SQL Server and Azure SQL Database GDPR Guidance

Microsoft

Table of Contents

01. Disclaimer

This white paper is a commentary on the General Data Protection Regulation (GDPR), as Microsoft interprets it, as of the date of publication. We've spent a lot of time with GDPR and like to think we've been thoughtful about its intent and meaning. But the application of GDPR is highly fact-specific, and not all aspects and interpretations of GDPR are well settled.

As a result, this white paper is provided for informational purposes only and should not be relied upon as legal advice or to determine how GDPR might apply to you and your organization. We encourage you to work with a legally qualified professional to discuss GDPR, how it applies specifically to your organization, and how best to ensure compliance.

MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS WHITE PAPER. This white paper is provided "as-is." Information and views expressed in this white paper, including URL and other Internet website references, may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this white paper for your internal, reference purposes only.

02.

Introduction

The General Data Protection Regulation (GDPR) applies to all businesses and organizations processing data of European Union residents and citizens or doing business in the European Union as of May 25, 2018.

GDPR is a significant piece of legislation whose impact on privacy and security practices of organizations and businesses around the world has already become visible. Organizations embracing the new GDPR privacy and security obligations must drive a coordinated effort among various disciplines and involve different organizational entities including legal, human resources, marketing, security, IT and others. The process starts with a gap analysis comparing the current state of privacy and data protection in the organization with the privacy and security requirements defined by GDPR. Addressing the identified gaps will almost always result in improved business processes that include data stewardship and enhanced data governance as well as a heightened organizational awareness of data privacy and a better understanding of current privacy concerns.

This paper summarizes the privacy and security principles that were relevant to our own work toward GDPR compliance. Following the security principle overview, we outline the security features available in SQL Server, Azure SQL Database and Azure SQL Data Warehouse that can be helpful when establishing the secure data handling processes and policies required to meet the enhanced security obligations stipulated by the GDPR. This includes features and capabilities of SQL products that can be used to enhance secure processing of sensitive data in accordance with industry best practices. Lastly, we give an overview of features that collect or save personal data, and how it can be removed in product. In the last section we introduce the GDPR and privacy options available to our customers in Windows and Azure.

03. Privacy principles

A compliant GDPR solution needs to address all privacy principles, which in turn will affect all business processes and workflows, IT systems, as well as all applications used and implemented by the organization. Addressing the security principle explained in section 4 of this document is particularly important and necessary, but not sufficient in terms of a GDPR-compliant solution. The key changes required to address GDPR are:

- Personal privacy
- Controls and notification
- Transparent policies
- IT security and training

The ability to impose significant fines for non-compliance may be the key driver for change as it creates pressure to comply with privacy principles and brings them into focus of every entity collecting or processing personal data in the European Union (EU) or pertaining to citizens of the EU. In addition, the regulation replaced the overly narrow NIST definition of "Personally Identifiable Information (PII)" with the European definition of "personal data," which is all data that relates to an identified or identifiable individual. A subset of personal data, called pseudonymous data, includes data sets that are not associated with specific individuals because obvious identifiers (name, social security number, or biometric data) have been removed. The data in pseudonymized data sets can be linked back to individuals, for example by using unique identifiers (device IDs, user IDs, IP addresses, or deterministic hashes of these IDs) that can be resolved by joining with other data sets. Storing data in pseudonymized data sets adds a layer of security, but it is important to remember that this data is still considered personal data under the law.

The following section outlines the privacy principles and the newly established rights that are awarded to the owners of personal data, known as "data subjects" in the legal language of the GDPR.

a. Lawfulness, fairness and transparency principle

The first question to answer about the collection or processing of personal data is if the owner of the data is aware of the collection and has given consent. Being able to establish that consent was obtained is essential. Here are some key points to consider:

- Confirm that legitimate business interest or instruction from the data controller to collect or process personal data exists.
- Share details of the data collection and related privacy practices openly and provide notice to customers when and why data is being collected (privacy statement).
- Decide when to do this in a clearly written privacy statement, and when to do this additionally at the time of data collection, for example directly in the user interface.

