete my information from your systems.

Note: AWS may retain certain of your information consistent with applicable law, rules, and regulations or our contractual obligations to you (if any).

Access and Portability – get a copy of your information that we process

- ☐ Yes, please send me a copy of the information AWS holds about me in relation to parts of AWS I have indicated above.

Note: AWS will provide you the information you have requested to your email address above.

In sum, sections 3.1 to 3.3 show that all Account, Diagnostic, Support and Website Data should be treated as personal data. Privacy Company was not allowed to inspect any of the network and webserver access logs with personal data Amazon processes about the use of AWS services and website. AWS did not grant access to these logs either in reply to the DSAR. It is however extremely plausible that these logs include personal data, also relating to visitors of applications and websites hosted on AWS VMs.

In reply to Part A of this DPIA, AWS objected that not all data mentioned in the report as 'missing' from the Data Subject Access Request (DSAR). Not all data are personal

data or should be provided, as "*disclosure in the context of a DSAR could result in unlawful disclosure of personal data and be antithetical to the objective of data protection.*"[93]

AWS provided three additional explanations for missing data in reply to DSAR requests:

1. AWS does not provide access to website logs, as the combination of a unique cookie ID and the IP address is insufficient to reliably identify a unique user.
2. AWS as processor provides its customers with access to Diagnostic Data via CloudTrail: this includes individual usage data that AWS also uses itself for invoices.
3. AWS does not provide access to security events from its own core security logs, because the data are 'not meaningful'. There are few directly identifiable personal data in these logs, only device and user identifiers plus IP addresses, and optionally sometimes configuration data, such as a file name.

These explanations will be assessed in Section 15 of this DPIA.

4. Purposes of the Processing

Under the GDPR, the principle of 'purpose limitation' dictates that personal data may only be collected for specified, explicit and legitimate purposes, and may not be further processed in a manner that is incompatible with the initial purpose. The purposes are qualified and assessed in part B of this DPIA. This Section provides a factual description of the purposes of the processing of Customer Data and Diagnostic Data by government organisations and AWS.

4.1 Purposes government organisations

The general interests government organisations may have to use Amazon Web Services are described in Section 6.1.

The purposes government organisations may have to process the Content Data are out of scope of this umbrella DPIA.

Government organisations have access to some Diagnostic Data collected by AWS about the individual behaviour of system administrators through the S3 access logs (Section 3.1.1), the CloudTrail logs (Section 3.1.2) and S3 bucket logs (Section 3.1.3).

Government organisations may need to process these data to comply with information security requirements, to verify access authorisations, to investigate and mitigate data security breaches and to comply with data subject right requests.

As data controllers, government organisations must determine when they need to access log files generated by AWS, what extra applications they want to use to get a better view of the available data and set security alerts, determine what retention periods are necessary to comply with security requirements while still complying with the data minimisation principle from the GDPR, and for what specific purposes specific

[93] AWS response to part A of the DPIA, 1 October 2021, Par. 6, p. 2.

personal data in the log files may be (further) processed and analysed. These specific purposes are not in scope of this umbrella DPIA.

4.2 Purposes AWS Content, Diagnostic, Account, Support and restricted Website Data (data processor)

Initially, AWS considered itself to be an independent data controller for the processing of personal data in Diagnostic Data, Website Data, Contact Data and Support Data. AWS continues to describe the purposes of its processing of personal data as a data controller in its Privacy Notice. Privacy Company identified other relevant purposes in other applicable contractual documentation. As a result of the discussions with the Dutch government, these findings are no longer relevant, and are largely removed from this DPIA.

Based on the contract with the Dutch government AWS acts as a data processor for the five categories of personal data (Content, Diagnostic, Account, Support and Restricted access Website Data). AWS may process these personal data to provide and maintain the services, secure the services and AWS network, provide customer-requested support, and perform basis troubleshooting, but only to the extent necessary to achieve these purposes. These purposes each have specified sub purposes, as shown in Table 1 below.

*Table 1: Overview of 3 data processing purposes with sub purposes [**Confidential**]*

Contractually AWS in its role as data processor may not determine any "further" or "compatible" purposes (within the meaning of Articles 5(1)(b) and 6(4) of the GDPR) other than the agreed compatible purposes. (See Section 4.3 below).

Additionally, AWS is prohibited from processing personal data for advertising purposes, or for profiling, data analytics and market research, except where such processing is explicitly allowed by the customer or if the individual admin has given valid consent.

4.3 Agreed compatible purposes

AWS is authorised by the Dutch government to process some personal data as data controller for a list of agreed compatible purposes, but solely when the processing is strictly necessary and proportionate [**Confidential**]. This includes billing and calculating employee compensation; complying with legal obligations; responding to data subject requests in AWS's role as controller; scanning to detect violations of the Acceptable Use Policy; combatting fraud, cybercrime and cyber-attacks; and analysing, improving and optimising the performance and core functionality of the Services.

Some of these purposes may not appear self-evident, and are explained in more detail below. Most importantly, AWS has committed to update an overview of the categories of personal data that AWS processes as a controller for agreed compatible purposes. Privacy Company has not yet seen that overview.

4.3.1 Compensation

AWS is permitted to use the necessary personal data to calculate compensation for sales staff, based on statics on infrastructure (not account!) usage data. If government organisations follow the recommendation to use pseudonymous accounts for admins, AWS promises never to re-identify these data.

4.3.2 Complying with legal obligations

AWS may be ordered to disclose personal data pursuant to a *valid and binding order of a governmental body (such as a subpoena or court order)*."[94]

According to the GDPR, only data controllers may take decisions to disclose personal data to governmental agencies outside of the EU. Article 48 of the GDPR creates an exception to this rule. This provision acknowledges that a data processor may sometimes be forced by order of a court or administrative authority in a third country, outside of the EU, to transfer or disclose personal data. Such orders may only be recognised or enforced in any manner if they are based on an international agreement such as a mutual legal assistance treaty. This exception is titled "*Transfers or disclosures not authorised by Union law*". This exception however does not change the main rule that only data controllers may take decisions whether to hand over personal data to governmental agencies outside of the EU. That is why data processors must redirect such orders to the data controllers.

[**Confidential**]

In February 2021 AWS published a Supplementary Addendum to its public DPA.[95] In this Addendum AWS provides guarantees with regard to its treatment of requests for Customer (Content) Data.[96] AWS promises to "*use every reasonable effort to redirect the requesting party*" to its Enterprise customer. If compelled to disclose Customer Data, AWS will notify its customer if permitted, or, if prohibited, AWS "*will use all reasonable and lawful efforts to obtain a waiver*."

AWS also commits to *challenge any overbroad or inappropriate Request (including where such Request conflicts with the law of the European Union or applicable Member State law).* If AWS is nonetheless compelled to disclose it will disclose "*only the minimum amount of Customer Data necessary to satisfy the Request*."[97]

AWS commits to follow these steps:

1. AWS will verify it is a lawful and binding order;
2. AWS will use every reasonable effort to redirect the authority to request the personal Data directly from the customer;
3. If AWS is compelled to disclose personal data, AWS will challenge any overbroad or inappropriate Request.
4. AWS will promptly notify the customer to allow the customer to seek a protective order or other appropriate remedy, if AWS is legally permitted to do so.

---

[94] In Section 3.2 of the Customer Agreement, AWS mentions the purposes of maintaining the Service and complying with the law, in the sentence: "We will not access or use Your Content except as necessary to maintain or provide the Service Offerings, or as necessary to comply with the law or a binding order of a governmental body."

[95] AWS Security Blog, AWS and EU data transfers: strengthened commitments to protect customer data, 17 February 2021, URL: https://aws.amazon.com/blogs/security/aws-and-eu-data-transfers-strengthened-commitments-to-protect-customer-data/

[96] AWS Supplementary Addendum to AWS GDPR DPA, February 2021, URL: https://d1.awsstatic.com/Supplementary_Addendum_to_the_AWS_GDPR_DPA.pdf.

[97] Idem, Sections 1.2 and 1.3.

5. If AWS is prohibited from notifying the customer, it will use all reasonable and lawful efforts to obtain a waiver of prohibition.
6. AWS will comply with Clause 14(e) of the SCC: "*The data importer agrees to **notify the data exporter promptly** if, after having agreed to these Clauses and for the duration of the contract, it has reason to believe that it is or has become subject to laws or practices not in line with the requirements under paragraph (a), including following a change in the laws of the third country or **a measure (such as a disclosure request) indicating an application of such laws in practice that is not in line with the requirements in paragraph (a).***"

AWS provides public information about its disclosure policy in a Help Q&A.[98] This URL lists bi-annual information request reports since 2015. These reports show the amount of subpoenas, requests and court orders received, country of origin of the requests processed by AWS, the percentage of requests relating to Content and Non-Content as well as the percentage of honoured requests for Content and Non-Content.

According to these reports none of the subpoenas, search warrants and court orders resulted in the disclosure to the U.S. government of enterprise or government Content Data located outside the United States. Since AWS included the metric in the reports (July 2020), the reports notes:

> "*How many requests resulted in the disclosure to the U.S. government of enterprise or government content data located outside the United States?*
>
> *None*."

AWS does not publish specific information if it has ever disclosed *Diagnostic Data* (Non-Content Data) from government or enterprise customers that host their data outside of the USA to law enforcement or security services.[99]

These numbers do not include possible National Security Requests. AWS describes that the range of National Security Requests since 2020 ranged between 0 and 249.

In the 2021 and 2022 reports, AWS does not publish detailed information about reviewing or fighting the legality of hand-over requests of Content Data based on the CLOUD Act. AWS only states: *"Amazon continues to object to overbroad or otherwise inappropriate requests as a matter of course regardless of where data is located."*[100] AWS has also published a blog about its commitment to protect Customer (Content) Data from these requests.[101]

---

[98] AWS Help & Customer Service, Law Enforcement Information Requests, URL: https://www.amazon.com/gp/help/customer/display.html?nodeId=GYSDRGWQ2C2CRYEF

[99] AWS transparency reports from 2015 through to 31 December 2022, URL: https://aws.amazon.com/compliance/amazon-information-requests/ .

[100] Amazon Information Request Report, 31 January 2021, URL: https://d1.awsstatic.com/certifications/Information_Request_Report_December_2020.pdf and Amazon Information Request Report, 31 June 2021, URL: https://d1.awsstatic.com/Information_Request_Report_June_2021_x.pdf, and the last report, from 1 July to 31 December 2022, URL: https://d1.awsstatic.com/Security/pdfs/Amazon_Information_Request_Report.pdf.

[101] AWS Security Blog, AWS and EU data transfers: strengthened commitments to protect customer data, 17 February 2021, URL: https://aws.amazon.com/blogs/security/aws-and-eu-data-transfers-strengthened-commitments-to-protect-customer-data/

In its Help Q&A AWS states it always notifies customers, unless it receives a gagging order, or has a clear indication of illegal conduct:

> "*Unless prohibited from doing so or there is clear indication of illegal conduct in connection with the use of Amazon products or services, Amazon notifies customers before disclosing content information.*"[102]

AWS also writes it takes a public stance on this matter.

> "*We have repeatedly challenged government demands for customer information that we believed were overbroad, winning decisions that have helped to set the legal standards for protecting customer speech and privacy interests. We also advocate in Congress to modernize outdated privacy laws to require law enforcement to obtain a search warrant from a court to get the content of customer communications. That's the appropriate standard, and it's the standard we follow.*"[103]

4.3.3 Scanning

Purpose no. 8 identifies scanning. The AWS Acceptable Use Policy (AUP) is part of the enrolment framework. In this document, AWS mentions processing of Content Data to detect violations of the AUP. AWS writes:

> *"We may investigate any suspected violation of this Policy, and remove or disable access to any content or resource that violates this Policy. You agree to cooperate with us to remedy any violation.*
>
> *When determining whether there has been a violation of this Policy, we may consider your ability and willingness to comply with this Policy, including the policies and processes you have in place to prevent or identify and remove any prohibited content or activity.*"[104]

To prevent any misunderstanding about the possibility of automated content scanning, AWS is contractually prohibited from automated scanning of Content Data to identify potentially abusive content or activity.

Scanning is only allowed under very limited circumstances, e.g. Amazon Simple Email Service scans a percentage of outgoing emails for SPAM and other types of email abuse in line with industry standards.

Additionally, AWS has explained that abuse complaints about a customer are processed by the Trust & Safety Team in the USA. This team always informs the customer, to enable the customer to take action, such as take-down. The T&S team cannot look inside the contents of EC2 instances and S3 buckets. AWS commits there will always be a human review to assess a possible conflict with the Acceptable Use Policy.

4.3.4 Cybercrime and cyber attacks

One of the essential purposes as a data controller from AWS's perspective is the right to process personal data about the core network services for its own security purposes. AWS identifies a shared responsibility between customers and AWS. In this

[102] AWS Help & Customer Service, Law Enforcement Information Requests.
[103] AWS Help & Customer Service, Law Enforcement Information Requests.
[104] AWS, Acceptable Use Policy, Last updated 1 July 2021, URL: https://aws.amazon.com/aup/.

approach AWS is solely responsible for protecting the global infrastructure that runs all of the AWS Cloud against cybercrime and cyber-attacks, while the customer is responsible for maintaining control over the content that is hosted on the AWS infrastructure.[105]

Customers are warned they are (exclusively) responsible for the management of the guest operating system (OS), application software and the configuration of the AWS provided security group firewall. [**Confidential**]. AWS tells customers that they must perform all necessary security configuration and management tasks to keep the data secure. AWS provides an illustration of these different responsibilities 'in' and 'of' the cloud.

Figure 29: AWS illustration of Shared Security Responsibility Model[106]

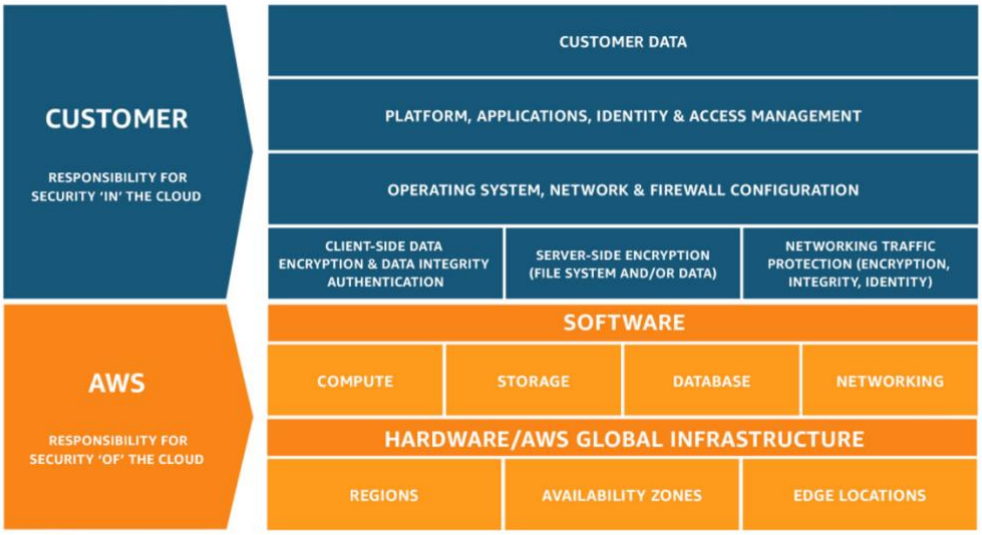

As described in Section 4.2 above, AWS acts as data processor when processing personal data across its customer database for security purposes. Every data processor is obliged under Article 28 of the GDPR to "*implement appropriate technical and organisational measures in such a manner that processing will meet the requirements of this Regulation and ensure the protection of the rights of the data subject*."

However, when AWS collects IP-addresses in its own central security logs from visitors to a VM, in a scenario where the government organisation has chosen to host a website on the VM, AWS processes these data as data controller.

After discussions with the researchers from Privacy Company about the specific purposes for which AWS processes data from its security log files, AWS provided the following limitation to the processing of website visitor data in its security logs:

> "*AWS has very limited data related to customers' end users. While there may be some discrete data related to end users that are accessing online resources hosted through AWS services (e.g., citizens accessing an application hosted by*

[105] AWS, Data protection in AWS Support, URL: https://docs.aws.amazon.com/awssupport/latest/user/data-protection.html.
[106] AWS, Shared Responsibility Model, URL: https://aws.amazon.com/compliance/shared-responsibility-model/

a government customer through AWS services), this data is not used by AWS for the purposes of profiling those end users, or advertising or marketing to them."

As described in Section 3.1.3 AWS did not provide access to these data to Privacy Company. In reply to this DPIA, AWS has pointed to the lack of non-compliance findings in its SOC-2 and C5_2020 audit reports. AWS has also suggested that government organisations that are very concerned about the logging of IP addresses of website visitors may deploy a proxy.

4.3.5 Improving and optimising

For a specific list of AI/ML (*Machine Learning*) services, AWS is permitted to use the necessary personal data for analysing, improving and optimising the performance and core functionality of these specific AI/ML services. AWS may only use anonymised data for Service Improvement, and only Content Data from a limitative list of services.

These AI/ML services are:

- Amazon Lex
- Amazon Transcribe
- Alexa for Business
- Amazon AppStream 2.0 User Pool
- Amazon CodeGuru Profiler
- Amazon Comprehend
- Amazon Connect Customer Profiles Identity Resolution
- Amazon Fraud Detector
- Amazon GuardDuty*
- Amazon Lex
- Amazon Polly
- Amazon Rekognition
- Amazon Textract
- Amazon Transcribe
- Amazon Translate
- Contact Lens for Amazon Connect107

Government organisations can opt-out of any reuse of their content and personal data for AI/ML from all of these services and all similar future services before implementation.[108]

4.4 Purposes AWS Commercial Contact and public Website Data (data controller)

AWS processes two categories of personal data as independent data controller: the Commercial Contact Data and the public Website Data. Pursuant to the AWS Privacy Notice and Cookie Notice, AWS processes these personal data for the following 20 purposes:

1. Provide and deliver AWS Offerings and [purpose split by Privacy Company];

[107] AWS, Sub-processors, AWS entities providing service improvement, URL: https://aws.amazon.com/compliance/sub-processors/.
[108] AWS, AI services opt-out policies, URL: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_ai-opt-out.html.