b. Purpose limitation principle

After consent has been obtained for the collection and processing of personal data, it is important to limit the further use of the data to the agreed purpose prior to collection. This, in GDPR terms, is a data controller. If the business is a data processor, the users issue instructions on how to process their data to the organization. These instructions are contained in a contract, such as terms & conditions. The purpose limitation principle is more difficult to implement, because once data has been collected, it is easily shared with other parts of the organization, who might use the data for purposes not initially disclosed. Implementing purpose limitation is where a privacy review process and the application of data sharing policies will be required. At time of data collection, the following needs to be implemented:

- State the purposes for which data is being collected publicly to the user (privacy statement).
- Implement data handling guidelines

and policies that limit the use of data throughout the organization.

c. Data minimization principle

This principle asks to collect only the data absolutely needed to fulfill the stated purpose. Don't over collect, and don't collect because it might be useful in the future. Any personal data collected that is not needed should be removed.

- Only collect the data that is required to fulfil the stated business purpose.
- Collect the least amount of data that you can operate with.

d. Accuracy principle

Personal data and especially contact information must be accurate to be useful. Customers and users need to be able to update their information. Keeping up with this requirement becomes more difficult if various copies of the data exist. Maintaining master copies and inventories will be required.

- Ensure personal data is accurate, complete and kept up-to-date.

e. Storage limitation principle

Personal data should only be stored for as long as it is needed to fulfill the purpose initially stated.

- Implement retention policies for all personal data that has been collected and stored.

f. Integrity and confidentiality principle

Personal data that has been collected lawfully and with the data owner's consent must be protected against potential risks such as unauthorized access, use or modification, loss or destruction, as well as disclosure (breach).

- Implement comprehensive security controls with the help of an established security framework.

GDPR requires that any security incidents where personal data has been lost, stolen or otherwise accessed by unauthorized third parties are reported within 72 hours of discovery.

The above privacy principles are specifically addressed in Chapter 2 Art. 5 GDPR Principles relating to processing of personal data.

g. Data transfer principle

GDPR limits the transfer of personal data to third countries. Organizations will have to create data maps and data flow diagrams to be able to demonstrate data flows are compliant.

- Inventory all data storage systems and their geographical location.
- Create data flow diagrams for all trans-national data flows.

Data transfer is addressed in Chapter 5 GDPR Transfers of personal data to third countries.

h. Accountability principle

GDPR requires that data controllers report any security incidents where personal data has been breached by unauthorized third parties to their data protection authority (DPA) within 72 hours of them becoming aware of it.

This principle requires implementation of robust breach detection, investigation and internal reporting procedures and puts additional pressure on organizations

processing personal data to implement security controls that protect data from unauthorized access and breaches.

These privacy and security requirements are reinforced by the risk of substantial fines. Poor privacy or security practices can turn into a large cost for the business—not only in terms of reputation, but now also in terms of significant monetary impact.

A last point to consider is that data controllers are responsible for data handling of their subcontractors or any other third parties they share the data with. This requires an audit of all relationships with third parties.

The accountability principle is addressed in Chapter 8 GDPR Remedies, liability and penalties.

i. Individual participation principle

After personal data has been collected lawfully and with consent, GDPR considers the person this data links to as the rightful owner of the data, and the owner is awarded the right to:

- access data held about them;
- request rectification of incomplete or inaccurate personal data;
- have their data deleted;
- restrict processing;
- data portability.

All these rights make it essential for organizations that process personal data to have systems and processes in place which enable them to identify, access, edit, delete, and export personal data—and be able to perform these operations quickly, or without undue delay and in any event within one month of individual rights requests. This is one of the most impactful requirements and the implementation details and how they apply to SQL Server and Azure services will be discussed in section 6 GDPR Data subject requests.

The following processes and workflows need to be implemented by the organization:

- Publish how the organization can be contacted regarding questions or concerns about data collection by the owner of the personal data (the data subject);
- Enable the data subject to access the data collected about them. A copy of the data must be provided to them free of charge, typically without undue delay and within one month of a request;
- Provide a process that allows data subjects to request changes or correct errors in their data;
- Allow the data subject to request for their data to be deleted, with some exemptions, such as for exercising freedom of expression and freedom of information;
- Enable the data subject to revoke consent to any further collection and processing;
- If applicable, allow for the data to be exported free of charge and in a structured, commonly used and machine-readable form.

These rights, referred to as “data subject rights” are outlined in Chapter 3 GDPR Rights of the data subject.