2. Process transactions related to AWS Offerings, including registrations, subscriptions, purchases, and payments;
3. Measure and technically improve AWS Offerings and [purpose split by Privacy Company];
4. Provide support to customers and [purpose split by Privacy Company];
5. Develop AWS Offerings;
6. Recommendations and Personalization "*We use your personal information to recommend AWS Offerings that might be of interest to you, identify your preferences, and personalize your experience with AWS Offerings.*";
7. Comply with legal obligations "*In certain cases, we have a legal obligation to collect, use, or retain your personal information. For example, we collect bank account information from AWS Marketplace sellers for identity verification + [Separately mentioned in rules for the EU customers] respond to lawful requests and orders*";
8. Communicate with you "*We use your personal information to communicate with you in relation to AWS Offerings via different channels (e.g., by phone, email, chat) and to respond to your requests*";
9. Marketing "*We use your personal information to market and promote AWS Offerings*";
10. Enrich data with information we receive from other sources
11. Fraud and Abuse Prevention and Credit Risk reduction "We release account and other personal information when we believe *release is appropriate to comply with the law, enforce or apply our terms and other agreements, or protect the rights, property, or security of AWS, our customers, or others. This includes exchanging information with other companies and organizations for fraud prevention and detection and credit risk reduction. [other sources: credit history information from credit bureaus]*";
12. Purposes for Which We Seek Your Consent "*We may also ask for your consent to use your personal information for a specific purpose that we communicate to you*";

[Cookie Notice]

13. Use cookies, pixels and other similar technologies to enable our systems to (i) recognize your browser or device (ii) learn more about your interests, (iii) provide you with essential features and services and (iv) for additional purposes, including:
14. Recognizing you when you sign in to use our offerings. This allows us to provide you with recommendations, display personalized content, and provide other customized features and services.
15. Keeping track of your specified preferences. This allows us to honor your likes and dislikes, such as your language and configuration preferences.
16. Conducting research and diagnostics to improve our offerings.
17. Preventing fraudulent activity
18. Improving security
19. Delivering content, including ads, relevant to your interests.

20. Reporting. This allows us to measure and analyze the performance of our offerings[109].

This is not a limitative list. AWS explains:

> *"AWS provides users with clear and comprehensive information in accordance with applicable law about the purposes of personal data processing when acting as a data controller in the AWS Privacy Notice. (…)*
> *Because the services and experiences offered by AWS are constantly evolving and customers use our services in different ways, we are not able to provide a limitative list."*[110]

Though AWS refuses to provide a limitative list of purposes, AWS has stated that it does not process these two categories of personal data for purposes AWS deems compatible with the purpose of 'providing the service'.

Privacy Company asked AWS to clarify some of these purposes, such as purpose no. 10 'Enrich data for marketing and sales generation with contact data from external sources'.

AWS replied that this is clear and comprehensive information in accordance with applicable law about the purposes of personal data processing. AWS confirmed all examples mentioned in the Privacy Notice also apply to a European public sector Customer, in this case: the procurement officers whose data are treated as Commercial Contact Data by AWS.

- "*marketing, sales generation, and recruitment information, including your name, email address, physical address, phone number, and other similar contact information;*
- *subscription, purchase, support, or other information about your interactions with products and services offered by us, our affiliates (such as AWS training courses), or third parties (such as products offered through the AWS Marketplace) in relation to AWS Offerings;*
- *search results and links, including paid listings (such as Sponsored Links) and;*
- *credit history information from credit bureaus*."[111]

AWS also replied that users have the right to request AWS to both restrict the processing of personal data where the processing is inappropriate and to object to the processing of personal data. AWS refers to the specific paragraph in the AWS Privacy Notice for data subjects in the EU that sums up the possible data subjects rights from articles 15 to 20 of the GDPR.[112]

5. Processor or controller

This section assesses the data protection roles of AWS and government organisations in the context of the tested Amazon Web Services.

[109] These purposes are mentioned as examples in the AWS Cookie Notice, 30 December 2022, URL: https://aws.amazon.com/legal/cookies/

[110] AWS response to DPIA questions, 17 July 2020, answer to Q3f.

[111] AWS Privacy Notice, Information from Other Sources.

[112] AWS Privacy Notice, the paragraph 'European Economic Area' in the section 'Additional Information for Certain Jurisdictions'.

5.1 Definitions

The GDPR contains definitions of the different roles of parties involved in the processing of data: (joint) controller, processor and sub-processor.
Article 4(7) of the GDPR defines the (joint) controller as:

> *"the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data; where the purposes and means of such processing are determined by Union or Member State law, the controller or the specific criteria for its nomination may be provided for by Union or Member State law."*

Article 26 of the GDPR stipulates that where two or more data controllers jointly determine the purposes and means of a processing, they are joint controllers. Joint controllers must determine their respective responsibilities for compliance with obligations under the GDPR in a transparent manner, especially towards data subjects, in an arrangement between them.

Article 4(8) of the GDPR defines a processor as:

> *"a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller."*

A subprocessor is another processor engaged by a processor that assists in the processing of personal data on behalf of a data controller.

Article 28 GDPR sets out various obligations of processors towards the controllers for whom they process data. Article 28(3) GDPR contains specific obligations for the processor. Such obligations include only processing personal data in accordance with documented instructions from the data controller and cooperating with audits by a data controller. Article 28(4) GDPR stipulates that a data processor may use subprocessors to perform specific tasks for the data controller, but only with the prior authorisation of the data controller.

When data protection roles are assessed, the formal contractual division of roles is not leading nor decisive. The actual role of a party must primarily be determined on the basis of factual circumstances.

5.2 Data processor and subprocessors

5.2.1 Assessment of AWS as a data processor

Contractually, AWS has become a data processor for all personal data in and about the use of Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon RDS. This applies to the processing of the following five categories of personal data: Content, Account, Diagnostic, Support and restricted access Website Data.

At the start of this DPIA project, AWS only acted as data processor for the Content Data, and qualified itself as independent data controller for the other categories of personal data.

As described in Section 4 above, the Dutch government and AWS have agreed on a limitative list of processing purposes for AWS as a data processor, and a list of agreed compatible purposes for which AWS may 'further' process some personal data as controller, when strictly necessary and proportionate.

Therefore the analysis of the purposes in the previous publicly available DPA and provisions from the Privacy Notice is no longer relevant for the assessment. These paragraphs have been deleted from this DPIA.

5.2.2 Subprocessors

In its publicly available list of sub-processors AWS describes three categories of sub-processors:

1. AWS affiliates that provide the AWS infrastructure
2. AWS service providers that provide specific services (for application/media services, for service improvement and for support)
3. third parties engaged as sub-processors[113]

The first category of sub-processors describes AWS-affiliates that manage infrastructure. For the Dutch government customers, the three most relevant AWS affiliates are A100 ROW GmbH in Frankfurt, Amazon Data Services France SAS and Amazon Data Services Ireland Limited. They are responsible for the AWS Region Europe.

The second category includes optional services, such as a video conferencing tool, not mandatory for customers of the tested EC2, S3 and RDS services. Even if a customer elects to use one of the services that build on user experiences, they can opt out of reuse of their data for any service improvement. See Figure 30 below. The second category also includes affiliates that provide support. See Figure 31 below.

Figure 30: Overview of AWS affiliates providing service improvement

b) AWS entities providing service improvement

Unless customer opts out of AWS using Customer Data for service improvement, where permitted in accordance with the AWS Service Terms, the AWS entities listed below provide service improvement for the applicable AWS services.

AWS entity	AWS service(s)	Processing location	Processing activity
Amazon Development Centre (India) Private Limited	Amazon Lex Amazon Transcribe	India	Service improvement
Amazon Web Services, Inc.	Alexa for Business Amazon AppStream 2.0 User Pool Amazon CodeGuru Profiler Amazon Comprehend Amazon Connect Customer Profiles Identity Resolution Amazon Fraud Detector Amazon GuardDuty* Amazon Lex Amazon Polly Amazon Rekognition Amazon Textract Amazon Transcribe Amazon Translate Contact Lens for Amazon Connect QuickSight Q for Amazon QuickSight	USA	Service improvement

*To the extent that the Amazon GuardDuty Malware Protection feature is enabled.

The list of support providers includes affiliates in countries outside of the EU with an adequate data protection regime, such as Japan and Canada. This is a relevant circumstance in the description of data transfers in Section 7 of this DPIA.

[113] AWS Sub-processors, list last revised 12 January 2023, URL: https://aws.amazon.com/compliance/sub-processors/.

AWS explained to SLM Rijk how the support employees work. When customers create a support case, AWS doesn't gain access to the customer's account. If necessary, support agents use a screen-sharing tool to view a customer's screen remotely and identify and troubleshoot problems. This tool is view-only. AWS specifically wrote: "*AWS personnel do not have the ability to log in to customer instances.*"[114]

Support agents cannot export any data from the customer, and cannot act for customers during the screen-share session. Customers must give consent to share a screen with a support agent.[115]

In reply to requests from the Dutch government to limit the processing of Support Data to EU-based employees, AWS has suggested an organisational measure. EU customers can ask their AWS account manager to flag all of their support requests with an internal contextual alert, in an existing system for support employees. Such an alert is specific to a customer. AWS Support Engineering and Customer Service will see these alerts displayed when accessing a customer case. Such an alert could warn employees that the customer only wants problems solved by EU-based employees, or for example only employees in a country with an adequate data protection regime, such as Japan.

Figure 31: Overview of AWS affiliates that provide support

c) AWS entities providing customer-initiated support

The AWS entities listed below provide customer-initiated support. These entities do not process Customer Data unless the customer agrees to share Customer Data in the course of requesting support.

AWS entity	Processing location (if applicable)
Amazon.com Services LLC	USA
Amazon Data Services SA (Pty) Ltd	South Africa
Amazon Development Centre (India) Private Limited	India
Amazon Development Centre Ireland Limited	Ireland
Amazon Development Centre (South Africa) (Proprietary) Limited	South Africa
Amazon Internet Services Private Limited	India
Amazon Support Services Costa Rica, SRL	Costa Rica
Amazon Web Services Australia Pty Ltd	Australia
Amazon Web Services Canada, Inc.	Canada
Amazon Web Services EMEA SARL	France and Ireland
Amazon Web Services Hong Kong Limited	Hong Kong
Amazon Web Services Japan G.K.	Japan
Amazon Web Services Korea LLC	Korea
Amazon Web Services Taiwan Ltd	Taiwan
Amazon Web Services, Inc.	USA
AWS India ProServe LLP	India
Elemental Technologies LLC	USA
Souq.com for E-Commerce LLC	Egypt
Twitch Interactive, Inc. (for Amazon Interactive Video Service only)	USA

The third category of third-party service providers, includes messaging services from application to person, and to provide geolocation services such as maps or points of

[114] Idem, answer to Q7a.

[115] AWS, Security for your AWS Support cases, URL: https://docs.aws.amazon.com/awssupport/latest/user/security-for-support-cases.html.

interest. AWS explains that data are processed in the customer's selected AWS Region(s).

In reply to DPIA questions AWS confirmed it has the necessary contractual terms with all affiliates and third party entities listed on the subprocessor page.[116]

AWS contractually commits that it will restrict the sub-processor's processing of personal data to only what is necessary for the sub-processor to perform the contracted obligations, and will enter into a written agreement with the sub-processor with the relevant contractual obligations that AWS has under the data processing agreement with its customers.

In spite of the obligation in Clause 9 sub c of Module Two (Controller to Processor) of the (new) Standard Contractual Clauses, AWS did not make copies available of the relevant privacy paragraphs in its contracts with subprocessors when so requested.[117]

> *"The data importer shall provide, at the data exporter's request, a copy of such a sub-processor agreement and any subsequent amendments to the data exporter. To the extent necessary to protect business secrets or other confidential information, including personal data, the data importer may redact the text of the agreement prior to sharing a copy."*[118]

AWS explains that it will inform its customers at least 30 days in advance prior to engaging a new subprocessor.

> *"At least thirty (30) days before AWS engages a new Sub-processor or replaces a Sub-processor (the "Objection Period"), AWS will update the applicable website (and provide Customer with a mechanism to obtain notice of that update) or otherwise notify Customer. If Customer does not object within the Objection Period, the new Sub-processor will be deemed authorised."*[119]

AWS has committed to offer several remedies in case a Dutch government customer objects to a new sub-processor. [**Confidential**]

Upon request, AWS explained that its agreements include comprehensive obligations on vendors to comply with Amazon's requirements to maintain the security of the data. AWS stated that the mechanisms to restrict unauthorised internal and external access to data and customer data are tested in the following available audits: "*ISO 27001, ISO 27017, and ISO 27018. Also the SOC 2 Type 1 report on AWS System relevant to privacy and SOC 2 Type 2 Report on AWS System relevant to security, availability, and confidentiality, and the AWS C5 Report."*[120]

AWS refers to the AWS SOC 2 report, which states:

---

[116] AWS response to DPIA questions, 17 July 2020, answer to Q5b.

[117] Clause 9 (c) of Module Two of the SCC specifies: "The data importer shall provide, at the data exporter's request, a copy of such a sub-processor agreement and any subsequent amendments to the data exporter. To the extent necessary to protect business secrets or other confidential information, including personal data, the data importer may redact the text of the agreement prior to sharing a copy."

[118] SCC, Clause 9, p. 23.

[119] AWS GDPR DPA, Article 6.1

[120] Idem.

"The delivery of services to customers does involve contracting with third parties to provide business functions, some of which include safeguarding and providing technical infrastructure for hosting a customer's content on AWS owned and managed compute equipment. AWS has business agreements with colocation services for providing compute infrastructure as well as business agreements with security services who provide physical security. Through contractual business agreements, AWS performs periodic reviews of colocation service providers to validate adherence with AWS security and operational standards (Control AWSCA-5.12). Through our contractual business agreements, third parties have a duty of confidentiality (Control AWSCA-11.1). AWS monitors the performance of third parties through periodic reviews, which evaluate performance against contractual obligations. AWS has a program in place to evaluate vendor performance against confidentiality and privacy commitments (Control AWSCA-11.2)."[121]

With regard to security measures in place to prevent unauthorized access by subprocessors to Content Data, AWS refers to the AWS SOC 2 report, which states:

"AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. User accounts are created to have minimal access. Access above these least privileges requires appropriate and separate authorization."

5.3 Assessment of AWS as data controller

As described in Section 4.3, the Dutch government has authorised AWS to 'further' process some personal data for a limitative list of agreed compatible purposes, when strictly necessary and proportionate.

Additionally, AWS acts as an independent data controller for the processing of personal data relating to the visits to its public website, and for the Commercial Contact Data. The purposes for these two categories of data are outlined in AWS's Privacy Notice.[122] Section 4.4 describes the 20 purposes that can be discerned in this notice.

6. Interests in the Data Processing

This paragraph outlines the different interests of AWS and of the Dutch government organisations in general. The interests of the Dutch governmental organisations may align with the interests of its employees, but this is not always the case. This section does not include an analysis of the fundamental data protection rights and interests of the employee-admins as data subjects. How their rights relate to the interests of AWS and the Dutch government organisations is analysed in part B of this DPIA.

6.1 Interests of AWS

As described in Section 1.1 of this report AWS is currently global market leader as provider of public cloud VMs, with a 32% global market share in the first quarter of

[121] AWS response to DPIA questions, 17 July 2020, answer to Q5d.

[122] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/.

2023, while Microsoft Azure climbed from a 21% market share in 2021 to a 23 percent market share in the first quarter of 2023, and Google Cloud from 10% to 11%.[123]

In its general competition with other big global players with similar offerings such as Microsoft, Google, IBM, SAP, Oracle and Alibaba, AWS distinguishes itself with its geographically wide spread. AWS has 76 availability zones in which its servers are located. These serviced regions are divided in order to allow users to set geographical limits on their services (if they so choose), but also to provide security by diversifying the physical locations in which data is held. Overall, AWS spans 245 countries and territories.[124] AWS is also renowned for its very wide portfolio of services. Computerworld writes:

> *"The key strength for the market leader continues to be the breadth and depth of its services, with more than 175 across compute, storage, database, analytics, networking, mobile, developer tools, management tools, IoT, security and enterprise applications, at last count."*[125]

The global public cloud providers do not seem to compete heavily on price. They all charge flexibly, by the second.[126]

To stay ahead of its competitors AWS has a financial (monetisation) interest in offering a high-quality service in terms of a wide choice of locations, services, scalability, flexibility and reliability.

Like its competitors, AWS has strong business ethical interests with regard to its compliance with international privacy and security standards and laws. In a world where many government organisations are still hesitant to entrust personal data to a cloud service provider, AWS has recently stepped up on its transparency on its disclose of personal data in response to government requests.

AWS provides a lot of documentation to show its compliance with the GDPR. AWS writes that its IT infrastructure is compliant with the following security standards:

1. SOC 1/ISAE 3402, SOC 2, SOC 3
2. FISMA, DIACAP, and FedRAMP
3. PCI DSS Level 1
4. ISO 9001, ISO 27001, ISO 27017, ISO 27018[127]

---

[123] Statista, Big Three Dominate the Global Cloud Market, quoting estimates from Synergy Research Group, 28 April 2023, URL: https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/.

[124] Investopedia, What Is Amazon Web Services and Why Is It so Successful?, last updated 6 November 2022, URL: https://www.investopedia.com/articles/investing/011316/what-amazon-web-services-and-why-it-so-successful.asp

[125] Computerworld, AWS vs Azure vs Google Cloud: What's the best cloud platform for Enterprise?, 23 January 2020, URL: https://www.computerworld.com/article/3429365/aws-vs-azure-vs-google-whats-the-best-cloud-platform-for-enterprise.html

[126] Idem. *"In general terms, prices are roughly comparable, especially since AWS shifted from by-the-hour to by-the-second pricing for its EC2 and EBS services in 2017, bringing it in line with Azure and Google."*

[127] AWS Security and Compliance, URL: https://docs.aws.amazon.com/whitepapers/latest/aws-overview/security-and-compliance.html.

6.2 Interests of government organisations

Dutch government organisations potentially have economic and financial interests in switching from hosting in local government datacentres to hosting of data on cloud servers. Some of the possible benefits of a switch to cloud VMs are:

1. Increased flexibility, because creating a VM and using a preconfigured database is faster and easier than installing an OS or database on a physical server. Admins can clone a VM with the OS already installed. Developers and software testers can create new environments on demand to handle new tasks as they arise.

2. Cost saving, because organisations may need less maintenance staff and are generally less effective in maximising the amount of VMs run on a single physical computer (bare metal server).

3. Increased reliability and availability, because it is easy to scale up computing capacity when necessary in the S3 bucket and the RDS, and to deploy multiple copies of the same virtual machine to better serve increases in load. Typically, it is much easier with virtualization software to reallocate hardware resources dynamically between one virtual machine and another.

4. Increased security and recovery, because the cloud VM provider has such a wide range of customers that it is able to faster detect new risks than engineers of local data centres. A globally operating cloud provider typically has more security engineers available, and 24/7, to mitigate risks. Additionally, the easy creation of VMs and databases also makes it possible to completely delete a compromised VM and then recreate it from a snapshot of the VM, hastening recovery from malware infections.

5. Easier GDPR compliance. As described in section 3.1.1 and 3.1.2 (Diagnostic Data), most cloud providers offer extensive logging possibilities that enable government organisations to regularly inspect the files for unauthorised access to personal data, and detect possible data breaches. Additionally, with VMs, because they are isolated from each other, it is easier to create independent user environments, for example a separate environment for software testing.

Government organisations have a strong interest in transparency, to be able to compare different services, forecast costs, to easily create an overview of different technical capacities and to assess the different privacy and security risks, including the risks of unlawful further processing through access by authorities in foreign countries, malicious state actors and hackers.

Additionally, government organisations need to prevent the possibility of 'legacy' and 'vendor lock-in'. Government organisations need to have the organisational and technical ability to easily transfer their data to other suppliers over time.

Finally, it is in the government organisations' interest that the platform is easy to use, so that employees who have little knowledge of IT can work with it. From a cost saving interest, it is important that the platform offers many good standard options, so that the government organisation does not need to develop much of its own functionality.

6.3 Joint interests

The interests of AWS and the Dutch government align when it comes to the protection of the integrity, availability, and reliability of Content Data in the cloud VM servers,

S3 buckets and RDS, as well as high scalability and flexibility of the deployment of the cloud services

The interests of AWS and the Dutch government also align with regard to the new Cloud Policy of the Dutch government. Based on this policy[128], and implementation guidance[129] Dutch government organisations are allowed to use public cloud services, based on a risk assessment, compliance with BIO and in accordance with C2000 criteria for risks regarding espionage, influencing or sabotage by state actors or other third parties.[130]

According to the Cyber Security Assessment Netherlands 2022 there are real risks related to unlawful access by state actors: "*The use of zero-day exploits by state actors against Dutch targets illustrates the structural and advanced state digital threat against Dutch economic and political security interests. Attackers are also increasingly focusing on exploiting the cloud. Cloud services have become crucial elements of many business processes over the past few years. Malicious actors see this dependence as a new opportunity to disrupt digital processes. More use of the cloud also means more potential victims. Outages or disruption of cloud services may have large-scale consequences for Dutch organisations and sectors.*"[131]

As part of the shared security interest, the creation of network and webserver access log files to detect security risks also concurs with the interests of the Dutch government organisations, provided that AWS performs the monitoring based on aggregated data, and does not store the personal data in these logs longer than strictly necessary for these security purposes.

The interests of AWS and Dutch government organisations do not necessarily align when it comes to the role of AWS as data controller for the Commercial Contact Data and the public Website Data, or the extra costs that may apply to encryption or the use of Nitro Systems.

7. Transfer of personal data outside of the EU

Even though customers can choose to host their Content Data within the EU availability zone, such a geolocation choice in the Admin Console is not available for the other categories of personal data (Account, Diagnostic, Support and restricted Website Data). AWS can transfer these personal data from its EU customers to the USA.

7.1 Available zones and regions for Content Data for Dutch government customers

AWS works with Regions, a physical location in a country where data centers are clustered. AWS has Regions in the EU. Each AWS Region consists of a minimum of

[128] Brief van de Staatssecretaris van Binnenlandse Zaken en Koninkrijksrelaties, Rijksbreed cloudbeleid 2022, 29 augustus 2022, URL: https://open.overheid.nl/documenten/ronl-a79331dc7c088f2cb6259f591c3b4f2fbcc9b5f1/pdf.

[129] Rijksoverheid, Implementatiekader risicoafweging cloudgebruik, versie 1.1, 5 januari 2023, URL: https://open.overheid.nl/documenten/ronl-734f947ec6465e4f75a56bed82fe64a1135f71a8/pdf.

[130] Kamerstukken II 2018/19, 25 124, nr. 96.

[131] NCTV, Cyber Security Assessment Netherlands, CSAN 2022, URL: https://english.nctv.nl/binaries/nctv-en/documenten/publications/2022/07/04/cyber-security-assessment-netherlands-2022/Cyber+Security+Assessment+Netherlands+2022.pdf.

three, isolated, and physically separate AZs within a geographic area. AWS calls each group of logical data centers an Availability Zone.

Customers of AWS EC2 VM, MySQL RDS and S3 bucket services can decide for themselves within which availability zone (geographic area) they want to store their Content Data. Figure 32 provides a map of geographic regions, Figure 33 sums up the available zones and regions.

Figure 32: AWS available geographic regions[132]

Figure 33: AWS tables of regions, countries and direct connect locations[133]

AWS explains:

"The AWS Cloud spans 99 Availability Zones within 31 geographic regions around the world, with announced plans for 15 more Availability Zones and 5 more AWS Regions in Canada, Israel, Malaysia, New Zealand, and Thailand."[134]

[132] Screenshot made on 23 May 2023 of AWS map, URL: https://aws.amazon.com/about-aws/global-infrastructure/

[133] Screenshot of AWS table, URL: https://aws.amazon.com/about-aws/global-infrastructure/

[134] Idem.

Figure 34: Available locations in test set-up (2020)

Region	Status	Action
Middle East (Bahrain)	**Disabled**	Enable
Africa (Cape Town)	**Disabled**	Enable
Asia Pacific (Hong Kong)	**Disabled**	Enable
Europe (Milan)	**Disabled**	Enable
Europe (Stockholm)	Enabled by default	
Asia Pacific (Mumbai)	Enabled by default	
Europe (Paris)	Enabled by default	
US East (Ohio)	Enabled by default	
Europe (Ireland)	Enabled by default	
Europe (Frankfurt)	Enabled by default	
South America (São Paulo)	Enabled by default	
US East (N. Virginia)	Enabled by default	
Asia Pacific (Seoul)	Enabled by default	
Asia Pacific (Osaka-Local)	Enabled by default	
Europe (London)	Enabled by default	
Asia Pacific (Tokyo)	Enabled by default	
US West (Oregon)	Enabled by default	
US West (N. California)	Enabled by default	
Asia Pacific (Singapore)	Enabled by default	
Asia Pacific (Sydney)	Enabled by default	
Canada (Central)	Enabled by default	

In the Management Console, under 'Preferences' AWS shows the available regions and specific locations. AWS explains that some regions, such as US East (N. Virginia) are enabled by default, and cannot be disabled.

AWS does not operate a data centre in the Netherlands, though AWS has 5 Edge locations in Amsterdam for faster network connectivity.[135] Privacy Company chose to set-up the S3 bucket in an AWS data centre in Frankfurt (eu-central-1). See Figure 35 and Figure 36 below.

[135] Amazon CloudFront Key Features, URL: https://aws.amazon.com/cloudfront/features/?p=ugi&l=emea&whats-new-cloudfront.sort-by=item.additionalFields.postDateTime&whats-new-cloudfront.sort-order=desc.

Figure 35: Location of EC2 Instance

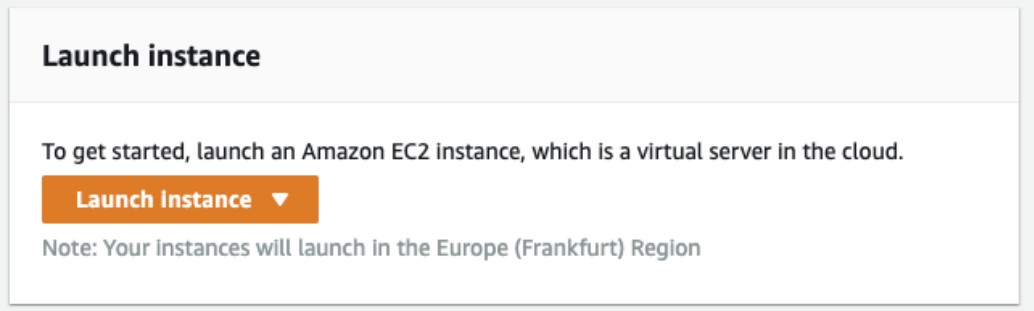

When the administrator selects a certain geolocation in the AWS console, the interface shows all instances in that specific geolocation and all instances that are created are created in that specific geolocation.

Figure 36: S3 bucket hosted in Frankfurt

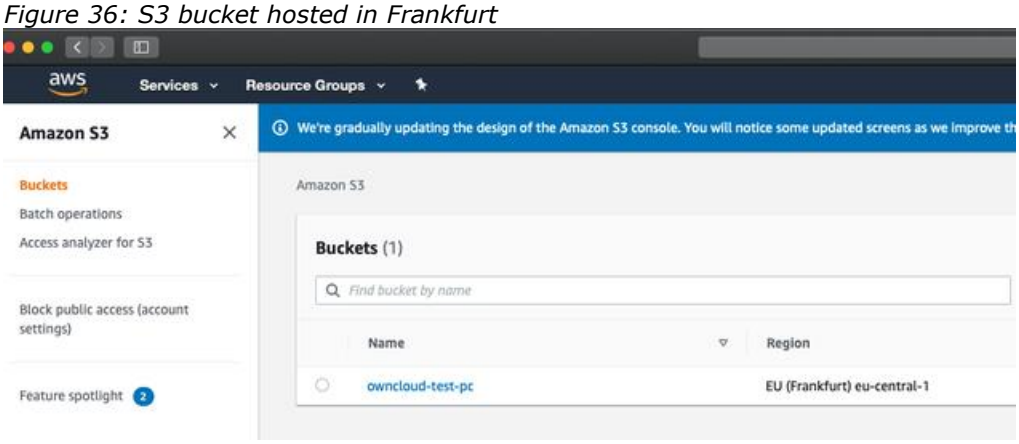

7.2 Transfers of other categories of personal data

In reply to requests from the Dutch government about the transfers of the Account, Diagnostic, Support and Restricted access Website Data to the USA, AWS produced two mitigating measures.

To restrict the processing of Support Data to EU-based employees, AWS has suggested an organisational measure. Dutch government customers can ask their AWS account manager to configure an alert for support employees that only EU based employees may respond to tickets, or employees in other countries with an adequate data protection regime, such as Canada or Japan.

Diagnostic Data are generated in the AWS Region where the service is used, and depending on the scope of the customer's interactions with AWS Offerings, may be stored in or accessed from multiple countries, including the United States. Website Data are generated in the USA, while Account Data are always transferred to the USA.

AWS advises Dutch government organisations to pseudonymise their Account Data (also collected as part of the Support, Diagnostic and restricted Website Data). AWS offers solutions to federate customer's employees, contractors, and partners (workforce) to AWS accounts and business applications, and offers federation support to customer's end-user-facing web and mobile applications. AWS supports commonly

used open identity standards, including Security Assertion Markup Language 2.0 (SAML 2.0), Open ID Connect (OIDC), and OAuth 2.0.[136]

As an additional mitigating measure AWS strongly recommends that customers never put confidential information or directly identifiable personal data, such as their email addresses, into tags or free-form text fields such as a Name field.[137]

7.3 GDPR rules for transfers of personal data

The GDPR contains specific rules for the transfer of personal data to countries outside the European Economic Area (EEA). In principle, personal data may only be systematically transferred to countries outside the EEA if the European Commission has taken a so-called adequacy decision or if the exporting organisation can provide appropriate safeguards.. Art. 46 of the GDPR offers multiple ways to provide such safeguards. For US based multinationals the two most commonly used instruments are Binding Corporate Rules or the EU Standard Contractual Clauses (SCC). Section 7.3.1 below addresses AWS's use of the (new) SCC.

The consequences and political follow-up of the jurisprudence in *Schrems-II* of the European Court of Justice are discussed in Section 7.3.2 through to 7.3.4 below, while Section 7.3.5 summarises the results of the separate Data Transfer Impact Assessment.

7.3.1 Standard Contractual Clauses

Personal data may be transferred from the EEA to third countries outside of the EEA using Standard Contractual Clauses (SCC, also known as EU model clauses). These clauses (hereinafter: SCC) contractually ensure a high level of protection.

On 16 September 2021, AWS published new SCC, based on the new SCC drafted by the European Commission in June 2021.[138] As part of the data processing agreement with the Dutch government the scope of the SCC is expanded to include the Account, Diagnostic Data, Support and restricted Website Data.

Dutch government customers (only) contract with Amazon Web Services EMEA SARL (AWS Europe), a Luxembourg-based AWS entity.

7.3.2 Consequences Schrems-2 ruling

On 16 July 2020, the European Court of Justice ruled that transfer of personal data based on the Privacy Shield is no longer valid with immediate effect.[139]

---

[136] AWS, Identity federation in AWS, URL: https://aws.amazon.com/identity/federation/ .

[137] AWS Service Terms, Paragraph 1.21, 30 May 2023, URL: https://aws.amazon.com/service-terms/

[138] AWS, New Standard Contractual Clauses now part of the AWS GDPR Data Processing Addendum for customers, 16 September 2021, URL: https://aws.amazon.com/blogs/security/new-standard-contractual-clauses-now-part-of-the-aws-gdpr-data-processing-addendum-for-customers/. Assuming Dutch government organisations are controllers, they can use AWS's Controller-to-Processor Clauses, as approved by the European Commission Implementing Decision (EU) 2021/914 of 4 June 2021. AWS publishes these SCC at the URL: https://d1.awsstatic.com/Controller_to_Processor_SCCs.pdf.

[139] European Court of Justice, C-311/18, Data Protection Commissioner versus Facebook Ireland Ltd and Maximillian Schrems (*Schrems-II*), 16 July 2020, ECLI:EU:C:2020:559, URL: http://curia.europa.eu/juris/document/document.jsf;jsessionid=CF8C3306269B9356ADF861B5

This judgement was the outcome of the lawsuit Max Schrems conducted against Facebook Ireland and the Irish Data Protection Commissioner. Earlier, in 2015, in another case instigated by Max Schrems, the European Court ruled the Safe Harbor agreement invalid, the predecessor of the Privacy Shield.

The Privacy Shield itself is since invalid as a legal basis for the transfer of personal data. The Court cited as the main reasons that the restrictions on privacy arising from the U.S. regulations are insufficiently defined and disproportionate and therefore constitute too great an invasion of privacy. Specifically, the Court describes the risks of mass surveillance (bulk data collection) by U.S. intelligence agencies under the surveillance programs PRISM and Upstream based on Section 702 FISA and based on E.O. 12333, and the lack of effective and enforceable rights for EU residents in the processing of those data by U.S. government agencies.
Although the European Court of Justice recognizes the validity of the decision of the European Commission with which it adopted the SCC, and data transfers on the basis of the SCC are therefore still permitted in principle, this validity cannot be assumed for systematic transfers of personal data to the United States.

The fact is that transfers via the SCC also require that the recipient country provides an adequate level of data protection as defined in EU law. Article 46(1) of the General Data Protection Regulation (GDPR) explains that this means that data subjects must have adequate safeguards, enforceable rights and effective legal remedies at their disposal. Whether this is the case, according to the Court, must be determined by the data controllers and cloud providers themselves.

The Court writes:

> *"The assessment required for that purpose in the context of such a transfer must, in particular, take into consideration both the contractual clauses agreed between the controller or processor established in the European Union and the recipient of the transfer established in the third country concerned and, as regards any access by the public authorities of that third country to the personal data transferred, the relevant aspects of the legal system of that third country. As regards the latter, the factors to be taken into consideration in the context of Article 46 of that regulation correspond to those set out, in a non-exhaustive manner, in Article 45(2) of that regulation."*[140]

7.3.3 US CLOUD Act and other applicable US law

In addition to these two specific surveillance powers, the USA legal regime enables law enforcement authorities and secret services to compel electronic communications services providers or remote computing service providers (such as cloud providers) that operate in the US to disclose personal data stored outside of the US. This includes disclosure of data from European customers stored in EU data centres.

The US CLOUD Act (*Clarifying Lawful Overseas Use of Data*) was specifically designed to obtain access to data stored in data centres in the EU. This act extends the jurisdiction of North American courts to all data under the control of companies operating in the USA, even if those data are stored in data centres outside the territory of the United States.

7785FDEE?text=&docid=228677&pageIndex=0&doclang=EN&mode=req&dir=&occ=first&part=1&cid=9812784 . See in particular par. 165 and 178-185.

[140] ECJ, Schrems-II, par. 104.

As explained by the EDPB and the European Data Protection Supervisor (EDPS) in their opinion on the CLOUD Act to the LIBE Committee of the European Parliament, transfers of personal data from the EU must comply with the Articles 6 (lawfulness of processing) and 49 (derogations for specific situations) of the GDPR. In case of an order based on the US CLOUD Act, the disclosure and transfer can only be valid if recognised by an international agreement between the EU and the USA.

The EDPB and EDPS write: "*Unless a US CLOUD Act warrant is recognised or made enforceable on the basis of an international agreement, and therefore can be recognised as a legal obligation, as per Article 6(1)(c) GDPR, the lawfulness of such processing cannot be ascertained, without prejudice to exceptional circumstances where processing is necessary in order to protect the vital interests of the data subject on the basis of Article 6(1)(d) read in conjunction with Article 49(1)(f).*" [141]

In their cover letter, the data protection authorities emphasise "*the urgent need for a new generation of MLATs to be implemented, allowing for a much faster and secure processing of requests in practice. In order to provide a much better level of data protection, such updated MLATs should contain relevant and strong data protection safeguards such as, for example, guarantees based on the principles of proportionality and data minimisation.*"[142]

Additionally, the data protection authorities refer to the ongoing negotiations since 2019 about an international agreement between the EU and the US on cross-border access to electronic evidence for judicial cooperation in criminal matters and negotiating directives. [143]

Only the UK has so far signed a specific agreement with the USA for the CLOUD Act. Negotiations between the EU and the US about updated MLATs did not produce any results yet.

7.3.4 Possible new adequacy decision companies USA

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.[144]

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

[141] Annex EDPB and EDPS joint response to US CLOUD Act, 10 July 2019, p. 8. URL: https://edpb.europa.eu/sites/default/files/files/file2/edpb_edps_joint_response_us_cloudact_annex.pdf .

[142] EDPB-EDPS Joint Response to the LIBE Committee on the impact of the US Cloud Act on the European legal framework for personal data protection, URL: https://edpb.europa.eu/sites/default/files/files/file1/edpb_edps_joint_response_us_cloudact_coverletter.pdf.