04. Security principle

The security principle relates to the collection of personal data: personal data must be processed in a manner that ensures appropriate security, including protection against unauthorized or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organizational measures to ensure "integrity and confidentiality" of personal data.

The legislation does not state any specific security requirements that must be implemented by the organization processing personal data, instead "adequate security" must be determined based on the sensitivity of data processed as established in the Data Processing Impact Assessment (DPIA) and the general risk profile. A practical way to meet the somewhat vague GDPR security obligations is to work with an existing security framework that has been created to assist with improving and operationalizing system security, such as security industry standards (ISO 27001, NIST 800-53) or industry-specific security certifications (HIPAA, SOC, PCI- DSS). The security policies and controls implemented by the organization must meet not only GDPR requirements but all industry-specific security or regulatory requirements that apply.

Therefore, we recommend defining a security policy for the organization that normalizes all applicable security requirements from diverse sources.

A common set of security control areas include:

- Information security policies (documented operating procedures)
- Human resource security (employee screening, training and awareness)
- Asset management (inventory)
- Information classification (classification, labeling and data handling policies)
- Media handling (secure disposal of media)
- Access control (user access control, implementation of "need to know based access" and "least privilege" principle, or "separation of duties")
- Cryptography (encryption requirements and key management)
- Malware protection
- Logging and monitoring (audit logs)
- Vulnerability management (patching)
- Secure development lifecycle (secure by

design and by default)
- Protection of test data (data masking, anonymization of data sets)
- Incident management and breach notification processes

It is important to reiterate that this list is not exhaustive, and we cannot determine if any security measure or set of measures either complies with, or does not comply with, GDPR. That determination can only be made by an organization in consultation with its legal and compliance advisors under consideration of their specific risk assessment.

Microsoft has decades-long experience building enterprise software and running some of the largest online services in the world. We use this experience to implement and continuously improve security-aware software development, operational management, and threat-mitigation practices that are essential to the strong protection of services and data. The security control areas above have been implemented for Azure SQL Database and Data Warehouse and the services comply with both international and industry-specific compliance standards and we participate in rigorous third-party audits that verify that our security controls have been implemented appropriately.

The Microsoft global incident response team works around the clock to mitigate the effects of any attack against our cloud services. And security is built into Microsoft business products and Azure services from the ground up, starting with the Security Development Lifecycle, a mandatory development process that embeds security requirements into every phase of the development process. Incident response processes and breach notification procedures are also part of our security fabric and incident response plans and drills get verified during third-party security audits.

In addition to the operational security discussed above, our products and services offer the following security features that can be used to implement security controls as defined by the following ISO 27001 security control families:

- Asset management
  - Data classification
  - TDE encrypted disks help meet the secure disposal requirement
- Access Control
  - Azure Role-Based Access (RBAC)
  - Granular permissions
  - Row-level security
  - Dynamic data masking
  - Azure Active Directory Support:
    - User name/ password
    - Integrated
    - Token-based
    - Multi-factor authentication
- Cryptography
  - Crypto functions (partial)
  - Encryption at rest: Transparent Data Encryption (TDE) - service-managed encryption keys and customer-managed keys with Azure Key Vault integration
  - Always encrypted (secure enclaves coming soon)
- Communications security
  - Service endpoints
  - full VNET coming soon
  - Encryption of client-server communication through TLS
- Operations security

- SQL Audit (database & server level)
- Vulnerability assessment
- SQL Threat Detection

- System acquisition, development and maintenance
  - Static data masking coming soon, available through partner offerings
- Business continuity management
  - High availability
  - Active geo-replication
  - Database copy
  - Point-in-time restore
  - Geo-restore
  - Automated backups
  - Long-term backup retention (preview)
  - Failover groups
  - Transactional replication
  - Zone redundant database
  - Zone redundant elastic pool
  - DataSync
  - Import-export (BCP)

This table below gives an overview of the most important features available for the different SQL products and services, abbreviated for the table headings as follows:

- SQL Server Enterprise Edition (SQL Server)
- Azure SQL Database (Azure SQLDB)
- Azure SQL Database Managed Instance (Azure SQLDB MI)
- Azure SQL Data Warehouse (Azure SQL DW)

The feature availability reflects the current status as of October 2018 and is subject to change. For a complete and updated list of features for Business Continuity Management refer to the documentation at: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity. For a complete and updated list of general security features, review the documentation at: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-security-overview.