[143] Council Decision authorising the opening of negotiations, 6 June 2019, URL: https://data.consilium.europa.eu/doc/document/ST-10128-2019-INIT/en/pdf and https://data.consilium.europa.eu/doc/document/ST-10128-2019-ADD-1/en/pdf.

[144] 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.[145] Following this EOP, the European Commission has prepared a new draft adequacy decision.[146]

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 the EOP, but expresses concerns, asks for clarifications from the Commission and calls on the Commission to monitor the implementation in future joint reviews.[147]

The LIBE committee of the European Parliament has taken a critical stance.[148] The plenary EP has adopted a similarly critical resolution on 11 May 2023.[149] Currently the ministers of the Member States are invited to agree (the Council). The European Commission cannot adopt a new adequacy decision before July 2023. The Commission has to verify that the USA have implemented the agreed policy measures, and the US first has to qualify the EU (or the individual member states) as *'qualifying state*', before inhabitants can invoke the protection measures from the EOP.

If the EC succeeds in adopting a new adequacy decision, companies such as AWS 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 SCC and BCR.[150]

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.*"[151]

7.3.5 Results of Data Transfer Impact Assessment

Based on the finalised guidelines of the EDPB on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data,[152] an

---

[145] 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/ .

[146] 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.

[147] 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.

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

[149] Resolution European Parliament on the adequacy of the protection afforded by the EU- U.S. Data Privacy Framework, 11 May 2023, URL: https://www.europarl.europa.eu/doceo/document/TA-9-2023-0204_EN.html

[150] 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.

[151] 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

[152] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, Version 2.0, Adopted on 18 June

umbrella Data Transfer Impact Assessment (DTIA) was performed. According to the EDPB exporting organisations have to perform an analysis of the risks of transfers of data outside of the European Economic Area (EEA). The DTIA is based on the model provided by the Swiss lawyer David Rosenthal, and modified by Privacy Company.

To perform a specific DTIA, the EDPB recommends taking six steps. [153] These steps are not repeated here, but some highlights are mentioned that are relevant to assess the risks for organisations that want to use the tested three AWS services. The risk assessment must include:

- The relevant laws
- The purposes for which the data are processed
- The categories of transferred data and their sensitiveness
- Whether the data will be stored in the third country or whether there is remote access to data stored within the EU/EEA
- Role of the parties (public/private, processor/controller)
- All actors, including subprocessors
- The format of the data
- Possibility of onward transfers[154]

Relevant USA legislation

In a whitepaper on Data Residency AWS give its views on data residency (in the Netherlands or in the EU).

> *"In today's complex computing environment, public sector organizations continue to have legitimate concerns about the security of their data. As a result, some governments have determined that mandating data residency – the requirement that all customer content processed and stored in an IT system remain within a specific country's borders – provides an extra layer of security. Data residency reflects a combination of issues primarily associated with perceived (and in some cases real) security risks around third-party access to data, including foreign law enforcement agencies. Public sector customers want the assurance that their data is protected from unwanted access from not only nefarious attackers, but also other governments."*[155]

In the paper, AWS attempts to de-bunk perceived security risks expressed by governments when they demand in-country data residency. AWS describes mechanisms such as Mutual Legal Assistance Treaties, Letters Rogatory and the US CLOUD Act as "*frameworks that strengthen legal due process for law enforcement requests*." AWS concludes that restricting a CSP to one jurisdiction does not better

---

2021, URL: https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf.

[153] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, Version 2.0, Adopted on 18 June 2021, URL: https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf

[154] Idem, p. 15.

[155] AWS Whitepaper, Data Residency, August 2020, URL: https://d1.awsstatic.com/whitepapers/compliance/Data_Residency_Whitepaper.pdf.

insulate data from governmental access. Because compelled access only occurs in a very limited number of cases, AWS concludes the risk is low.

Though this whitepaper was updated after the Schrems-II ruling, AWS does not mention mass surveillance, orders from intelligence or security services, or FISA and/or the Executive Order 12 333. AWS mentions the existence of a (generic) EU US Mutual Lateral Assistance Treaty (MLAT), but omits to explain that negotiations about a specific MLAT legitimising law enforcement access to data and access to e-evidence are dragging on since September 2019.[156]

It is possible that AWS qualifies as communication service provider as defined in 50 USC § 1881(b)(4). This article defines:

> *The term "electronic communication service provider" means:*
> *A. a telecommunications carrier, as that term is defined in section 153 of title 47;*
> *B. a provider of electronic communication service, as that term is defined in section 2510 of title 18;*
> *C. a provider of a **remote computing service**, as that term is defined in section 2711 of title 18;*
> *D. any other communication service provider who has access to wire or electronic communications either as such communications are transmitted or as such communications are stored; or*
> *E. an officer, employee, or agent of an entity described in subparagraph (A), (B), (C), or (D).*

As *electronic communication service provider* AWS not only falls under FISA 702, but also under the US CLOUD Act, and other surveillance powers. Orders from law enforcement authorities and secret services may be accompanied by a gagging order.

Categories of transferred data
As explained in the introduction, it is up to the government organisations that purchase AWS services to determine the nature of the data they allow to be processed by AWS and to customize their choices accordingly. Organisations should not only consider the Content Data they actively upload, but also the four other categories of personal data that can be processed by AWS.

AWS explained in May 2023 that both its operational logs and security logs are processed separately by region, in the customer's region, and in separate partitions. This regional distinction is in the 'engineering standard'. If there is a security incident in Ireland, the log entry is created and retained in Ireland. There is a distinction between the general logging of operational diagnostic data (about usage infrastructure) and the separate security event log.

Sections 2.3 and 2.5 contain examples of the possible data that may be processed, and the possible types of data subjects that may be affected by the data processing. Additionally, organisations must take into account that the Support Data may include snippets of Content Data (when included by a customer), to illustrate possible bugs or errors.

---

[156] See for example: Theodore Christakis, Fabien Terpan, EU–US negotiations on law enforcement access to data: divergences, challenges and EU law procedures and options, 12 February 2021, in: International Data Privacy Law, Volume 11, Issue 2, April 2021, Pages 81–106, URL: https://academic.oup.com/idpl/article/11/2/81/6133744 .

Applicable transfer instrument
The data processing agreement between AWS and the Dutch government entities includes the new SCCs.

Likelihood that personal data are handed over or accessed to authorities
Providers such as AWS cannot in all circumstances provide effective measures against the interference created by the law of the U.S. with the fundamental rights of persons whose data are transferred to the U.S.A.

Organisations therefore need to take three general risks for unlawful further processing of personal data into account when enabling AWS to process personal data:

1. orders to AWS from US law enforcement authorities, security agencies and secret services;
2. rogue administrators at AWS and at subprocessors, and;
3. hostile state actors.

AWS takes a number of different technical and organisational measures to protect personal data against the risks of rogue administrators and against attacks from hostile state actors. AWS encrypts the data in transit and offers options to encrypt Diagnostic Data in CloudTrails and to encrypt Content Data in block storage and S3.

The encryption options are detailed in Section 8.1 of this DPIA. As explained in the introduction, in the section about the reply from AWS, AWS offers extra protection against tampering with encryption with Nitro System for EC2 instances. The Key Management System is "*designed so that no one, including AWS employees, can retrieve customer plaintext KMS keys from the service.*" Based on the review of the design the risk of forced decryption is now assessed to be near zero. Even though the probability of decryption by AWS is likely to be extremely low, the solution does not meet the letter of the three possible guarantees provided by the EDPB: "*the keys are retained (i) solely under the control of the data exporter, or (ii) by an entity trusted by the exporter in the EEA or (iii) under a jurisdiction offering an essentially equivalent level of protection to that guaranteed within the EEA [numbering added by Privacy Company].*"[157]

When compared with local, on-premises hosting, AWS as a cloud provider offers better guarantees for the timely detection of risks, and for the implementation and monitoring of up-to-date security measures.

However, given the legal circumstances in the USA, AWS cannot offer an absolute guarantee that it will never be compelled to give access to, or disclose, the personal data processed by Dutch government organisations. As explained above, even if Content Data are exclusively stored in the EU, the applicability of the US CLOUD Act potentially enables US law enforcement authorities to force cloud providers to hand-over of such personal data.

[157] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, Version 2.0, Adopted on 18 June 2021, URL: https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf.

In the final version of the EDPB Recommendations on measures to supplement transfer tools, the EU supervisory authorities and the EDPS allow for a risk-based approach in a DTIA.

> *"You may decide to proceed with the transfer without being required to implement supplementary measures, if you consider that you have no reason to believe that relevant and problematic legislation will be applied, in practice, to your transferred data and/or importer".*[158]

This means data controller can take into account the amount of received and/or honored requests and orders. The EDPB recommends to read the transparency reports from the importing organisation. However, "*Documented practical experience of the importer with relevant prior instances of requests*" alone cannot be relied on.[159] Annex 3 of the EDPB guidelines contains a non-exhaustive list of sources to consult, such as reports from various credible organisations and warrants from other entities.

As explained in Section 5.2.3 of this DPIA, AWS publishes transparency reports about the amount of requests it receives from law enforcement and the range of orders from security services for personal data relating to its Enterprise customers. In its four most recent reports, about the 2021 and 2022, AWS reports a range of requests for security services between 0 and 249. AWS notes that none of the requests resulted in the disclosure to the U.S. government of enterprise or government content data located outside the United States (which would concern requests under the CLOUD Act).

According to AWS, because compelled access only occurs in a very limited number of cases, or not all, the risk of unlawful further processing for data subjects is low.[160] However, the potential impact of such disclosure may be high, depending on the nature of the data.

Since February 2021 AWS provides additional guarantees in the Supplementary Addendum to the DPA, that it will challenge *any overbroad or inappropriate* requests or gagging orders. As quoted there, AWS writes that it has *repeatedly challenged government demands for customer information that we believed were overbroad, winning decisions that have helped to set the legal standards for protecting customer speech and privacy interests."*[161] However, different from providers such as Microsoft and Google, AWS does not make any information available about such challenges and court cases. The six steps AWS has committed to follow are listed in Section 4.3.2.

Another important assurance is AWS's guarantees in Clause 14 of the SCC that it has no reason to believe that it cannot fulfill its obligations under the clauses due to lawful access orders and requests. Clause 14 says:

> "*The Parties warrant that they have no reason to believe that the laws and practices in the third country of destination applicable to the processing of the personal data by the data importer, including any requirements to disclose personal data or measures authorising access by public authorities, prevent the data importer from fulfilling its obligations under these Clauses. This is based*

---

[158] EDPB, Recommendations 01/2020, Version 2.0, par. 43.3.

[159] Idem, para 47.

[160] AWS Whitepaper, Data Residency, p. 7. "*The reality is that such compelled access occurs in a very limited number of cases, and generally only where there is an extreme need for information (e.g., to prevent terror- related events).*"

[161] AWS Help & Customer Service, Law Enforcement Information Requests.

on the understanding that laws and practices that respect the essence of the fundamental rights and freedoms and do not exceed what is necessary and proportionate in a democratic society to safeguard one of the objectives listed in Article 23(1) of Regulation (EU) 2016/679, are not in contradiction with these Clauses."

Supplementary measures that can be applied
The EDPB recommends strong encryption as the best supplementary measure to mitigate the risks of unlawful processing when the data are transferred outside of the EEA. However, the EDPB also acknowledges that end-to-end encryption with customer-held keys is not always possible for cloud providers. In certain cases (e.g., where access to data in the clear is required) the cloud provider may need access to unencrypted data during the processing to provide certain services.

As AWS has pointed out in reply to Part A of this DPIA, government organisations can minimise the risks of bulk searches for identifying data by U.S. government agencies of the Diagnostic and admin Account Data by using AWS federated Identity and Access Management (IAM). This will ensure that these AWS logs generated in the USA only contain pseudonymised personal data of government employees.

These two possible technical measures are described in more detail in Section 8, about encryption and pseudonymisation.

In sum, government organisations must assess on a case by case basis if they can apply technical measures such as encryption and pseudonymisation to the different categories of personal data to lower the likelihood of unlawful further processing.

8. Data minimisation: encryption and pseudonymisation

AWS offers its customers different encryption options, both for the Content Data and for the customer-managed log files. Enterprise customers may also choose to pseudonymise the Admin Account Data.

8.1 Encryption

The question of adequate encryption of data has become very topical with the finalised guidelines from the European Data Protection Board (EDPB) on technical measures that data controllers must take when transferring personal data to a country with a non-adequate level of data protection.[162]

AWS provides HTTPS endpoints using the TLS protocol for communication.[163] This provides encryption of the data-in-transit when customers use AWS APIs. AWS recommends customers should use TLS 1.2 or later.[164]

[162] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, 18 June 2021, and: EDPB Recommendations 02/2020 on the European Essential Guarantees for surveillance measures, Adopted on 10 November 2020, URL: https://edpb.europa.eu/our-work-tools/our-documents/recommendations/edpb-recommendations-022020-european-essential_en .
[163] AWS, AWS service endpoints, undated, URL: https://docs.aws.amazon.com/general/latest/gr/rande.html.
[164] AWS Security blog, TLS 1.2 to become the minimum TLS protocol level for all AWS API endpoints, 28 June 2022, URL: https://aws.amazon.com/blogs/security/tls-1-2-required-for-aws-endpoints//

Customers can use the AWS Certificate Manager (ACM) service to generate, manage, and deploy the private and public certificates they use to establish encrypted transport between systems for their workloads. The ACM supports the importing of third party certificates. The AWS Key Management Service (KMS) and the ACM both support the hybrid post-quantum TLS ciphers.[165]

AWS provides customers with different options to protect the security of personal data in Content Data. AWS writes:

> "*AWS offers customers options to implement strong encryption for their customer content in transit or at rest, and also provides customers with the option to manage their own encryption keys or use third party encryption mechanisms of their choice. The AWS Well-Architected Framework 'security' pillar offers best practices in this regard.*"[166]

AWS provides a long list of services which customers can choose to use encryption. The list includes Amazon EC2, Amazon S3 and Amazon RDS.[167]

To encrypt the images on a VM, AWS offers EBS encryption (Elastic Block Store).[168] Customers can use EBS encryption to encrypt both boot and data volumes of Virtual Machines. This means that the data-at-rest inside the volume are encrypted, as well as the data-in-transit between storage and the VM instance and all snapshots and volumes created from the disk images.

AWS VM instances always need to have access to the key to decrypt, because AWS needs to be able to read the contents and boot the operating system inside the images. It is technically not possible to completely 'hide' the decryption key from AWS. However, AWS offers CloudHSM to allow the customer to manage the keys and (further) limit access by AWS to the Content Data. AWS explains:

> *"This is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys on the AWS Cloud. With CloudHSM, you can manage your own encryption keys using FIPS 140-2 Level 3 validated HSMs."*[169]

Elastic Load Balancing is integrated with ACM and is used to support HTTPS protocols. If Content Data are distributed through Amazon CloudFront, it supports encrypted endpoints.

AWS explains in its Whitepaper on GDPR compliance that customers may also encrypt customer managed logs, such as the S3 access logs and CloudTrail that are stored in S3.

> "*Logs can be encrypted at rest by configuring default object encryption in the destination bucket. The objects are encrypted using server-side encryption*

---

[165] AWS Cloud Security, Post-Quantum Cryptography, undated, URL: https://aws.amazon.com/security/post-quantum-cryptography/

[166] AWS response to DPIA questions, 17 July 2020, answer to Q2f. AWS refers to AWS Well-Architected Framework: https://aws.amazon.com/architecture/well-architected/

[167] AWS Service Capabilities for Privacy Considerations, URL: https://aws.amazon.com/compliance/data-privacy/service-capabilities/

[168] AWS EBS encryption, URL: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html.

[169] AWS CloudHSM, URL: https://aws.amazon.com/cloudhsm/

> *with either Amazon S3-managed keys (SSE-S3) or customer master keys (CMKs) stored in AWS Key Management Service (AWS KMS).”*[170]

Customers can use S3 buckets to store data that are encrypted with their own key. This can be useful for archived data, but not for logs that are stored (dynamically) by AWS in S3. This solution also doesn’t work if S3 is used to store static files for publicly accessible websites, which need to be readable by design.

Technically, customers could also decide to encrypt parts of the database, for example with MySQL enterprise encryption.[171] In the test set-up, this approach was not used because it requires modifications to the application to use these features. Anyway, such encryption would be of limited use to protect against the risk of unlawful access by AWS, because the keys to decrypt the data are still present in the VM. Otherwise, the application could not read or write the data, and with a web based application, such read and write access is continuously necessary.

This does not mean that disk encryption is not a useful security measure, or that access to the encrypted content or key material is readily available to AWS staff. However, according to the EDPB Recommendations on additional measures for transfers, disk encryption can only be considered a sufficient measure if the importer does not have access to the keys.[172] The EDPB writes that it considers the encryption performed on storage of personal data by a hoster in a third country an effective supplementary measure if six criteria are met. Here especially the criterion is relevant that the importer does not have access to the key.

> *“the keys are retained solely under the control of the data exporter, or by an entity trusted by the exporter in the EEA or under a jurisdiction offering an essentially equivalent level of protection to that guaranteed within the EEA”*[173]

The EDPB similarly explains in Use case 6 of the Guidelines that transport encryption and data-at-rest encryption are not sufficient to ensure an essentially equivalent level of protection if the data importer (in this case AWS) is in possession of the cryptographic keys.[174]

However limited the exposure, in 2021 Privacy Company concluded that the key material was not solely under customer or trusted third party control, the absolute norm imposed by the EDPB.

Since, AWS published a third-party verification of an even stronger supplemental measure, its Nitro System by the NCC Group.[175] As AWS explains, Nitro system is a collection of security measures that allows hardware based VM encryption and key

---

[170] AWS whitepaper, Navigating GDPR Compliance on AWS, p. 15.
[171] MySQL Enterprise Encryption, URL: https://www.mysql.com/products/enterprise/encryption.html
[172] EDPB, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, 18 June 2021, URL: https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf.
[173] Idem, par. 84, p. 30.
[174] Idem, par. 95, p. 35.
[175] NCC Group, Public Report – AWS Nitro System API & Security Claims, URL: https://research.nccgroup.com/2023/05/03/public-report-aws-nitro-system-api-security-claims/

management for customer virtual machines.[176] AWS adds that the Nitro System was designed to provide confidentiality from the provider administrators and software. AWS writes that Nitro System's security model is locked down and prohibits administrative access, eliminating the possibility of human error and tampering. AWS also points to other security capacities of the Nitro System on this webpage.