ISO Security Control Families	Features	SQL Server	Azure SQLDB	Azure SQLDB MI	Azure SQL DW
Asset Management	Data Classification	✓	✓	Planned	Planned
	TDE	✓	✓	✓	✓
Access Control	Role-Based Access (RBAC)		✓	✓	✓
	Granular permissions	✓	✓	✓	✓
	Row-level Security	✓	✓	✓	✓
	Dynamic Data Masking	✓	✓	✓	NO
	Azure Active Directory Support:	N/A	✓	✓	✓
	User name/password	N/A	✓	✓	✓
	Integrated	N/A	✓	✓	✓
	Token-based	N/A	✓	✓	✓
	Multi-factor Authentication	N/A	✓	✓	✓
	Windows Integrated Authentication	✓	N/A	Planned	N/A
Cryptography	Crypto Functions	✓	Partial	Partial	Partial
	Encryption at rest: TDE	✓	✓	✓	✓
	Customer Managed Keys	EKM* and HSM**	Azure Key Vault	Azure Key Vault	Azure Key Vault
	Always Encrypted				
	Software-based	✓	✓	✓	Planned
	With secure enclaves	vNext	Planned	Planned	Planned
Communications Security	Service Endpoints	N/A	✓	NO	✓
	VNET	N/A	✓	✓	✓
	Encryption of Client-server communication	✓	✓	✓	✓
Operations Security	SQL Audit (Database & Server level)	✓	✓	✓	✓
	Vulnerability Assessment	✓	✓	✓	✓
	SQL Threat Detection	Planned	✓	✓	✓
System Acquisition, Development and Maintenance	Static Data Masking	Planned	Planned	Planned	NO
Business Continuity Management	High Availability	Always On	Built-in	Built-in	Built-in
	Disaster Recovery	Availability Groups	Geo-DR & Active Geo-Replication	Geo-DR & Active Geo- Replication	Geo-DR
	Database Copy	N/A	✓	✓	NO
	Automated Backups	N/A	✓	✓	✓
	Point-in-time Restore	N/A	✓	✓	✓
	Long Term Backup Retention	N/A	✓	✓	NO
	Failover Groups	N/A	✓	✓	NO
	Transactional replication	YES	✓	✓	NO
	Zone redundant database	N/A	✓	✓	✓
	DataSync	✓	✓	✓	NO
	Always On Availability Groups	✓	N/A	N/A	N/A
	Always On Failover Clusters	✓	N/A	N/A	N/A
	Fast Recovery	✓	✓	✓	NO
	Peer-to peer replication	✓	N/A	N/A	N/A
	Merge replication	✓	N/A	N/A	N/A
	Backup/Restore	✓	NO	✓	NO
	Change data capture	✓	NO	✓	NO
	Change tracking	✓	✓	✓	NO
	Log shipping	✓	NO	Planned	NO

*EKM: Extensible Key Management **HSM: Hardware Security Module

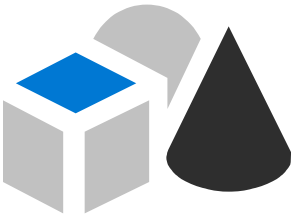

05.

GDPR Impact on SQL features

In this section we will outline features that are either impacted by GDPR or can help with maintaining data, such as implementing data retention policies. Furthermore, we will list features that may store personal data, for example data from database users' or administrators' accounts or logins, that need to be removed or updated as needed to meet GDPR compliance. This list is not comprehensive and will get updated over time.

SQL-based applications must be reviewed and analyzed for their GDPR compliance. Here we will discuss the following features:

- Backups and long-term backups
- Data expiration
- Azure Active Directory Authentication
- SQL Server Metadata incl. Reporting Services and SQL Server Integration Services
- Threat detection and SQL Audit logs
- Email notification features

Backups and long-term backups:

Personal data in backups is in scope of GDPR, which means any personal data contained in backups must be made available to data subjects upon request, and in some cases data subjects can request for data to be deleted. In addition, there may be conflicting data retention requirements that require certain records or transactions to be retained for a certain timeframe.