On 4 May 2023, AWS published a blog about the role of Nitro systems in its Digital Sovereignty Pledge.[177] AWS writes:

> *"With the AWS Nitro System, which is the foundation of AWS computing service Amazon EC2, we designed and delivered first-of-a-kind innovation by eliminating any mechanism AWS personnel have to access customer data on Nitro. Our removal of an operator access mechanism was unique in 2017 when we first launched the Nitro System."*[178]

AWS summarises the findings of a 2023 independent audit on Nitro System:

> *"This report confirms that the AWS Nitro System, by design, has no mechanism for anyone at AWS to access your data on Nitro hosts. The report evaluates the architecture of the Nitro System and our claims about operator access. It concludes that "As a matter of design, NCC Group found no gaps in the Nitro System that would compromise these security claims." It also goes on to state, "NCC Group finds…there is no indication that a cloud service provider employee can obtain such access…to any host."*[179]

Privacy Company has not tested the Nitro System itself, as this measure was proposed to the Dutch government by AWS in 2023 as an additional measure for government organisations, after the testing was completed.

AWS commits it has not built in any *backdoors* or similar programming in the services that could be used by AWS or by third parties to obtain unauthorised access to the system.

8.2 Pseudonymisation

AWS enables organisations to limit the exposure of directly identifiable information from admins by using AWS federated Identity and Access Management. What that means in practice is that the administrator user-accounts can be stored in a user database outside of AWS and AWS can be instructed to authenticate users against that external database through the OpenID Connect or SAML 2.0 protocols.[180]

To implement AWS IAM, the customer must first provide the unique identity of each authorised user to the (local or external) identity provider. The customer has more freedom to determine how the users are identified to AWS, for example only with pseudonymous account identifiers that are not easily identifiable by AWS. The

---

[176] AWS Nitro System, URL: https://aws.amazon.com/ec2/nitro/.

[177] AWS, Delivering on the AWS Digital Sovereignty Pledge: Control without compromise, 4 May 2023, URL: https://aws.amazon.com/blogs/security/delivering-on-the-aws-digital-sovereignty-pledge-control-without-compromise/.

[178] Idem.

[179] Idem.

[180] AWS, Identity providers and federation, undated, URL: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers.html.

customer is free to host its own identity provider or choose available commercial offerings.

As quoted in Section 5.1 of this report, pseudonymised data are still personal data. The 'exporter' (Dutch government organisation) has a realistic possibility to re-identify individuals based on a look-up table of the original identifying data stored with the identity provider and the aliased data provided to AWS.

While the use of AWS IAM does not prevent the processing of personal data of admins in AWS logs, the pseudonymisation can lower the data protection risks for the system admins. As described in Use case 2 of the EDPB Recommendations on measures that supplement transfer tools, pseudonymisation may be an effective supplementary measures to lower transfer risks, provided that the exporting organisation can also ensure that the pseudonymised data cannot be reidentified:

> *"by means of a thorough analysis of the data in question - taking into account any information that the public authorities of the recipient country may be expected to possess and use - that the pseudonymised personal data cannot be attributed to an identified or identifiable natural person even if cross-referenced with such information."*[181]

9. Additional legal obligations: e-Privacy Directive

This section only describes the additional obligations arising from the current ePrivacy Directive and (possible) future e-Privacy Regulation. In view of the limited scope of this DPIA, other legal obligations or frameworks (for example in the area of information security, such as BIO) are not included in this report.

The act of reading or placing information (through cookies or similar technology), or enabling third parties to read information from the devices of end users triggers the applicability of Article 5(3) of the ePrivacy Directive, regardless of who places or reads the information, and regardless of whether the content is personal data or not. Consent is required prior to the retrieval or storage of information on the devices or browsers of end users, unless one of the exceptions applies, such as the necessity to deliver a requested service, or necessity for the technical transmission of information.

Based on article 3(1) of the GDPR, because the data processing takes place in the context of the activities of data controllers (Dutch government organisations), the GDPR applies to all phases of the processing of these data.

Applicability of the GDPR rules does not exclude applicability of the ePrivacy rules or vice versa. The European Data Protection Board writes:

> *"Case law of the Court of Justice of the European Union (CJEU) confirms that it is possible for processing to fall within the material scope of both the ePrivacy Directive and the GDPR at the same time. In Wirtschaftsakademie, the CJEU applied Directive 95/46/EC notwithstanding the fact that the underlying processing also involved processing operations falling into the material scope of the ePrivacy Directive. In the pending Fashion ID case, the Advocate General*

[181] EDPB Recommendations, Use case 2, p. 31.

expressed the view that both set of rules may be applicable in a case involving social plug-ins and cookies.[182]

Article 5(3) of the ePrivacy Directive was transposed in article 11.7a of the Dutch Telecommunications Act. The consequences of the cookie provision are far-reaching, since it requires clear and complete information to be provided *prior* to the data processing, and it requires consent from the user, unless one of the legal exceptions applies. The consent is identical to the consent defined in the GDPR.

The most frequently used exception in the Netherlands is the processing of such information for analytical purposes, literally:

> *"to obtain information on the quality or effectiveness of a delivered information society service provided that it has no or little impact on the privacy of the subscriber or user concerned."*

As described in Section 3.2, AWS does not seem to collect (read) or set any information through cookies, pixels or comparable technologies on its restricted access websites (Admin Console and Support Portal) that would require informed consent from visitors (government admins). However, as described in Section 3.2.1, AWS does collect analytical data through browser telemetry, without providing any option to opt-out. Regardless if the admin chooses the second option in the cookie banner *Continue without accepting* or the third option *Customize cookies,* AWS reads analytical data with the full account name in the browser telemetry. AWS is still studying this traffic, and will discuss this with SLM Rijk in the ongoing dialogue.

The fact that AWS collects the account name, that may contain directly identifying personal data if the admins do not use pseudonymised accounts, from the browser on the end-user device, seems excessive for analytical purposes. It is plausible that this type of data collection does not comply with the specific consent-exception for the collection of analytical data in the Dutch Telecommunications Act mentioned above.

If AWS wishes to successfully invoke this Dutch analytical exception, it should offer users an option to refuse this data collection, refrain from collecting the user name, and based on guidance from the EDPS, Austrian, French, Italian and Danish DPA[183], stop transferring these analytical data to the USA.

---

[182] EDPB, Opinion 5/2019 on the interplay between the ePrivacy Directive and the GDPR, in particular regarding the competence, tasks and powers of data protection authorities, adopted on 12 March 2019, Paragraph 30. URL: https://edpb.europa.eu/sites/edpb/files/files/file1/201905_edpb_opinion_eprivacydir_gdpr_interplay_en_0.pdf In footnotes the EDPB refers to: CJEU, C-210/16, 5 June 2018, C-210/16, ECLI:EU:C:2018:388. See in particular paragraphs 33-34 and the Opinion of Advocate General Bobek in Fashion ID, C-40/17, 19 December 2018, ECLI:EU:C:2018:1039. See in particular paragraphs 111-115.

[183] See for example IAPP, Garante orders stop on Google Analytics transfers, 23 June 2022, URL: https://iapp.org/news/a/italian-dpa-orders-stop-on-google-analytics-transfers/ (last viewed 10 August 2022).

Figure 37: Timeline decision making ePrivacy Regulation

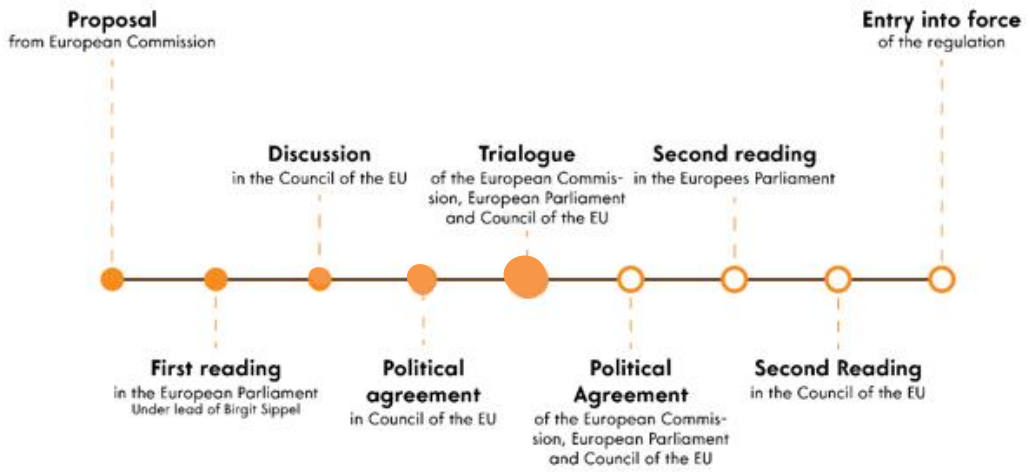

On 10 January 2017, the European Commission published a proposal for a new ePrivacy Regulation.[184] This was followed by an intense political debate the last five and a half years. The European Parliament responded quickly and positively[185], but it has taken the representatives of the EU Member States three years to draft a compromise about the proposed ePrivacy Regulation. The Council sent its agreed position[186] to COREPER to start the trialogue on 10 February 2021.[187] The trilogues began on 20 May 2021. The last publicly available update from the Council dates from 28 March 2022, in which the proposed compromises are all blacked out.[188] Figure 37 above shows the required legislative steps for adoption of the ePrivacy Regulation.

The points of view of the European Parliament and the European Council are widely diverging. Therefore, it is not likely that the ePrivacy Regulation will enter into force anytime soon, and AWS will have to comply with the current ePrivacy rules in the next few years.

That means that if AWS wants to continue to use 'performance' cookies when a user selects *'continue without accepting*', and continues to read telemetry data from the browser on the end user device without prior consent, it needs to become more transparent, and offer an opt-out to admins, to ensure the impact on data subjects (government admins visiting its website) is minimal.

[184] European Commission, Proposal for a Regulation on Privacy and Electronic Communications, 10.1.2017 COM(2017) 10 final, URL: https://ec.europa.eu/digital-single-market/en/proposal-eprivacy-regulation

[185] On 23 October 2017, the EP adopted the report from rapporteur Birgit Sippel. URL: https://www.europarl.europa.eu/doceo/document/A-8-2017-0324_EN.html.

[186] The Council mandate (its agreed integrated text version of the ePrivacy Regulation), 10 February 2021, URL: https://data.consilium.europa.eu/doc/document/ST-6087-2021-INIT/en/pdf.

[187] Press release Council of the European Union, Confidentiality of electronic communications: Council agrees its position on ePrivacy rules, 10 February 2021, URL: https://www.consilium.europa.eu/en/press/press-releases/2021/02/10/confidentiality-of-electronic-communications-council-agrees-its-position-on-eprivacy-rules/

[188] French presidency, preparation for trialogue, 7458/22, 28 March 2022, URL: https://data.consilium.europa.eu/doc/document/ST-7458-2022-INIT/x/pdf.

10. Retention Periods

In its Privacy Notice, AWS describes some criteria to retain data. AWS writes: "*We keep your personal information to enable your continued use of AWS Offerings, for as long as it is required in order to fulfil the relevant purposes described in this Privacy Notice, as may be required by law (...).*"[189]

However, AWS does not provide details about the retention periods of the Account, Diagnostic, Support and Website Data. During a conference call on 3 November 2020 AWS explained its security logs are stored for a retention period of 10 years. In May 2023 AWS explained the minimum retention period for security events is 12 months, while the maximum remains 10 years.

As processor for the Content Data, AWS explains that the Customer determines the retention periods of the Content Data. Most cloud providers distinguish between active deletion of Content Data (including backups) by customers, and passive deletion, when Content Data are deleted after a certain period of time. In reply to this DPIA, AWS explained that customers may actively delete their content at any time. After an AWS account is closed, any content still stored in the customer's account will be deleted after 90 days.[190]

In reply to a question about the retention period of (centralised) backups, AWS only referred to different (paid) options for the Customer to make backups of Content Data.[191] In reply to this DPIA AWS referred to information about backups of relational databases: "*If you don't set the backup retention period, the default backup retention period is one day if you create the DB instance using the Amazon RDS API or the AWS CLI. The default backup retention period is seven days if you create the DB instance using the console.*"[192]

In the dialogue with SLM Rijk in May 2023, AWS explained that for all personal data except for the Content Data, the retention period is 90 days, unless one of four criteria necessitate or allow for a longer retention period:

1. As long as required to sustain ongoing use of the service
2. As long as necessary for security investigation – if related to security events
3. As long as required to generate invoices and to comply with tax legislation

AWS gave the example of Luxembourg tax law which requires the retention of billing data for 10 years.[193]

---

[189] AWS Privacy Notice, last updated 5 May 2023, URL: https://aws.amazon.com/privacy/

[190] AWS response to Part A of the DPIA, 1 October 2021, Par. 20.

[191] AWS refers to its (paid) AWS Backup services, URL: https://aws.amazon.com/backup/ . AWS explains: "*AWS Backup is AWS' native data protection platform that offers centralized, fully managed, and policy-based service to protect customer data and ensure compliance and business continuity across AWS service.*" In: AWS Storage blog, AWS Backup provides centralized data protection across your AWS resources, 9 November 2020, URL: https://aws.amazon.com/blogs/storage/aws-backup-provides-centralized-data-protection-across-your-aws-resources/

[192] AWS, Working with backups, URL: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html#USER_WorkingWithAutomatedBackups.BackupRetention.

[193] Luxembourg commercial code retention period. URL: https://legilux.public.lu/eli/etat/leg/loi/1931/05/22/n1/jo .

AWS additionally explained it ensures lawful use by policies that only allow for access when necessary.

Part B. Lawfulness of the data processing

This second part of the DPIA assesses the lawfulness of the data processing. This part contains a discussion of the legal grounds, an assessment of the necessity and proportionality of the processing, and of the compatibility of the processing in relation to the purposes.

11. Legal Grounds

To be permissible under the GDPR, processing of personal data must be based on one of the grounds mentioned in Article 6 (1) GDPR. Essentially, for processing to be lawful, this article demands that the data controller bases the processing on the consent of the user, or on a legally defined necessity to process the personal data.

The grounds mentioned in Article 6 (1) GDPR are as follows:

a) the data subject has given consent to the processing of his or her personal data for one or more specific purposes.

b) processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract.

c) processing is necessary for compliance with a legal obligation to which the controller is subjected.

d) processing is necessary in order to protect the vital interests of the data subject or of another natural person.

e) processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller.

f) processing is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.

The appropriate legal ground depends on AWS's role as processor, or as independent data controller. As explained in Section 5, AWS can formally and factually be qualified as data processor for the personal data in and about the use of Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon RDS. This applies to the processing of the following five categories of personal data:

1. Content Data (customer uploaded Content Data in the VMs and storage spaces)
2. Account Data (including Contact Data)
3. Diagnostic Data (including Configuration and Security Data)
4. Support Data
5. Website Data (the restricted access Admin Console)

In addition, AWS is authorised to 'further' process some personal data as independent data controller for a list of agreed compatible purposes, when strictly necessary and proportionate. The legal ground will be discussed in Section 11.2 below.

The assessment of available legal grounds (sometimes called 'lawful bases') for the five types of data processing by AWS (Content, Account, Diagnostic, Support and Website Data) is tied closely to the principle of purpose limitation. The EDPB notes:

"The identification of the appropriate lawful basis is tied to principles of fairness and purpose limitation. […] When controllers set out to identify the appropriate legal basis in line with the fairness principle, this will be difficult to achieve if they have not first clearly identified the purposes of processing, or if processing personal data goes beyond what is necessary for the specified purposes."[194]

Thus, in order to determine whether a legal ground is available for a specific processing operation, it is necessary to determine for what purpose, or what purposes, the data were or are collected and will be (further) processed. There must be a legal ground for each of these purposes.

As data processor AWS may only process the personal data for a limited set of specific and legitimate purposes. This makes it possible to identify appropriate legal grounds for each of the five types of data processing identified in this report.

Below four of the different possible legal grounds are briefly discussed for the different categories of personal data for the different purposes AWS is instructed to process the personal data when the Dutch government organisation is the sole data controller. Since this report is an umbrella DPIA, it cannot identify all the specific purposes a government organisation may want to serve through the use of the tested AWS services. As mentioned in Section 5.1 government organisations must identify at least one security purpose for the data they export from the different logs. Based on security requirements in art. 32 of the GDPR, and national security standards such as BIO, government organisations are legally required to retain and process these logs for some time for the purpose of testing, assessing and evaluating the effectiveness of their information security policy. The legal ground for this processing is not discussed separately in this report, as this report focusses on the data processing risks of remote data processing by AWS.

Section 11.2 below (AWS as independent data controller) discusses the agreed compatible purposes for which AWS may further process a limited set of personal data if strictly necessary and proportionate.

Additionally, AWS acts as an independent data controller for the processing of personal data relating to the visits to its public website, and for the Commercial Contact Data. The legal ground for the processing of personal data from these two sources is not assessed below, as this is the exclusive responsibility of AWS. The contract excludes possible joint controllership.

11.1 Legal grounds government organisations

The four relevant legal grounds are, in short: consent, contract with the data subject, public interest and legitimate interest. The legal ground of vital interest is not discussed, since nor AWS nor the Dutch government organisations have a vital (lifesaving) interest in processing personal data via AWS's cloud services. Additionally, there is no legal obligation for the Dutch government to use AWS (or any comparable cloud services).

[194] EDPB, Guidelines 2/2019 on the processing of personal data under Article 6(1)(b) GDPR in the context of the provision of online services to data subjects - version adopted after public consultation, 16 October 2019, URL: https://edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-22019-processing-personal-data-under-article-61b_en.

11.1.1 Consent

Article 6 (1) (a) GDPR reads: "*the data subject has given **consent** to the processing of his or her personal data for one or more specific purposes*"

Article 4(11) GDPR defines consent as "*consent of the data subject means any freely given, specific, informed and unambiguous indication of the data subject's wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her.*"

Consent generally is not available as legal ground for government organisations in relation to data processing of their employees, in view of the power imbalance. The only exception is the use of cookies. For the consent to be valid, the default settings must be privacy protective, and the website must effectively enable visitors to give free, specific and informed consent.

AWS uses a confusing cookie banner on its restricted access websites, as described in Section 3.2.1 of this report. Admins can proceed to only accept the strictly necessary 'Essential' cookies by selecting the third option, to *'Customize cookies'*. However, if an admin chooses this third option, and accepts the default setting in this pop-up screen, AWS still collects analytical personal data through website telemetry. These data include the full user account name. This type of data collection is not compatible with the specific Dutch consent exception described in Section 9 of this report. Hence, AWS should stop collecting this name, or ask for specific consent, and inform the admins about this specific data processing for analytical purposes. If AWS does not stop collecting the user name, government organisations may have to rely on the legal ground of necessity for their legitimate interest for this data processing. To rely on that ground, it is crucial that they inform admins, and pseudonymise the admin account data.

11.1.2 Contract

Article 6 (1) (b) GDPR reads: "*processing is **necessary for the performance of a contract** to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract.*"

Government organisations frequently rely on the legal ground of necessity for the performance of a (labour) contract with their employees, certainly when the processing involves personal data from admins. This legal ground can only be used when the processing is strictly necessary to perform the contract with these individuals. The EDPB explains:

> "*the controller should be able to demonstrate how the main subject-matter of the specific contract with the data subject cannot, as a matter of fact, be performed if the specific processing of the personal data in question does not occur. The important issue here is the nexus between the personal data and processing operations concerned, and the performance or non-performance of the service provided under the contract*."[195]

In order to manage the tested AWS services, admins must create an account with AWS. Hence the processing of these personal data by AWS (that can be

[195] EDPB, Guidelines 2/2019 on the processing of personal data under Article 6(1)(b) GDPR in the context of the provision of online services to data subjects, version 2.0, 8 October 2019, par. 30, p. 9-10, URL: https://edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines-art_6-1-b-adopted_after_public_consultation_en.pdf. .

pseudonymised to minimise data protection risks) is necessary for admins to perform their work tasks.

Similarly, admins working for government organisations must necessarily access the restricted access websites (Admin Console and Support Centre) to manage the AWS services. They must file a support request when they can't solve an issue with the 3 tested services. They cannot fulfil the main object of their contract with the Dutch government organisation without visiting the restricted access websites and filing support requests when necessary. These activities inevitably results in the processing of their personal data in AWS's webserver access log and support ticketing system.

As described in Section 4.2, the Dutch government has instructed AWS as processor to process the personal data from the admins for the following purposes: to provide and maintain the services, secure the services and AWS network, provide customer-requested support, and perform basis troubleshooting.

These purposes, with their specific sub-purposes, ensure that the data processing is limited to the necessary operations to deliver the requested services in a well-functioning and secure manner to the Dutch government. This means government organisations can "*justify the necessity of its processing by reference to the fundamental and mutually understood contractual purpose*."[196]

Based on security requirements in art. 32 of the GDPR, and national security standards such as BIO, government organisations are legally required to retain and process logs with Diagnostic Data about access to the 3 tested AWS services for some time for the purpose of testing, assessing and evaluating the effectiveness of their information security policy. For the processing of these exported personal data, government organisations can also plausibly rely on the legal ground of contract.

In sum, it is likely that government organisations can rely on the legal ground of contract for the processing of Account, Diagnostic, Support and restricted access Website Data by AWS as data processor, for the agreed purposes.

11.1.3 Necessity for a public interest

Article 6 (1) (e) GDPR reads: "*processing is necessary for the **performance of a task carried out in the public interest** or in the exercise of official authority vested in the controller.*"

The use of cloud storage and database tools can be necessary to perform tasks carried out by government organisations in the public interest, certainly if it involves direct interaction with members of the public/inhabitants. As the EDPB notes in a letter to the European Commission on the use of contract tracing apps relating to the COVID-19 pandemic,

> "*When public authorities provide a service, based on a mandate assigned by and in line with requirements laid down in law, it appears that the most relevant legal basis for the processing is the necessity for the performance of a task for public interest.*"[197]

The legal ground of public interest should be used in combination with the necessity to perform the (employment)contract and, as will be assessed below, in limited cases, also for the necessity for the legitimate interest of the government organisations.

[196] Idem, par 32, p. 10.

[197] https://edpb.europa.eu/sites/edpb/files/files/file1/edpbletterecadvisecodiv-appguidance_final.pdf

11.1.4 Legitimate interest

Article 6 (1) (f) GDPR reads: "*processing is **necessary for the purposes of the legitimate interests pursued by the controller** or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.*"

The last sentence of Article 6(1) of the GDPR adds: "*Point (f) of the first subparagraph shall not apply to processing carried out by public authorities in the performance of their tasks.*"

The last sentence of Article 6(1) of the GDPR excludes the application of the legitimate interest ground for processing carried out by public authorities in the performance of their tasks. However, government organisations may still invoke this legal ground for necessary and proportionate data processing by cloud providers. Public sector organisations may process personal data in a different role, outside of the tasks they carry out in the public interest, for example, when they hire office space or pay salaries to employees. The choice to use certain cloud services is secondary to the performance of public tasks and can be considered as a task primarily exercised under private law.

Recital 47 of the GDPR explains:

> "*Such legitimate interest could exist for example where there is a relevant and appropriate relationship between the data subject and the controller in situations such as where the data subject is a client or **in the service of the controller**.*"

Whether a government organisation can successfully invoke the legal ground of necessity for its legitimate interest, must be assessed on a case-by-case basis. The proportionality of the processing plays a crucial role. This will be elaborated in Section 14.2 of this report. In any event, government organisations must ensure "*that the interests or the fundamental rights and freedoms of the data subjects are not overriding, taking into account the reasonable expectations of data subjects based on their relationship with the controller*" (Recital 47 GDPR).

There is one specific example of data processing, related to the restricted access Websites, that requires government organisations to rely on the necessity for their own legitimate interest. As described in Section 3.2.1, AWS does not set cookies on its restricted access websites (Admin Console and Support Centre) that require consent. However, the restricted access Websites do read (collect) telemetry data from the website, visible in the browser of the admins, with their directly identifying account name. The purpose seems to be analytical, to examine the use and performance of these websites in relation to different browsers and settings. Though there is an exception on the consent requirement for analytical cookies in the Dutch implementation of the ePrivacy Directive, the collection of directly identifiable data does not comply with this exception. AWS does not provide any public explanation about this behaviour, and does not offer an opt-out to admins. To mitigate the risk of this excessive data processing, government organisations must pseudonymise the admin accounts.

Dutch government organisations also have to rely on the necessity for their legitimate interest to allow AWS to further process some personal data for its own controller purposes. This will be discussed below, in Section 11.2 of this DPIA.

In sum, it appears Dutch government organisations generally can rely on the legal ground of necessity to perform a contract and sometimes on the legal ground of public interest for the data processing by AWS as processor. Government organisations generally should not rely on the legal ground of necessity for a legitimate interest. However, if AWS does not ask for consent for the reading of website telemetry data, or stop collecting the directly identifiable data, government organisations have to pseudonymise the admin data, and invoke this legal ground.

11.2 Legal grounds AWS as independent data controller

As explained in Sections 4.3 and 5.3 the Dutch government has authorised AWS to 'further' process some personal data for a list of agreed compatible purposes, when strictly necessary and proportionate. These purposes range from billing and calculating employee compensation to combatting fraud, and from responding to data subject access requests for personal data in AWS's controller role to improving the performance and core functionality of the services.

The Dutch government generally relies on the legal ground of article 6(4) GDPR: the compatibility test for data processing by AWS for purposes other than those for which the personal data are collected. This legal ground is only available if the processing is not based on the data subject's consent or on a Union or Member State law.

As explained in Section 4.3, AWS is bound contractually to strict limits to this data processing, by ensuring AWS generally only processes aggregated data above tenant level for these purposes, and only individual personal data when strictly necessary and proportionate. The processing for most of the purposes is strictly necessary for AWS's existence as a commercial company. AWS may seek to rely on consent as legal ground for subscriptions of admins to professional newsletters.

There is only one situation in which the Dutch government does not have a legal ground for the further processing by AWS: if AWS is compelled to disclose customer data to a government authority and is prohibited from informing and redirecting the order to its customer. AWS has committed to challenge any overbroad or inappropriate request, and has agreed to follow a strict procedure to minimise the occurrence of this risk. It follows from AWS's public transparency reporting and audit reports that the probability of occurrence of this scenario is extremely slim with regard to Dutch government personal data. See Section 4.3.2 for more details.

*Table 2: Overview of agreed compatible purposes and legal grounds AWS and Dutch government [**Confidential**]*

12. Special categories of personal data

As explained in section 2.3.1 of this DPIA (*Content Data*), it is up to the individual government organisations to determine if they process special categories of data on a VM, in an S3 bucket or in a RDS database.

Special categories of data are

> *data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health, data concerning a natural person's sex life or sexual orientation or data relating to criminal convictions and offences.*

As described in Section 7.3.5 a separate Data Transfer Impact Assessment (DTIA) was performed to determine if the specific data protection risks associated with the storing of special categories of data on AWS's cloud computers can be sufficiently mitigated through additional protection measures.

Government organisations must take all the circumstances of the transfer into account, such as the probability risk of AWS providing compelled access to these data when they are exclusively stored in datacentres in the EU, in view of AWS's transparency reports and contractual guarantees.

According to Use case 6 of the final version of the EDPB Recommendations on measures that supplement transfer tools, encryption cannot be considered a sufficient measure if the encryption keys are accessible to AWS. However, as described in Section 8.1, with the use of AWS Nitro System, the access by AWS employees to the customer keys to encrypt the Content Data is nearly impossible. Auditor NCC group recently completed an audit on the system design and concluded "*there is no indication that a cloud service provider employee can obtain such access…to any host.*"[198]

The data protection risks for data subjects are not limited to the processing of special categories of data. Similar risks may apply to other categories of personal data of a sensitive nature, classified or secret data. The EDPS explains in its guidelines on the use of cloud computing services by European institutions that special categories of data should be interpreted broadly when interpreting the risks for data subjects.

The EDPS writes:

> "*Nevertheless, this is not the only factor determining the level of risk. Personal data that do not fall under the mentioned categories might lead to high levels of risk for the rights and freedoms of natural persons under certain circumstances, in particular when the processing operation includes the scoring or evaluation of individuals with an impact on their life such as in a work or financial context, automated decision making with legal effect, or systematic monitoring, e.g. through CCTV.*" [199]

The EDPS also refers to the criteria provided by the Article 29 Working Party when a Data Protection Impact Assessment (DPIA) is required. [200]

Government organisations must also consider the risk that special categories of data (or otherwise sensitive, confidential or secret data) could end up in Support Requests filed by system administrators. Similarly, some Admin Contact and Diagnostic Data may be confidential, if it is strictly confidential what employees work for a specific government organisation.

---

[198] Idem.

[199] EDPS, Guidelines on the use of cloud computing services by the European institutions and bodies, 10 March 2018, URL: https://edps.europa.eu/sites/edp/files/publication/18-03-16_cloud_computing_guidelines_en.pdf

[200] The EDPB has adopted the Article 29 Working Party guidelines WP 248 rev.01, Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679.

The DTIA specifies two mitigating measures government organisations can take to lower the probability of compelled disclosure of such special categories of data in the Account and Support Data:

1. Pseudonymise the admin account data
2. Ask the AWS sales manager to issue an internal alert to have support tickets only dealt with by employees in countries with an adequate data protection regime (the European Economic Area, Japan and Canada).

13. Purpose limitation

Article 5(1) (b) of the GDPR obliges data controllers to comply with the principle of purpose limitation. Data may only be "*collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, in accordance with Article 89(1), not be considered to be incompatible with the initial purposes.*"

Essentially, the above means that the controller must have a specified purpose for which he collects personal data and can only process these data for purposes compatible with that original purpose, unless the controller can meet additional requirements, that have to be met before any further processing can take place.

Data controllers must be able to prove based on Article 5(2) of the GDPR that they comply with this purpose limitation principle (accountability). They cannot comply if they engage a cloud provider as data processor, but ignore that that provider also reserves the contractual right to process the personal data for its own purposes, in its own commercial business interest. As described in the introduction, and in Section 4.2, initially AWS reserved the right to process all personal data for its own purposes, in a role as independent controller. The contract with the Dutch government includes a limitative list of 3 main purposes, with specific sub-purposes for which AWS as processor may process the personal data relating to the 3 tested services. These purposes apply to the five categories of personal data, and not just to the Content Data. Additionally, AWS is authorised to further process personal data as controller for a second list of compatible purposes, when strictly necessary and proportionate.

The principle of purpose limitation is closely linked to transparency and 'fairness' of the processing. Together, implementation of these principles should lead to surprise minimisation for the government employees. As the Article 29 Working Party wrote in its guidelines on transparency (adopted by the EDPB):

> *"A central consideration of the principle of transparency outlined in these provisions is that the data subject should be able to determine in advance what the scope and consequences of the processing entails and that they should not be taken by surprise at a later point about the ways in which their personal data has been used. This is also an important aspect of the principle of fairness under Article 5.1 of the GDPR and indeed is linked to Recital 39 which states that "[n]atural persons should be made aware of risks, rules, safeguards and rights in relation to the processing of personal data(...)."*[201]

[201] The EDPB has adopted the Article 29 Working Party guidelines WP 260 rev 1, Guidelines on transparency under Regulation 2016/679, adopted on 29 November 2017, as last Revised and Adopted on 11 April 2018, Par. 41.

It follows that government organisations must be able to inform their admin employees for what specific purposes their personal data will be processed, and be able to inform other persons whose personal data may be processed as part of the Content Data about the agreed purposes, both when AWS qualifies as processor, and as data controller.

Initially, at the start of this DPIA project, it required close reading of the contractual framework to understand for what purposes AWS processed the different categories of personal data. The contract now contains strong purpose limitation, with two lists of well-defined processor and 'further processing' purposes. AWS has additionally committed to update an overview of the categories of personal data it may 'further' process as controller. Government organisations can use this future documentation to inform their employees and update their records of processing activities (ROPA).

Though AWS provides some access to the Diagnostic Data in the S3 Access Logs and Cloud Trail logs, AWS does not publish any information about the Diagnostic Data it collects on its own servers and has only provided limited access to very limited contents of its security logs for this DPIA. AWS demonstrated that the personal data in the log files is limited and can be further limited by the government organisations (see below).

In follow-up discussions in May 2023 AWS also pointed to a recently completed new audit, against the German C5:2020 standard. The audit confirms that AWS solely collects and uses the Diagnostic Data for the three purposes of billing, incident management and security management.[202] The report also states that AWS complies with other tested criteria, such as use of solely anonymous metadata to deploy and enhance the cloud service, no commercial use of these data and automated removal of personal data from log data as far as technically possible. As explained in Section 3.1.3, to Privacy Company, none of the Diagnostic Data seem excessive for security purposes. The collection rules seem in line with cloud provider industry practices. However, since the security logs may also include IP addresses of visitors to Dutch government applications or websites (if they are hosted on AWS), Dutch government organisations should consider using a proxy, if they conclude the processing by AWS posed a high data protection risk for data subjects. For example, if AWS would collect IP addresses from visitors to the website https://www.meldmisdaadanoniem.nl/[203] and would be compelled to disclose these data to Dutch or foreign law enforcement authorities.

The Dutch government will use its audit right to verify AWS's data processing in this respect.

14. Necessity and proportionality

Article 5(1)c GDPR requires that every processing has to be limited to what is necessary to achieve the set purpose(s). It is therefore important to examine whether

[202] EY describes the inspection method as follows: "*Inquired of an AWS Security Assurance Program Manager to ascertain formal security policies existed, included designation of responsibility and accountability for managing the system and controls, and provided guidance for information security within the organization and the supporting IT environment. No deviations noted.*"

[203] M. 'Guaranteed anonymity', URL: https://www.meldmisdaadanoniem.nl/english/

every processing is in fact necessary for the purposes for which the data controllers process personal data.

The concept of necessity is made up of two related concepts, namely proportionality and subsidiarity. The personal data which are processed must be necessary for the purpose pursued by the processing activity. First, it has to be assessed whether the same purpose can reasonably be achieved with other, less invasive means. If so, these alternatives have to be used.

Second, proportionality demands a balancing act between the interests of the data subject and the data controller. Proportionate data processing means that the amount of data processed is not excessive in relation to the purpose of the processing. If the purpose can be achieved by processing fewer personal data, then the amount of personal data processed should be decreased to what is necessary. Therefore, essentially, the data controller may process personal data insofar as is necessary to achieve the purpose but may not process personal data he or she may do without. The application of the principle of proportionality is thus closely related to the principles of data protection from Article 5 GDPR.

14.1 Assessment of the proportionality

The key questions are: are the interests properly balanced? And, does the processing not go further than what is necessary?

To assess whether the processing is proportionate to the interest pursued by the data controller(s), the processing must first meet the principles of Article 5 of the GDPR. As legal conditions they have to be complied with in order to make the data protection legitimate.[204]

Data must be '*processed lawfully, fairly and in a transparent manner in relation to the data subject*' (Article 5 (1) (a) GDPR). This means that data subjects must be informed about the processing of their data, that all the legal conditions for data processing are adhered to, and that the principle of proportionality is respected.

There is no public, centrally accessible source of information from AWS that describes the personal data collected in the different categories of Diagnostic Data identified in this DPIA. System administrators can view some of the system generated server logs in the S3 Access and the CloudTrail logs, but none of the data generated through website visits (the publicly accessible website, the Admin Console and the Support Portal). AWS does not publish information about the personal data that it may process in logs that it generates about the use of all its services by system administrators (system-generated server logs), its SIEM logs about activities by admins and external visitors to information hosted on its network, VMs and databases or its webserver access logs. These logs include data about access attempts by third parties. In view of the contractual purpose limitation, the outcomes of the recent C5:2020 audit report, and the right of the Dutch government to conduct its own audit on AWS's

[204] See for example CJEU, C-131/12, Google Spain SL, Google Inc. v Agencia Española de Protección de Datos (AEPD), Mario Costeja González, ECLI:EU:C:2014:317. Paragraph 71: *In this connection, it should be noted that, subject to the exceptions permitted under Article 13 of Directive 95/46, all processing of personal data must comply, first, with the principles relating to data quality set out in Article 6 of the directive and, secondly, with one of the criteria for making data processing legitimate listed in Article 7 of the directive (see Österreichischer Rundfunk and Others EU:C:2003:294, paragraph 65; Joined Cases C-468/10 and C-469/10 ASNEF and FECEMD EU:C:2011:777, paragraph 26; and Case C-342/12 Worten EU:C:2013:355, paragraph 33).*

compliance, this lack of transparency does not make the data processing inherently disproportionate.