For short-term backups, some of the extra work required by GDPR can be avoided by limiting the backup timeframe to shorter periods. Backup retention for premium databases is 35 days by default but can be adjusted to shorter periods. Standard databases only provide 7 days of backup retention. If a premium database with a customized backup retention is changed to standard, and then changed back to premium, the backup retention will fall back to the default of 35 days.

Long-term backups that are available for periods of up to 10 years must meet GDPR compliance as well. Here it is important to determine if the organization has data retention requirements—including for personal data—that do not allow for personal data to be deleted, for example transactions that need to be retained for tax purposes for at least 6 years. That determination can only be made by an organization in consultation with its legal and compliance advisors under consideration of their specific data retention requirements.

Minimizing personal data needed in long-term backups, pseudonymizing data sets where possible or applying technologies such as tokenization or data masking might be helpful when considering long-term backups.

Data expiration:

Data minimization can be accomplished by implementing data retention policies. For SQL Server (on-prem or in a VM) SQL Server Agent can run a scheduled job that deletes data based on any query saved as a stored procedure. This can be used to delete records automatically after a specified time.

For Azure SQL Database, the Elastic Jobs feature provides comparable capabilities in Azure:

- Jobs can target databases in one or more Azure SQL elastic pools, logical servers, shard maps, and across multiple subscriptions.
- Jobs can be composed of multiple steps to customize the execution sequence.
- List of target databases in the target groups are dynamically enumerated. Any databases added or dropped from the target group are automatically picked up without having to make explicit job script changes.
- Jobs can be configured to limit the number of databases a job runs against in parallel to ensure resource optimization.
- Target groups can be customized with "include" or "exclude" references for specific databases, pools, or servers.
- T-SQL, PowerShell, REST APIs, and Azure portal support.

Azure Active Directory Authentication:

Azure Active Directory (AAD) is the place where SQL Database and DW administrators can delete identities belonging to users, which will result in a purge of their personal identifier and associated data like telephone numbers, email addresses, etc. throughout the Azure system. When using Azure Active Directory Authentication with Azure SQL Database, AAD user accounts may also be duplicated, either as contained database users or, in the case of Azure SQL Database Managed Instance, as logins. Contained users and logins in SQL DB/DW that are associated with AAD identities have to be removed separately from specific databases or MasterDB for the Managed Instance respectively.

To avoid having to find all database objects a user/login has permissions for, we recommend the use of groups. Here, users are added to groups and groups have specific permissions to objects. This allows for users/logins to be removed from a group in a single action.

SQL Server Metadata incl. Reporting Services and SQL Server Integration Services:

SQL Server stores metadata about users in system tables. Here is an overview of what is stored and how it can be removed, including the retention period if applicable.

SQL Server Reporting Services

This is execution log storage which is automatically aged out after 14 days and is configurable through SSMS:

Server	Database	Table	Query
*	ReportServer	ExecutionLogStorage	Possible ID

Detailed documentation can be found at:

https://docs.microsoft.com/en-us/sql/reporting-services/tools/server-properties-advanced-page-reporting-services?view=sql-server-2017

https://blogs.msdn.microsoft.com/shiyangqiu/2016/09/12/reporting-service-execution-log-retention/

Subscription definitions are user's content and will need to be either deleted manually or edited to remove the specific users. This can be done through UI/tools or scripted through Report Server APIs. Schedule owners can be updated through UI or server APIs as well. We do not recommend setting values to NULL directly in the catalog as that will probably cause subscription failures.

Server	Database	Table	Query
*	ReportServer	Subscriptions	Possible email

Similar to subscriptions, schedules are user content. These will need to be manually updated or edited. Recommend this be done through the product UI or Server APIs. Manually editing the schedule content could lead to unexpected errors.

Server	Database	Table	Query
*	ReportServer	Schedule	Possible ID

UserName can be manually set to NULL, however note the user must be first removed from Active Directory or it will be refreshed on a daily schedule.

Table	Query
Users	User ID

SQL Server Integration Services

The following columns are only used for logging purposes, no functionality depends on these columns, so they can be set to a different value by a simple update statement.

created_by_sid	0x0105000000000000515000000FD3742403DE3084D828BA628D4340500
created_by_name	domain\userid
created_by	domain \userid
deployed_by_sid	0x0105000000000000515000000FD3742403DE3084D828BA628D4340500
deployed_by_name	domain \userid

Threat Detection and SQL Audit:

SQL Audit collects data in scope for GDPR, mainly client IP addresses and user logins. A complete list of data points collected by the audit feature can be found at: https://go.microsoft.com/fwlink/?linkid=829599

Collecting data for security purposes is generally allowed, as securing a system and protecting from malicious actors is in the best interest of all users. Personal data collected for security purposes is still in scope of GDPR requirements though, which means purpose limitation and data subject rights still apply. Although in most cases data subjects cannot request to be deleted from audit logs, but they may still request to view the data that was collected.

Email notifications:

Some of the advanced security features include optional email notifications. If email addresses were submitted to Microsoft to facilitate the email notifications, they can be managed, viewed, removed directly in the Azure Portal through the user interface.

06.

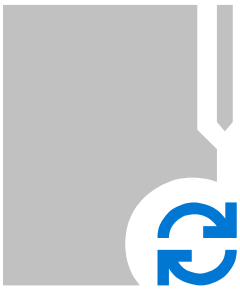

GDPR Data subject requests

The Individual participation principle discussed in section 3.i addresses the rights that are awarded to the owners of personal data by GDPR, who are "data subjects" in the language of the legislation. This section shows how the Azure platform gives customers access to personal data collected through the use of the service. Microsoft's GDPR capabilities will enable enterprise customers to fulfill the GDPR requirement to give data subjects access to the data collected about them upon request. These requests are referred to as Data Subject Rights requests (DSRs). In the context of Azure, we distinguish between two categories of data:

1. Customer data
   Data owned and authored by our enterprise customers, such as:
   - Contents of a database
2. Service data
   Data generated by Microsoft when operating products & services:
   - Pseudonymized information
   - Support data
   - Administrator data

Both Azure SQL Database and SQL Server provide capabilities to enable DSR requests for Customer Data, which customers can perform in-product.

DSR Requests for Customer Data include:

- Discover – find
- Access – retrieve
- Rectify – make changes or perform other actions
- Restrict – restrict processing via license revocation or removal of data from the Microsoft Cloud to another location