Customers can mitigate the risks of AWS's lack of publicly accessible documentation by pseudonymising the admin accounts. As mentioned above in Section 13, if they conclude that AWS's recording of IP addresses of visitors to Dutch government applications or websites causes a high data protection risk, they can either host their website somewhere else, or use a proxy before redirecting visitors to AWS.

The principles of data minimisation and privacy by default demand that the processing of personal data is limited to what is necessary: Data must be "*adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed*" (article 5 (1) (c) GDPR). This means that government organisations may not collect and store data (or allow AWS to collect and store data) which are not directly related to a legitimate purpose. According to this principle, the default settings for the data collection should be set in such a way as to minimise data collection by using the most privacy friendly settings. AWS is commendable for having taken this principle to heart for cookies used on its websites, but unfortunately, AWS does collect analytical data through website telemetry, without informing the admins, or giving them an option to opt-out from this data processing.

The principle of storage limitation demands that personal data are only retained as long as necessary for the purpose in question. Data must be "*kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed*" (article 5 (1) (e), first sentence GDPR). This principle therefore demands that personal data are deleted as soon as they are no longer necessary to achieve the purpose pursued by the controller. The text of this provision goes on to clarify that "*personal data may be stored for longer periods insofar as the personal data will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes in accordance with Article 89(1) subject to implementation of the appropriate technical and organisational measures required by this Regulation in order to safeguard the rights and freedoms of the data subject*" (article 5 (1) (e), second sentence, GDPR).

As described in Section 10 of this report, AWS has informed SLM Rijk that it generally retains personal data for a period of 90 days, and longer if necessary for one of the four criteria:

1. As long as required to sustain ongoing use of the service
2. As long as necessary for security investigation – if related to security events
3. As long as required to generate invoices and to comply with tax legislation

Additionally, AWS has explained the minimum retention period for security events is 12 months, while the maximum retention period is 10 years, for example when AWS is compelled to retain billing data for a period of 10 years based on Luxembourg tax law.

In view of the strict purpose limitation for AWS (both as processor, and as authorised further processing controller) and the confirmation in the recent C5:2020 audit that AWS does not use these data for any commercial purposes, the long retention periods are not per definition disproportionate.

14.2 Assessment of the subsidiarity

The key question is whether the same goals can be reached with less intrusive means.

As described in Section 1, AWS competes with global players with similar infrastructure and database offerings. Like its competitors, AWS offers a high-quality service at a reasonable cost, in terms of scalability, flexibility and reliability.

There is no publicly available information that compares the compliance of all of these competing suppliers of cloud services with privacy laws and regulations. The Dutch government has negotiated or aims to negotiate GDPR-compliant contracts for the use of cloud services from Microsoft, AWS, Google, Oracle, Salesforce, SAP and IBM. These contracts ensure that government organisations can choose between multiple cloud providers.

As described in Section 8 (*Data minimisation: encryption and pseudonymisation*) AWS offers privacy configuration controls that allow government organisations to influence the impact of the data processing. These controls are important, but they are not sufficient by themselves to mitigate all high risks. The only alternative for individual employees is not to reject using AWS's services, but that is only possible if their employer has a realistic option to switch to another cloud provider.

Most cloud providers are USA based companies, which at least in part cause the same risks for data subjects, for example with regard to the transfer of personal data to the United States. Of the companies mentioned above, only SAP is exclusively EU based. As mentioned in Section 7.2.3, Microsoft has announced it will keep all the data processing within the EU by the end of 2024 the latest. Oracle has announced the creation of an exclusive EU company, without any involvement from the USA.[205]

A relatively new development is the provision of cloud services via GAIA-X, an collaborative EU cloud infrastructure initiated with French and German government support that focusses on data sovereignty.[206]

Even in the event of a potential switch to another cloud provider, the government organisations must first identify the privacy and security risks, and the question whether the competitor offers the necessary functionalities. There may also be costs for migration to new systems and redevelopment of specific applications that government organisations already use. This situation is also referred to as *vendor lock-in*. Though one of GAIA-X's aims is to liberate organisations from vendor lock-in, it is not clear what alternatives the members of the Dutch Gaia-X Hub currently offer, and if a switch from a hyperscaler to a Dutch member is a viable option.[207]

In sum, government organisations should consider if competing providers can offer the desired services with a higher level of privacy protection. SLM Rijk will publish updates and DPIAs at https://slmmicrosoftrijk.nl/downloads-dpias/ .

---

[205] Oracle Cloud Infrastructure blog, Introducing Oracle's sovereign cloud regions for the European Union, 11 July 2022, URL: https://blogs.oracle.com/cloud-infrastructure/post/introducing-oracles-sovereign-cloud-regions-for-the-european-union.

[206] Gaia-x homepage, URL: https://www.data-infrastructure.eu/GAIAX/Navigation/EN/Home/home.html

[207] Press release TNO (in Dutch), Nederlandse hub gaia-x voor invulling en bijdrage Europese data- en cloudinfrastructuur, 29 June 2021, URL: https://www.tno.nl/nl/over-tno/nieuws/2021/7/nederlandse-hub-gaia-x/

15. Rights of data subjects

The GDPR grants data subjects the right to information, access, rectification and erasure, object to profiling, data portability and file a complaint. It is the data controller's obligation to provide information and to duly and timely address these requests. If the data controller has engaged a data processor, the GDPR requires the data processing agreement to include that the data processor will assist the data controller in complying with data subject rights requests.

As discussed in Section 5, AWS qualifies as data processor for almost all data processing in the context of the three tested cloud services and will assist the controller with any data subject access requests.

15.1 Right to information

Data subjects have a right to information. This means that data controllers must provide people with easily accessible, comprehensible and concise information in clear language about, inter alia, their identity as data controller, the purposes of the data processing, the intended duration of the storage and the rights of data subjects. As quoted above in Section 13, the EDPB explains in its Guidelines on transparency that controllers should clearly explain the most important consequences of the processing.

> *"[n]atural persons should be made aware of risks, rules, safeguards and rights in relation to the processing of personal data."*[208]

With the help of this umbrella DPIA, government organisations that wish to use the 3 tested AWS services can inform their employees about the agreed scope and purposes of the data processing. Additionally, AWS has committed to update an overview of the categories of Personal Data that AWS processes as a controller for agreed compatible purposes.

15.2 Right to access

Secondly, data subjects have a right to access personal data concerning them. Upon request, data controllers must inform data subjects whether they are processing personal data about them. If this is the case, data subjects should be provided with a copy of the personal data processed, together with information about the purposes of processing, recipients to whom data have been transmitted, the period for which personal data are to be stored, and information on their further rights as data subjects, such as filing a complaint with the Data Protection Authority.

As a data processor AWS enables government organisations (as customers) to upload, change and delete the Content Data. AWS also offers some tools to government organisations to get access to some personal data that are part of the Account Data, and/or work with federated identity to pseudonymise the Account Data. Admins can access, view and edit the personal information they have actively provided to AWS via the AWS Management Console, can access payment information, charges and account activity, and security credentials via the AWS Account pages, and access and update communications preferences by visiting the AWS Communications Center.

Additionally, the root admin can grant admins access to personal data in the S3 Access logs and the CloudTrail. AWS does not offer a specific tool for the root admin to selectively download all personal data relating to a specific admin, including

[208] The EDPB has adopted the Article 29 Working Party guidelines WP 260 rev 1, Guidelines on transparency under Regulation 2016/679, adopted on 29 November 2017, as last Revised and Adopted on 11 April 2018, Par. 41.

Diagnostic, Support and Website Data. AWS considers that admins have access to all necessary data through the Service Controls. However, AWS does have a separate Data Request Form that admins can use to ask for a copy of their Console usage data.[209] The customer may request further assistance from AWS in exceptional circumstances. AWS may charge additional costs for such requests.

As described in Section 3.3. AWS did not provide a complete overview of personal data when Privacy Company filed a data subject access request for the data relating to its test user. The 'missing' data from the DSAR response include:

1. use of the AWS Admin and Root Account outside of available logs in CloudTrail;
2. information recorded in the webserver access logs with information about IP address, end user, device and activities;
3. information about filed Support Requests;
4. information collected by AWS in its network logs, and;
5. the necessary extra information elaborated in Article 15 (1) GDPR, in particular the provisions c (recipients), d (retention period), (f) the right to lodge a complaint with the DPA, (g) information about external sources and (h) automated decision-making, including profiling, in particular with regard to credit-scoring.

In reply to the final DPIA, AWS explained that all information in the fifth bullet point is (currently) available in its Privacy Notice, and therefore not 'missing'. However, as the EDPB explains, a reference to a Privacy Notice is not personalised enough, as information on recipients, on categories and on the source of the data may vary depending on who makes the request and what the scope of the request is.[210]

In reply to Part A of this DPIA, AWS objected that not all data mentioned in the list above as 'missing' from the Data Subject Access Request (DSAR) are personal data or should be provided, as "*disclosure in the context of a DSAR could result in unlawful disclosure of personal data and be antithetical to the objective of data protection.*"[211]

AWS did not elaborate when certain data are not personal data, or when providing access would be 'unlawful' or 'antithetical'.

Privacy Company did not perform a re-test of the Data Subject Access Request in the new situation with AWS as processor, using AWS's Service Controls, Therefore this DPIA cannot draw any conclusions about the quality of AWS's (future) responses. If admins are unable to obtain access to some personal data, they can ask AWS for assistance.

If AWS, in such an exceptional assistance procedure, would want to argue that disclosing certain personal data would harm its company confidential objectives, such as its exact security posture, it should provide detailed reasoning. As the EDPB explains in its Guidelines on restrictions under Article 23:

---

[209] AWS Data Request Form, URL: https://s3-us-west-2.amazonaws.com/aws-support-documents/Forms/AWSDataRequestForm.pdf.

[210] EDPB. Adopted Guidelines 01/2022 on data subject rights - Right of access, Version 2.0, Adopted on 28 March 2023, par. 113, URL: https://edpb.europa.eu/system/files/2023-04/edpb_guidelines_202201_data_subject_rights_access_v2_en.pdf .

[211] AWS response to part A of the DPIA, 1 October 2021, Par. 6, p. 2.

"*Any restriction shall respect the essence of the right that is being restricted. This means that restrictions that are extensive and intrusive to the extent that they void a fundamental right of its basic content, cannot be justified. In any case, a general exclusion of data subjects' rights with regard to all or specific data processing operations or with regard to specific controllers would not respect the essence of the fundamental right to the protection of personal data, as enshrined in the Charter. If the essence of the right is compromised, the restriction shall be considered unlawful, without the need to further assess whether it serves an objective of general interest or satisfies the necessity and proportionality criteria.*"[212]

The Dutch implementation law implements the same exceptions as Art. 15(4) and Art. 23 of the GDPR, with the legal explanation that data controllers may only invoke these exceptions when it is strictly necessary, as foreseen in Art. 8 (2) of the ECHR and Art. 51 of the Charter of Fundamental Rights.[213]

In sum, based on the role of AWS as processor for almost all of the data processing, system administrators should be able to access the available Account, Support and Diagnostic Data, as well as the restricted access Website Data during the time period that the admin can reliably identify him/herself or colleagues. If AWS would refuse access to some personal data, and the admin is not convinced by the arguments, he/she can file a complaint with the Dutch Data Protection Authority. It is then up to the Dutch DPA to assess the validity of such a refusal.

15.3 Right of rectification and erasure

Thirdly, data subjects have the right to have inaccurate or outdated information corrected, incomplete information completed and - under certain circumstances - personal information deleted or the processing of personal data restricted.

Customers can delete Content Data, and can access and change some of the Account and Support Data. However, as noted in Section 10 of this DPIA, AWS applies a grace period of 90 days before deleting the Content Data. This period allows customers to either retrieve or delete the data. Customers cannot delete any data retained up to 10 years for purposes such as security and fiscal obligations. It is not clear if AWS is obliged under the very broad Luxembourg fiscal retention obligation to keep copies of Account and Support Data. However, as a result of purpose limitation such personal data are only retained and accessed for that specific fiscal audit purpose, and not for any other purpose.

AWS has recently confirmed to Privacy Company that it is possible for a customer to request deletion of Support Tickets. This procedure was not tested. Though the information about the retention periods is minimal, and the retention period of 10 years for security data seems excessively long, AWS seems to meet the minimum requirements of Article 17(1)(a) and Article 17(1)(d) of the GDPR. These provisions require a data controller to delete personal data without undue delay upon request of a data subject if they are no longer needed for the purposes for which they were

---

[212] EDPB , par. 14

[213] Memorie van toelichting bij de UAVG Art. 41: "*Gelet op het belang van de rechten van betrokkene, de meldplicht en de beginselen dienen verwerkingsverantwoordelijken alleen van de bevoegdheid om af te wijken gebruik te maken indien dit strikt noodzakelijk is en op proportionele wijze gebeurt. Net als onder artikel 43 van de Wbp geldt voor de toepasselijkheid van deze gronden dus een strikt noodzakelijkheidscriterium (vergelijk artikel 8, tweede lid, van het EVRM en artikel 52, eerste lid, van het Handvest).*"

collected or otherwise processed, or when the personal data have been unlawfully processed.

15.4 Right to object to profiling

Fourthly, data subjects have the right to object to an exclusively automated decision if it has legal effects.

AWS contractually guarantees that it does not use the personal data from its customers (sales contacts or admins) for profiling purposes, unless the admin has provided valid consent for example for tracking cookies.

Therefore, this specific right of objection does not apply in this case.

15.5 Right to data portability

Employees have a right to data portability if the processing of their personal data is carried out by automated means and is based on their consent or on the necessity of a contract. As explained in Sections 11.1 and 11.2 of this report, the processing by AWS on behalf of government organisations is generally based on the necessity of performing a (employment) contract with the admin. The exercise of the right to data portability is not realistically possible for admins: as they would copy confidential data and personal data from the government organisation.

The individual right to data portability is independent of the situation where government organisations themselves would want to move their processing and files collectively to another provider, to escape from vendor lock-in (see Section 14.2 above).

15.6 Right to file a complaint

Finally, government organisations as controllers must inform their employees about their right to complain, internally to their Data Protection Officer (DPO), and externally, to the Dutch Data Protection Authority (Autoriteit Persoonsgegevens).

Part C. Discussion and Assessment of the Risks

This part concerns the description and assessment of the risks for data subjects. This part starts with an overall identification of the risks to the rights and freedoms of data subjects as a result of the processing of the Content, Account, Diagnostic, Support and restricted access Website Data. The risks will subsequently be classified according to the likelihood they might occur, and the impact on the rights and freedoms of the data subjects when they do.

16. Risks

16.1 Identification of risks

Below, a general distinction is made between the risks of the processing of metadata on the one hand, and the Content Data on the other hand. Subsequently, 9 specific data protection risks are described, 8 of which apply to the Metadata and 2 to the Content Data. One risks overlaps: of the lack of control over the use of subprocessors.

The data protection risks can be grouped in the following categories:

- inability to exercise rights (including but not limited to privacy rights)
- inability to access services or opportunities
- loss of control over the use of personal data
- discrimination
- identity theft or fraud
- financial loss
- reputational damage
- physical harm
- loss of confidentiality
- re-identification of pseudonymised data or
- any other significant economic or social disadvantage[214]

These risks have to be assessed against the likelihood of the occurrence of these risks and the severity of the impact.

The UK data protection commission ICO provides the following guidance: *Harm does not have to be inevitable to qualify as a risk or a high risk. It must be more than remote, but any significant possibility of very serious harm may still be enough to qualify as a high risk. Equally, a high probability of widespread but more minor harm might still count as high risk.* [215]

In order to weigh the severity of the impact, and the likelihood of the harm for these generic risks, this report combines a list of specific risks with specific circumstances of the specific investigated data processing.

---

[214] List provided by the ICO, https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/data-protection-impact-assessments-dpias/how-do-we-do-a-dpia/

[215] ICO, How do we do a DPIA?, URL: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/data-protection-impact-assessments-dpias/how-do-we-do-a-dpia/.

16.1.1 Metadata

Because the level of data protection in the USA is not considered adequate (yet), there are inherent risks related to the transfer of personal data to the USA.

The processing by AWS and its affiliates as sub-processors in case of Support Requests of the real names and email addresses of certain government employees can lead to additional data protection risks if such organisations work with highly confidential data, or publicly controversial databases.

Government organisations that use the AWS cloud services are technically able to use metadata about the use of the service by admins (use of the Admin account, filing of Support Requests) to gain some insight in the work patterns of admins with Cloud Services accounts. The system administrators may experience a (small) chilling effect as a result of the monitoring of their behavioural data.

The risks mentioned above are inherent to the use of any IT service. These risks are not specific for the use of AWS's cloud storage and database services.

Government organisations can mitigate these risks by pseudonymising the admin account data, and instructing admins never to upload personal data as part of support tickets.

AWS requires actively given consent from admins before subscribing them to commercial mails. AWS may send unsolicited mail to its Commercial Contacts, such as the procurement officers of Dutch government organisations. They can opt-out from such newsletters and mails by visiting the Marketing Preference Centre. Privacy Company did not re-test the configuration of the Marketing Preference Centre for admins under the new conditions, but assumes AWS no longer automatically subscribes admins to its commercial communications.

Overall, the processing of metadata relating to the use of AWS's cloud services results in the following 8 data protection risks.

1. Loss of control, loss of confidentiality through disclosure or unlawful access to Account, Diagnostic, Support and Website Data as a result of transfer to the USA
2. Loss of control through lack of legal ground for cookies and website telemetry restricted access Website Data
3. Loss of control over processing by subprocessors
4. Loss of control through lack of transparency Diagnostic, Support and Website Data
5. Incomplete exercise of data subjects rights to some Account, Diagnostic, Support, and Website Data
6. Chilling effects employee monitoring system
7. Loss of control over personal data in inaccessible AWS security logs
8. The sending of unsolicited marketing mail to procurement officers

16.1.2 Content Data

In its role as data processor for the Content Data, AWS can process highly sensitive or confidential/secret content data from Dutch government organisations. Organisations can use different encryption modules to the data at rest. As explained

in Section 8.1, with AWS Nitro System the probability of unlawful access to the key or decrypted data has become extremely small.

Similar to the transfer risks described above for the metadata, the use of cloud storage and database services entails the same risks of unlawful access by a rogue system administrator or hostile state actor, and valid orders from law enforcement and secret services/security agencies to hand-over content data, even when the Content Data are exclusively processed within the EU. Unlawful disclosure of the contents of data hosted on an AWS VM, bucket or database could breach government and professional confidentiality and secrecy classifications.