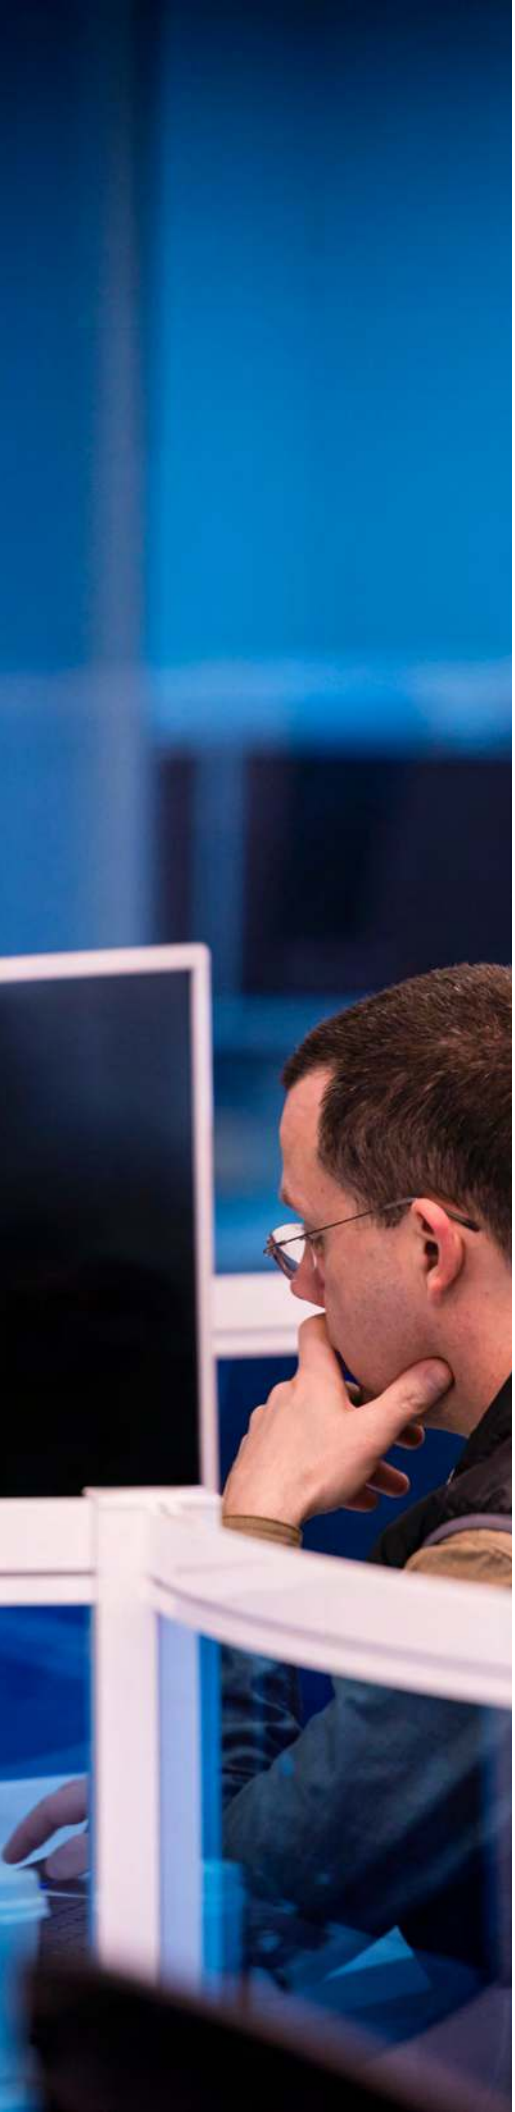

- Delete – permanently remove data
- Export – provide an electronic copy

In addition, Microsoft provides capabilities to enable DSR requests for service data, specifically data generated as a by-product of using the products & services. DSR requests for service data in Azure include the following capabilities:

- Access—find data that resides in the Microsoft cloud and allow to export the data.
- Rectify (not offered)—service data constitutes factual actions conducted within the Microsoft cloud (e.g., a user logging into a service). Rectification is not offered as it would compromise the historical record and increase fraud and security risks.
- Restrict—restriction of processing results in the removal of data from the Microsoft cloud for storage at a user-specified location.
- Delete—permanently remove personal data from the Microsoft cloud.
- Export—provide an electronic copy (in a native format) of personal data of a given data subject.

Using the Azure Active Directory Admin Center, a tenant administrator can permanently delete a data subject from Azure Active Directory and related services: Azure AD Admin Center

For SQL Server, customer's DSR requests will go through Windows, but generally there are three places for customers to go to exercise DSRs:

- On Device (AAD, MSA, unauthenticated users)
  - Windows Diagnostics Viewer and Delete
  - Delete Activity History (incl. SQL Server)
- Consumer Privacy Dashboard (for MSA accounts only)
- Privacy Response Center (AAD and MSA accounts)

The SQL Server Privacy Addendum is available as a supplement to the Microsoft Privacy Statement.

07. Summary

To meet GDPR requirements we have implemented technical (tooling) and organizational measures (processes) to ensure that data processing is in accordance with GDPR privacy regulations and produced documentation to demonstrate how tools and processes enable us to meet the new privacy obligations. The actual path to GDPR compliance will be different for every business or organization, and Microsoft has made general resources available in the new Service Trust Portal to help customers get started. This document is supposed to be helpful for customers who are evaluating Azure SQL Database or SQL Server for GDPR compliance. SQL Server and Azure SQL Database offer a strong suite of security features that are useful when implementing security controls based on a comprehensive security framework such as ISO 27001.

Privacy principles are more difficult to implement as they will require changes in applications, business processes and data handling practices and less guidance exists, but privacy principles are now costly to ignore. “Privacy by design” is no longer just a nice idea; privacy by design and privacy by default are now firm legal requirements. This means minimizing the processing of personal data wherever possible—limiting activity to only what is necessary for a specific purpose, carrying out privacy impact assessments and maintaining up-to-date records to prove compliance—is advised. Lastly, consent requirements for processing personal data are considerably strengthened under GDPR, requiring solutions for consent management, such as tracking consent and allowing consumers to revoke it later.

The maximum fine that organizations can receive for the most serious infringements of the regulation is 4% of their global annual turnover (or €20M, whichever is greater). GDPR has significantly increased the financial risk associated with handling or mishandling personal data, which has resulted in much needed executive attention. This in turn has improved funding for innovative solutions that help to address the privacy problems and strengthen data handling processes and security standards, which will benefit customers and users of software and services worldwide.