Overall, the processing of Content Data in AWS VMs, buckets and databases results in the following two data protection risks:

1. Unlawful disclosure to, or access of, government authorities to Content Data
2. Loss of control processing of Content Data in Support Requests by AWS's affiliate sub-processors.

16.2 Assessment of risks

16.2.1 Disclosure or access to Content Data as result of transfer to the USA

As described in Section 7.3.5, a separate Data Transfer Impact Assessment (DTIA) was performed for the different categories of personal data processed by AWS. It follows from the Schrems II ruling from the ECJ and the analysis of the EDPB essential guarantees that the legal regime of the USA does not offer an essentially equivalent level of protection.

Based on AWS's transparency reporting, audited security measures, contractual assurances that it will resist government orders for Content Data, the commitment that it has not built in any *backdoors* or similar programming in the services that could be used by AWS or third parties to obtain unauthorised access to the system, as well as transparency reporting from other cloud service providers, there are very few requests (if at all) for data from EU public sector organisations. The DTIA shows that the probability is extremely low that AWS is compelled to provide unlawful access to Content Data (including Content Data in support tickets) to US authorities, even though such disclosure may also be ordered when government organisations have chosen to exclusively store the data in the Netherlands.

As quoted in Section 4.3.2 AWS was not compelled to provide access to Content Data from any Enterprise or government customer located outside of the US, as stated by AWS in each of its twice-yearly Information Request Reports since July 2020.The reports all note:

> "*How many requests resulted in the disclosure to the U.S. government of enterprise or government content data located outside the United States?*
>
> *None*."

AWS only mentions a range between 0 and 249 to describe how often it has received orders from security services for these data, the only information AWS is permitted to disclose by law. Therefore it is unknown if AWS has ever been compelled by security services.

In spite of all these assurances and protections, unlawful disclosure or access to the Content Data cannot be completely ruled out. There are different possibilities for Dutch government organisations to apply encryption to the Content Data. As auditor NCC group states: there is *no mechanism for anyone at AWS to access your data on Nitro hosts.* Alternatively, in some use cases governments can use S3 buckets to store data that are encrypted with keys that are exclusively under the customer's control.

The DTIA has assessed the risks in relation to the nature of the data. If the data would for example include a secret list of suspected terrorists, or the detailed communications of participants involved in multinational serious crime, the likelihood is much higher that a foreign government compels access, or hostile state actors try to obtain access in an illegal way. In that case, based on the national cloud policy, governments have to generally refrain from using public cloud services for the processing, unless they can guarantee that the encryption cannot be broken.

Regardless of the nature of the data, the probability of access in plain text is very small. Even though the consequences for data subjects can vary from low to very serious, this results in a low risk for data subjects.

16.2.2 Disclosure or access to Account, Diagnostic, Support and Website Data as result of transfer to the USA

As described above, the DTIA was also performed for the transfer of Account, Diagnostic, Support and Website Data.

Some government organisations work with highly confidential data, or publicly controversial databases. The processing by AWS (and by affiliate subprocessors in case of Support Requests) of the real names and email addresses of certain government employees could lead to additional data protection risks for these individuals. The same logic applies to the collection of the admin user name in the website telemetry when an admin access the restricted access Admin Console or Support Centre.

These employees could become the targets of spear phishing, social engineering and blackmailing by foreign law enforcement authorities if their identity as employees of these particular organisations was breached, or if AWS, or a processor of AWS that processes Support Requests, was ordered to disclose such data to law enforcement authorities, security services or intelligence agencies.

Government organisations also have to take into account that Content Data may be sometimes become part of Support Data, if the organisation provides such data as part of a support request. These data may include sensitive or confidential (state secret) information, and sensitive and special categories of data of all kinds of data subjects, not just government employees.

In discussions with SLM Rijk, AWS has explained that support employees may use a screen sharing tool to view a customer's screen remotely and identify and troubleshoot problems, but this is not a standard procedure. This tool is view-only. Support agents cannot export any data from the customer, and cannot act for customers during the screen-share session. Additionally, customers must give consent to share a screen with a support agent. These measures do not prevent customers from uploading attachments with personal data with a support ticket. Government organisations cannot encrypt these data, but mitigate this risk by instructing admins never to upload personal data as attachments with support tickets. According to recent information from AWS, admins can also request AWS to delete support tickets.

Additionally, AWS has suggested an extra organisational measure to lower the probability of transfer. Dutch government organisations can ask their AWS account manager to configure the alert for support employees that only EU based employees may respond to tickets, or employees in other countries with an adequate data protection regime, such as Canada or Japan.

Government organisations can mitigate the transfer risks for the other three categories of personal data by pseudonymising the admin account data.

Assuming government organisations comply with these recommendations, the chance of reidentification is very small and the risk may be assessed as low.

16.2.3 Loss of control cookies and telemetry restricted access Website Data

As described in Section 3.2.1 AWS uses a confusing cookie banner. This does not result in surprising data processing: AWS does not appear to set and read any analytical or advertising cookies, regardless of the admin choice. However, AWS does collect analytical data via the website telemetry even when an admin has opted out from all but essential cookies. This type of data processing is more difficult to observe than use of cookies. The lack of information, and the absence of an opt-out for the website telemetry lead to loss of control over the personal data of the government admins.

As assessed in Section 11.1.1, the processing of the observed website telemetry data may be exempted from consent based on the Dutch legal exception for analytical cookies, but this exception does not apply to the collection of the directly identifiable account names. The Dutch government can mitigate this risk by using pseudonymous admin accounts.

When the admin accounts are pseudonymised, even though the probability of collection by AWS is 100% (as observed by Privacy Company in repeated testing), the impact on the admins is low, and hence this can be qualified as a low risk for the admins. SLM Microsoft, Google and Amazon Web Services Rijk will address this issue in the ongoing dialogue with AWS.

16.2.4 Loss of transparency Diagnostic, Support and Website Data

As described in Sections 2.3.2 through to 2.2.5 of this report, AWS does not publish sufficient documentation about the contents of the Account Data, Diagnostic Data, the metadata collected through Support Requests and the restricted Website Data. The only available documentation is in the Privacy Notice, which only applies to the Commercial Contact Data and the public Website Data.

AWS does not provide a public, centrally accessible source of information that describes the personal data collected in the different categories of Diagnostic Data identified in this DPIA. AWS makes some of the Diagnostic Data available through the S3 Access Logs and the CloudTrail log, but AWS considers all information about its own logging activities confidential, including its security logs (See for the security logs, the separate risk 16.2.8 below).

The same lack of transparency applies to AWS's processing of restricted access Website Data (logs with personal data about visits to the Admin Console and the Support Portal) and Support Data. The retention periods of these data are unknown, but according to the four retention criteria provided by AWS, they could be retained as long as 10 years. AWS's publicly available audit reports do not specify if AWS

creates profiles of specific admins based on their support requests, for example through a comments section in the user interface for support employees.

In view of AWS's processor role, and the limited purposes for which AWS is contractually allowed to process these personal data, SLM Rijk can largely mitigate the risks from the lack of transparency by organising audits to verify AWS's compliance with the guarantees in the enrolment framework.

Government organisations can further lower the impact of this risk for employees by forcing them to use pseudonymous accounts.

The probability of occurrence of this risk is high, while the impact for data subjects can be lowered to minimal impact. Therefore this lack of transparency qualifies as a low risk for the admins.

16.2.5 No access for data subjects to some Account, Diagnostic, Support, and Website Data

As described in Section 3.3 and assessed in Section 15.3 AWS did not provide the required overview of personal data when Privacy Company as data subject (admin) filed a Data Subject Access Request with AWS, through its DSAR form.

Privacy Company did not perform a re-test of the Data Subject Access Request in the new situation with AWS as processor, using AWS's Service Controls. Therefore this DPIA cannot draw any conclusions about the quality of AWS's (future) responses. If admins are unable to obtain access to some personal data, they can ask AWS for assistance.

With regard to the collection of the directly identifying user account name with analytical data through the website telemetry, the lack of public documentation prevents government organisations as controllers to assist employees with the exercise of their data subject access rights.

The probability of occurrence of this risk is 100%, as evidenced in this report. However, if the admins use pseudonymous accounts, and inform admins about this processing, the impact of this lack of transparency can be minimised, and qualified as a low risk for data subjects.

16.2.6 Loss of control subprocessors

During this DPIA process, AWS has updated its information about subprocessors it engages for the Content Data. Most of the subprocessors are AWS affiliates. A few are external third parties, but they only process data from specific services or limited to customers in specific countries.

Contractually AWS imposes relevant contractual obligations on its sub-processors. AWS has committed to offer several remedies in case a Dutch government customer objects to a new sub-processor. In spite of the obligation in Clause 9 sub c of Module Two (Controller to Processor) of the (new) EC Standard Contractual Clauses, AWS did not make copies available of the relevant privacy paragraphs in its contracts with subprocessors when so requested. Even though the contracts could not be verified, AWS has greatly improved transparency about its subprocessors, and what categories of personal data they may process.
This transparency, together with the contractual guarantees, greatly reduce the likelihood that the data protection risk of a loss of control factually occurs.

Government organisations can avoid the AI/ML-services that reuse Content Data for service improvement. Even though the impact on data subjects (both admins and individuals whose data may be processed on the VM, in the bucket or in the RDS) can be high, AWS's use of subprocessors qualifies as a low risk for data subjects.

16.2.7 Chilling effects employee monitoring system

Government organisations that use AWS's hosting and database services are technically able to use the Diagnostic Data and other metadata about the use of the Admin account and filing of Support Requests to gain some insight in the work patterns of admins with AWS accounts. As employers, government organisations can use the Diagnostic Data in the S3 Access logs and CloudTrail logs to distil a (limited) picture/create a profile of a person. The system administrators may experience a *chilling effect* as a result of the monitoring of their behavioural data when they undertake processing operations on the VMs, buckets and RDS. A *chilling effect* is the feeling of pressure someone can experience through the monitoring of his or her behavioural data, discouraging this person from exercising their rights, such as accessing certain content. [216]

System admins may also feel observed knowing that AWS keeps track in webserver access logs when they access the Admin Console or the Support Portal.

Government organisations already have other ways of logging the behaviour of their employees in digital systems, through the use of log-in/sign-up logs and logs of behaviour in central work applications. The AWS S3 Access and CloudTrail logs, as well as the information the AWS Support Portal, provide an additional source that can be combined with existing logs to reconstruct a pattern of effective working hours, from first log-in to last log-out, and time spent on the different VMs. Even though the available logs provide limited information, and log data are unreliable for this purpose, government organisations could (theoretically) use this information for a negative performance assessment, if such use of these data is not excluded in an internal privacy policy.

Based on the case law of the European Court of Human Rights , government organisations need to expand on their internal privacy policies, and in particular disclose to employees under which circumstances and for which specific purposes these behavioural data may be processed. This includes listing the specific risks against which the logs will be checked, and what measures the organisations will take to ensure purpose limitation. Even though these AWS logs by itself only provide limited information, government organisations have to take the risk into account of the combination of these logs with other logs from work tools. It is likely that government organisations already have such rules. Therefore, the probability that these risks will occur, can be estimated as very low. Because of this remote chance, even though the impact may be very high, the data protection risks for the employees are low.

16.2.8 Loss of control over personal data in inaccessible AWS security logs

AWS has explained it retains its security logs for a period of maximum 10 years. These logs may contain personal data, for example IP addresses from visitors to websites hosted on AWS.

---

[216] Merriam-Webster Online Dictionary, "chilling effect", URL: https://www.merriam-webster.com/legal/chilling effect.

Privacy Company has not been allowed to see any of the contents of the logs created by AWS about the test set-up of this DPIA. AWS did - [**Confidential**] confirm that it collects log files with personal data as part of its security program. To Privacy Company, none of the categories of collected data mentioned in this AWS internal policy seem excessive for security purposes, but this could not be verified.

AWS's lack of transparency, combined with the long retention period of 10 years, leads to a risk for government employees of unlawful (further) processing of their personal data. The chance that a privacy risk occurs is per definition higher with a long retention period, due to an increased risk of unlawful processing, data becoming inaccurate/outdated and data breaches. Government organisations can mitigate this risk by forcing employees to use pseudonyms when logging in, for example with federated identity. This lowers the impact on the admins. AWS guarantees it will not process these security logs for any other purposes. These guarantees lower the probability of the loss of control, resulting in a low risk for the admins (8a).

Another high risk could occur if visitors of Dutch government websites hosted on AWS would be identified through their IP addresses, in spite of anonymity assurances. For example, if AWS would collect IP addresses from visitors to the website https://www.meldmisdaadanoniem.nl/[217] and would be compelled to disclose these data to Dutch or foreign law enforcement authorities. Government organisations can mitigate this risk by not hosting websites on AWS VMs that require full control over the registration of visitor data. By applying that measure the likelihood of occurrence of this risk is highly unlikely. Therefore, even though the impact on the data subjects may be very high, the risk can also be qualified as low (8b).

16.2.9 The sending of unsolicited marketing mail to procurement officers

AWS requires actively given consent from admins before subscribing them to commercial mails. AWS may send unsolicited mail to its Commercial Contacts, such as the procurement officers of Dutch government organisations. Such e-mails have a low impact (perhaps annoyance, and perhaps some loss of worktime). Commercial Contacts can opt-out from such newsletters and mails by visiting the Marketing Preference Centre. Privacy Company did not re-test the configuration of the Marketing Preference Centre for admins under the new conditions, but assumes AWS no longer automatically subscribes admins to its commercial communications.

Because commercial contacts such as procurement officers can easily opt-out and the impact of this data processing is low, this can be classified as a low risk.

The risks are plotted in the matrix created by the ICO[218] in Table 3 below.

[217] M. 'Guaranteed anonymity', URL: https://www.meldmisdaadanoniem.nl/english/

[218] Copied from the DPIA guidance from the UK data protection commission, the ICO. URL: https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulationgdpr/data-protection-impact-assessments-dpias/how-do-we-carry-out-a-dpia/

Table 3: Risk assessment

Severity of impact	**Serious harm**	**Low risk** **1, 6, 7, 8b**	**High risk**	**High risk**
	Some impact	**Low risk**	**Medium risk**	**High risk**
	Minimal impact	**Low risk** **2, 8a**	**Low risk**	**Low risk** **3, 4, 5, 9**
		Remote	Reasonable possibility	More likely than not
		Likelihood of harm (occurrence)		

Part D. Description of risk mitigating measures

Part D describes the proposed (counter-)measures that are necessary to mitigate the remaining low risks found in Part C.

17. Risk mitigating measures

The following section contains a table of the mitigating technical, organisational and legal measures that need to be taken by the governmental organisation or by the supplier in order to reduce the remaining 9 low risks for the data subjects.

17.1 Measures to be taken to mitigate low risks

No.	Low risks	Recommended measures government organisations	Recommended measures AWS
1.	**Disclosure or access to Content Data as a result of transfer to the USA**	Apply encryption. For S3: Encrypt files stored in S3 with keys outside of AWS's control if the files contain personal data AWS may not access. For EC2/RDS: AWS Nitro is designed to prevent AWS from accessing the Content Data inside of the VM.	Continue to organise external audits on compliance with access policies and disclosure to authorities and continue to disclose the findings via Artifakt.
			Continue to publish transparency reports about requests and disclosures.
2.	**Disclosure or access to Account, Diagnostic, Support and Website Data as a result of transfer to the USA**	Pseudonymise admin employee accounts, for example using identity federation.	Increase transparency about government access to personal data other than Content Data
		Do not host a website on an EC2 instance if the identifiability of visitors through their IP addresses is sensitive, or use a proxy.	
		Ask the AWS account manager to configure the alert for support employees so that only employees based in EU Member States and in countries for which an adequacy decision is available such as Canada or Japan may respond to tickets.	Ensure that the account managers are able to set the alert per customer.
3.	**Loss of control cookies and website telemetry restricted access Website Data**	Pseudonymise admin employee accounts, for example using identity federation.	Ensure that no analytical website telemetry data with user account names or account identifiers are being sent from the website when an admin selects the *'Customize cookies*' option in the cookie consent banner.
		Select the third option *'Customize cookies'* in the cookie banner on the restricted access websites (Admin and Support).	
4.	**Loss of control subprocessors**	Where available, opt-out of Service Improvement for Services.	Continue to organise external audits on compliance of subprocessors with the agreed data protection guarantees.
5.	**Loss of transparency Diagnostic, Support and Website Data**	Advise admins never to upload personal data in Support Requests.	Publish more detailed and up-to-date documentation, including essential information such as the processing of IP addresses of visitors to AWS hosted websites / applications.

			Update an overview of the categories of Personal Data that AWS processes as a controller for the agreed compatible purposes.
		Pseudonymise admin employee accounts, for example using identity federation.	Warn admins not to upload personal data in attachments to Support Requests, encourage and enable masking of personal data in screenshots.
6.	**No access for data subjects to some Account, Diagnostic, Support, and Website Data**	Inform employees about access to the data in the available admin log files and in the Support Centre.	Assist admins as controllers to honour data subject access rights for all personal data in Diagnostic, Website (Admin Console and Support Centre), Support and Account Data and explain to admins when such access is denied on a case by case basis.
7.	**Chilling effects employee monitoring system**	Complement internal privacy policy for the processing of employee personal data with rules for what specific purposes specific personal data in the log files may be (further) processed and analysed. This includes listing the specific risks against which the logs will be checked, and which measures the organisations will take to ensure purpose limitation.	-no measures necessary-
8.	**Loss of control over personal data in inaccessible AWS security logs**	(SLM Rijk) conduct audits on compliance with purpose limitation, data minimisation and retention periods.	Continue to organise relevant audits on compliance with purpose limitation, data minimisation and retention periods.
9.	**The sending of unsolicited marketing mail to procurement officers**	Instruct procurement officials (Commercial Contacts) to opt-out from marketing communications through the AWS Email Preference Center.	-no more measures necessary, AWS will ask admins for consent for commercial newsletters and mails (no more opt-out).

Conclusion

As a result of the negotiations with the Dutch government, AWS has become a data processor for all personal data in and about the use of Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3) and Amazon RDS. If Dutch government organisations follow the recommended measures from this DPIA, they can use these 3 AWS services without any known high data protection risks.

If government organisations encrypt the Content Data with self-managed keys and apply the other risk mitigating measures such as the use of pseudonymous account data for the admins, the transfer risks are no longer qualified as high.